

UNCONSCIOUS PERCEPTION RECONSIDERED

IAN PHILLIPS

University of Birmingham and Princeton University

Most contemporary theorists regard the traditional thesis that perception is essentially conscious as just another armchair edict to be abandoned in the wake of empirical discovery.¹ Here I reconsider this dramatic departure from tradition. My aim is not to recapture our prelapsarian confidence that perception is inevitably conscious (though much I say might be recruited to that cause). Instead, I want to problematize the now ubiquitous belief in unconscious perception. The paper divides into two parts. Part One is more purely philosophical. It explains how standard arguments for unconscious perception rely on contentious background assumptions concerning the relation between ordinary perception and the explanatory constructs of scientific psychology. Part Two, in contrast, offers detailed engagement with relevant empirical work. It exposes how, even setting aside the concerns identified in Part One, a dilemma confronts the believer in unconscious perception. This dilemma arises because ordinary perception is an individual-level state or occurrence, yet criteria sufficiently stringent to guarantee that a putatively perceptual state is unconscious vitiate the grounds for its attribution to the individual. The dilemma foments a hypothesis, namely that the conditions for genuine, individual-level perception coincide with the

¹ For this narrative, see Prinz 2010. A good overview of the contemporary consensus and its empirical basis can be found in Merikle et al. 2001. Further recent defences of unconscious perception include Dretske 2006, Burge 2010: 374-6 (who significantly follows Palmer 1999: chpt. 13), Brogaard 2011, Prinz 2015, Block 2016, Block in Phillips and Block 2016, and Block and Kentridge in Peters et al. 2017.

conditions for perceptual consciousness. The viability of this hypothesis should unnerve anyone who thinks unconscious perception is simply an empirical given.

PART ONE

1. What is Perception?

To say that perception can occur unconsciously is to say something about perception. But what is perception? In Part One I explain why this question is of paramount importance in thinking about unconscious perception. I begin by showing why a traditional, relationalist conception of perception can make unconscious perception seem incoherent (§1). I then criticize attempts to define perception in some neutral way to avoid closing-off the issue (§2). A more promising suggestion underpins many recent arguments for unconscious perception. This is that perception is a natural kind whose nature and independence from consciousness is settled by psychological investigation. I first explore this suggestion (§3) and then challenge it (§4). I do so in two ways. First, I explain how a constitutionalist approach to the relation between the kinds identified by psychology and the manifest kind *perception* allows us to take perceptual psychology seriously without identifying perception with a psychological kind. Adopting this constitutionalist approach topples arguments for unconscious perception at the first hurdle. Second, I explain why, even if we do assume that perception is a psychological kind, it will always be open to think of this kind as identical with conscious perception.

To begin, consider the traditional, literally Moorean view that perceiving is simply a way of being conscious, just as being red is a way of being colored (Moore 1925: 46-7). On this picture, talk of unconscious perception makes no more sense than talk of uncolored red

things, or shapeless squares. In the background here is Moore's subscription to a relational account of perception. On such an account, perception involves a subject standing in relations of conscious acquaintance or awareness to various presented elements: for Moore, non-physical sense-data; for contemporary naïve realists, aspects of mind-independent reality.² The appeal to such presented elements is intended to answer a question we can each pose for ourselves: *what is it like to be me, from my present perceptual perspective?* We answer this question (at least in part) by indicating which items we are acquainted with from our perspective, as well as how they are arranged and qualified. The explanatory force of such views derives from the fact that they treat objects of awareness as genuine constituents of perception. As such, the objects of awareness literally shape the contours of consciousness.

On any such relationalist picture, the idea of unconscious perception can seem incoherent. For it would seem to involve being acquainted with some element, and yet that element making no contribution to the subject's conscious perspective on the world. This is commonly regarded as a serious difficulty for relationalism. Thus, various theorists have argued that the alleged wealth of empirical evidence for unconscious perception shows that relationalist accounts must be misguided. Block, for example, targeting contemporary relationalist views, finds "it difficult to see what their account of unconscious perception could be" (2010: 49) and, failing to discover any "direct realist [i.e. relationalist] discussion of this issue, even a brief one" laments "one more sign of the profound disconnect between direct realism and the science of perception" (ibid: 30).³ Yet there is something puzzling

² For sense-datum theories see, in addition to Moore, e.g. Price 1932 and O'Shaughnessy 2000; for naïve realist theories see, e.g. Martin 2009 [1997], 2006, Campbell 2002, and Brewer 2011.

³ See also Block's opening statement in Phillips and Block 2016: 169, and Berger and Nanay 2016. As discussed in Anaya and Clarke 2017 (see also Phillips in Phillips and Block 2016), such arguments can be resisted in various ways (e.g. by exploiting the so-called "third-relatum" in Campbell and Brewer's version of relationalism). The present discussion suggests a stronger line: the theorist who celebrates the fact that their

about this complaint. For, insofar as unconscious perception is incoherent on a relationalist approach, only a theorist who had already rejected that approach to perception could possibly interpret a given empirical case in such terms. But then it is unclear whether the issue is, after all, empirical or, instead, a disagreement concerning which conception of perception one ought to adopt in the first place. This issue concerns not just relationalist views, but any *experientialist* view, that is any view on which perception is treated as constitutively connected with conscious experience (O'Shaughnessy 2000: chpt. 15).

2. A Neutral Definition?

Can the impasse just identified be avoided? Attempts to defend unconscious perception often begin by offering a “neutral” definition of perception: one deemed acceptable to all parties and which prescind from theoretical controversies concerning the precise nature of perception. For example, in his recent defence of unconscious perception, Kentridge (in Peters et al. 2017) looks to the *Oxford English Dictionary* (Simpson and Weiner 1989) for a “simple ... ‘experience-neutral’ definition of perception”. He selects the following: “The process of becoming aware of physical objects, phenomena, etc., through the senses”, before going on to “adopt a working definition of visual perception simply as the process through which we become acquainted with the visual properties of objects in the world (i.e. their distal properties)”. However, these definitions appeal respectively to *awareness* and *acquaintance*, two notions which arguably imply consciousness.⁴ At the very least then, some

view *can* accommodate unconscious perception is simply exposing what should be an embarrassing fact, namely that their account of the intrinsic nature of perception (e.g. in terms of representational contents) fails to explain why perception has a phenomenal nature.

⁴ Given Kentridge's appeal to the *OED*, it is perhaps worth noting the *OED* entry for “awareness” which runs: “The quality or state of being aware, consciousness”. “Acquaintance” is a technical term but in its technical sense, the *OED* offers: “Direct or immediate experience or awareness of anything”.

further “experience-neutral” definition of these notions is required. None being given, Kentridge’s approach threatens to serve quite the contrary purpose to that intended.

A related approach is to define perception so as explicitly to avoid implying any necessary connection with awareness.⁵ Consider the following passages from Kanwisher and Palmer respectively.

If the scientific investigation of awareness is different from the scientific investigation of perception, then the two phenomena must not be identical. (In keeping with the possibility that they are distinct, the word ‘perception’ will be used ... to refer to the extraction and/or representation of perceptual information from a stimulus, without any assumption that such information is necessarily experienced consciously.) So the most basic question is whether all perception is accompanied by awareness, or whether the two phenomena can be uncoupled. (Kanwisher 2001: 89-90)

Until now, I have been taking for granted that you know what I mean by “visual perception.” I do so in large part because I assume that you are reading the words on this page using your own eyes and therefore know what visual experiences are like. Before we go any further, however, we ought to have an explicit definition. ... **visual perception** will be defined as the process of acquiring knowledge about environmental objects and events by extracting information from the light they emit or reflect. ... One interesting feature of this definition is that it does not contain any reference to visual awareness of experience. This might seem like an oversight, but it was not. Conscious visual experience was left out because it is logically possible for vision to occur in the absence of awareness. (Palmer 1999: 5 and 630-1)

⁵ Compare discussions of quasi-memory in Shoemaker 1970 and Parfit 1984. For criticism, including the worry that negative characterisations may simply misfire, see Evans 1982: chpt. 7 and Wiggins 2001: chpt. 7.

Reading these passages, it is easy to feel victim to a slight-of-hand. With Palmer, one wonders how the possibility of unconscious vision could have been established *prior* to his definition of perception, not least given the acknowledged fact that our intuitive grip on perception is bound up with visual experience. On the other hand, if the possibility of unconscious vision only follows *given* his definition of perception, one wonders why we should think it a good one.⁶ Similarly, Kanwisher offers us a definition of perception disconnected from experience in order to open up an empirical issue. But why suppose that there is an open empirical issue here? Put another way: what reason have we for thinking that the notion she defines is not simply a technical one when what we wanted to know was whether perception *in the ordinary sense* inevitably involves awareness?

There is absolutely nothing wrong with introducing a technical or scientific notion of “perception” appropriate for certain theoretical purposes. However, advocates of unconscious perception are often adamant that they are not simply making a claim about perception in some potentially distinct, technical or scientific sense. This is most obvious when it comes to those philosophers who cite unconscious perception as an objection to relationalism since relationalism is explicitly concerned with perception in the ordinary sense (Campbell 2010: 202, 210). According to a currently popular viewpoint, the concerns voiced in this section are avoided by recognizing that perception is a natural, specifically psychological, kind. I turn to this approach in the next section.

3. Perception as a Natural Kind

According to a view much in vogue, perception (and/or its determinates such as seeing) is a natural kind, whose nature and independence from consciousness is established by

⁶ Dretske (2006: 90) and Burge (2010: 417, fn. 56) also criticise these respective definitions, though neither on the ground that perception requires consciousness.

psychological research. Thus, Burge (2010) “elaborate[s] a conception of sense perception as a distinct psychological kind” (367) which “grounds scientific as well as commonsense explanation.” (ibid.) and “whose instances are often but not always conscious” (362-3, fn. 97) Similarly Block (in Phillips and Block 2016: 169) holds: “Seeing is a single fundamental natural kind of which conscious and unconscious seeing are sub-kinds.”⁷

Thinking of perception as a natural kind allows theorists to avoid the unpromising project of defining the term in some neutral way. Instead, they can think of “perception” as picking out states *such as these*, where the demonstrative “these” picks out familiar conscious perceptual episodes and the “such as” functions to generalize to all states of the same fundamental nature. The experientialist, who conceives of perception as constitutively conscious, need not disagree. What they will likely take issue with is the further contention (explicit in Prinz 2015: 371-2) that consciousness can then be treated merely as part of the prototype used to identify instances of perception, leaving it as an open scientific question whether consciousness forms part of the essence of the kind. Against this, the experientialist may insist that perception wears its fundamental nature *qua* conscious episode on its sleeve.

The idea that perception wears its nature on its sleeve does not conflict with the unelaborated thought that perception is a natural phenomenon; one, for example, about which inductive generalizations can be made.⁸ For one, it is not obvious that everything in the natural world has a hidden nature, revealed only by scientific inquiry. As O’Shaughnessy writes: “simple

⁷ Block here focuses specifically on seeing as opposed to perception. I focus on perception in general as Block does on other occasions. I assume that a commitment to unconscious perception is a commitment to unconscious perception in some recognized perceptual modality, and so to unconscious seeing, hearing, or tasting, for example.

⁸ The connection between induction and natural kind-hood is in any case opaque. Some of our most secure inductive generalizations concern artefacts which are precisely built to function alike (Hacking 2007).

artefacts like knives and chairs, simple pre-scientific items like lakes and mountains. Such objects have *no* hidden nature!" (2000: 424) But even if most natural phenomena do have a hidden nature, perception may be special. In Moore, we encountered the idea that perception is simply a way of being conscious. Combined with the traditional thought that there cannot be anything more to a conscious episode than is apparent to its subject upon first-person reflection, it follows that perception lacks a hidden, underlying nature. There can be no fool's perception.⁹

Resistance to this idea comes not from a commitment to perception being a natural phenomenon, but rather from two more specific ideas. First, that perception is a *psychological* kind whose fundamental nature it is up to psychological science to establish. Second that this psychological kind is distinct from the kind, *conscious perception*.¹⁰ As I explore in the next section, both of these claims can be resisted. It is not beyond doubt that perception is a psychological kind. Nor is it beyond doubt that, if it is a psychological kind, that psychological kind is not a conscious kind.

4. Perception as a Conscious Kind

Burge struggles to understand how one could deny that perception is a psychological kind:

The psychology of perception centers on explaining perception, as ordinarily conceived. It does not merely explain enabling conditions of perception in something like the way neuro-

⁹ Hallucinations are not fool's perception in the sense at issue. This is true even for the naïve realist, at least if they adopt Martin's view on which hallucinations lack any positive nature other than being indiscriminable from veridical perceptions of a certain kind (Martin 2004, 2006). Hallucinations, so conceived, *do* wear their nature on their sleeves. They mislead because their subjects cannot tell that they lack a further nature: namely that of actually being a veridical perception of a certain kind.

¹⁰ For a vigorous defence of both claims see Burge 2010: chpt. 9. See also Block 2012, 2016, and in Phillips and Block 2016 and Peters et al. 2017.

physiology explains the underlying neural enabling conditions for perception. Perceptual psychology, strange to have to say it, theorizes about perception. (2005: 46)

It is, however, possible to agree with much of what Burge wishes to insist on whilst still denying that perception is a psychological kind. In particular, it is possible to agree (a) that empirical psychology theorizes about perception in no less a way than chemistry theorizes about ordinary kinds such as gold or water, and (b) that scientific psychology will identify a single, fundamental kind which is naturally labelled “perception”, whilst denying that this kind should be identified with the ordinary kind, *perception*. In turn, this opens up the possibility that unconscious instances of the scientific kind are entirely consistent with experientialism, the view that perception is constitutively linked to consciousness.

To approach the view I have in mind, consider a somewhat arch question, posed by Johnston (1997), namely: is all gold, golden? Take this to mean: does all gold have the manifest form of gold (ibid: 577)? In favour of a negative answer, one might argue as follows: gold (as we now know) is simply the chemical element, Au; not every quantity of Au is golden (a single atom arguably lacks any manifest form; it certainly is not golden); so not all gold is golden. Against this, Johnston argues¹¹ that we should instead hold that the chemical kind, Au, constitutes the *manifest kind* gold. A manifest kind being “a kind whose instances we identify and re-identify on the basis of their manifest properties” (565). Moreover, just as some lumps do not constitute statues, not all quantities of Au constitute gold. For a quantity of Au to

¹¹ Johnston's argument for the case of water and H₂O runs as follows. If water = H₂O, then water vapour = H₂O and snow = H₂O, so by the transitivity of identity, snow = water vapour. This being absurd, Johnston argues we should replace the initial identity claim with a constitution claim. A similar argument can easily be given for gold and Au. Although Johnston's argument deserves much greater attention than I can give it here, I do not think it is decisive. An alternative reaction is to deny the conditional premise and hold that whilst water = H₂O, water vapour and snow are merely constituted by H₂O. Then since constitution is not symmetric, no absurd conclusion follows. Here and elsewhere in this section I'm grateful to Harvey Lederman for discussion.

constitute gold, that quantity must meet certain requirements of *manifest form* which a single atom of gold (or a scattering of such atoms) fails to meet.¹² On this view, the suggestion that all gold is golden can be reconciled with the existence of a chemical science of gold. This chemical science is a science of gold's constituent basis. Chemistry tells us that Au is what constitutes gold. But such constitution only occurs when certain formal constraints are met. Confusingly, the term "gold" is used to refer to both chemical and manifest kinds.

Now consider our central question: is all perception conscious? Perception's manifest form is plausibly its phenomenal nature, hence we can think of this question as equivalent to the question whether all perception has the manifest form of perception. Against this, one might argue as follows (again Burge 2010: 374-6 is exemplary): empirical psychology tells us that perception is a psychological kind; it also tells us that not every instance of this kind is conscious; so: not all perception is conscious. A Johnstonian constitutionalist can reply to this argument as follows. Call the kind identified by perceptual psychology, P. P should not be thought of as identical with, but as constituting, perception. Perception is rather a manifest kind. It is constituted by P only when certain requirements of manifest form are met. Such requirements involve having a phenomenal nature. Hence all perception is conscious. But this is quite consistent with the existence of a psychological science of perception whose kinds can occur unconsciously.

It is important to note that nothing in the above argument requires that psychology will uncover a single, fundamental kind which constitutes ordinary perception in the way, say,

¹² Temporal form, i.e. stability, is plausibly also a constraint. One might reasonably doubt that all the various radioisotopes of Au constitute the manifest kind gold.

that Au might be thought of as constituting gold.¹³ Concerning the kinds adduced by psychologists, the experientialist should be equally open to eliminativism and pluralism—eliminativism being the denial that any psychological kind deserves the appellation “perception,” and pluralism being the insistence that more than one kind deserves the appellation (perhaps depending on one’s explanatory interests).¹⁴ Strictly, all the experientialist requires is that ordinary perception is not *identifiable* with a psychological kind. Which kind or kinds constitute perception can be left as a complex and unsettled question.¹⁵

This constitutionalist picture just sketched shows that we can concede a great deal concerning the outcome of scientific psychology without being forced to acknowledge unconscious perception. As such it shows that the contention that ordinary perception can occur unconsciously involves a substantive metaphysical commitment. It is not simply an empirical datum. In what follows, however, I propose to grant this assumption and assume that perception is indeed a psychological kind and ask whether even then we are forced to acknowledge unconscious perception. There is a simple reason to be suspicious of any quick argument from perceptual psychology to such a conclusion. To see this notice that a *complete* psychological science will not simply offer a theory of perception but also a theory of

¹³ Even the case of Au and gold is more complex than it first seems. For, as Johnston points out, “almost all the gold we have ever encountered is a mixture of Au and other ingredients. These mixtures can vary significantly without there being any question but that what we have before us is a lump or nugget or coin of gold.” (577-8)

¹⁴ Nanay 2015 defends a version of pluralism drawing on precedent views in philosophy of biology concerning traits (e.g. Dupré 1993 and Hacking 2007). See also Matthen 1998. Note though that Nanay does not distinguish between manifest and scientific kinds and so does not consider the possibility of monism about the manifest kind and pluralism about the scientific kind(s).

¹⁵ The issues here closely connect to long-standing disputes as to where to draw “the” perception/cognition border (e.g. Firestone and Scholl 2016 and B. Phillips forthcoming).

consciousness.¹⁶ As a result any such complete science will uncover a scientific kind which correlates with *conscious perception*. This raises an obvious question: why should this correlate not be identified with perception? It cannot be enough to point out that psychology has also uncovered some other kind which does not correlate with conscious perception. What we need to know is why we should opt to identify this kind with perception instead of something else, say, perceptual or pre-perceptual processing. Thus, even granting that perception is a psychological kind, we need to justify the claim that this kind is not the very same kind as corresponds to conscious perception.

Put so bluntly, it is natural to object that when we actually do the science, many reasons will emerge for distinguishing perception from conscious perception. This is presumably why Burge takes the view he does in the parenthetical remark in the following notorious passage:

Perceptual psychology as it now stands does not attempt to give a complete theory of the essence of all perceptual states. For example, it is possible that consciousness is an aspect of the essence of some perceptual states. (It is almost surely, however, not an essential feature of all perception.) The psychological theories that I have discussed do not attempt to explain consciousness. There is, currently, no scientific theory of consciousness. (2005: 46)

In Part Two, I explore the empirical case for unconscious perception granting the assumption that perception is a psychological kind. For quite principled reasons, this case transpires to be extremely fraught. This shows that the simple concern just raised cannot be so easily dismissed.

¹⁶ I set scepticism about the very possibility of science of consciousness aside here. In the present dialectical context, such scepticism would naturally raise the prospect of related scepticism about a science of perception.

Before turning to these issues, one further point. Those who argue that water and gold are chemical kinds identical to H₂O and Au respectively will commonly point to the deferential behaviour of ordinary subjects when informed by experts that various substances are instances of the relevant kinds despite lacking their familiar, superficial characteristics. In contrast, the folk seem much more reluctant to accept alleged cases of perception without consciousness as genuine cases of perception. As Farah puts it, “Most people would say that one has not perceived something if one is not consciously aware of that thing.” (1994: 203) This judgement is even stronger, I submit, in relation to specific modes of perception such as hearing, seeing, and tasting. Furthermore, such reluctance is implicit amongst experts themselves. Thus scientists commonly talk of suppression techniques such as masking or flash-suppression (often used in purported demonstrations of unconscious perception) as rendering stimuli *invisible* (e.g. Song and Yao 2016). And when talking of unconscious perception or seeing, scare-quotes or qualifiers such as “in some sense” are commonly exploited.¹⁷ This suggests that even experts rely on a special, nonliteral or extended sense of “seeing” or “perceiving”. This may also help explain oxymorons such as blindsight, numb-sense, deaf-hearing as well as Goodale and Milner’s title *Sight Unseen* (2004). These considerations further motivate taking seriously the hypothesis that perception is a conscious kind, be it manifest or psychological.

¹⁷ An example of both: “There are also weird phenomena such as ‘blindsight’ (in which people who are consciously ‘blind’ behave in ways which clearly show they are, in some sense, able to ‘see’...)” (Black 2011: 15).

PART TWO

5. Contemporary Orthodoxy

According to contemporary orthodoxy, perception occurs unconsciously. This claim is rooted largely in alleged dissociations of perception and consciousness in clinical syndromes such as blindsight and neglect; in neurotypical subjects under conditions of inattention or suppression; and finally in arthropods such as bees and spiders.¹⁸ Contemporary orthodoxy is typically representationalist, often explicitly on the ground that only a representationalist approach to perception is consistent with vision science (e.g. Nanay 2014). This makes sense given the orthodox identification of perception with a psychological kind, since psychological science traffics in representational states within an information processing approach to the mind.

A standard bearer for the orthodox approach is Burge who, drawing on his reading of vision science, argues that perception is specifically *objective sensory representation by the individual* (2010: esp. chpt. 9). Burge's approach is notable for his focus on the necessity of perceptual constancies in achieving objective representation.¹⁹ As a result, for Burge, a key test for whether we have unconscious perception is whether constancy-implicating representation can occur unconsciously. But Burge also draws attention to a point which

¹⁸ See references in note 1.

¹⁹ For Burge, perception is *sensory* in that it begins with the registration of proprietary stimulation by a specialized system. For processing of this initial stimulation to produce perception, however, it must yield *objective representation*: “content that is *as of* a subject matter beyond idiosyncratic, proximal or subjective features of the individual” and instead of “entities in the physical environment” (2010: 397). Perception therefore requires capacities for distinguishing between “what concerns the individual's receptors and receptor-independent reality” (398). Specifically: “Perception requires perceptual constancies” (399)—“capacities to represent environmental attributes, or environmental particulars, as the same, despite radically different proximal stimulations” (114). For Burge, then, perceptual states must have an objective representational content accrued by dint of constancy mechanisms.

applies to any plausible explication of perception, though one rarely made explicit. This is that perception is *by the individual*.²⁰ This condition applies to any plausible explication because a representation not attributable to the individual could not possibly be identified with perception in the ordinary sense which is plainly a personal or individual-level achievement. It follows that a very general test for whether we have unconscious perception in some putative instance is whether the relevant representation is genuinely attributable to the individual. This point lies at the heart of the critique which follows where I argue that attempts to decouple perception and consciousness stumble at precisely this juncture.

My argument proceeds by partitioning studies of unconscious perception into two classes. I begin with the first class which purports to show discriminative responding in relation to a feature in the absence of consciousness of that feature. I argue that all such studies confront the so-called problem of the criterion and, as a result, cannot provide compelling evidence of unconscious perception. I illustrate this first in relation to blindsight (§6), before generalizing the point to a range of other paradigms (§7). I then turn to the second class of studies which avoids the problem of the criterion by attempting to show unconscious perception of a feature in the absence of discriminative responding in relation to that feature. I argue that such studies face what I call the problem of attribution: they fail to show that the states they implicate are attributable to the individual. I illustrate this with respect to priming studies at the so-called objective threshold (§8). Following this, I extend the worry to studies of

²⁰ Burge denies that all *perceptual representations* are attributable to individuals (i.e. whole animal organisms); some, he thinks, are only attributable “to their brains or other subsystems” (190). Thus some perceptual representations do not constitute perception proper. Block (personal communication) rejects the idea of genuinely perceptual representations in subsystems, and so insists that attribution to the individual is a constitutive requirement on perceptual representation. I take no stand on this issue. The critical point on which Block, Burge and other theorists agree is that *genuine perception* is constitutively by the individual. Such theorists also insist that individual-level perception can occur unconsciously (e.g. Burge 2010: 374).

cognitive control at the objective threshold and to work on vision-for-action (§9). I then complete my critique of the orthodox case for unconscious perception by considering arguments based on studies of perception in lower animals such as bees and spiders (§10). Finally, I conclude by responding to recent objections to my view levelled by Block (2016) (§11).

Before proceeding, a word about consciousness. Throughout, I use “consciousness” to refer to phenomenal consciousness without prejudice to its putative dissociation from what Block calls access consciousness (Block 1995 and 2005; see also Burge 2010: 188, fn. 90 and 190, fn. 95). Little consensus exists with respect to (phenomenal) consciousness. In what follows, I aim to remain as neutral as possible. I make one exception which is that I set aside actualist higher-order thought theories of consciousness since these threaten a trivial positive answer concerning the existence of unconscious perception. According to such approaches, a first-order perceptual state is conscious only if it is the object of an actual higher-order thought whose relationship to its target first-order state is merely causal (e.g. Rosenthal 1986, 2005). Such a view effectively entails the possibility of unconscious perceptual states, since for a perceptual state to be unconscious is simply for it to fail to be accompanied by a higher-order thought, and there is no necessity in the causal connection between lower and higher-order states.²¹

²¹ Not all higher-order theories need setting aside. For instance, whether a *dispositional* higher-order theory requires a commitment to unconscious perception will depend in part on precisely which stimulus conditions are treated as relevant to the manifestation of the higher-order disposition (i.e. under which counterfactual circumstances it is claimed that the state needs to be the object of a higher-order representation for it to be conscious in actuality). If such counterfactual circumstances include situations with different response criteria, then as will be clear from what follows, dispositional theories are potentially consistent with the claim that perception is always conscious. For higher-order theories of different varieties, see e.g. Armstrong 1968; Lycan 1987, 1996; Weiskrantz 1997: esp. 71-76; Carruthers 2005; and Weisberg 2011.

Actualist higher-order thought theories are highly controversial and much controverted (e.g. Dretske 1993, Byrne 1997, Block 2011). In contrast, the existence of unconscious perception is common ground not only amongst theorists who deny that consciousness requires higher-order representation (e.g. Dretske 2006, Prinz 2010, Brogaard 2011), but even amongst those who believe that an episode can be conscious even though its content is unavailable for use in reasoning or for the rational control of speech and action (Block 2012: 11-12, Burge 2010: 374-5). This common ground is built upon a shared empirical foundation. That foundation is my focus here.

6. Blindsight and the Problem of the Criterion

Blindsight is a condition least prejudicially defined as “residual visual processing after destruction of primary visual cortex [= V1 or striate cortex]” (Cowey 2010: 3). The standard view is that blindsight provides a clear case of unconscious perception. Thus according to Burge: “blindsight patients perceive environmental conditions. The perception involves perceptual constancies—including motion, location, and size constancies. The perception guides action. There is strong reason to believe that some of these patients lack phenomenal consciousness in the relevant perceptions.” (2010: 374) Here Burge is insisting that blindsight constitutes objective sensory representation by the individual without consciousness and therefore constitutes unconscious perception. Specifically, Burge holds that the representations in blindsight are objective since they involve various constancies, and they are individually attributable since they guide action.

Burge takes his claims about the preservation of constancies in blindsight as uncontested, principally citing the first-edition of Weiskrantz's canonical *Blindsight* (2009 [1986]). However, writing sixteen years later, Weiskrantz notes “that size constancy, or in fact any of

the visual constancies, has never been addressed in any blindsight studies of which I am aware" (2002: 572).²² Moreover, there are both theoretical and experimental grounds to doubt that constancies are preserved in blindsight. In terms of theory, we know that monkeys with extra-striate (V2) lesions respond to retinal image size not object size, arguably because of problems coding distance (Ungerleider et al. 1977, Humphrey and Weiskrantz 1969). V1 provides the primary feed-forward projections to V2, and its responses are highly attenuated if V1 is lesioned (Schmid et al. 2009). Furthermore, "V1 is adapted to the specific requirements of depth perception, so as to perform essential pre-processing of the signals it receives from the retinae" (Read 2005: 90). Essentially, V1 contains "a kind of 'cyclopean' retina" (ibid: 102) integrating monocular inputs ahead of further processing to yield depth. Given this, V1 lesions *predict* size constancy failures since they abolish pre-processing of signals arguably essential for depth perception (see further Sperandio and Chouinard 2015).

In terms of direct evidence, Alexander and Cowey (2010) find only evidence of sensory capacities to detect sharp luminance contours and/or stimulus transients in two patients (GY and MS); Azzopardi and Hock (2011) show that motion discrimination in GY is limited to "objectless" first-order motion energy (i.e. spatiotemporal changes in luminance) as opposed to changes in position or shape; and Kentridge et al. (2007) show that DB matches colored stimuli purely on the basis of wavelength and so lacks even the rudiments of color constancy mechanisms (see also Alexander and Cowey 2013, and Kentridge 2015). None of these

²² In the relevant paper Weiskrantz reports afterimages obeying Emmert's law in DB, the first intensively studied blindsight patient. Emmert's law implies constancy mechanisms (see Phillips 2013: §6). However, the "images" obeying Emmert's law in DB are *conscious* and so do not reveal a dissociation between constancies and consciousness. Indeed, Cowey (suggesting that DB's afterimages may be due to islands of preserved V1) comments: "How ironic if the discovery of blindsight proves to be based on a patient who does not possess it!" (2010: 7) Note that because of the metal clips implanted in DB's brain, imaging techniques cannot be used to establish the completeness of DB's lesion. Furthermore, whilst DB's more recent visual capacities have been impressive (Trevethan et al. 2007), it is quite possible that he has recovered some conscious vision.

capacities implies objective environmental representation. Thus, such evidence raises serious doubts as to whether blindsight constitutes genuine perception.²³

Rather than pursue that issue here, however, I want to focus on the question of whether blindsight patients do in fact lack consciousness in respect of their residual vision. Elsewhere, I have argued at length that, despite appearances, patients with blindsight and other clinical conditions such as unilateral neglect may well be conscious, albeit in a degraded or highly abnormal way.²⁴ At the crux of that argument is the notorious problem of the criterion. Here I briefly review this problem and its relevance to clinical cases like blindsight. I then extend my earlier work to show how this same problem applies to a range of other non-clinical paradigms alleged to demonstrate unconscious perception.

The problem of the criterion arises because tasks standardly used to assess awareness are subject to conservative “bias” and so risk underestimating awareness, especially at the limits of perception. In essence, these tasks are subject to bias because they require the subject themselves to decide whether a stimulus was present or seen, and subjects are often cautious in making such decisions. In contrast, tasks standardly used to establish residual sensitivity (i.e. perception) are naturally “unbiased”. In effect, they relieve the subject of responsibility in deciding whether a stimulus was present or seen. This leads to an apparent but potentially wholly artefactual dissociation between performance (i.e. perception) and consciousness.

²³ Burge also appeals to cases of action-blindsight (Danckert and Rosetti 2005) and evidence of unconscious attention in blindsight (Kentridge et al. 1999, 2004). I respond to these appeals in discussing the problem of attribution (§§8-9).

²⁴ See Phillips 2016. For a classic early statement of this concern in relation to blindsight see Campion et al. 1983.

The standard framework for modelling these ideas is signal detection theory (SDT) (Tanner and Swets 1954; Green and Swets 1966). To illustrate, consider the kind of “yes/no” (yn) task typically used to assess awareness, wherein a subject must say whether or not a stimulus has been presented on a given trial. Or, equally, whether or not they saw, or were aware of, a stimulus. SDT models a perceiver’s sensitivity to stimulus presence in terms of the distance between the means of two distributions of sensory responses—one associated with noise (in their sensory system and environment), the other with stimulus presence together with omnipresent noise. Making the large assumption that these distributions are normal and equivariant, this distance is given, in units of their standard deviation, by the parameter d' .

However, how often a subject responds “yes” is not settled by her sensitivity alone. This further requires knowing her criterion—the variable threshold which a sensory response must reach to generate a positive judgement.

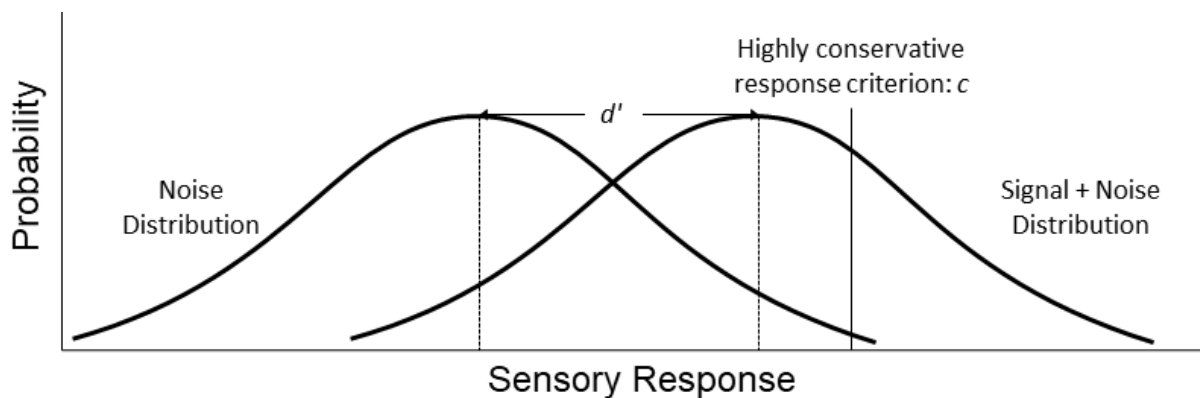


Fig. 1: SDT analysis of a simple yn task showing a highly conservative response criterion.

As can be seen from Fig. 1, a subject with reasonable perceptual sensitivity to stimulus presence ($d' \gg 0$) may repeatedly deny seeing a stimulus if her criterion is sufficiently conservative (far to the right). In this case, most sensory responses associated with stimulus-

presence will fall short of such a criterion and so go unreported. Consequently, knowing how often a subject correctly judges whether or not a stimulus is present is insufficient to determine her underlying perceptual sensitivity (Azzopardi and Cowey 1998). To determine sensitivity, the experimenter must either manipulate the subject's criterion, plotting a "receiver operating characteristic" (ROC) curve whose shape suffices to calculate d' , or exploit a forced-choice task. The distinctive advantage of a forced choice task is that it is naturally unbiased. For example, in a classic two-alternative forced-choice (2afc) task in which a subject must say in which of two intervals a stimulus is presented, subjects naturally adopt a symmetrical criterion, simply choosing whichever interval corresponds to the strongest sensory stimulation. As a result such tasks directly reveal perceptual sensitivity whenever it is present (Green and Swets 1966: 107-8).²⁵

A failure to report a stimulus in a biased yn task together with the presence of perceptual sensitivity ($d' > 0$) does not suffice to show unconscious perception. It is perfectly consistent with the subject being conscious (albeit perhaps in a dim, distorted or degraded manner) but operating with a conservative response criterion in relation to her conscious experience. Such a possibility is not outlandish. Ordinary subjects are naturally and systematically conservative in yn tasks (Björkman et al. 1993) and such biases have especially pronounced effects at the limits of perception—precisely where studies of unconscious perception probe.

²⁵ 2afc tasks are also (in a sense) *easier* than yn tasks since subjects in 2afc tasks are getting two bites at the cherry, one for each interval. On this point Macmillan and Creelman comment: "The relative ease of 2AFC has an impact on some aspects of subjective experience: Observers often report surprise that they can perform above chance with small stimulus differences, which they might be unwilling to report as above a yes-no criterion." (2005 [1991]: 179) This difference between the tasks can, of course, be corrected in the mathematical analysis of sensitivity. Note that whilst 2afc tasks are prized in empirical work because they are naturally unbiased, they are not problem free (see, e.g. Lin and Murray 2014).

To see how these points apply to blindsight, consider that “blindness” (i.e. lack of consciousness) is typically established by static perimetry—i.e. mapping the subject’s field defect by asking them whether or not they can see a light at various intensities and locations in their visual field. This is a biased yn task. Similarly, consider the so-called “commentary key” responses introduced by Weiskrantz to establish the absence of consciousness in a given task: one key indicating that the subject was aware of some “visual aspect of the stimulus presentation” (Weiskrantz et al. 1995: 6122), the other indicating that they were not. Again this is a biased yn task. In contrast, “sight” is typically established by asking the subject to indicate in which of two temporal or spatial intervals a stimulus is presented, i.e. a naturally unbiased 2afc task.²⁶

Rigorous psychophysical investigations of GY (Azzopardi and Cowey 1997, 1998, see also Azzopardi and Cowey 2001) show that his residual capacities are consistent with a detection theoretic model on which blindsight arises from conservative (and, in the case of static stimuli, unstable) criteria setting combined with residual perceptual sensitivity. Though such psychophysical work has not been conducted across the gamut of clinical conditions in which performance/awareness dissociations are observed, the same theoretical concerns apply. For example, the large majority of studies of unconscious perception in unilateral neglect—“the failure to report, respond, or orient to novel or meaningful stimuli presented to the side opposite a brain lesion, when this failure cannot be attributed to either [elementary] sensory

²⁶ Dissociations can also be found even when the “performance” task is not 2afc. For example, subjects may be able to say whether a line is horizontal or vertical, or a shape an “X” or an “O” (Weiskrantz et al. 1974). They may also be able to do much better than chance in a detection task if encouraged to guess. However, whilst potentially subject to bias, these tasks are much less likely to be biased than the yn tasks used to assess “awareness”, so again such dissociations can be readily explained in terms of differences in criteria across tasks.

or motor defects” (Heilman et al. 1993: 279)—can be interpreted as cases of residual sensitivity unreported due to conservative biases.²⁷

The implications of such concerns are delicate. As Macmillan (1986: 39) notes: “SDT takes no stand on whether below-criterion stimuli are consciously perceived.” SDT simply clarifies two thresholds: an “objective” threshold above which stimuli are *discriminable* above chance; and a “subjective” threshold above which a subject will respond positively in an appropriate task.²⁸ However, the only immediate way to block the concerns above is to insist that only stimuli above the subjective threshold are conscious, i.e. insist that an explicit report is necessary for consciousness. This would be anathema to theorists such as Block and Burge who hold that phenomenal consciousness dissociates from cognitive accessibility. It would also be anathema to theorists who claim that phenomenal consciousness dissociates from cognitive *access* even if not from accessibility (Snodgrass and Shevrin 2006: 75; cf. Block 2005 and discussion in Phillips 2016). The association of consciousness with the subjective threshold also ignores the many factors—“demand characteristics”, task design, experimental instructions, implicit and explicit pay-offs, prior probabilities, preconceptions, and natural propensities—which impact a subject’s judgements over-and-above the simple fact of their awareness. As Draine and Greenwald write: “it is well known that the stimulus-presentation conditions at which any perceiver places the boundary between judged presence and absence of a stimulus can be influenced by instructional or motivational variations. It is difficult to accept a subject’s assertion of subjective absence of a stimulus at face value when it is known

²⁷ For fuller discussion of all these points, especially in relation to neglect, see Phillips 2016.

²⁸ For this terminology see Cheesman and Merikle 1986 and Merikle and Cheesman 1986.

that the subject, with somewhat different instructions, might have indicated presence" (1998: 287).²⁹

This leaves us with two alternatives: embrace an objective threshold, i.e. treat all stimuli for which $d' > 0$ as conscious (e.g. Eriksen 1960, Holender 1986), or eschew any simple association between detection theoretic thresholds and consciousness.³⁰ Either way, our earlier concern stands: studies of patients with blindsight (and other clinical syndromes) fail to provide straightforward evidence of unconscious perception, since they are consistent with a simple, alternative interpretation which appeals to degraded conscious perception unreported due to conservative biases.

7. The Problem of the Criterion Extended

Such concerns are not restricted to studies of clinical patients. They extend to many paradigms standardly alleged to show unconscious perception in neurotypical individuals. Consider, for example, inattentional blindness (IB; Mack and Rock 1998, Most et al. 2001, 2005). In IB (see Fig. 2) naïve subjects fixate on a central cross for 1500msec before being presented with a larger cross in their parafovea for 200msec followed by a 500ms pattern mask. Subjects must say which arm of this cross is longer: a difficult task, demanding covert attention (i.e. attention unaccompanied by overt eye-movements and so foveation). On critical trials an unexpected critical stimulus is presented in whichever of the cross's quadrants corresponds to fixation (i.e. foveally). Only 20-40% of subjects report this stimulus on subsequent questioning. Nonetheless, such unreported critical stimuli can produce high-

²⁹ For examples of blindsight subjects reporting awareness under different instructions see Stoerig and Barth 2001, Overgaard et al. 2008 and Mazzi et al. 2016. See also the striking variation in awareness judgements suggestive of criterion shifts offered by GY under identical task conditions in Zeki and ffytche 1998 (see esp. Fig. 2).

³⁰ For an excellent discussion of this issue from a philosophy of science perspective see Irvine 2013.

level priming effects. For example, if the unreported stimuli are words, they can dramatically affect performance on later stem-completion tasks.

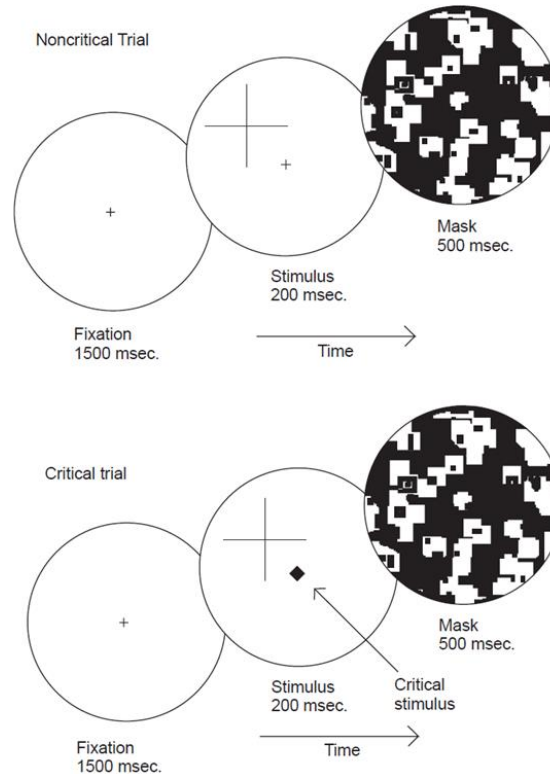


Fig. 2: Inattention Blindness (Mack and Rock 1998: 16). Copyright © MIT Press. Permission Pending.

On these grounds, many theorists claim that the missed stimuli in IB are unconsciously perceived (e.g. Mack and Rock 1998, Prinz 2015). This interpretation raises many concerns.³¹ However, a large but largely neglected issue is that the measure of awareness used is inadequate since patently subject to conservative bias. As Dulany complains:

³¹ The most common (Wolfe 1999, Moore 2001, Driver et al. 2001) is that the methodology fails to rule out the possibility that subjects consciously see the unattended stimulus but fail to encode it in such a way as to be able to report it. Ward and Scholl (2015) provide evidence against an inattentional amnesia version of this response but not an inattentional inaccessibility version (Block 2001). Given the use of priming to establish perception, such paradigms also face the problem of attribution discussed in §§8-9.

[Mack and Rock's] detection measure is only a subjective report of presence or identification subject to well-known criterion and bias effects (...). The experimenter's tone, or even inadequately phrased writing, can suggest that the real business of science was the evaluation of cross arms, and whether something else was observed is only incidental—and perhaps not to be mentioned if it would reveal that they hadn't kept their attention where they were told to (2001: §4).

Dulany also points out that Mack and Rock's awareness question was, "Did you see anything on the screen on this trial that had not been there on previous trials?" As a result, a subject might see the critical stimulus yet not report doing so because unsure whether or not it had been present on a *previous* trial.

Similar points apply to attentional blink (AB) paradigms in which a second target stimulus is "missed" in a roughly 400-600ms window following presentation of a first target in a rapid serial visual display (Chun and Potter 1995, Raymond et al. 1992, Luck et al. 1996, Marois et al. 2004, Pesciarelli et al. 2007). Such paradigms are again often interpreted as involving unconscious perception of the missed stimuli (e.g. Burge 2010: 375, fn. 12). Yet they also standardly use a biased detection measure (without calculating d') to argue that there is no conscious awareness of the "missed" target. Consider Luck et al. 1996 whose "awareness" measure is a plainly bias-prone task in which subjects have to say whether a word is semantically related to a target "context" word or not.³²

³² As with IB, one might also worry about whether such tasks measure awareness or memory. Thus, Luck et al. (1996: 617) write: "it is difficult to determine whether the probe words were identified without reaching awareness or if they momentarily reached awareness and were then rapidly forgotten".

As before, it is only if one insists that explicit report is necessary for consciousness that such paradigms can be directly interpreted as evidencing unconscious perception. Otherwise, even if one holds that consciousness requires actual access, it is possible to maintain (as Dulany suggests) that subjects *do* have access to the unreported stimulus but simply fail to report it. Alternatively, if one instead holds that consciousness requires only *accessibility*, then, if the scenarios relevant to determining accessibility include scenarios varying in response criterion and/or attentional distribution, one might agree with Lamme (2006: 496) that in IB and AB “the unattended information is not inaccessible ... just not currently accessed”. Lastly, one might agree with Block (2001) who suggests that the critical stimuli in IB are outright inaccessible to cognition due to the removal of attention *but yet phenomenally conscious*. For all these reasons, scepticism is warranted concerning the evidential force of attentional paradigms as regards unconscious perception.

A compelling case of unconscious perception requires that we address the problem of the criterion. Since the problem will arise in any case where a subject exhibits discriminative sensitivity, e.g. as witnessed by better than chance performance in a forced-choice task, the clearest way to avoid the problem is to turn to cases where discriminative sensitivity is at chance, i.e. $d' = 0$, the so-called “objective threshold”. I now turn to such studies.

8. The Objective Threshold and The Problem of Attribution

Can perception occur at the *objective* threshold? If so, we could straight-forwardly avoid the problem of the criterion. The demonstration of perception at the objective threshold is precisely the goal of much psychophysical work over the last several decades. A dominant paradigm is masked priming (Kouider and Dehaene 2007). In a typical such paradigm, the subject is first presented with a *subliminal* prime, i.e. a stimulus whose presence and/or

properties have been rendered invisible by masking. The subject is then presented with a *supraliminal* target (in some cases the mask itself) which the subject must respond to in some way (e.g. identify or categorize). Priming occurs when responses to the supraliminal target are modulated differentially (e.g. are faster or more accurate) when preceded by a congruent subliminal prime, as compared either to an unrelated or absent prime.

A vast literature debates the complex experimental and statistical issues which arise in evidencing genuine chance-level performance in such paradigms. The naïve reader beware: many well-known studies fail to establish that their effects are truly at the objective threshold.³³ Here, however, I want to consider the implications of what I take to be the consensus view in the field, namely that priming effects can be elicited at the objective threshold.

In many priming paradigms it is impossible to say whether the priming effect is mediated merely by sensory information which in some way correlates with an objective environmental feature, or whether it involves representation of the objective feature itself. To block the Burgean concern that mere sensory information would not constitute perception proper, I focus on a study by Norman et al. (2014) which aims to establish constancy-implicating priming effects—and thus genuinely objective representation—at the objective threshold. Norman et al. presented subjects with colored disc primes, followed by colored annulus targets which acted as metacontrast masks, suppressing awareness of the discs (see Fig. 3). Subjects were tasked to identify the color of the annulus as quickly as possible. By altering the illumination under which the annulus was presented and then using two different annuli, one green (matching in reflected color, i.e. wavelength), the other blue (matching in surface

³³ An excellent critical review can be found in Sand 2016.

color), Norman et al. were able to show that color identification responses were slightly but significantly faster when preceded by a prime which matched in surface as opposed to reflected color. This was true even when the prime was undetectable. Since surface color representations require the operation of color constancy mechanisms to discount the changing illumination, this indicates that “color constancy can occur in the absence of color experience” (2014: 2826).

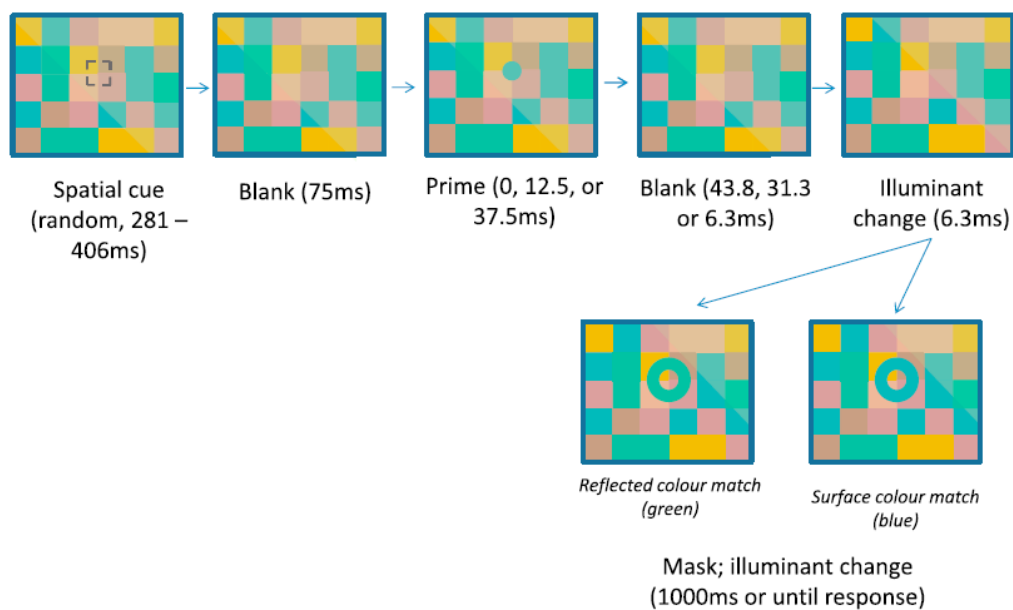


Fig. 3: Trial sequence used in Norman et al. 2014: 2824. Copyright © 2014 Elsevier Ltd.

Reprinted with permission.

Concerns can inevitably be raised whether awareness was completely absent.³⁴ I ignore these. Instead, I argue that the relevant priming data do not suffice to demonstrate genuine perception by the individual. This introduces what I call the problem of attribution.

³⁴ The awareness measure employed in Norman et al. 2014 is a confidence-scale measure of prime presence and not a test of prime-color discrimination (the feature driving the effect). Consequently, the awareness task does not assess awareness of the very same information which drives the priming effect (cf. Reingold and Merikle 1988, Schmidt and Vorberg 2006, and Schmidt 2007). Nonetheless, as Kentridge et al. write elsewhere, “it is

Genuine perception is an individual-level phenomenon. Above we saw Burge's insistence on this point. But it is equally a claim endorsed by psychologists. Thus consider Klotz and Neumann discussing evidence of motor activation at the objective threshold:

...the term *perception* seems logically inappropriate in this context (...). In ordinary usage, perceiving is something that a person or an animal does, not something that can be properly ascribed to stages, sub-systems, brains areas, or the like. The triggering of a sneeze by an external stimulus does not imply that the reflex center that controls it "perceives" the stimulus. Similarly, the activation of a manual response by a stimulus that cannot be consciously discriminated should not be called perception (1999: 976)

But when are states or occurrences attributable to the individual? On some answers to this question, unconscious perception is ruled out *a priori*. This is most obvious on the view (arguably suggested by Klotz and Neumann) that representations are properly attributable to individuals *just when they are conscious*.³⁵ Interestingly, Burge expresses occasional sympathy towards such a view. He writes at one point: "I do believe that there are certain notions of proprietary ownership of psychological states that hinge on consciousness. If bees are not conscious, they lack a certain *type* of ownership of their visual states." (2010: 190, his emphasis) And later he appears to connect attribution to the individual with accessibility to consciousness.

extremely unlikely that a subject is having any experience of ... stimuli whose presence or absence he or she cannot discriminate [better than chance]" (2008: 866; cf. Snodgrass et al. 2004a).

³⁵ Cf. the parable of John in King and Carruthers (2012: esp. 216-17).

Many processes that occur in perceptual systems ... are not attributable to individuals. Transformations of sensory information into perceptions and transformations among perceptions are almost never attributable to the individual. The individual does not make them occur; *they are not conscious or accessible to consciousness*; they are not exercises of the individual's central capabilities. But, necessarily and constitutively, individuals perceive. (369; my emphasis)

This is not Burge's official view, however. Burge joins contemporary orthodoxy in holding that "consciousness is not the basic factor in determining what in a perceptual system is an individual's and what is merely a subsystem's" (2010: 374).

Another traditional idea is that attributable states must be cognitively accessible. Kant famously declares in the *Deduction* (2003: B131–32) that, "It must be possible for the 'I think' to accompany all my representations; for otherwise something would be represented in me which could not be thought at all, and that is equivalent to saying that the representation would be impossible, or at least would be nothing to me." To Kant's dictum, Pereboom (2014) objects: "Plausibly, some of my representations are so thoroughly subconscious that I cannot attribute them to myself, while they are nevertheless mine due to the causal relations they bear to other representations and to actions that are paradigmatically mine." However, great care is needed here. Given the causal dependency of the individual on their subsystems, a *merely* causal condition on ownership will threaten the very idea of a distinction between attribution to the individual and attribution merely to a subsystem. Thus, even if Pereboom is right, we remain in need of a better developed criterion of individual attributability to replace Kant's.

More recent authors have tended to insist, as Evans puts it, that “it is not thoughts about the experience that matter, but thoughts about the world” (1982: 158). Thus, for Evans, only “when sensory input ... serves as the input to a *thinking, concept-applying, and reasoning system* ... [that] we can say that the person, rather than just some part of his brain, receives and processes the information” (ibid.). Both Kantian and Evansian criteria raise a serious difficulty for establishing unconscious perception.³⁶ It is plausible to think that only representations above the objective threshold are cognitively accessible in any sense. Hence, all stimuli below the objective threshold will fail to be attributable to the individual, and so objective threshold paradigms will be inherently incapable of establishing unconscious perception given the problem of the criterion discussed above.³⁷

If they are to avoid the problem of the criterion, the defender of unconscious perception must, it seems, defend a weaker condition for individual-attribution. Burge offers the following “key”: “Where a sensory state, non-perceptual or perceptual, can initiate action by an individual, it is attributable to the individual.... Sensory states that are integral to accounts of the initiation of such actions [eating, navigating, mating, etc.] are attributable to the individual.” (373) Burge also suggests that attributable perceptual representations “single out particulars that action aims for or aims to avoid” (370) and figure “directly in guiding action” (375). These are plausibly generic formulations, best read as claiming that the *paradigm* of an individually attributable representation is one which is “available to central coordinating agency” (333). In consequence, we cannot infer from the fact that a representation is

³⁶ The fuller quotation from Evans is: “we arrive at *conscious* perceptual experience when sensory input ... serves as an input to a *thinking, concept-applying, and reasoning system*” (my emphasis). In this way the condition for attribution is directly tied to the arrival of consciousness.

³⁷ Snodgrass and Shevrin suggest that being above the objective threshold is not merely necessary but *sufficient* for cognitive accessibility on the ground that “when sufficiently motivated, participants can lower their criterion and include such stimuli in reflective [i.e. access] consciousness” (2006: 75).

unavailable to central agency that it is not an individual's. Nonetheless, a representation's availability to central agency comprises our best evidence for attributing a representation to an individual. Indeed, it provides our basic grip on idea. Furthermore, where a representation is not so available, a question will naturally arise as to what grounds we have for thinking of the representation as the individual's.³⁸

This question arises forcefully in respect of the representations implicated in priming paradigms such as Norman et al.'s. Subjects in Norman et al.'s study respond to the colored annulus whose color they are told to identify. This annulus is consciously seen. The datum of interest is that their identification judgements are faster when the annulus is preceded by a surface color matched prime as opposed to a reflected color matched prime (or by no prime at all). How should we understand this effect? One natural and extremely generic understanding is that the prime generates activity in the perceptual system, up to and including surface color representations. In turn, this activation of color representations means that the perceptual system is better able to process subsequent stimuli which match in surface color. Subsequent processing is more "fluent". Annuli matching in surface color are thus perceived more quickly (and potentially more accurately), and the subject is able to identify them faster. Though skeletal and speculative, this story suffices to show why Norman et al.'s paradigm provides no grounds for attributing representations of the prime's color to the individual. For the only role that such representations are required to play is to speed subsequent processing *within the perceptual system*. There is no reason to think that, in order to explain the facilitation effect, such representations must be available to central coordinating agency, let alone to thought or reasoning.

³⁸ For a non-generic formulation of a similar idea see Dretske's test for perception in his 2006 where he insists that the information in genuinely perceptual states must be directly available for the control and guidance of action.

By way of analogy, consider a much simpler case. Certain stimuli cause pupillary dilation despite not being consciously perceived. It is easier to see, and so respond to, various things with dilated pupils. Thus, a stimulus which elicits pupillary dilation will facilitate responses to various subsequent stimuli. But such facilitation effects are clearly insufficient grounds for claiming that the dilation provoking stimulus was itself perceived, or even registered, by the individual. The effects of the stimulus may well be limited to the relevant circuitry governing the pupillary response (i.e., retina, optic nerve, and pretectum).

The priming effect in Norman et al.'s study may also arise because the prime automatically activates associated responses in the *motor* system (Schmidt et al. 2006, Kentridge in Peters et al. 2017). However, such "response" priming does no better at demonstrating that the relevant information is available for individual-level action control and guidance. Such priming plausibly activates the motor system in a way which bypasses central agency (Kunde et al. 2003, Ansorge et al. 2011, and discussion in §9). The problem of attribution thus casts the methodology of using priming to establish perception into grave doubt. And this remains true even if we adopt Burge's relatively weak criterion for attribution to the individual.

9. The Problem of Attribution Extended

The concerns of the last section generalize beyond simple perceptual and motor priming to a wide range of popular paradigms. Consider a currently much-discussed paradigm often claimed to reveal unconscious perception is continuous flash suppression (CFS; Tsuchiya and Koch 2005). This is a technique in which distinct stimuli are presented to each eye, one a flashing "Mondrian" pattern which is thought to suppress awareness of the other for several seconds. CFS experiments typically rely on an indirect measure to establish that the suppressed

stimulus is genuinely perceived.³⁹ However, even making the large assumption that CFS entirely abolishes consciousness of the suppressed stimulus, these indirect measures fail to implicate individual-level perception. Thus, for example, Raio et al. (2012) used acquired skin conductance responses as an index of fearful face perception. Yet such responses are manifestations of the autonomic nervous system, not of central agency. Similarly, Jiang et al. (2006) show that gendered nudes presented under suppression can differentially elicit reflexive spatial attention. But why think that stimulus-driven, reflex-like attentional responses count as manifestations of central agency, and so witness individual-level perception?

The same basic complaint can be made about the (nonetheless extremely important) demonstration of unconscious object-based attention in Norman et al. 2013. What this study shows is that *attentional processing* can be elicited by unseen objects. However, to the extent that subjects cannot themselves use representations of the attended but unseen objects to guide their responses, such representations do not witness genuine perception. The differential processing they produce is instead akin to a stimulus-driven reflex, operating entirely outside of voluntary, agentic control.⁴⁰

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

³⁹ I here leave aside so-called “breaking CFS” studies (e.g. Yang et al. 2007, Mudrik et al. 2011) which arguably simply reveal differences in conscious detectability (Stein et al. 2011). For a much fuller discussion of this issue and flash suppression paradigms in relation to unconscious perception see Phillips and Block 2016.

⁴⁰ Cf. Buehler (unpublished manuscript) who comments: “Individuals typically cannot suppress exogenous orientation to a stimulus, even if they know that the stimulus interferes with their ongoing actions, and even if they try to suppress the reflex.” Buehler cites Giordano et al. 2009: 8 and Carrasco 2011: 1488.

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. In what remains of this section I consider two challenges to this verdict. First, the existence of priming effects apparently involving intentional action or cognitive control. Second, evidence of unconscious action-guiding representations associated with the dorsal stream.

As Kentridge (in Peters et al. 2017) argues, “Unseen primes can do much more than elicit motor responses. They can modulate switching between ‘task-sets’ (e.g. Lau & Passingham, 2007), they can slow or completely inhibit responses by priming ‘no-go’ signals (e.g. van Gaal et al, 2009) and even modify task goals in masked semantic priming (e.g. Fitzsimons & Bargh, 2003).” Do such cases not witness perception by the individual, as Kentridge suggests?

Many studies in this area (including those Kentridge cites) face serious concerns about whether the stimuli were genuinely subliminal. (For extensive discussion here see Sand 2016, esp. chpt. 8.) A reasonable suspicion is that many notable studies simply illustrate criterion effects *within conscious perception* (i.e. subjects doing one thing when a stimulus falls above a subjective criterion, another when it falls below). With that concern in mind, I focus here on a body of work by Snodgrass, Shevrin and colleagues which stands out in its ambition to establish that the primes are truly unconscious. I then explain why, contrary to the view of its authors, the evidence from this paradigm still does not show perception by the individual. I then generalize this claim to other studies (relaxing concerns about awareness).

Across a series of experiments, Snodgrass, Shevrin and colleagues have sought to show that intentional judgements can be mediated by unconscious stimuli.⁴¹ Their basic paradigm involves the tachistoscopic presentation of one of four possible emotionally valenced words (e.g. Pain, Rose) for 1ms against a field of uniform luminance which renders them undetectable ($d' = 0$) due to so-called “energy masking” (Turvey 1973). Under these conditions, subjects are asked to identify the presented words. Unsurprisingly, *overall* identification performance is at chance (i.e. 25%). However, subjects are instructed to adopt one of two strategies: a “look” strategy in which they are urged to rely on “on any available conscious perception” (i.e. attend carefully to any partially perceived aspects of the stimulus) and a “pop” strategy in which they are told “to respond with the first word that [comes] to mind” (2004a: 858). Subjects are also asked which strategy they prefer and performance is analysed for each preference group when exploiting their preferred and non-preferred strategies. The striking outcome of this analysis is that when “lookers” (subjects who preferred the “look” strategy) exploit their favoured “look” strategy their identification performance increases to slightly, but significantly, above chance (~28% accurate). In contrast, when lookers adopt a “pop” strategy, their performance falls slightly, but significantly, *below* chance (~22-3% accurate).⁴²

There are a number of questions one might press here. However, let us suppose that the finding is robust. Minimally this indicates that identification judgements can be *affected* by subliminal stimuli. But does it provide evidence of genuine individual-level perception? In particular, are the relevant representations available to central agency, and so exploitable by,

⁴¹ See Snodgrass et al. 1993a (substantially replicated by Van Selst and Merikle 1993 on which see Snodgrass et al. 1993b), Snodgrass et al. 2004a, b, Snodgrass and Shevrin 2006.

⁴² I base this summary on the meta-analyses and large scale-replication in Snodgrass and Shevrin 2006: §§8-9 and 13. These suggests slightly different reliable interaction effects than their original studies.

or attributable to, subjects? The notion of central agency is a placeholder for whichever systems subserve an agent's capacity for genuine, individual-level action. In his classic discussion, Frankfurt argues that action (so understood) involves *guidance* by the agent.

When we act, our movements are purposive ... their course is guided.... The dilation of the pupils ... does not mark the performance of an action by the person; his pupils dilate, but he does not dilate them. This is because the course of the movement is not under his guidance.

The guidance in this case is attributable only to the operation of some mechanism with which he cannot be identified. (1988 [1978]: 159)⁴³

Do the effects found by Snodgrass and colleagues count as cases where the course of an agent's action is purposively guided or steered by the agent in the relevant respect? No. Plausibly, they should rather be regarded as cases where an intentional act is, to paraphrase Frankfurt, *affected* by the operation of some mechanism with which the agent themselves cannot be identified.⁴⁴ This is consistent with Snodgrass and Shevrin's own assessment. For in their own view the effects of the subliminal words in their paradigm are "radically uncontrollable" (2006: 73 and §17), operating quite independently of the subject's own intentions. This is most evident in the case of the looker under "pop" instructions. Here the subject's intentions are to identify the word correctly, yet their unconscious processing of the word actually *impairs* their performance, bringing it below chance. True, in the case of the looker adopting a "look" strategy the subject's performance is facilitated in line with their subject-level intentions. However, this is plausibly mere happy coincidence given the counter-volitional effects found under different instructions. In consequence, though the

⁴³ Buehler 2014 offers a substantive account of central agency based on this Frankfurtian starting point.

⁴⁴ An analogy: imagine that whilst shooting baskets someone uses targeted TMS on your motor cortex to interfere with your performance.

effects of the words may accidentally coincide with the subject's intentions, their effects are non-volitional. Subjects cannot exploit the words to guide their behaviour. The words merely affect their behaviour outside their direct control. If perception is by the individual and such attribution requires availability for action guidance, then this is not perception.

As mentioned, much other work on unconscious cognitive control is problematized by problems with the assessment of consciousness. Nonetheless, even if we bracket such concerns, a plausible reaction to this literature is to concede that it shows significant effects of stimuli which are not consciously perceived on high-level cognitive processes, but to insist that such effects are, in Snodgrass and Shevrin's terms, intrinsically and radically uncontrollable.⁴⁵ In short: such stimuli cannot be exploited by subjects to guide and control their actions, and so fall foul of the problem of attribution. To illustrate, consider an exemplary study by Cressman et al. (2013). In it subjects were asked to make a pointing movement from a "home" position towards a central square in a display of three squares. As Fig. 4 shows, in 70% of trials subjects were then presented with a neutral (star) prime followed by a neutral metacontrast mask (Box A). In these trials the subject was simply to continue with their pointing movement towards the central square target. In 30% of trials, however, participants saw a *non-neutral* (arrow) mask which pointed one way in 80% of such trials and the other way in 20% of such trials (Box B). Subjects were instructed to alter their pointing movement quickly in response to such masking arrows. These masking arrows were preceded either by neutral (star) or directional (arrow) primes.

⁴⁵ Snodgrass and Shevrin relate this to the familiar line of thought (e.g. Cheesman and Merikle 1986, Merikle and Joordens 1997, Merikle et al. 1995) that stimuli below the *subjective* threshold are intrinsically uncontrollable and so unconscious. However, as Snodgrass and Shevrin go on to argue, there are good reasons to think that "subjective threshold effects are at least potentially controllable" (2006: 70). Instead, then it is the *objective* threshold which is most plausibly associated with *intrinsic* uncontrollability.

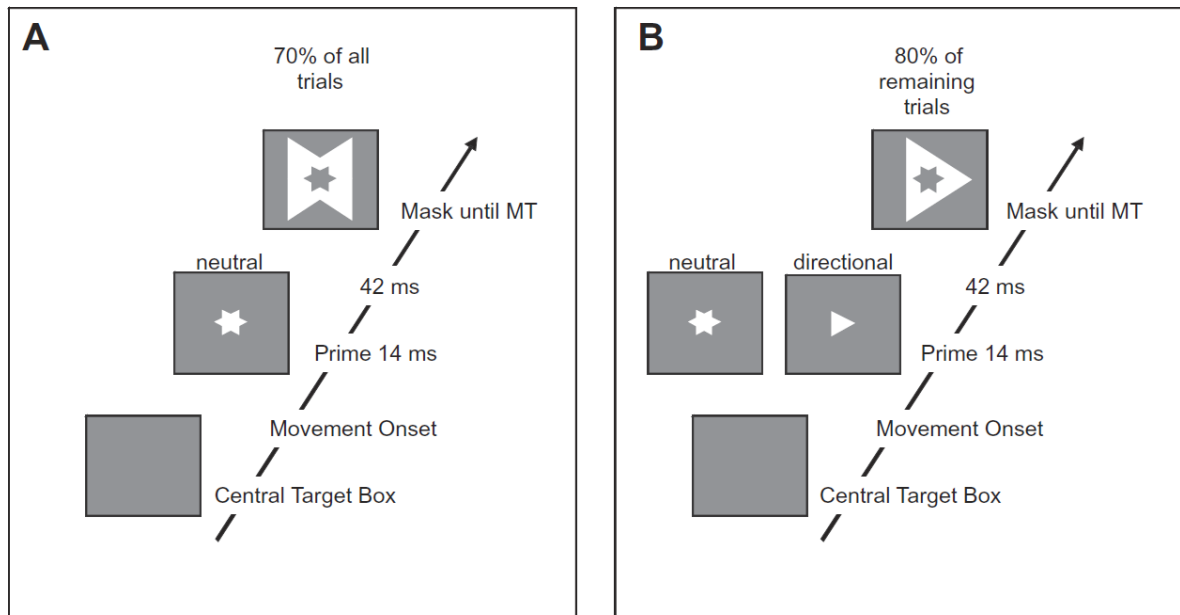


Fig. 4: Temporal sequence of priming from Cressman et al. 2013: 718. Copyright © 2013

Elsevier Inc. Reprinted with permission.

Subjects knew that the masking arrows were strongly biased in one direction (e.g. 80% were right-arrows). As a result they exhibited intentional response bias, responding significantly faster to these high-probability masks than their low-probability counterparts. Subjects' responses also revealed effects of the primes. (I assume, purely for the sake of discussion, that these were not consciously perceived.) Responses were faster to masks preceded by a congruent prime as compared to an incongruent or neutral prime. Critically, however, "the influence of the invisible prime was not affected by the probability expectation associated with the visible mask" (720). In other words, whereas the *visible* masks modulated subjects' performance in line with the known bias in mask probabilities, the invisible prime directly activated its associated response regardless of the subjects' expectations and intentions.

These results fits nicely with the discussion of control and guidance above. They suggest that invisible primes can *affect* response selection, directly activating the motor system in relation

to a pre-learned response. But they also suggest that such activation bypasses the agent's own control and guidance as revealed by the fact that such activation occurs quite independently of the subject's knowledge and intentions. In short: invisible primes are not useable or exploitable by the individual to guide their actions. They are not then the objects of genuine perception even given a minimal criterion for individual attribution.⁴⁶

A very different kind of case often appealed to in defence of unconscious perception concerns so-called vision-for-action representations associated with the dorsal stream. These are most familiar from studies of patients with visual form agnosia (Milner and Goodale 2006 [1995]). But similar issues arise in connection with so-called "action-blindsight" (Dankert and Rossetti 2005, Whitwell et al. 2011). A classic paradigm used to study such patients involves asking the patient to judge the width of a circular disc either verbally or using their thumb and forefinger. The inaccuracy of such judgements suggest that such patients lack constancy-implicating perception. However, when asked to reach out and grasp the discs in question, accurate grip scaling implicates (at least vergence based) size-constancy mechanisms (Sperandio et al. 2012; Marotta et al. 1997; see also Servos 2006; Mon-Williams et al. 2001).

⁴⁶ It is doubtful that Burge would accept the link here insisted on between availability to central coordinating agency and guidance or control. He claims that perceptually-guided ducking "can be an action even if it is against one's own attempt to inhibit the ducking" (2010: 334), and on this ground denies a control or guidance condition on agency. However, it is obscure how Burge proposes to distinguish between ducking (here presumably conceived of as an escape reflex) and saccadic eye movements which he claims are "normally not imputed to individuals" (333). Neither behaviour is straightforwardly a reflex "in the classical sense", and both occur involuntarily in some but not all situations. I suggest it is much more attractive to retain the link between guidance and agency in line with Frankfurt. That said, the proper development of an account of agency and guidance is an enormous task. The claim here is only that a natural central conceit about agency animates scepticism about unconscious perception in light of extant empirical evidence. For a substantially stronger account of individual attributability in terms of cognitive integration which would rule out many representations—including even those implicated in familiar cases of blindsight—from counting as individually attributable see Bayne 2013.

A standard interpretation is that subjects' actions are guided by visual representations which are not available for offline report and judgement, and are not associated with consciousness.

Do such representations constitute unconscious perception? Two now familiar and interlocking issues arise here. Suppose we think that the fine-grained modulation of grip-size constitutes the guidance of action by the individual. Then we need to ask why this does not suggest that the relevant representations are conscious?⁴⁷ The more promising response, I suggest, however, is to deny that vision-for-action constitutes genuine perception on the grounds that the pertinent modulations of behaviour (e.g. grasp aperture) do not witness genuine control and guidance *by the individual*, and so fail to meet relevant conditions for perception proper.⁴⁸ This view is strongly suggested by the metaphors which theorists use to describe the dorsal system. It is an automatic pilot (Pisella et al. 2000), a tele-assisted semi-autonomous robot (Goodale and Humphrey 1998: §9, Goodale and Milner 2004: 98-103, Milner and Goodale 2006 [1995]: §8.2.3) or a heat-seeking missile (Campbell 2002: 56).

One way of understanding the force of such metaphors is as claiming that representations involved solely in fine-grained motor programming are not individually attributable. They are confined to the autonomous robot—a sub-system of the individual. In contrast, the only representations attributable to the individual are those associated with target and action-type

⁴⁷ Returning to previous discussion of the problem of the criterion, we also need to ask whether the measures of awareness are suitably sensitive and unbiased. Here consider Whitwell et al. who write of their patient, SJ: “It is important to note, however, [that her] failure to show a target redundancy effect in our experiment does not mean that she is completely incapable of detecting targets in her blind field (using a button press). Had we used a forced-choice variant of this task she may have very well exhibited better-than-chance levels of performance.” (2011: 915)

⁴⁸ Whilst Milner and Goodale themselves do think that perception can occur unconsciously, they insist that “[t]he visual information used by the dorsal stream for programming and on-line control ... is not perceptual in nature” (2008: 776; cf. 2006 [1995]: 2). However, this is because they take potential for consciousness to be criterial of the perceptual.

selection which are associated with consciousness (cf. Clark 2007: 576). This claim closely echoes discussion of unconscious priming above. For example, Danckert and Rossetti report how the dorsal-parietal system “often functions automatically, rapidly modifying visually guided hand movements ... in contradiction to conscious commands” (2005: 1042, see also Pisella et al. 2000). Such representations arguably thus affect and modulate behaviour despite not being useable or exploitable by the individual themselves to guide their actions.

In summary, the problem of attribution threatens a host of different paradigms which have been interpreted in terms of unconscious perception. In combination with the problem of the criterion, a hypothesis suggests itself, namely that the conditions for individual attributability coincide with those for perceptual consciousness. I have not argued directly in favour of this hypothesis. Nonetheless, its simplicity and consistency with the data, mean that we must take it seriously. And that of course is to take seriously the claim that there is no such thing as (individual-level) perception without consciousness.

10. Lower Animals

In this penultimate section, I consider whether a case for unconscious perception can be mounted by looking to evidence of perception in lower animals. Here is Burge mounting that case:

A ... set of considerations that strongly suggests that perception by individuals need not be conscious derives from what is known about animal perception.... some arthropods clearly have perceptual capacities.... certain spiders visually perceive color, shape, motion, spatial location, and so on. They exhibit associated perceptual constancies. Whether ... spiders are phenomenally conscious is unknown. These cases are not known to illustrate individual

perception without consciousness. But the epistemic situation supports not taking consciousness to be constitutive of individual perception (375; cf. Block 2012: 11-12)

In short, according to Burge, we are in the following “epistemic situation”: we know that jumping spiders can perceive, we do not know whether they are conscious. This is said to support us not only in “not taking consciousness to be constitutive of individual perception” but in thinking “that perception by individuals need not be conscious”, i.e. taking consciousness not to be constitutive of individual perception.

Even granting that we do know that such animals “illustrate individual perception” in Burge’s sense, this form of argument is doubly problematic. First, constitution can plausibly be a posteriori (Kripke 1980). Thus, the fact that we know a state to be perceptual without knowing whether it is conscious no more shows that perception is not constitutively conscious than the fact that a young Lavoisier knew rain to be water without (yet) knowing it to be H₂O shows that water is not constitutively H₂O. Furthermore, even if it is a priori *knowable* that all perceptual states are conscious, it does not follow that it is *known*. Perhaps we are beguiled by erroneous theoretical beliefs concerning perception or consciousness or both, and so fail to exploit our epistemic position.⁴⁹ Either way, the appeal to arthropods falls short of establishing a dissociation between perception and consciousness.

⁴⁹ Cf. Williamson (2006: §2) on Peter who denies that all vixens are female foxes despite fully understanding the relevant concepts because he: (a) falsely believes that for a claim of the form “All *F*s are *G*s” to be true at least one *F* must exist; and (b) has the “weird belief” that there are no vixens after gullibly reading a conspiracy theory website. Williamson uses Peter to try to show that there are no conceptual truths. An alternative reaction is to think that Peter is, in virtue of his semantic competence, *in a position to know* that all vixens are female foxes, but, because of his misguided beliefs, is unable to exploit his epistemic position.

Things would be quite different if we had good evidence for thinking that arthropods exhibiting perceptual constancies were *not* conscious. This appears to be Block's view:

We have many theories of what consciousness is in the brain and none of those that are taken seriously by substantial numbers of working neuroscientists apply to bees or spiders. (You can see what working scientists think of panpsychism here: (Block et al., 2014)). For example there is no evidence of anything approximating a “global neuronal workspace” in arthropods. So we have some—far from decisive—scientific reason to believe that spiders and bees have no conscious states. (2016: 453)

This is not the place for a proper evaluation of the complex theoretical and methodological issues surrounding the assessment of cognition and consciousness in lower animals.⁵⁰ However, it is worth emphasising three points. First, Block's attempt to tar the hypothesis that spiders are conscious with the same brush as panpsychism should be ignored. What working scientists ridicule in the letter which Block cites (and indeed co-authors) is “the view that electrons are conscious” (Block et al. 2014: 557). Spiders are not electrons. Second, it is obscure exactly what argumentative force the alleged inapplicability of “global neuronal workspace” models to spiders has, given that Block himself denies that such models capture *phenomenal* consciousness (e.g. Block 2005).⁵¹

A third point is the most telling. One of the “working scientists” whom Block invokes is Christoph Koch. Yet, writing in 2008, Koch highlights experiments demonstrating flexible

⁵⁰ For an excellent philosophical introduction see Allen and Trestman 2014, especially §6.6.

⁵¹ Relatedly, Block elsewhere objects to integrated information theory (Tononi and Edelman 1998) on the grounds that it fails to distinguish intelligence from consciousness, arguing that “on the face of it, mice or even lower animals might have phenomenal consciousness without much intelligence” (2009: 1112).

and sophisticated working memory (and other cognitive capacities) in bees (Giurfa et al. 2001; Brandt et al. 2005). Koch comments:

Although these experiments do not tell us that bees are conscious, they caution us that we have no principled reason at this point to reject this assertion. Bees are highly adaptive and sophisticated creatures ... Given all of this ability, why does almost everybody instinctively reject the idea that bees or other insects might be conscious? The trouble is that bees are so different from us and our ilk that our insights fail us. But just because they are small and live in colonies does not mean that they can't have subjective states.... I am not a mystic. I am not arguing for pan-psychism..... What this dilemma highlights is that there is no accepted theory of consciousness, no principled theory that would tell us which systems, organic or artificial, are conscious and why. In the absence of such a theory, we must at the very least remain agnostic about consciousness in these creatures.

Koch does not simply exemplify a major figure in the neuroscience of consciousness who takes the possibility of consciousness in bees seriously. His remarks also raise a question about Block's contention there is "no evidence of anything approximating a 'global neuronal workspace' in arthropods". The idea of a global workspace is closely linked to the classical concept of working memory—Block (2008: 306) talks of "the 'working memory' system—the 'global workspace'". Thus, one might wonder whether the evidence of flexible working memory in bees which Koch points to is not precisely evidence of something approximating a "global neuronal workspace". Moreover, although, as Barron and Klein (2016) note, global workspace theories exhibit a "strongly cortical bias" which might appear to exclude lower animals, theorists who take insect consciousness seriously will point to evidence that midbrain mechanisms play a crucial role in information integration and behavioural control—

a role analogous to (or perhaps even involved in) our global neuronal workspace (see also Merker 2007).

As with bees, it was long supposed that spiders were all-alike “simple, instinct-driven automatons” (Jackson and Cross 2011: 115). However, as Jackson and Cross continue: “research on spider biology is revealing increasing evidence of their cognitive abilities”. In particular, work on certain species of jumping spider (*Salticidae*) demonstrates that these spiders possess (in rudimentary form) many of the capacities standardly appealed to in theories of consciousness, e.g. selective attention (Jackson and Cross 2011), working memory (Cross and Jackson 2014), and complex and flexible “problem solving, decision making and forward planning” abilities (Jackson et al. 2006: 290). Indeed, so impressed are Jackson and colleagues with these spiders’ sophisticated approach to predation that they refer to them as “eight-legged cats” (Harland and Jackson 2000; also Stimson Wilcox and Jackson 1998; Cross and Jackson 2006). In line with Koch’s comments on bees, none of this shows that jumping spiders are conscious (nor indeed would the absence of such capacities conclusively show that they were not). It does, however, suggest that the hypothesis should not be dismissed out of hand. We should remain agnostic. And agnosticism does not support an argument for unconscious perception.

11. Block’s Objections and Concluding Remarks

Block has recently objected to my earlier criticisms of alleged cases of unconscious perception, claiming that they “are an ad hoc group with no unity—except the superficial unity of ‘not perception’ or ‘not unconscious’” (2016: 452). He continues:

Breitmeyer (2015) describes 24 methods of producing unconscious perception ... what is the likelihood that each of the 24 paradigms is subject to its own fatal flaw? All of the 24 paradigms have passed the test of peer review, and in many cases have been subject to many years of intense scrutiny and subsequent refinement.... Without some unified reason for skepticism, the plausibility that something different is wrong with each of the 24 methods is not high. (ibid.)

Block relates this to an “Anna Karenina Principle” regarding conscious perception: “All conscious perceptions are alike but each unconscious perception is unconscious in its own way.” (ibid.)

In response to this, first note that Block's appeal to Breitmeyer is not entirely candid. For Breitmeyer does not describe 24 methods of “producing unconscious perception”, as Block claims, but “24 ways to noninvasively suppress the conscious *report* of visual stimuli” (Breitmeyer 2015: 235, my emphasis). This distinction matters: suppressing reports is one thing, abolishing phenomenal consciousness another—a point which Block himself has repeatedly made in relation to many of the 24 methods which Breitmeyer lists. Consider inattentional and change blindness, which Block (2011) argues involves inaccessible phenomenal states, or crowding, at least special cases of which Block (2012) argues involves phenomenal consciousness outside attention. Breitmeyer shares Block's views on these matters, commenting:

Similar to what happens with visual crowding and Gestalt-switching methods, with any of these [attentional] methods [of suppression], most likely here again all stimuli do register in what Block (2011) refers to as perceptual or phenomenal consciousness but, for lack of focused attention, do not register in access consciousness and thus go unnoticed. (2015: 239)

Thus, Block and Breitmeyer themselves both deny that seven or eight of the peer-reviewed and intensely scrutinized “methods of producing unconscious perception” produce unconscious perception!

Like conscious perception, conscious report is plausibly subject to its own “Anna Karenina Principle”: many things need to go “right” in order for a subject to issue a conscious report, and many different things can go “wrong” to explain why no report is issued: failures of attention, memory, confidence, understanding and motivation, not to mention basic physical capacity (absent e.g. in extreme cases of locked-in syndrome). Block can hardly disagree that these motley failures are all legitimate concerns to raise in relation to a putative finding of unconscious perception. Many of them echo points he makes himself.⁵² Moreover, many of these potential failures are compressed by SDT into a single parameter, viz. response criterion (Green and Swets 1966: 118-9). SDT thereby abstracts a powerful and unified criticism of many studies of unconscious perception: the problem of the criterion encountered above.

More generally, the criticisms of this paper at least are not simply a piecemeal exercise in devil’s advocacy. They can be read as taking seriously a unified hypothesis, namely that the conditions for genuine individual-level perception coincide with the conditions for perceptual consciousness. A prediction of this hypothesis is that suppression methods which eliminate phenomenal consciousness will thereby eliminate genuine perception. This is entirely

⁵² Regarding brain damage, for example, Block acknowledges that the damage may have affected “the cognitive processing underlying the subjects’ reports” (2016: 453) as opposed to abolishing phenomenal consciousness. Here “cognitive processing” is presumably a catch-all for the many processes required for subjective report.

consistent with there being many ways in which consciousness may be suppressed, for there are doubtless many ways in which perception may not be achieved.

The hypothesis that the conditions for genuine perception coincide with the conditions for perceptual consciousness returns us to a traditional conception of perception as a determinate of consciousness encountered at the outset of this paper. It also accords with one natural reading of detection theory as applied to individual subjects on which discriminative sensitivity affords a measure both of perception and of perceptual consciousness. We should not cleave dogmatically to such views if they are empirically untenable. However, as we have seen, the empirical case in favour of unconscious perception, whilst superficially overwhelming, is on proper consideration eminently questionable. As argued in Part One, perception in its ordinary sense may be essentially conscious even if a related scientific kind is not—either because perception in its ordinary sense is a manifest kind, or because perception should be identified with the scientific kind which correlates with conscious perception. Moreover, as argued in Part Two, even setting aside such concerns, when properly recognized as an individual-level phenomenon, it is far from obvious that there is (or could be) evidence for unconscious perception given the dilemma posed by the problem of the criterion and the problem of attribution.

Acknowledgments

Ideas in this paper have been presented at Antwerp, Barnard-Columbia, Cambridge, Edinburgh, Manchester, NYU, Oxford, Princeton, Warwick, York, a Rethinking the Senses workshop at Centre L'Oubradou in France, and at the Association for the Scientific Study of Consciousness in Buenos Aires. I cannot begin to thank everyone for the many stimulating

Forthcoming in *Analytic Philosophy*. Author's Final Draft.

questions and challenges I received on those occasions but I would like to record special gratitude to Paul Azzopardi, Tim Bayne, Ned Block, Denis Buehler, David Chalmers, Sam Clarke, Martin Davies, Will Davies, Robert Foley, Anil Gomes, Chris Hill, Mark Johnston, Bob Kentridge, Hakwan Lau, Harvey Lederman, Hemdat Lerman, Eric Mandelbaum, Brian McLaughlin, John Morrison, Bence Nanay, Chris Peacocke, Megan Peters, Jake Quilty-Dunn, Louise Richardson, Nick Shea, Josh Shepherd, Wayne Wu, and as always to Hanna Pickard. I also gratefully acknowledge the support of the Leverhulme Trust under grant number RF-2013-278 during the early days of this project.

References

- Alexander, I. and Cowey A. 2010. "Edges, color and awareness in blindsight" *Consciousness and Cognition* 19: 520–533.
- Alexander, I. and Cowey, A. 2013. "Isoluminant colored stimuli are undetectable in blindsight even when they move" *Experimental Brain Research* 225: 147–152.
- Allen, C. and Trestman, M. 2014. "Animal Consciousness" In E. N. Zalta (ed.) *The Stanford Encyclopedia of Philosophy (Summer 2014 Edition)*.
- Anaya, A. and Clarke, S. 2017. "Naïve realism and unconscious perception: A reply to Berger and Nanay" *Analysis* 77(2): 267–273.
- Ansorge, U., Fuchs, I., Khalid, S. and Kunde, W. 2011. "No conflict control in the absence of awareness" *Psychological Research* 75: 351–365.
- Armstrong, D. M. 1968. *A Materialist Theory of the Mind*. London: Routledge.
- Azzopardi, P. and Cowey, A. 1997. "Is blindsight like normal, near-threshold vision?" *Proceedings of the National Academy of Sciences* 94(25): 14190–14194.

Forthcoming in *Analytic Philosophy*. Author's Final Draft.

Azzopardi, P. and Cowey, A. 1998. "Blindsight and visual awareness" *Consciousness & Cognition* 7(3): 292–311.

Azzopardi, P. and Cowey, A. 2001. "Why is Blindsight Blind?" In De Gelder, B., De Haan, E. H. F. and Heywood, C. A. (eds) *Out of Mind: Varieties of Unconscious Processes*. Oxford University Press, pp. 3–19.

Azzopardi, P. and Hock, H. S. 2011. "Illusory motion perception in blindsight" *Proceedings of the National Academy of Sciences USA* 108: 876–881.

Barron, A. and Klein, C. 2016. "What insects can tell us about the origins of consciousness" *Proceedings of the National Academy of Sciences* 113(18): 4900–4908.

Bayne, T. 2013. "Agency as a marker of consciousness" In T. Vierkant, J. Kiverstein and A. Clark (eds) *Decomposing the Will*. Oxford: OUP, pp. 160–182.

Berger, J. and Nanay, B. 2016. "Relationalism and unconscious perception" *Analysis* 76(4): 426–433.

Björkman, M., Juslin, P. and Winman, A. 1993. "Realism of confidence in sensory discrimination: The underconfidence phenomenon" *Perception & Psychophysics* 54: 75–81.

Black, D. M. 2011. *Why things matter: the place of values in science, psychoanalysis and religion*. Routledge.

Block, N. 1995. "On a Confusion about a Function of Consciousness" *Behavioral and Brain Sciences* 18(2): 227–287.

Block, N. 2001. "Paradox and Cross Purposes in Recent Work on Consciousness" *Cognition* 79(1-2): 197–219.

Block, N. 2005. "Two Neural Correlates of Consciousness" *Trends in Cognitive Sciences* 9(2): 46–52.

Forthcoming in *Analytic Philosophy*. Author's Final Draft.

Block, N. 2008. "Consciousness and cognitive access" *Proceedings of the Aristotelian Society* 108(3): 289–317.

Block, N. 2009. "Comparing the Major Theories of Consciousness" In M. Gazzaniga (ed.) *The Cognitive Neurosciences IV*, Cambridge: MIT Press, pp. 1111–1122.

Block, N. 2010. "Attention and Mental Paint" *Philosophical Issues* 20: 23–63.

Block, N. 2011. "The higher order approach to consciousness is defunct" *Analysis* 71(3): 419–431.

Block, N. 2012. "The Grain of Vision and the Grain of Attention" *Thought: A Journal of Philosophy* 1(3): 170–84.

Block, N. 2016. "The Anna Karenina Principle and Skepticism about Unconscious Perception" *Philosophy and Phenomenological Research* 93(2): 452–459.

Block, N., Carmel, D., Fleming, S. M., Kentridge, R. W., Koch, C., Lamme, V. A. F., Lau, H. and Rosenthal, D. 2014. "Consciousness science: real progress and lingering misconceptions" *Trends in Cognitive Sciences* 18(11): 556–7.

Brandt, R., Rohlfing, T., Rybak, J., Krofczik, S., Maye, A., Westerhoff, M., Hege, H. C. and Menzel, R. 2005 "Three-Dimensional Average-Shape Atlas of the Honeybee Brain and Its Applications" *Journal of Comparative Neurology* 492(1): 1–19.

Breitmeyer, B. 2015. "Psychophysical 'blinding' methods reveal a functional hierarchy of unconscious visual processing" *Consciousness and Cognition* 35: 234–50.

Brewer, B. 2011. *Perception and its Objects*. Oxford: OUP.

Brogaard, B. 2011. "Are There Unconscious Perceptual Processes?" *Consciousness & Cognition* 20(2): 449–63.

Forthcoming in *Analytic Philosophy*. Author's Final Draft.

Buehler, D. 2014. *Psychological Agency: Guidance of Visual Attention*. Doctoral Thesis, University of California, Los Angeles.

Buehler, D. Unpublished Manuscript. "A dilemma for 'selection-for-action'".

Burge, T. 2005. "Disjunctivism and Perceptual Psychology" *Philosophical Topics* 33: 1–78.

Burge, T. 2010. *The Origins of Objectivity*. Oxford: OUP.

Byrne, A. 1997. Some like it HOT: Consciousness and higher-order thoughts. *Philosophical Studies* 2(2): 103–29.

Campbell, J. 2002. *Reference and Consciousness*. Oxford: OUP.

Campbell, J. 2010. "Demonstrative Reference, the Relational View of Experience, and the Proximity Principle" In R. Jeshion (ed.) *New Essays on Singular Thought*. Oxford: OUP, pp. 193–212.

Campion, J. Latto, R. and Smith, Y. M. 1983. "Is blindsight an effect of scattered light, spared cortex, and near-threshold vision?" *Behavioural and Brain Sciences* 3: 423–486.

Carrasco, M. 2011. "Visual attention: The past 25 years" *Vision Research* 51(13): 1484–1525.

Carruthers, P. 2005. *Consciousness: Essays from a Higher-Order Perspective*. Oxford: OUP.

Cheesman, J. and Merikle, P. M. 1986. "Distinguishing conscious from unconscious perceptual processes" *Canadian Journal of Psychology* 40(4): 343–367.

Chun, M. M. and Potter, M. C. 1995. "A two-stage model for multiple target detection in rapid serial visual presentation" *Journal of Experimental Psychology: Human Perception and Performance* 21: 109–127.

Forthcoming in *Analytic Philosophy*. Author's Final Draft.

Clark, A. 2007. "What Reaching Teaches: Consciousness, Control, and the Inner Zombie" *British Journal for the Philosophy of Science* 58: 563–594.

Cowey, A. 2010. "The blindsight saga" *Experimental Brain Research* 200: 3–23.

Cressman, E. K., Lam, M. Y., Frank, I. M., Enns J. T. and Chua, R. 2013. "Unconscious and out of control: subliminal priming is insensitive to observer expectations" *Consciousness & Cognition* 22(3): 716–28.

Cross, F. R. and Jackson, R. R. 2006. "From eight-legged automatons to thinking spiders" In K. Fujita and S. Itakura (eds) *Diversity of Cognition: Evolution, Development, Domestication and Pathology*. Kyoto: Kyoto University Press, pp. 188–215.

Cross, F. R. and Jackson, R. R. 2014. "Specialised use of working memory by *Portia africana*, a spider-eating salticid" *Animal Cognition* 17(2): 435–44.

Danckert, J. and Rossetti, Y. 2005. "Blindsight in action: What can the different subtypes of blindsight tell us about the control of visually guided actions?" *Neuroscience and Biobehavioral Reviews* 29(7): 1035–1046.

Draine, S. C. and Greenwald, A. G. 1998. "Replicable unconscious semantic priming" *Journal of Experimental Psychology: General* 127: 286–303.

Dretske, F. 1993. "Conscious experience" *Mind* 102(406): 263–283.

Dretske, F. 2006. "Perception without awareness" In Gendler, T. S. and Hawthorne, J. (eds) *Perceptual Experience*. Oxford: OUP.

Driver, J., Davis, G., Russell, C., Turatto, M. and Freeman, E. 2001. "Segmentation, attention and phenomenal visual objects" *Cognition* 80: 61–95.

Dulany, D. E. 2001. "Inattentional awareness" *Psyche: An Interdisciplinary Journal of Research on Consciousness* 7(5).

Forthcoming in *Analytic Philosophy*. Author's Final Draft.

Dupré, J. 1993. *The Disorder of Things. Metaphysical foundations of the disunity of science*. Cambridge, MA: Harvard University Press.

Eriksen, C. W. 1960. "Discrimination and learning without awareness: a methodological survey and evaluation" *Psychological Review* 67: 279–300.

Evans, G. 1982. *The Varieties of Reference*. Oxford: OUP.

Farah, M. J. 1994. "Visual perception and visual awareness after brain damage: A tutorial review" In M. Moscovitch and C. Umiltà (eds) *Conscious and Unconscious Information Processing: Attention and Performance XV* Cambridge: MIT Press, pp. 37–76. Reprinted in N. Block, O. Flanagan and G. Guzeldere (eds) *The Nature of Consciousness: Philosophical Debates*. Cambridge: MIT Press, 1997, pp. 203–236. (Page references are to this reprint.)

Firestone, C., & Scholl, B. J. 2016. Cognition does not affect perception: Evaluating the evidence for 'top-down' effects. *Behavioral and Brain Sciences* 39, doi:10.1017/S0140525X15000965

Fitzsimons, G. M. and Bargh, J. A. 2003. "Thinking of you: Nonconscious pursuit of interpersonal goals associated with relationship partners" *Journal of Personality and Social Psychology* 84: 148–164.

Frankfurt, H. 1988 [1978]. "The Problem of Action" In his *The Importance of What We Care About*, New York: CUP.

Giordano, A. M., McElree, B., Carrasco, M. 2009. "On the automaticity and flexibility of covert attention: A speed-accuracy trade-off analysis" *Journal of Vision* 9(3): 1–10.

Giurfa, M., Zhang, S., Jenett, A., Menzel, R. and Srinivasan, M. V. 2001. "The Concepts of 'Sameness' and 'Difference' in an Insect" *Nature* 410: 930–933.

Goodale, M. A. and Humphrey, G. K. 1998. "The objects of action and perception" *Cognition* 67(12): 181–207.

Forthcoming in *Analytic Philosophy*. Author's Final Draft.

Goodale, M. A. and Milner, D. 2004. *Sight Unseen: An Exploration of Conscious and Unconscious Vision*. Oxford: Oxford University Press.

Green, D. M. and Swets J. A. 1966. *Signal Detection Theory and Psychophysics*. New York: Wiley.

Hacking, I. 2007. "Natural kinds: Rosy Dawn, Scholastic Twilight" *Royal Institute of Philosophy Supplement* 82(61): 203–239.

Harland, D. P. and Jackson, R. R. 2000. "'Eight-legged cats' and how they see—a review of recent research on jumping spiders (*Araneae: Salticidae*)" *Cimbebasia* 16: 231–240.

Heilman, K. H., Watson, T. W. and Valenstein, E. 1993. "Neglect and related disorders" In Heilman, K. H. and Valenstein, E. (eds) *Clinical Neuropsychology*. New York: OUP, pp. 279–336.

Holender, D. 1986. "Semantic activation without conscious identification in dichotic listening, parafoveal vision, and visual masking: A survey and appraisal" *Behavioral and Brain Sciences* 9: 1–66.

Humphrey, N. K. and Weiskrantz, L. 1969 "Size Constancy In Monkeys With Inferotemporal Lesions" *Quarterly Journal of Experimental Psychology* 21: 225–38.

Irvine, E. 2013. *Consciousness as a Scientific Concept*. Dordrecht: Springer.

Jackson, R. R. and Cross, F. R. 2011. "Spider cognition" *Advances in Insect Physiology* 41: 115–174.

Jackson, R. R., Cross, F. R. and Carter, C. M. 2006. "Geographic variation in a spider's ability to solve a confinement problem by trial and error" *International Journal of Comparative Psychology* 19(3): 282-296.

Forthcoming in *Analytic Philosophy*. Author's Final Draft.

Jiang, Y. V., Costello, P., Fang, F., Huang, M. and He, S. 2006. "A gender- and sexual orientation dependent spatial attentional effect of invisible images" *Proceedings of the National Academy of Sciences* 103(45): 17048–17052.

Johnston, M. 1997. "Manifest Kinds" *Journal of Philosophy* 94(11): 564–583.

Kant, I. 2003. *Critique of Pure Reason*, trans. N. Kemp Smith. New York: Palgrave Macmillan.

Kanwisher, N. 2001 "Neural Events and Perceptual Awareness" *Cognition* 79: 89-113.

Kentridge, R. W. 2015. "What is it like to have type-2 blindsight? Drawing inferences from residual function in type-1 blindsight" *Consciousness & Cognition* 32: 41–4.

Kentridge, R. W., Heywood, C. A. and Weiskrantz, L. 1999. "Attention without awareness in blindsight" *Proceedings of the Royal Society of London B* 266: 1805–1811.

Kentridge, R. W., Heywood, C. A. and Weiskrantz, L. 2007. "Color-contrast processing in human striate cortex" *Proceedings of the National Academy of Sciences (USA)* 104: 15129–15131.

Kentridge, R. W., Heywood, C. A., and Weiskrantz, L. 2004. "Spatial attention speeds discrimination without awareness in blindsight" *Neuropsychologia* 42: 831–835.

Kentridge, R. W., Nijboer, T. C. W. and Heywood, C. A. 2008. "Attended but unseen: Visual attention is not sufficient for visual awareness" *Neuropsychologia* 46: 864–869.

King, M. and Carruthers, P. 2012. "Moral responsibility and consciousness" *Journal of Moral Philosophy* 9: 200–228.

Klotz, W. and Neumann, O. 1999. "Motor activation without conscious discrimination in metacontrast masking" *Journal of Experimental Psychology: Human Perception and Performance* 25(4): 976–992.

Forthcoming in *Analytic Philosophy*. Author's Final Draft.

Koch, C. 2008. "Exploring consciousness through the study of bees" *Scientific American* November 26, 2008. <http://www.scientificamerican.com/article/exploring-consciousness/>

Kouider, S. and Dehaene, S. 2007. "Levels of processing during non-conscious perception: a critical review of visual masking" *Philosophical Transactions of the Royal Society of London, Series B, Biological Sciences* 362(1481): 857–875.

Kripke, S. 1980. *Naming and Necessity*. Cambridge: Harvard University Press.

Kunde, W., Kiesel, A. and Hoffman, J. 2003. "Conscious control over the content of unconscious cognition" *Cognition* 88(2): 223–42.

Lamme, V. A. F. 2006. "Towards a true neural stance on consciousness" *Trends in Cognitive Sciences* 10(11): 494–501.

Lau, H. C. and Passingham, R. E. 2007. "Unconscious activation of the cognitive control system in the human prefrontal cortex" *The Journal of Neuroscience* 27: 5805–5811.

Lin, Z. and Murray, S. O. 2014. "Priming of awareness or how not to measure visual awareness" *Journal of Vision* 14(1): 1–17.

Luck, S. J., Vogel, E. K., and Shapiro, K. L. 1996. "Word meanings can be accessed but not reported during the attentional blink" *Nature* 383: 616–618.

Lycan, W. G. 1987. *Consciousness*. Cambridge, MA: MIT Press.

Lycan, W. G. 1996. *Consciousness and Experience*. Cambridge, MA: MIT Press.

Mack, A. and Rock, I. 1998. *Inattention blindness*. Cambridge, MA: MIT Press.

Macmillan, N. A. 1986. "The psychophysics of subliminal perception" *Behavioural and Brain Sciences* 9(1): 38–39.

Macmillan, N. A. and Creelman, C. D. 2005 [1st edition 1991]. *Detection Theory: A User's guide*. Cambridge: Cambridge University Press.

Forthcoming in *Analytic Philosophy*. Author's Final Draft.

Marois, R., Yi, D. J. and Chun, M. M. 2004. "The neural fate of consciously perceived and missed events in the attentional blink" *Neuron* 41: 465–472.

Marotta, J. J., Behrmann, M. and Goodale, M. A. 1997. "The removal of binocular cues disrupts the calibration of grasping in patients with visual form agnosia" *Experimental Brain Research* 116: 113–121.

Martin, M. G. F. 2009 [1997]. "The Reality of Appearances" In H. Logue and A. Byrne (eds), *Disjunctivism: Contemporary Readings*. Cambridge, MA: MIT Press, pp. 91–115.

Martin, M. G. F. 2004. "The limits of self-awareness" *Philosophical Studies* 120(1-3): 37–89.

Martin, M. G. F. 2006. "On being alienated" In T. Gendler & J. Hawthorne (eds) *Perceptual Experience*. Oxford: Oxford University Press, pp. 354–410.

Matthen, M. 1998. "Biological Universals and the Nature of Fear" *Journal of Philosophy* 95: 105–32.

Mazzi, C., Bagattini, C. and Savazzi, S. 2016. "Blind-sight vs. degraded-sight: different measures tell a different story" *Frontiers in Psychology*. 7: <http://dx.doi.org/10.3389/fpsyg.2016.00901>.

Merikle, P. M. and Cheesman, J. 1986. "Consciousness is a 'Subjective' State" *Behavioral and Brain Sciences* 9: 42–43.

Merikle, P. M. and Joordens, S. 1997. "Parallels between perception without attention and perception without awareness" *Consciousness & Cognition* 6(2-3): 219–36.

Merikle, P. M., Joordens, S., and Stolz, J. 1995. "Measuring the relative magnitude of unconscious influences" *Consciousness & Cognition* 4: 422–439.

Merikle, P. M., Smilek, D., and Eastwood, J. D. 2001. "Perception without awareness: perspectives from cognitive psychology" *Cognition* 79: 115–34.

Forthcoming in *Analytic Philosophy*. Author's Final Draft.

Merker, B. 2007. "Consciousness without a cerebral cortex: A challenge for neuroscience and medicine" *Behavioral and Brain Sciences* 30(1): 63–81.

Milner, A. D. and Goodale, M. A. 2008. "Two visual systems re-viewed" *Neuropsychologia* 46(3): 774–85.

Milner, D. and Goodale, M. A. 2006 [1st edition 1995]. *The Visual Brain in Action*. Oxford: OUP. (Page references are to the second edition.)

Mon-Williams, M., Tresilian, J. R., McIntosh R. D. and Milner, A. D. 2001. "Monocular and binocular distance cues: Insights from visual form agnosia I (of III)" *Experimental Brain Research* 139: 127–136.

Moore, C. M. 2001. "Inattentional Blindness: Perception or memory and what does it matter?" *Psyche: An Interdisciplinary Journal of Research on Consciousness*, 7.

Moore, G. E. 1925. "A Defence of Common-Sense" in Muirhead, J. H. (ed.) *Contemporary British Philosophy* (2nd series) London: Allen and Unwin. Reprinted in his *Philosophical Papers* London: George, Allen and Unwin, 1959. (Page references are to this reprint.)

Most, S. B., Scholl, B. J., Clifford, E. and Simons, D. J. 2005. "What you see is what you set: Sustained inattentional blindness and the capture of awareness" *Psychological Review* 112: 217–242.

Most, S. B., Simons, D. J., Scholl, B. J., Jimenez, R., Clifford, E. and Chabris, C. F. 2001. "How not to be seen: The contribution of similarity and selective ignoring to sustained inattentional blindness" *Psychological Science* 12(1): 9–17.

Mudrik, L., Breska, A., Lamy, D. and Deouell, L. 2011. "Integration without awareness: expanding the limits of unconscious processing" *Psychological Science* 22(6): 764–770.

Forthcoming in *Analytic Philosophy*. Author's Final Draft.

Nanay, B. 2014. "Empirical problems with anti-representationalism" In B. Brogaard (ed.) *Does Perception have Content?* Oxford: OUP, pp. 39–50.

Nanay, B. 2015. "The Representationalism versus Relationalism Debate: Explanatory Contextualism about Perception" *European Journal of Philosophy* 23(2): 321–336.

Norman, L. J., Akins, K., Heywood, C. A., and Kentridge, R. W. 2014. "Color constancy for an unseen surface" *Current Biology* 24(23): 2822–6.

Norman, L. J., Heywood, C. A., and Kentridge, R. W. 2013. "Object-based attention without awareness" *Psychological Science* 24(6): 836–43.

O'Shaughnessy, B. 2000. *Consciousness and the World*. Oxford: OUP.

Overgaard, M., FehI, K., Mouridsen, K., Bergholt, B., and Cleeremans, A. 2008. "Seeing without seeing? Degraded conscious vision in a blindsight patient" *PLoS ONE* 3(8): e3028. doi:10.1371/journal.pone.0003028

Palmer, S. E. 1999. *Vision science: Photons to phenomenology*. Cambridge, MA: Bradford Books/MIT Press.

Parfit, D. 1984. *Reasons and Persons*. Oxford: Clarendon Press.

Pereboom, D. 2014. "Kant's Transcendental Arguments" In E. N. Zalta (ed.) *The Stanford Encyclopedia of Philosophy (Fall 2014 Edition)*.

Pesciarelli, F., Kutas, M., Dell'Acqua, R., Peressotti, F., Job, R. and Urbach, T. P. 2007. "Semantic and repetition priming within the attentional blink: An event-related brain potential (ERP) investigation study" *Biological Psychology* 76: 21–30.

Peters, M. A. K., Kentridge, R. W., Phillips, I. B. and Block, N. 2017. "Does unconscious perception really exist? Continuing the ASSC 20 Debate" *Neuroscience of Consciousness* 3(1): 11–1.

Forthcoming in *Analytic Philosophy*. Author's Final Draft.

Phillips, B. Forthcoming. "The Shifting Border Between Perception and Cognition" *Noûs*
doi: 10.1111/nous.12218

Phillips, I. 2013. "Afterimages and sensation" *Philosophy and Phenomenological Research*
87(2): 417–453.

Phillips, I. 2016. "Consciousness and criterion: on Block's case for unconscious seeing"
Philosophy & Phenomenological Research 93(2): 419–451.

Phillips, I. and N. Block. 2016. "Debate on unconscious perception" In In Nanay, B. (ed.)
Current Controversies in Philosophy of Perception New York: Routledge, pp. 165–192.

Pisella, L., Gréa, H., Tilikete, C., Vighetto, A., Desmurget, M., Rode, G., Boisson, D. and
Rossetti, Y. 2000. "An 'automatic pilot' for the hand in human posterior parietal cortex:
toward reinterpreting optic ataxia" *Nature Neuroscience* 3(7): 729–36.

Price, H. H. 1932. *Perception*. London: Methuen & Co. Ltd.

Prinz, J. J. 2010. "When is Perception Conscious?" In Nanay, B. (ed.) *Perceiving the World*
Oxford: OUP, pp. 310–332.

Prinz, J. J. 2015. "Unconscious Perception" in M. Matthen (ed.) *Oxford Handbook of
Philosophy of Perception*. Oxford: OUP, pp. 371–389.

Raio, C. M., Carmel, D., Carrasco, M. and Phelps, E. A. 2012. "Nonconscious fear is quickly
acquired but swiftly forgotten" *Current Biology* 22(12): R477–R479.

Raymond, J. E., Kimron, L. S. and Arnell, K. M. 1992. "Temporary suppression of visual
processing in an RSVP task: An attentional Blink" *Journal of Experimental Psychology:
Human Perception and Performance* 18(3): 849–860.

Read, J. 2005. "Early computational processing in binocular vision and depth perception"
Progress in Biophysics and Molecular Biology 87: 77–108.

Forthcoming in *Analytic Philosophy*. Author's Final Draft.

Reingold, E. M. and Merikle, P. M. 1988. "Using direct and indirect measures to study perception without awareness" *Perception & Psychophysics* 44: 563–575.

Rosenthal, D. 1986. "Two Concepts of Consciousness" *Philosophical Studies* 49(3): 329–359.

Rosenthal, D. 2005. *Consciousness and Mind*. Oxford: Clarendon Press.

Sand, A. 2016. *Subliminal or not?* PhD Thesis, Stockholm University.

Schmid, M. C., Panagiotaropoulos, T., Augath, M. A., Logothetis, N. K., Smirnakis, S. M. 2009 "Visually driven activation in macaque areas V2 and V3 without input from the primary visual cortex" *PLoS One* 4(5): e5527.

Schmidt, T. 2007. "Measuring unconscious cognition: Beyond the zero-awareness criterion" *Advances in Cognitive Psychology* 3(1): 275–287.

Schmidt, T. and Vorberg, D. 2006. "Criteria for unconscious cognition: Three types of dissociation" *Perception & Psychophysics* 68(3): 489–504.

Schmidt, T., Niehaus, S. and Nagel, A. 2006. "Primes and targets in rapid chases: Tracing sequential waves of motor activation" *Behavioural Neuroscience* 120: 1005–1016.

Servos, P. 2006. "Preservation of Emmert's law in a visual form agnostic" *Neurocase* 12(2): 122–6.

Shoemaker, S. 1970. "Persons and Their Pasts" *American Philosophical Quarterly* 7(4): 269–285.

Simpson, J. and Weiner, E. 1989. *Oxford English Dictionary* (2nd ed.) Oxford: Clarendon Press.

Snodgrass, M. and Shevrin, H. 2006. "Unconscious inhibition and facilitation at the objective detection threshold" *Cognition* 101(1): 43–79.

Forthcoming in *Analytic Philosophy*. Author's Final Draft.

Snodgrass, M., Bernat, E. and Shevrin, H. 2004a. "Unconscious perception: A model-based approach to method and evidence" *Perception & Psychophysics* 66: 846–867.

Snodgrass, M., Bernat, E. and Shevrin, H. 2004b. "Unconscious perception at the objective detection threshold exists" *Perception & Psychophysics* 66(5): 888–895.

Snodgrass, M., Shevrin, H. and Kopka, M. 1993a. "The Mediation of Intentional Judgments by Unconscious Perceptions: The Influences of Task Strategy, Task Preference, Word Meaning, and Motivation" *Consciousness & Cognition* 2: 169–193.

Snodgrass, M., Shevrin, H. and Kopka, M. 1993b. "Absolute Inhibition is Incompatible with Conscious Perception" *Consciousness & Cognition* 2: 204–209.

Snowdon, P. 1990. "The Objects of Perceptual Experience" *Proceedings of the Aristotelian Society* 64: 121–50.

Song, C. and Yao, H. 2016. "Unconscious processing of invisible visual stimuli" *Scientific Reports* 6(38917): 1–6.

Sperandio, I. and Chouinard, P. A. 2015. "The Mechanisms of Size Constancy" *Multisensory Research* 28: 253–283.

Sperandio, I., Whitwell, R. L., Chouinard, P. A. and Goodale, M. A. 2012. "Preservation of size constancy for action, but not perception, in a patient with bilateral occipital lesions" *Journal of Vision* 12(9): 837.

Stein, T., Hebart, M. and Sterzer, P. 2011. "Breaking continuous flash suppression: a new measure of unconscious processing during interocular suppression?" *Frontiers in Human Neuroscience* 5(167): 1–17.

Forthcoming in *Analytic Philosophy*. Author's Final Draft.

Stimson Wilcox, R. and Jackson, R. R. 1998. "Cognitive abilities of araneophagic jumping spiders" In I. Pepperberg, A. Kamil and R. Balda (eds) *Animal Cognition in Nature*. San Diego: Academic Press, pp. 411–433.

Stoerig, P. and Barth, E. 2001. "Low-level phenomenal vision despite unilateral destruction of primary visual cortex" *Consciousness & Cognition* 10(4): 574–587.

Tanner Jr., W. P. and Swets, J. A. 1954. "A decision-making theory of visual detection" *Psychological Review* 61(6): 401–409.

Tononi, G. and Edelman, G. M. 1998. "Consciousness and complexity" *Science* 282(5395): 1846–51.

Trevelan, C. T. Sahraie, A. and Weiskrantz, L. 2007. "Can blindsight be superior to 'sighted-sight'?" *Cognition* 103: 491–501.

Tsuchiya, N. and Koch, C. 2005. "Continuous flash suppression reduces negative afterimages" *Nature Neuroscience* 8: 1096–1101.

Turvey M. 1973. "On peripheral and central processes in vision: inferences from an information-processing analysis of masking with patterned stimuli" *Psychological Review* 80: 1–52.

Ungerleider, L., Ganz, L. and Pribram, K. H. 1977. "Size Constancy in Rhesus Monkeys: Effects of Pulvinar, Prestriate and Inferotemporal Lesions" *Experimental Brain Research* 27: 251–269.

Van Gaal, S., Ridderinkhof, K. R., van den Wildenberg, W. P. M., Lamme, V. A. F. 2009. "Dissociating consciousness from inhibitory control: Evidence for unconsciously triggered response inhibition in the stop-signal task" *Journal of Experimental Psychology: Human Perception and Performance* 35(4): 1129–1139.

Forthcoming in *Analytic Philosophy*. Author's Final Draft.

Van Selst, M. and Merikle, P. M. 1993. "Perception below the Objective Threshold?" *Consciousness and Cognition* 2(3): 194–203.

Ward, E. J. and Scholl, B. J. 2015. "Inattention blindness reflects limitations on perception, not memory: Evidence from repeated failures of awareness" *Psychonomic Bulletin & Review* 22(3): 694–700.

Weisberg, J. 2011. "Abusing the Notion of What-It's-Like-Ness: A Response to Block" *Analysis* 71(3): 438–443.

Weiskrantz, L. 1997. *Consciousness lost and found*. Oxford: Oxford University Press.

Weiskrantz, L. 2002. "Prime-Sight and Blindsight" *Consciousness & Cognition* 11(4): 568–581.

Weiskrantz, L. 2009 [1986]. *Blindsight: A case study spanning 35 years and new developments*. Oxford: OUP.

Weiskrantz, L., Barbur, J. L., Sahraie, A. 1995. "Parameters affecting conscious versus unconscious visual discrimination with damage to the visual cortex (V1)" *Proceedings of the National Academy of Sciences* 92: 6122–6.

Weiskrantz, L., Warrington, E. K., Sanders, M. D. and Marshall, J. 1974. "Visual Capacity in the Hemianopic Field Following a Restricted Occipital Ablation" *Brain* 97: 709–728.

Whitwell, R. L., Striemer, D. L., Nicolle, D. A. and Goodale, M. A. 2011. "Grasping the nonconscious: preserved grip scaling to unseen objects for immediate but not delayed grasping following a unilateral lesions to primary visual cortex" *Vision Research* 51(8): 908–24.

Wiggins, D. 2001. *Sameness and Substance Renewed*. Oxford: Oxford University Press.

Williamson, T. 2006. "Conceptual Truth" *Aristotelian Society Supplementary Volume* 80(1): 1–41.

Forthcoming in *Analytic Philosophy*. Author's Final Draft.

Wolfe, J. M. 1999. "Inattentional amnesia" In V. Coltheart (ed.) *Fleeting Memories: Cognition of Brief Visual Stimuli*. Cambridge, MA: MIT Press, pp. 71–94.

Yang, E., Zald, D. and Blake, R. 2007. "Fearful expressions gain preferential access to awareness during continuous flash suppression" *Emotion* 7: 882–886.

Zeki, S. and ffytche, D. H. 1998. "The Riddoch syndrome: insights into the neurobiology of conscious vision" *Brain* 121: 25–45.