

- Pearce, D. 2009. Reprogramming predators, last updated 2015. <<https://www.hedweb.com/abolitionistproject/reprogramming-predators.html>>, last accessed 3 February 2019.
- Pearce, D. 2015 [1995]. The hedonistic imperative. <<https://www.hedweb.com/>>, last accessed 3 February 2019.
- Tomasik, B. 2009. How many wild animals are there? last updated 2018. <<http://reducing-suffering.org/how-many-wild-animals-are-there/>>, last accessed 3 February 2019.
- Unger, P. 2000. The survival of the sentient. *Noûs* 34: 325–48.

## *Object files and unconscious perception: a reply to Quilty-Dunn*

IAN PHILLIPS 

### *1. Introduction*

A wealth of cases – most notably blindsight and priming under inattention or suppression – have convinced philosophers and scientists alike that perception occurs outside awareness. In recent work (Phillips 2016a, 2018, Phillips and Block 2017, Peters et al. 2017), I dispute this consensus, arguing that any putative case of unconscious perception faces a dilemma. The dilemma divides over how absence of awareness is established. If subjective reports are used, we face the problem of the criterion: the concern that such reports underestimate conscious experience (Eriksen 1960, Holender 1986, Peters and Lau 2015). If objective measures are used, we face the problem of attribution: the concern that the case does not involve genuine individual-level perception. Quilty-Dunn (2019) presents an apparently compelling example of unconscious perception due to Mitroff et al. (2005) which, he contends, evades this dilemma. The case is fascinating. However, as I here argue, it does not escape the dilemma’s clutches.

### *2. The case*

In Mitroff et al.’s experiment, subjects see two stationary circles at the upper corners of their screen.<sup>1</sup> In each, a single letter is displayed. The circles then move diagonally towards the screen’s centre, subjects sometimes taking them to ‘stream’ through each other, sometimes to ‘bounce’ off one another (Figure 1). The circles then stop, and a single letter appears in one.

1 For ease of exposition, I focus only on their first experiment. Readers are encouraged to view an example trial sequence here: <http://perception.yale.edu/Brian/demos/OF-BounceStream.html> (Demonstration #1) (last accessed 29 October 2018).




Figure 1. (a) Ambiguous display used in Mitroff et al. (2005, exp. 1); (b) representation of ‘bouncing’; (c) representation of ‘streaming’. Reprinted with permission. Copyright © 2004 Elsevier B.V. All rights reserved.

Subjects must judge as quickly as they can whether it matches either original letter. After that, they make a ‘streaming’ or ‘bouncing’ judgement.

By using smooth, fairly rapid motion, Mitroff et al. induced subjects to report ‘streaming’ on 95% of trials. Despite this, subjects’ responses were faster when the letter matched that shown in the circle directly above, as opposed to that in the opposite corner, where a streaming circle would have come from (compare Figure 1(b) and (c)). Carefully ruling out alternative interpretations, Mitroff et al. argue that this reflects an ‘object-specific preview benefit’: a priming effect due to the letter appearing in the ‘same’ object. This suggests that subjects’ ‘object-file’ systems represent the circles as bouncing not streaming. ‘Thus’, Mitroff et al. conclude, ‘object files and conscious perception can diverge’ (76).

Following Mitroff et al., Quilty-Dunn (2019) argues that this implicates unconscious perception. His reasoning is simple: object-file representations are perceptual, and their divergence from subjects’ conscious percepts demonstrate that they are unconscious. Quilty-Dunn further alleges that the dilemma noted at the outset is avoided. Here I argue that this is doubly mistaken. In the next section, I explain why the problem of the criterion does arise and, as a result, why we cannot be confident that the relevant object-file representations are unconscious. In §4, I provide two reasons why the relevant object-file representations do not clearly reveal individual-level perception. First, such representations belong in the first instance to working memory not perception. Second, the problem of attribution is harder to elude than Quilty-Dunn supposes.

### 3. *Are object-file representations unconscious?*

How can we be sure that subjects are not conscious of bouncing? Quilty-Dunn’s answer is that ‘the [streaming] experience [which subjects] describe is plainly at odds with their [bouncing] object files’. Quilty-Dunn further claims that the problem of the criterion does not arise. This is mistaken. Concerns about response criteria are not, as Quilty-Dunn suggests, limited to

'circumstances where subjects fail to indicate that they have seen a stimulus though they may have had dim conscious awareness of it'. They arise whenever a subject's performance in a task is the joint product of their underlying perceptual sensitivity and their response criterion. In any such case, performance may misrepresent awareness owing to response biases.

This is not merely a theoretical point. Consider the sound-induced visual bounce effect whereby playing a sound at the point at which the dots coincide causes a dramatic jump in the proportion of 'bounce' versus 'stream' judgements from roughly 25% to 75% (Sekuler et al. 1997). Grove et al. (2012) show that this effect is exclusively due to the 'strong influence of sound on criterion' (8), i.e. an increased tendency to say 'bounce' as opposed to 'stream'.<sup>2</sup> Evidently then, 'stream' judgements cannot simply be assumed to be a direct read-out of what subjects see.<sup>3</sup>

If stream judgements in Mitroff et al.'s study reflect response biases, they would potentially be consistent with subjects consciously seeing bouncing, and so with object files reflecting the contents of consciousness after all. Such a combination of perceived bouncing and judged streaming would occur if subjects had a strong bias towards 'streaming' responses in the relevant experimental conditions. Empirical assessment of this hypothesis is greatly complicated by the fact that Mitroff et al.'s display is ambiguous. This means that there is no straightforwardly correct answer to the 'stream/bounce' question, and so no straightforward way of applying the usual tools of detection theory which rely on classifying responses as hits and false alarms to prise apart sensitivity and bias (Green and Swets 1966, Macmillan and Creelman 2005). This is why Grove et al. turn to (what they take to be) non-ambiguous displays in their study. Nonetheless, Grove et al.'s findings show that an explanation of divergence between judgements and perception in terms of response bias must be taken seriously. As a result, the problem of the criterion remains very much with us.

The hypothesis just canvassed offers a way of developing Block's avowedly speculative suggestion that Mitroff et al.'s finding might be due to the 'conscious perceptual representation of a bounce [being] mis-conceptualized as a case of streaming' (forthcoming). However, whilst this is the starkest way to understand a divergence between what is reported and what is seen, subtler possibilities also need considering. Such possibilities arise because the phenomenological situation in Mitroff et al.'s experiment is significantly more complex than it may first appear.

2 See further Grove et al. 2016 for evidence of intra-modal 'cognitive' influences on 'stream/bounce' judgements.

3 Grove et al. make the orthodox assumption that criterion effects all alike reflect a post-perceptual 'decisional process' (2012: 9). This is controversial (see Morgan et al. 1990, Witt et al. 2015 and discussion in Phillips 2016b). Nonetheless, there is no doubt that some such effects are post-perceptual. This suffices for the present point.

Quilty-Dunn insists that subjects (predominantly) ‘had a clear, unambiguous percept of streaming’ (2019). In this he follows Mitroff et al. who note: ‘Conscious percepts were reported via a dichotomous keypress simply because all percepts fell into one of these categories. . . . each individual event is always clearly seen as one [streaming] or the other [bouncing]’ (75). However, even if subjects find it easy, under instruction, to classify trials, it does not follow that their percepts are either binary or unidimensional. Nor does it follow from the fact that it requires ‘other cues, such as disambiguating featural information’ (or sounds) for subjects to *report* substantial numbers of ‘bouncing’ trials, that *experiencing* bouncing requires such cues, as Quilty-Dunn contends. There are plausibly many aspects to subjects’ experiences in Mitroff et al.’s experiment. In consequence, response bias need not be construed as a tendency to mis-conceptualize experience. It might equally reflect a tendency to report one aspect of experience over another. Here are three ways in which this might occur.

- (a) It is quite possible to see something as *F* despite it *looking G*, even when being *F* is incompatible with being *G*. Consider the tilted penny which is typically seen as circular (due to constancy mechanisms) even whilst naturally being said to look elliptical.<sup>4</sup> Martin (2010) treats such looks reports as comparative, reporting the existence of a contextually salient similarity between the look of an *F* object (for Martin a construction from its basic visible qualities) and the paradigm look of a *G* object. Mitroff et al.’s dots do have a look which is saliently similar to that of streaming dots: their pattern of movement violates the solidity constraint that ‘no parts of two distinct objects coincide in space and time’ (Spelke et al. 1992: 606, quoted in Mitroff et al. 2005: 89). Combined with their rapid, smooth motion, the dots might thus *look* like they are streaming even whilst being seen as bouncing. These factors may likewise explain the hypothesized bias towards a streaming response.<sup>5</sup>
- (b) Mitroff et al.’s display is a depiction. As a result, subjects might see the dots as bouncing whilst simultaneously seeing them as *depicting* streaming (or vice-versa). Another possibility is that the dots themselves are neither seen to bounce nor stream (being just images on a screen) but are nonetheless seen to depict both.

4 Quite how to characterize the case is controversial. See Schroer 2017 for a recent discussion. Schroer’s starting point is the ‘popular claim’ that the case involves ‘an experiential doubleness of shape’ (211). Such ‘doubleness’ – however elucidated – is all that is required to make the present point.

5 One could also suppose that the dots are seen as streaming but look like they are bouncing. Either way, the bouncing object-files would align with an aspect of phenomenology.

- (c) Subjects may see *both* bouncing and streaming. This need not necessarily involve incoherent contents. Numerous possibilities arise. For instance, bouncing and streaming would be coherent if a second dot, initially hidden by the first, was seen to diverge post-coincidence, or if one of the dots was seen to undergo ‘fission’ at the point of coincidence.<sup>6</sup> If subjects see both bouncing and streaming, the case for unconscious perception plainly collapses, since object-file representations will then partially reflect a complex conscious percept. Again, the fact that subjects reliably classify trials as ‘streaming’ given a dichotomous question, does not entail that their percepts are themselves binary. Subjects may simply have a strong bias to report one aspect of the display over the other.

These points all bring out ways in which subjects’ responses may fail fully to capture their conscious experience. They thereby stress the continuing presence of the problem of the criterion. They also all bear on Quilty-Dunn’s central methodological proposal:

If (i) a subject visually experiences *F* and (ii) also visually represents *G* and (iii) *F* and *G* are clearly incompatible, then, *ceteris paribus*, there is good reason to think that the representation of *G* is unconscious.

Even if we agree with the proposal, its application requires us to be sure that two incompatible features are both represented visually as opposed to simply in judgement. Furthermore, its application requires us to be sure that the presence of both features in experience would involve incompatibility or incoherence. Yet, as the titled penny shows, there is no incompatibility in seeing an object as *F* but yet it looking incompatible *G*. Likewise, as two-dimensional paintings of three-dimensional scenes show, there is no incompatibility in seeing an object as *F* and yet it depicting incompatible *G*. Finally, we need to be open to ways in which seeing both bouncing and streaming would be perfectly coherent. For all these reasons, there is substantial room for doubt that Mitroff et al.’s study involves object-file representations disassociated from conscious experience.

#### 4. *Are object-file representations perceptual?*

Suppose we are convinced, despite the various concerns of the last section, that the object-file representations evidenced by Mitroff et al.’s study do not correspond to conscious experience. Is this tantamount to embracing unconscious perception? There are two reasons to think not.

First, as Quilty-Dunn acknowledges, the behavioural effects in Mitroff et al.’s study (as with object-specific preview benefits in general) are ‘achieved

6 Quilty-Dunn holds that incoherence would inevitably involve ‘jarring, even disorienting’ phenomenology and so would be implausible in the present context. Someone unconvinced will want to take seriously the more obvious possibility of seeing the dots as both (perhaps partially) bouncing and streaming.

by object files stored in working memory'. Yet working memory representations cannot be assumed to be identical to perceptual representations. Certainly, the process of storing and maintaining representations in working memory is partly independent of perception: object representations in working memory can last up to eight seconds (Noles et al. 2005) and do not seem to compete for the same resources as object representations in perception (Zhang et al. 2010). Nonetheless, one might think that working memory representations are simply stored perceptual representations and thus provide a clear window onto perception.

In recent work, Block (forthcoming) rejects this assumption, arguing for a sharp distinction between object representations in mid-level vision, and object representations in working memory. For Block, working memory differs in capacity (having low as opposed to high capacity limits), content (being conceptual as opposed to non-conceptual) and format (being discursive as opposed to iconic). If Block is right, it is far from obvious what object-specific preview benefits tell us about perceptual representation. Block's view is controversial (see Quilty-Dunn forthcoming: §5 for a trenchant critique). Fortunately, for present purposes, a much more minimal idea suffices, namely that working memory representations can in some cases diverge from perceptual representations. This is quite consistent with it generally being true that working memory representations are simply stored perceptual representations.

This minimal idea of occasional divergence between perception and working memory should not be controversial given Mitroff et al.'s results (setting aside the points of §3). On Quilty-Dunn's interpretation, subjects' initial perception of bouncing comes apart from their subsequently reported conscious experience of streaming (cf. discussion of 'downstream overriding' in Mitroff et al. 2005: §7). Since their experiential reports are presumably also grounded in working memory, this interpretation requires a divergence between perceptual representations and certain working memory representations. Once this is seen, however, we must consider an alternative proposal on which the divergence occurs, not between initial perception and reported experience, but between initial perception and preview benefit-supporting object-file representations in working memory. This hypothesis denies any conflict between subjects' reports and their initial perception. Its availability means that Quilty-Dunn has at most made the case for unconscious working memory not unconscious perception.

Second, and perhaps most importantly, we must consider whether the object-file representations implicated in Mitroff et al.'s study are individual-level. This is the problem of attribution. The reason for positing such representations is the object-specific preview benefits which they provide. However, in line with the discussion of priming in my 2018 (and in Phillips and Block 2017 and Peters et al. 2017), the mere fact that a representation affects (e.g. speeds) the way that a subject responds is not a

sufficient ground for attributing that representation to the individual as opposed to their brain or some psychological sub-system. The key (at least evidential) criterion for such individual attribution is, I propose, availability of the representation for guidance of intentional action. Thus, the most powerful consideration which Quilty-Dunn (2019) gives in favour of such representations being individual-level is that they are routinely exploited in intentionally guided saccades.<sup>7</sup> The problem, however, is that Mitroff et al.'s study is not routine. Rather (again setting aside the concerns of §3) it is unique in teasing apart object-file representations in working memory from the representations underlying consciousness and explicit report when normally the two march in step. Given these special circumstances, it cannot be assumed that the specific object-file representations evinced by Mitroff et al. remain available for the guidance of intentional action (and so count as individual-level), disconnected as they are from consciousness and report. Indeed, the very fact that they are not reportable shows that in one respect they are not so available. Direct evidence of availability is thus needed. In its absence, the object-file representations in question continue to confront the problem of attribution.

Quilty-Dunn (2019) insists that an object file is a 'personal-level type of mental representation' and that 'object files are a paradigm case of personal-level perceptual representation'. Understood generically I am happy to concede these claims. However, it is obscure why we should accept the much stronger theses that *all* object files are personal-level. Distinguishing between generic and universal claims reveals the invalidity of Quilty-Dunn's argument, which he summarizes thus: 'Object files are genuine personal-level perceptual representations. They can also occur unconsciously. Therefore, there is unconscious perception.' Compare: Birds fly. Some birds are flightless. Therefore, there are flightless birds which fly. The object-file representations in Mitroff et al.'s study may be the cognitive equivalents of flightless birds.

In summary: Quilty-Dunn highlights a fascinating putative case of unconscious perception. Ultimately, however, it fails to avoid the problems of criterion and attribution which beset more familiar cases. The discussion underscores just how deep the dilemma facing any alleged example of unconscious perception runs.<sup>8</sup>

7 Quilty-Dunn also points to the fact that such representations are the outputs of constancy mechanisms and are available to working memory. However, I deny that constancies are sufficient for individual-level perception (even if they are a requirement on genuinely perceptual representation as proposed by Burge 2010). I also see no reason to assume that all working memory representations are individual-level including those which fail to meet the agency criterion discussed in the main text.

8 I am grateful to an anonymous referee for very helpful comments, to Ned Block and Sam Clarke for discussion, and especially to Jake Quilty-Dunn for continuing correspondence about these issues.

*Department of Philosophy and Department of  
Psychological and Brain Sciences  
Johns Hopkins University  
Baltimore, MD 21218, USA  
ianbphillips@jhu.edu*

### References

- Block, N. forthcoming. *The Border between Seeing and Thinking*. Cambridge, MA: MIT Press.
- Burge, T. 2010. *Origins of Objectivity*. Oxford: Oxford University Press.
- Eriksen, C.W. 1960. Discrimination and learning without awareness: a methodological survey and evaluation. *Psychological Review* 67: 279–300.
- Green, D.M. and J.A. Swets. 1966. *Signal Detection Theory and Psychophysics*. New York: Wiley.
- Grove, P.M., J. Ashton, Y. Kawachi and K. Sakurai. 2012. Auditory transients do not affect visual sensitivity in discriminating between objective streaming and bouncing events. *Journal of Vision* 12: 1–11.
- Grove, P.M., C. Robertson and L.R. Harris. 2016. Disambiguating the stream/bounce illusion with inference. *Multisensory Research* 29: 453–64.
- 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.
- Macmillan, N.A. and C.D. Creelman. 2005 [1st edition 1991]. *Detection Theory: A User's Guide*. Cambridge: Cambridge University Press.
- Martin, M.G.F. 2010. What's in a look? In *Perceiving the World*, ed. B. Nanay, 160–225. Oxford: Oxford University Press.
- Mitroff, S.R., B.J. Scholl and K. Wynn. 2005. The relationship between object files and conscious perception. *Cognition* 96: 67–92.
- Morgan, M.J., G.J. Hole and A. Glennerster. 1990. Biases and sensitivities in geometric illusions. *Vision Research* 30: 1793–810.
- Noles, N.S., B.J. Scholl and S.R. Mitroff. 2005. The persistence of object file representations. *Perception & Psychophysics* 67: 324–34.
- Peters, M.A.K., R.W. Kentridge, I. Phillips and N. Block. 2017. Does unconscious perception really exist? Continuing the ASSC20 debate. *Neuroscience of Consciousness* 3: 1–11.
- Peters, M.A.K. and H. Lau. 2015. Human observers have optimal introspective access to perceptual processes even for visually masked stimuli. *eLife* 4: 1–30.
- Phillips, I. 2016a. Consciousness and criterion: on Block's case for unconscious seeing. *Philosophy and Phenomenological Research* 93: 419–51.
- Phillips, I. 2016b. Naïve realism and the science of (some) illusions. *Philosophical Topics* 44: 353–80.
- Phillips, I. 2018. Unconscious perception reconsidered. *Analytic Philosophy* 59: 471–514.



- Phillips, I. and N. Block. 2017. Debate on unconscious perception. In *Current Controversies in Philosophy of Perception*, ed. B. Nanay, 165–92. New York, NY: Routledge.
- Quilty-Dunn, J. 2019. Unconscious perception and phenomenal coherence. *Analysis* 79: 461–9.
- Quilty-Dunn, J. forthcoming. Perceptual pluralism. *Nous*. doi:10.1111/nous.12285.
- Sekuler, R., A.B. Sekuler and R. Lau. 1997. Sound alters visual motion perception. *Nature* 385: 308.
- Schroer, R. 2017. Hume's table, Peacocke's trees, the tilted penny and the reversed seeing-in account. *Mind & Language* 32: 209–30.
- Spelke, E.S., K. Breinlinger, J. Macomber and K. Jacobson. 1992. Origins of knowledge. *Psychological Review* 99: 605–32.
- Witt, J.K., J.E.T. Taylor, M. Sugovic and J.T. Wixted. 2015. Signal detection measures cannot distinguish perceptual biases from response biases. *Perception* 44: 289–300.
- Zhang, H., Y. Xuan, Z. Fu and Z.W. Pylyshyn. 2010. Do objects in working memory compete with objects in perception? *Visual Cognition* 18: 617–40.

## *Inanimation: A network of feeling and perception*

MATTEO RAVASIO

Consider the Elton John song 'I guess that's why they call it the blues'. The narrator describes the experience of longing for the presence of one's partner when away from home. The memory and anticipation of happy moments with his loved one are contrasted with the emptiness and boredom that characterize the time away from the object of his desire. The narrator refers to these negative feelings as 'the blues'.

Why do people call it 'the blues'? The etymology of the expression is debated, but it seems plausible to assume that the connection between the colour blue and a state of sadness or melancholy is part of the reason why this expression emerged and survived: the colour blue is expressive of melancholy, sadness and more generally of a subdued psychological state.

There are two questions other than the etymological one that need an answer.

First, how can we meaningfully ascribe expressive properties to a sensation such as that of colour, and how should its connection to the experience of a mental state be fleshed out? This is the familiar problem of expressiveness. Second, and relatedly, what is it in the experience of a mental state that invites its description in terms of a sensory property – in the above case, blue?

This paper seeks to elucidate the latter issue, although a connection to the former will emerge. We often use terms primarily concerned with the description of inanimate objects in order to characterize psychological states or