

Fuzziness in the mind

Taylor, John

DOI:

[10.1111/phpr.12592](https://doi.org/10.1111/phpr.12592)

License:

None: All rights reserved

Document Version

Peer reviewed version

Citation for published version (Harvard):

Taylor, J 2019, 'Fuzziness in the mind: can perception be unconscious?', *Philosophy and Phenomenological Research*. <https://doi.org/10.1111/phpr.12592>

[Link to publication on Research at Birmingham portal](#)

Publisher Rights Statement:

Checked for eligibility: 26/03/2019

This is the peer reviewed version of the following article: Taylor, H. (2019), Fuzziness in the Mind: Can Perception be Unconscious?. *Philosophy and Phenomenological Research*. doi:10.1111/phpr.12592, which has been published in final form at <https://doi.org/10.1111/phpr.12592>. This article may be used for non-commercial purposes in accordance with Wiley Terms and Conditions for Use of Self-Archived Versions

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

FUZZINESS IN THE MIND: CAN PERCEPTION BE UNCONSCIOUS?

Henry Taylor
University of Birmingham

Forthcoming in *Philosophy and Phenomenological Research*
Please cite published version.

This is the peer reviewed version of the following article: Fuzziness in the mind: can perception be unconscious? which has been published in final form at: [DOI unavailable at time of upload]. This article may be used for non-commercial purposes in accordance with Wiley Terms and Conditions for Use of Self-Archived Versions.

Abstract

Recently, a new movement has arisen in the philosophy of perception: one that views perception as a natural kind. Strangely, this movement has neglected the extensive work in philosophy of science on natural kinds. The present paper remedies this. I start by isolating a widespread and influential assumption, which is that we can give necessary and sufficient conditions for perception. I show that this assumption is radically at odds with current philosophy of science work on natural kinds. I then develop an alternative, new view of perception. This new view takes as its starting point the dominant position on kinds in the life sciences: that kinds are homeostatic property clusters. I show that, if you accept this view, then all of the putative cases of unconscious perception are more plausibly seen as cases where it is *indeterminate* whether the mental episode in question is an instance of *perception*.

1. Perception as a psychological natural kind

We are witnessing the rise of a new way of thinking about perception. The new way conceives of perception as a natural kind. It hopes to find a theory of perception that *carves nature at the joints*. The new way has much to recommend it. If we have an account of perception that carves nature at the joints, then we have a principled reason to prefer this view over others. Surprisingly, this movement has proceeded largely without input from the philosophy of science work on natural kinds. This needs fixing. I identify an assumption at the heart of the new way, which is commitment to an essentialist view of perception. I show that this is radically at odds with current philosophy of science work on biological and psychological kinds (§2). I then introduce the dominant approach to natural kinds in the life sciences, which is the homeostatic property cluster account, and explore the consequences of applying this account to perception (§§3-4). This generates a radically new view of perception, which is currently entirely unexplored in philosophy. I argue that this new view has important consequences for the debate over unconscious perception. Specifically, if you accept this view, then all of the putative cases of unconscious perception are more plausibly seen as cases where it is indeterminate whether the mental episode in question is an instance of perception (§5).

The purpose of this paper is exploratory. It is not primarily concerned with resolving issues in the philosophy of science about the nature of kinds. Rather, the aim is to *apply* the dominant account of kinds from the philosophy of science to the new movement in the philosophy of perception.

2. Essentialism

The prime proponent of the new way of thinking about perception is Tyler Burge (2005, 2010). He writes that perception consists in ‘*objective sensory representation by the individual*’ (2010, p.368). On this view, there are three features that a mental episode must have to be perceptual. First, objectivity. To Burge, representing something objectively requires representing ‘environmental attributes, or environmental particulars, as the same, despite radically different proximal stimulations’ (2010, p.114). In this way the notion of objectivity is closely associated with constancy. Second, it must be sensory: it must use information gathered from the sense organs. Third, perception is an *individual*-level phenomenon, rather than something at the sub-personal level.

For decades, it was received wisdom that perception can be unconscious. Recently, this has been questioned, and Burge’s criteria for perception have taken centre stage in this debate. Thus, the debate over unconscious perception centres on whether there is a mental episode which:

- (i) is unconscious
- (ii) is sensory
- (iii) involves constancy
- (iv) is individual-level

The issue is *unconscious* perception, so clearly (i) needs to be fulfilled. But (i) is not part of what makes something perceptual.¹ Furthermore, everyone agrees that some sensory episodes can be unconscious (transduction from the retina to the optic nerve presumably counts as sensory). The debate thus centres on criteria (iii-iv), which give us the two properties that are required for a sensory episode to be an instance of perception.

The interlocutors in the debate accept Burge’s criteria for perception, at least for the sake of argument. For example, Phillips considers a range of experimental paradigms, and argues that in all cases, we have

¹ At one point, Phillips suggests that it may be (forthcoming, pp.8-11). He does not pursue this, and none of his arguments rely on it, so I set it aside here.

reason to doubt that the mental episodes in question fulfil all of (i-iv) (2016, p.420; forthcoming; Block and Phillips 2016). Given that no episodes have all the properties, we lack good reason to believe that perception occurs unconsciously. So Phillips claims. Those who defend unconscious perception reply by claiming that at least some sensory episodes fulfil all the criteria (Block and Phillips 2016, pp.182-183; Quilty-Dunn forthcoming).

Some background assumptions about kindhood are implicit in this debate. The first is that properties (ii-iv) are *necessary* for an episode to be perceptual. To see this assumption, consider that the opponents of unconscious perception argue that no unconscious episodes instantiate (ii-iv). The friends of unconscious perception argue that some unconscious episodes fulfil all of them. These arguments would not hit their intended targets unless we assume that *all* of the properties are required to have an instance of perception. We can also see the assumption that they are *sufficient* for a mental episode to be perceptual. Otherwise, giving an example of an unconscious episode that fulfils them would not licence us to conclude that we have an instance of unconscious perception, as we find in, e.g. Quilty-Dunn (forthcoming). In other words, it is implicitly assumed that (ii-iv) are *necessary and sufficient* for perception.²

This, in turn, embodies a more overarching assumption, which I call ‘perceptual kind essentialism’ (PKE):

PKE: There exists a set of properties that are necessary and sufficient for a mental episode to be a member of the kind ‘perception’.

Whether Burge himself commits to (PKE) is less clear. He says:

To be an instance of a kind... something must meet certain collateral *constitutive conditions*. These are conditions that are necessary, sufficient, or necessary and sufficient to be something of that kind (2010, p.58).³

Burge suggests that the constitutive conditions on kindhood need not be necessary and sufficient, but could be necessary *or* sufficient. However, he does not elaborate on these claims. In any case, Burge’s

² I am not saying that the interlocutors explicitly agree with this. Rather, it is implicit in the way that the debate proceeds. It may be that they only accept it for the sake of argument. Phillips, for example, mentions that these criteria can be questioned (Phillips 2016, p.420; Block and Phillips 2016, pp.185-186). Block also accepts the Burgean criteria for the sake of argument but elsewhere argues that perception has properties such as iconicity and adaptation effects (2014). However, it’s unclear whether he thinks these are *definitive* of perception as a kind.

³ Thanks to an anonymous referee for drawing my attention to this quotation.

conditions have been interpreted by others as necessary and sufficient for a mental episode to be an instance of perception. I take this idea as my target.

Let essentialism be the view that natural kinds are defined by sets of properties, which are necessary and sufficient for membership in the kind. (PKE) is an instance of essentialism. The overwhelming consensus in the philosophy of science is that essentialism about biological and psychological kinds is fundamentally wrong. In philosophy of biology, essentialism is almost universally seen as a damaging relic of pre-Darwinian thinking (Hull 1965, Sober 1980, Dupré 1981).⁴ Things are even clearer in the philosophy of psychology. I do not know of a single thinker who explicitly defends essentialism about psychological kinds (Griffiths 1997, Machery 2009). Dupré sums up the general view when he says that essentialism diverges from ‘some actual biological facts and theories’ (1981, p.66). Very rarely in philosophy do we see such overwhelming agreement that a particular view is wrong.

I do not have space for a full analysis of essentialism’s failure. However, I will summarise some of the main problems with the view, and show how they apply to perception. Then, I outline the homeostatic property cluster (HPC) view, which is now the dominant position in the philosophy of the life sciences. I also explain how the HPC view resolves the problems that plague essentialism. The purpose of this paper is not to defend the HPC view exhaustively, so I restrict myself to summarising the core issues. The primary aim of the paper is to *apply* this dominant view to perception, and draw out the consequences for whether perception can occur unconsciously.

One prominent problem for essentialism involves vagueness in biological and psychological kinds. It is widely held that evolution does not produce kinds with clear boundaries that are neatly defined by necessary and sufficient conditions. Rather, as a new kind emerges throughout evolutionary history, there will always be a long period during which some properties are lost, and others emerge. These new properties may themselves be replaced with others on the road to establishment of the new kind. This long period of gradual evolutionary refinement is all part of how a biological or psychological kind emerges (Hull 1965, Sober 1980). Given this, there is no one set of properties, such that lacking one of them determinately puts one outside of a particular kind.⁵ We would expect something similar to arise with perception, since it is an

⁴ I say *almost* universally (Devitt 2008).

⁵ The essentialist can accept some vagueness. Take the properties that the essentialist says are necessary and sufficient for membership in a kind. Presumably, it can be vague whether some entity instantiates those properties (Devitt 2008).

evolved faculty. In the context of a philosophical programme that aims to understand the gradual emergence of perception from its evolutionary predecessors (Burge 2010), we would expect a period when some properties definitive of perception were present and others were not. This period inevitably brings with it vagueness in the extension of the kind, which sits poorly with the essentialist picture.

There are further difficulties for essentialism. The view doesn't have the resources to distinguish genuine natural kinds from groups of entities that share properties by coincidence (Griffiths 1997). There is also nothing in the essentialist picture that *explains* why certain entities share the properties in question. This also applies to perception. We do not only want to know which properties are present in perceptual episodes. It would be preferable if our view can help us understand *why* perceptual episodes instantiate these properties.

In the wake of the rejection of essentialism, by far the dominant account of biological and psychological kinds is the HPC view, from Richard Boyd (1989, 1991; Wilson *et.al.* 2007). Samuels and Ferreira say 'philosophers of science have, in recent years, reached a consensus-or as close to consensus as philosophers ever get-according to which Natural Kinds are *Homeostatic Property Clusters*' (2010, p.222). The account has been described as the 'received' view (Ereshefsky and Reydon 2015, p.969). The HPC view has even filtered into biology itself (Assis and Brigandt 2009). Essentialism is now almost dead as a view of biological and psychological kinds, and the HPC view is the dominant approach (Machery 2009, Griffiths 1997, Taylor 2018a).

On the HPC view, kinds are (partially) defined by a cluster of properties, but instantiating this cluster is not sufficient for kind membership. The properties must also be supported by common homeostatic mechanisms. The claim about mechanisms can cause confusion. It is not that all the properties in the cluster must be caused by one individual mechanism. Rather, there can be multiple mechanisms that work in concert to produce the clustering of properties. It can even be that some of the properties in the cluster are causally sustained by other properties in the cluster, rather than by an underlying mechanism (Boyd 1989,

I do not have space to fully address this idea here, but note that it still commits the essentialist to the claim that there is some set of properties, such that if an entity determinately lacks any one of those properties, it is determinately not a member of the kind. This is generally thought to be implausible for biological and psychological kinds, given the gradual nature of evolution.

p.16; Slater 2015). The point is just that different instances of the kind should tend to share properties *because they share underlying causal factors*.

On the HPC view, if some entities share properties by coincidence, then they're not a natural kind. Entities must share properties because of similar underlying mechanisms. The mechanisms also explain *why* these properties tend to be found together. It is because the same underlying causal factors will produce the same properties in a range of cases (Boyd 1989). With perception, the underlying mechanisms will include neural factors, DNA, transcription and translation mechanisms that underpin the expression of DNA into phenotypic traits, and so on. When we talk about 'the' common mechanism that underpins perception, we are talking about the systematic interaction of all these factors, since this is what ensures the clustering together of the properties.⁶

In these ways, the HPC view allows us to distinguish natural kinds from cases where entities coincidentally share properties, and it can invoke homeostatic mechanisms to explain why the properties occur together. Recall that it is also desirable if our account of natural kinds can accommodate vagueness in biological and psychological kinds. In relation to this, Boyd says:

There will be many cases of extensional "vagueness"... There will be things that display some but not all of the properties [in the cluster] (and/or in which some but not all of the relevant homeostatic mechanisms operate) such that no rational considerations dictate whether or not they are to be classed [as members of the kind] (1989, p.17)

We have a spectrum. At one end are individuals that are clearly members of a kind. We then pass through a fuzzy area when certain properties aren't present, or where the homeostatic mechanisms are only partially operative. At the other end of the spectrum are cases that are clearly not members of the kind. Note that the fuzziness can occur across evolutionary time, where the relevant properties and mechanisms have not yet evolved fully. It can also occur later. Even when the evolutionary process has produced the kind in question, it might be that in certain cases, the mechanism partially breaks down, or some of the properties aren't present.

⁶ Proponents of the argument from multiple realizability may claim that the neural mechanisms underpinning perception are not shared across species. However, even proponents of this argument agree that such mechanisms are unified at the *functional* level, and it is at this level that neural mechanisms are individuated (Piccinini and Craver 2011). If it does transpire that different instances of perception are supported by radically different mechanisms, which display no unity at all (even at the functional level) then the claim that perception is a natural kind must be revised. But we can relax, because the empirical evidence for this pessimistic conclusion is weak (Bechtel and Mundale 1999).

We have seen how the HPC view can resolve certain issues that render essentialism implausible. The virtues of the HPC position do not end here. It also explains the success of scientific inferences. Suppose we find that certain members of a kind have two particular properties, which are caused by common underlying mechanisms. That gives us good (not infallible) reason to infer that many members of the kind will also have those properties. This is because the same mechanisms will produce similar properties in a range of cases. We routinely see this inference pattern in perceptual psychology. When a feature of perception is discovered in a group of 30 subjects, we think we have good reason to infer that the feature is widespread throughout the population. The HPC view explains what underpins the success of such inferences.

Of course, not everyone agrees with the HPC position. Here is not the place for a full defence of the view. However, I will note that many of the objections to it do not apply to perception, so we need not worry about them here. For example, Ereshefsky and Reydon (2015) argue that it doesn't apply to microbiological kinds, but perception is not one of those. Ereshefsky and Matthen (2005) claim that it cannot account for stable polymorphisms in biological species, but there is no reason to think that similar problems arise for psychological kinds.

This is only a sketch of the virtues the HPC view. In the rest of this paper, I will apply the view to perception. The main conclusion of this paper is conditional: *if* you reject essentialism in favour of this alternative picture, then a new view of perception emerges. On this view, all of the purported instances of unconscious perception are more plausibly seen as cases where it is indeterminate whether the mental episode is perceptual. This draws an important link between the dominant view of kindhood in the philosophy of science and an increasingly prominent and exciting tradition in the philosophy of perception. It also has dramatic upshots for debates over the possibility of unconscious perception. For these reasons, the conclusion is a very substantial one.

3. Perception as a homeostatic property cluster kind

3.1 Which properties?

On the view we are considering, perception is defined by a cluster of properties, supported by common mechanisms. Not any old cluster of properties will do though, and we don't want our choice of which properties to include in the cluster to be arbitrary. The view should not be anything goes. To implement this restriction, proponents of the HPC view claim that the properties in the cluster must be

‘methodologically important’ (Boyd 1991, p.141) or ‘scientifically relevant’ (Machery 2009, p.232). What does this amount to? We can make progress on this question by returning to Burge, who emphasises explanation: ‘Explanation guides us to what kinds there are’ (2010, p.315, cf. 2010 pp.57-60). Burge himself does not have the HPC view in mind of course, but his arguments for singling out particular properties as definitive of perception can be integrated into an HPC view of perception as a kind, to guide us in how to populate the cluster with properties.

For Burge, the primary explanatory task for perceptual psychology is to provide an explanation of how proximal input (such as retinal stimulation) is converted into fully representational states (2010, p.344). This is based on the claim that there is a fundamental distinction between creatures like bacteria and paramecia (the sensory systems of which can be exhaustively described using biological concepts) and metazoans such as animals, for which we must invoke a distinctively psychological notion of representation (Burge 2010, pp.339-342). Representational states, in this sense, are those that are characterised using notions such as ‘veridicality conditions’, rather than biological notions like (mal)function. One way that perceptual psychology performs this explanatory task is to ascribe *constancy* to perceptual systems. That is a crucial theoretical posit in explaining how we move from the relatively poor resources afforded us by the information from our sense organs to full-blown representational states (2010, pp.379-388). That’s why we should include constancy in the properties definitive of perception. It is easy to construct a similar argument for inclusion of the sensory property. A core resource in explaining how we get from proximal stimulation of sense organs to objective representation is recognising that the input to the perceptual system is delivered by the sense organs.

The justification for the ‘individual-level’ property is different, and stems from *central agency*. This is the capacity that underpins organism-level actions such as mating, navigating and eating (2010, p.370). Burge says that perception is a core resource in explaining the organism’s individual-level actions (2010, p.371). He then says that ‘If perception is to ground this explanatory role, it must be attributable to individuals’ (2010, p.371).⁷ It does not follow that *all* individual-level perceptions are available to central agency.

⁷ Burge also claims that the individual-level nature of perception is a priori (2010, p.369). For those who accept the HPC view, this isn’t a good reason for us to take the individual-level property as definitive of perception. On the HPC view, such properties must be determined a posteriori by empirical investigation, not by a priori reflection (Boyd 1989, p.16).

However, given that being individual-level is ascribed to perceptions precisely because this is essential for explaining how perceptual episodes interface with individual-level central agency, it is generally accepted that the best evidence we could have for a perceptual episode's being individual-level is that it is available to central agency systems (Burge 2010, p.333; Block and Phillips 2016, p.178; p.181). Interestingly, Burge accepts that agency by individuals is explicable by biology (2010, p.370).

This neatly addresses the 'anything goes' worry. Not just any old property gets into the cluster. Rather, it must be justified by its importance within psychology, and one of the best ways of doing this is by showing that attributing it to perception is a core part of psychology's explanatory structure. This picks up on a strong tradition that takes scientific success to be the best indicator of kindhood (Magnus 2012, Taylor 2018a). This isn't to say that explanatory success is the only thing that can qualify a property for inclusion in the cluster. It could be that usefulness for other scientific goals (such as prediction and intervention) may also qualify. Note also that not all the properties in the cluster are *unique* to perception. We aren't saying that perception is the only mental faculty that is individual-level, or that it is the only faculty that consumes sensory information. It is the cluster and the mechanisms as a whole that define the kind. It isn't the job of each individual property to do that.

3.2 *Expanding the Cluster*

Further developing the view of perception as a HPC kind, I shall now argue that we have good reason to extend the list by at least one further property, involving attention. I do not claim that this is the only property that should be added to the cluster, just that we should add *at least* this additional property.⁸

Recall that constancy is a crucial resource in explaining how we get from proximal stimulation of our sense organs to representation of the outside world. The faculty of attention does something similar. The influence of attention on a mental episode is invoked to explain how a perceptual state is constructed. It is used to explain which properties are attributed to a distal object in perception (Carrasco et al. 2004, Tse 2005). It also explains the selection and modulation of sensory information, which is of course crucial for explaining how we get from proximal input to representational states (Chun et al. 2011, pp.77-80). Attention is also used in explaining how we perceive properties as bound to objects (Treisman 1996). Furthermore,

⁸ Note that one could be an essentialist and agree that the cluster should be bigger. This section is not meant to be unique to those who hold the HPC view. Rather, it is a general point that contributes to the overall picture of perception that I am developing.

recall that the property of being individual-level is justified based on the claim that it is crucial for explaining the influence of perception on central agency. Clearly attention is also important for this explanatory role. Attention serves to distinguish relevant from irrelevant information, which is hugely important in explaining how perception links up with the organism's action (Wu 2014). Attention clearly and obviously nestles at the core of a huge number of explanations in perceptual psychology. If anything counts as explanatorily important in the relevant ways, the availability of a sensory episode to attention does. So availability to attention should qualify for inclusion in the cluster of properties that is definitive of perception.⁹

Someone may object that availability to attention is 'only' a *dispositional* property of a mental episode. This is no problem. Science is full of reference to dispositional properties (Bird 2007). Thinking of them as less real is just metaphysical prejudice, not philosophy of science. Another potential objection is that primitive organisms do not have attentional systems, and yet they can still perceive things, so it cannot be one of the properties in the cluster definitive of perception. But this is false. We find attention across biological taxa, including in bees (Moranetz and Spaethe 2012). Another objection could be that we can imagine creatures with perceptual capabilities, but which do not have any attentional mechanisms. But for those attracted to the HPC view, what we can imagine is irrelevant. The properties in the cluster are decided a posteriori by what scientifically relevant properties cluster together as a result of homeostatic mechanisms. They are not decided by what we can imagine (Boyd 1989, p.16).

I must reiterate that I do *not* claim that availability to attention is the only property that should be added to the cluster. Determining precisely which properties fall into the cluster is a long, difficult theoretical process, and it would be silly to attempt it in full here. I claim only that the list of properties should be extended to include *at least* the attentional one.

4. Fuzziness

As noted above, if you embrace the HPC view of perception that I am exploring, there will be many cases of extensional vagueness.¹⁰ For the mental episodes that fall within this penumbral zone, it will be

⁹ It is controversial whether attention can be entirely explicated in neurobiological terms (Fazekas and Nanay forthcoming). But whether it is or not, it forms a core part of perceptual psychology's explanatory framework, which is all that is needed.

¹⁰ Note that it is not the case that if a mental episode instantiates only *one* property in the cluster, then it becomes indeterminate whether it is perceptual. Since some of the properties are not unique to perception, there may be cases where mental episodes have one of the properties, but are clearly not perception (where the fuzzy zone begins and ends may also be indeterminate).

indeterminate whether they are perceptual or not. Within the penumbral zone there will be cases that are more or less close to paradigmatic instances of perception, just as there is a spectrum between being very hirsute and being totally bald. Does all this fuzziness threaten the natural kind status of perception? No. As I emphasised above, many (perhaps all) kinds in the life sciences are fuzzy.

Before we turn to concrete examples of fuzziness, I will highlight another way that fuzziness can arise on this HPC picture of perception. Many of the properties in the cluster constitutively rely on the presence of other mental faculties. For example, for a perceptual episode to be available to attention, attention itself must be operative. But if such faculties are themselves HPC kinds, they too would be associated with a cluster of properties sustained by a common mechanism. For example, ‘availability to attention’ is a property composed of the ability to *trigger* the attraction of attentional systems, as well as to be *modulated* by those attentional systems. Attentional systems themselves are composed of a cluster of faculties such as spatial and object attention, orienting, executive and alerting attention (Carrasco 2011, Taylor 2018b). If some but not all of these faculties are present, then it will be indeterminate whether the faculty of attention is present, in which case it will be indeterminate whether a particular mental episode instantiates the property of availability to attention. Since availability to attention is one property definitive of a perceptual episode (I have argued) it would then be indeterminate whether the episode is perceptual. Here, the indeterminacy flows up, from its being indeterminate whether attention is present, to its being indeterminate whether an episode is perceptual.¹¹

Let’s summarise. I have argued that if we replace an essentialist view of perception with one inspired by the HPC account, then we can offer a new picture of perception. On this picture, perception is a HPC natural kind. The properties in the cluster definitive of this kind include *at least* being individual-level, being sensory, being available to attention, and involving constancy. There may be other properties. These properties are allowed entry into the cluster because of their stellar explanatory credentials. The mechanisms underpinning the clustering include at least DNA, gene expression mechanisms and various neural mechanisms. What this amounts to is a new view of perception.

¹¹ Block briefly mentions a similar idea. He says that being engaged with the subject’s preferences or needs, being available to central agency, and engaging the subject’s understanding are all potential markers of being individual-level (Peters et.al. 2017, p.8). He also suggests that there is no determinate boundary between the personal-level and the subpersonal-level (Peters et al. 2017, p.8).

If we accept this picture, then there are at least three ways that indeterminacy can arise:

- 1) Some but not all of the properties in the cluster are present.
- 2) The underlying mechanisms are partially inoperative.
- 3) It is indeterminate whether a certain property in the cluster is present.

5. Unconscious perception?

I now turn to unconscious perception. I will survey the main proposed examples of unconscious perception. I show that, if we accept the view outlined above, they are more plausibly regarded as cases where it is indeterminate whether the mental episode in question is an instance of perception. I cannot discuss every supposed instance of unconscious perception, so I constrain myself to the ones that are most prominently discussed. Extension to others will be straightforward.

I start by setting two issues aside. First, to establish unconscious perception, we must show that the item was not perceived *consciously*. Phillips uses signal detection theory to address this issue (Phillips 2016). This is important, but not the primary focus of this paper. This paper is interested in whether the stimuli in question were *perceived*, not whether they were conscious. Second, the focus of the debate has been on whether certain mental episodes involve constancy and the individual-level property. It is not debated that at least some unconscious mental episodes are *sensory*. All of the examples I consider are sensory. The indeterminacy interpretation that I'll argue for comes from elsewhere.

5.1 Blindsight

In one experiment, a blindsight subject (GY) was presented with either a cue that appeared at the same location as a subsequently presented target in his blindfield (exogenous condition); or he was presented with an arrow in his healthy field that pointed toward the location of a subsequently presented target in his blindfield (endogenous condition).¹² In both of these conditions, GY's discrimination of the target was faster than in cases where the target was absent, or the cue was misleading (Kentridge *et.al.* 1999). Such effects are taken to be indicative of attention to the stimulus.¹³ So it is the case that the stimuli could have

¹² Blindsight: a condition arising from damage to visual cortex. Subjects deny awareness of objects placed in a certain area of their visual field, but can still use visual information from these stimuli to perform experimental tasks (Weiskrantz 2002).

¹³ Attention was measured using the Posner paradigm, the gold standard for measuring attention effects in perceptual psychology (Posner 1980). Some may claim that other varieties of attention were not present. Prinz, for example, understands attention in terms of availability to working memory (2010). If we accept that this is one expression of attention, then this will only be more grist to my mill: some but not all of the markers of attention were present, so the indeterminacy arises here as well.

attention directed upon them (endogenous condition) and could themselves trigger attention (exogenous condition). Given that I have argued for the inclusion of attention into the cluster of properties definitive of perception, this pushes us toward the claim that the stimulus was perceived.

Take the property of being available to central agency, which is indicative of being individual-level. There are some features of blindsight that point toward the claim that accessibility to central agency is present. GY is pressing a button, which he has been instructed to do by the experimenters. The stimuli are having a causal impact on his personal-level agency, which feeds into the completion of instructions that he has been consciously given. In some conditions, blindsight subjects can respond to stimuli without prompts (Storeig 2010). Blindsight subject TN was able to spontaneously navigate a hallway full of obstacles (De Gelder *et.al.* 2008). Since availability to central agency is indicative of being individual-level, this pushes us toward the conclusion that the individual-level property is present, which in turn contributes to the claim that blindsight subjects' sensory episodes are perceptual.

But other things pull us in the opposite direction. As Phillips argues, evidence suggests that constancy is abolished in blindsight (forthcoming, pp.17-19). For example, Kentridge *et.al.* (2007) found that blindsight subject DB lacks even the basic processing required for colour constancy. So, *at the very least*, we have strong reason to doubt that blindsight patients are capable of *colour* constancy. Therefore, we have good reason to think that not all of the mechanisms associated with constancy are operative. This pushes us away from the conclusion that the stimuli are perceived. Furthermore, not all of the faculties distinctive of availability to central agency are present. In many studies, GY persistently denied seeing anything, claiming that he did not even have a 'feeling that something was there' (Kentridge et al. 1999, p.1806; 2004, p.833). He even accused the experimenters of running a control experiment in which there really were no cues (Kentridge 2011, p.239). This indicates that the stimuli were not available for voluntary inspection or use in inference.

Blindsight is also clearly a case where the mechanisms that underpin perception are only partially operative.¹⁴ Blindsight results from damage to visual cortex, and since the neural mechanisms underpinning visual perception are partially located there, these mechanisms were partially malfunctioning. Of course, it

¹⁴ Similar things go for the case of DF, discussed by Block (Peters et al. 2017, p.8), and cases of unilateral neglect, discussed by Phillips (2016).

is not plausible to claim that all of the mechanisms underpinning perception were destroyed. For example, GY has only suffered damage to his left striate cortex (Kentridge et al. 1999, p.1805). Blindsight subjects still have much of their visual system intact, including (for example) the lateral geniculate nuclei.

So, the perceptual effects show availability to engage endogenous and exogenous attention, but there is no evidence of constancy (at least, not all of the features definitive of constancy are present) and there are some but not all of the markers definitive of availability to central agency. The mechanisms underpinning perception are also partially inoperative. If you hold the HPC view of perception as a kind, the verdict is clear: it is indeterminate whether blindsight subjects are perceiving.¹⁵

5.2 *Meta-contrast masking*

Meta-contrast masking is a process where a prime (e.g. a disk) is presented, followed by an annulus (e.g. a ring), the inside edge of which coincides with the outside edge of the prime. The visual processing of the prime is rendered unconscious, whilst the annulus is consciously perceived. If the annulus is the same colour as the prime, subjects will be faster at discriminating the colour of the annulus (Breitmeyer *et.al.* 2004). Norman *et.al.* (2014) showed that subjects are faster at distinguishing the colour of the annulus when it matched the prime's surface colour, as opposed to the wavelength at which light was reflected from it. The ability to distinguish surface colour from wavelength is pretty much definitional of colour constancy, and so (since subjects aren't conscious of the disc) it was concluded that unconscious stimuli can exhibit colour constancy. This pushes us toward the claim that they were perceived.

Phillips doesn't deny that these effects illustrate unconscious constancy, but he objects that meta-contrast masking only demonstrates priming (Block and Phillips 2016, p.4). And priming effects (says Phillips) are not individual-level, i.e. priming only facilitates sub-personal processing. This pushes us further from the claim that meta-contrast masked stimuli are perceived. Of course, it might be claimed that priming exhibits accessibility to central agency of at least some kind. But even if we accept this, it clearly falls short of the full range of abilities distinctive of availability to central agency. So, constancy is present but we lack reason to claim that the property of being individual-level is entirely there. So some but not all of the properties definitive of perception are present. Indeterminacy.

¹⁵ Hyman (1991) argues for an interpretation of blindsight that is in some ways similar from a Wittgensteinian point of view.

5.3 *Continuous flash suppression*

In continuous flash suppression (CFS), the eyes are presented with two distinct conflicting images, typically a high-contrast set of colours to one eye (the ‘noise patch’), and a normal picture to the other. When the high-contrast coloured image is presented at a very high frequency, subjects deny seeing the image presented to the other eye (Tsuchiya and Koch 2005, Fang and He 2005). In one study (Jiang *et.al.* 2006), a similar technique was used to render perception of nude images unconscious.¹⁶ Gender and sexual-orientation specific effects with relation to the nudes were found. For example, heterosexual males were faster to identify the orientation of a Gabor patch (a rippled texture) that appeared at the same location as a previously presented female nude.

These experiments have been subject to much debate (Block and Phillips 2016). At the very least, we should accept that the nudes attracted attention.¹⁷ So at least some of the features indicative of availability to attention are present. This pushes us further toward the claim that they are perceived. One hotly contested issue is whether the stimuli in question affect central agency (Block and Phillips 2016, p.178, p.181). Recall that central agency is the system that controls individual-level actions such as feeding, fighting, mating etc. The contribution of sexual orientation to the deployment of attention is part of mate-recognition behaviour, and so it should be admitted that the nudes were available to the central agency system in at least some way. But we should deny that the stimuli were fully available to the central agency system. Though there is a reaction time advantage engendered by the nude, which is linked to sexual orientation, there is no evidence that the nude elicits the full range of reactions associated with sexual attraction. In summary, with the nudes study, some put not all of the features definitive of perception are present. We have indeterminacy again.¹⁸

¹⁶ One eye received a picture and a scrambled version of that picture, and the other received two scrambled black and white pictures.

¹⁷ We may doubt whether the nudes themselves were *subject to* attention, but we must accept at least that they triggered the deployment of attention to the subsequently presented Gabor patch (as reflected in the reaction time advantage). So at least the ability to *trigger* attention was present. Phillips (Block and Phillips 2016, p.174) cites Prinz (2010), arguing that the nudes study can be accounted for using eye saccades, rather than attention. Prinz does deny that the nude was subject to attention, but he doesn’t deny that the Gabor was.

¹⁸ I take the nudes study as my main example of CFS. Block and Phillips (2016) discuss other examples, including one in which images trigger changes in skin conductance, which Block argues are indicative of a fear response (Raio *et.al.* 2012). There is debate over whether these responses are cortical or merely physiological so it’s not clear whether they are individual-level (Block and Phillips 2016, pp.181-182). Drawing on Mudrik et al. (2011) Block also claims that high-level perceptual categorisation is present in cases of CFS, a point that Phillips denies (Block and Phillips 2016, pp.171-176). Because of the difficulty of interpreting these results, I will not take a stand on them here. At the very least, they do not display the full range of features distinctive of being available to central agency. So at best, they are also cases of indeterminacy.

5.4 Object files

Jake Quilty-Dunn (forthcoming) uses what he calls ‘phenomenal coherence’ to argue for the presence of unconscious perception. He points to studies in which some objects are displayed, and a feature appears inside each object. The features disappear, the objects move, and then another feature appears. Subjects are asked if this feature matches any of the previously presented ones. They are faster if the feature appears in the same object in which it appeared initially. This is the object-specific preview benefit (OSPB). In Mitroff *et.al.* (2005), the visual appearance of the objects’ movements was ambiguous between the two objects ‘bouncing off’ one another, and ‘streaming through’ one another. Subjects claimed to see the objects as streaming, but they had an OSPB that indicated they had tracked the objects as though they had bounced. Quilty-Dunn claims that the ‘bouncing’ representations are instances of unconscious perception.

Attention is standardly invoked to explain how we visually track objects (Rensink 2013), so we can accept that attention was paid to the objects as though they had bounced. What about the individual-level criterion? The representations certainly have an impact on behaviour: they decrease the reaction times of subjects in performing actions that they have been instructed to perform. But this is not indicative of the full range of capacities constitutive of central agency. Presumably, the conscious planning and implementation of actions would only make use of the representations of streaming, since these are the only representations that subjects are aware of (to assume otherwise would be to directly contradict their reports). So the ‘bouncing’ object files do not demonstrate the full range of activities indicative of central agency.

Quilty-Dunn gives other reasons to think that the object files are individual-level, but they are inconclusive. He takes working memory to be an individual-level capacity (forthcoming, p.7). He then notes that to fulfil the task, subjects had to attend to the objects in question, and attention is the capacity that decides which information is encoded in working memory. Of course, we may accept that *one* function of attention is to oversee access to working memory, but it doesn’t follow that everything that is attended is encoded in working memory. Blindsight subjects pass experimental tests for attention to stimuli that no one thinks are encoded in working memory. Quilty-Dunn also notes that we can keep track of about four object files, and that the capacity of working memory is about four. But this view of working memory as having four ‘slots’ is now widely rejected, in favour of a view in which working memory does not have

slots, but is a flexible resource that can be widely spread across many or few items (Ma *et.al.* 2014). In any case, it is far from clear that being encoded in working memory is sufficient for being individual-level (Soto *et.al.* 2011).

I have surveyed the range of experiments at the centre of the debate over unconscious perception. I have argued that, if we embrace the HPC view of perception suggested above, they are more plausibly seen as cases where it is indeterminate whether the mental episodes in question are perceptual.

6. Conclusion

I have argued for three things:

- i) We should extend the list of properties definitive of perception at least to include availability to attention.
- ii) If we accept a view of perception that is in line with the HPC view (which is the overwhelmingly dominant position on kinds in the life sciences), then we can develop a new view of perception.
- iii) On this view, all of the putative cases of unconscious perception are more plausibly seen as cases where it is indeterminate whether the mental episode in question is perceptual.

These conclusions have widespread consequences. If we accept this view, it provides a new way of thinking about perception. On such a view, perception is a messy cluster, supported by mechanisms that don't always work. It comes along with a fuzzy transition between the non-perceptual and the perceptual. This is an entirely new way of thinking about perception, hitherto completely unexplored in philosophy.

If this view is accepted, then there will be important upshots for philosophical theories of consciousness. Such theories usually assume that perception can be both conscious and unconscious, and then attempt to slot consciousness into the picture (Prinz 2010). But if the view given above is accepted, then we have a complex interplay of different cases between the perceptual and the non-perceptual, and our theory has to knit consciousness in amongst this complexity somehow. There are also dramatic upshots for philosophical theories of perception such as naïve realism and representationalism. Philosophers of perception have paid some attention to the possibility of unconscious perception (Berger and Nanay 2016), but there has been no discussion of cases where it is indeterminate whether a mental episode is perceptual

or not. If the view offered above is correct, it will be incumbent on such theories to accommodate these cases. How they might do this is anything but clear.

There is a much more general lesson to be drawn here, which is that we must not cut off philosophy of perception from the valuable insights about kindhood from philosophy of science. When we recognise this, things are going to get very complicated. But then, no one ever said it would be easy.¹⁹

¹⁹ I would like to thank Alexandria Boyle, Ian Phillips and an anonymous referee for comments on a previous draft of this paper. Thanks to Ori Beck, Will Davies, Bob Kentridge and an audience at *The European Society for Philosophy and Psychology* in Hertfordshire for helpful discussion.

References

- Assis, L. and Brigandt, I. 2009. Homology: homeostatic property cluster kinds in systematics and evolution. *Evolutionary Biology*. 36: 248-255.
- Bechtel, W. and Mundale, J. 1999. Multiple realizability revisited: linking cognitive and neural states. *Philosophy of Science*. 66: 175-207
- Berger, J. and Nanay, B. 2016. Naïve realism and unconscious perception. *Analysis*. 76: 426-433.
- Bird, A. 2007. *Nature's Metaphysics*. New York: OUP.
- Block, N. 2014. Seeing-as in the light of vision science. *Philosophy and Phenomenological Research*. 89: 560-572.
- Block, N. 2016. The Anna Karenina principle and skepticism about unconscious perception. *Philosophy and Phenomenological Research*. 93: 452-459
- Block, N. and Phillips, I. 2016. Block/Phillips Debate on Unconscious Perception. In B. Nanay (ed.) *Current Controversies in Philosophy of Perception* New York: Routledge.
- Boyd, R. 1989. What realism implies and what it does not. *Dialectica*. 43, 5-29.
- Boyd, R. 1991. Realism, anti-foundationalism and the enthusiasm for natural kinds. *Philosophical Studies*. 61: 127-148.
- Breitmeyer, B et.al. 2004. Unconscious color priming occurs at stimulus-not percept-dependent levels of processing. *Psychological Science*. 15: 198-202.
- Burge, T. 2005. Disjunctivism and Perceptual Psychology. *Philosophical Topics*. 33: 1-78.
- Burge, T. 2010. *Origins of Objectivity*. New York: OUP
- Carrasco, M. 2011. Visual attention: the last 25 years. *Vision Research*. 51: 1484-1525.
- Carrasco, M. et.al. (2004). "Attention alters appearance." *Nature*. 7: 308-313.
- Chun, M. et.al. 2011. "A taxonomy of internal and external attention." *Annual Review of Psychology*. 62: 73-101.
- Craver, C. 2009. Mechanisms and natural kinds. *Philosophical Psychology* 22: 575-594.
- Devitt, M. 2008. Resurrecting Biological Essentialism. *Philosophy of Science*. 75: 344-382.
- Dupré, J. 1981. Natural kinds and biological taxa. *The Philosophical Review*. 90: 66-90.
- Ereshefsky, M. and Matthen, M. 2005. Taxonomy, polymorphism and history. *Philosophy of Science*. 72: 1-21.
- Ereshefsky, M. and Reydon, T. 2015. Scientific Kinds, *Philosophical Studies*. 172: 969-986
- Fang, F. and He, S. 2005. Cortical responses to invisible objects in the human dorsal and ventral pathways. *Nature Neuroscience*. 8: 1380-1385.
- Fazekas, P. and Nanay, B. 2018. Attention is amplification, not selection. *British Journal for the Philosophy of Science*. doi.org/10.1093/bjps/axy065
- De Gelder, B. et.al. 2008. Intact navigation skills after bilateral loss of striate cortex. *Current Biology*. 18: R1128-R1129.
- Griffiths, P. E. 1997. *What Emotions Really Are*. Chicago: University of Chicago Press.
- Hull, D. 1965. The Effect of Essentialism on Taxonomy: Two Thousand Years of Stasis (I). *British Journal for the Philosophy of Science*. 15: 314-326
- Hyman, J. 1991. Visual experience and blindsight. In Hyman, J. (ed.) *Investigating Psychology*. London: Routledge.
- Jiang, Y. et.al. 2006. A gender- and sexual orientation-dependent spatial attentional effect of invisible images. *Proceedings of the National Academy of Sciences*. 103: 17048-17052.
- Kentridge, R. 2011. Attention without awareness: a brief review. In Mole, C., Wu, W. and Smithies, D. (eds.) *Attention: Philosophical and Psychological Essays*. New York: Oxford University Press.
- Kentridge, R. et.al. 1999. Attention without awareness in blindsight. *Proceedings of the Royal Society (London) Series B*. 266: 1805-1811.
- Kentridge, R. et.al. 2004. Spatial attention speeds discrimination without awareness in blindsight. *Neuropsychologia*. 42: 831-5.
- Kentridge, R. et.al. 2007. Colour contrast processing in human striate cortex. *Proceedings of the National Academy of Sciences*. 104: 15129-15131.
- Ma, W. et al. 2014. Changing concepts of working memory. *Nature Neuroscience*. 17: 347-356.
- Machery, E. 2009. *Doing Without Concepts*. New York: OUP.
- Magnus, P.D. 2012. *Scientific Enquiry and Natural Kinds*. New York: Palgrave Macmillan.
- Mitroff, S. et.al. 2005. The relationship between object files and conscious perception. *Cognition*. 96: 67-92.

- Morawetz, L. and Spaethe, J. 2012. Visual attention in a complex search task differs between honeybees and bumblebees. *The Journal of Experimental Biology*. 215: 2515-2523.
- Norman, L. *et.al.* 2014. Color constancy for an unseen surface. *Current Biology* 24: 3822-3826
- Peters, M. *et.al.* 2017. Does Unconscious Perception Really Exist? *Neuroscience of Consciousness*. 3: 1-11.
- Piccinini, G and Craver, C. 2011. Integrating psychology and neuroscience: functional analyses as mechanism sketches. *Synthese*. 183: 283-311.
- Phillips, I. 2016. Consciousness and Criterion: On Block's Case for Unconscious Seeing. *Philosophy and Phenomenological Research*. 93: 419-451.
- Phillips, I. forthcoming. Unconscious Perception Reconsidered. *Analytic Philosophy*.
<https://www.ianbphillips.com/papers.html>
- Posner, M. 1980. "Orienting of attention." *Quarterly Journal of Experimental Psychology*. 32: 3-25.
- Prinz, J. 2010. When is perception conscious? In B. Nanay (ed.) *Perceiving the World*. Oxford: OUP.
- Quilty-Dunn, J. forthcoming. Unconscious Perception and Phenomenal Coherence. *Analysis*.
<https://sites.google.com/site/jakequiltydunn/home>
- Raio, C.*et.al.* 2012. Nonconscious fear is quickly acquired but swiftly forgotten. *Current Biology*. 22: 477-479.
- Rensink, R. 2013. "Perception and attention." In D. Reisberg (ed.) *The Oxford Handbook of Cognitive Psychology*. New York: Oxford University Press.
- Samuels, R. and Ferreira, M. 2010. Why *don't* concepts constitute a natural kind? *Behavioral and Brain Sciences* 33: 222-223.
- Slater, M. 2015. Natural Kindness. *British Journal for the Philosophy of Science*. 66: 375-411.
- Sober, E. 1980. Evolution, population thinking and essentialism. *Philosophy of Science*. 47: 350-383.
- Soto, D. *et.al.* 2011. Working memory without consciousness. *Current Biology*. 21: R912-R193.
- Storeig, P. 2010. Cueless Blindsight. *Frontiers in Human Neuroscience*. 3: 1-8.
- Taylor, H. 2018a. Emotions, Concepts and the Indeterminacy of Natural Kinds. *Synthese*.
<https://doi.org/10.1007/s11229-018-1783-y>
- Taylor, H. 2018b. Attention, psychology and pluralism. *British Journal for the Philosophy of Science*.
doi.org/10.1093/bjps/axx030
- Treisman, A. 1996. "The binding problem." *Current Opinion in Neurobiology*. 6: 171-178.
- Tse, P. 2005. "Voluntary attention modulates the brightness of overlapping transparent surfaces." *Vision Research*. 45: 1095-1098.
- Tsuchiya, N. and Koch, C. 2005. Continuous flash suppression reduces negative afterimages. *Nature Neuroscience* 8: 1096-1101.
- Weiskrantz, L. 2002. Prime-Sight and Blindsight. *Consciousness and Cognition*. 11: 568-581.
- Wilson, R. *et.al.* 2007. When traditional essentialism fails. *Philosophical Topics*. 35: 189-215.
- Wu, W. 2014. *Attention*. New York: Routledge.