

ON THE EPISTEMIC SIGNIFICANCE OF PERCEPTUAL STRUCTURE

BY DOMINIC ALFORD-DUGUID

Our awareness of the boundedness of the spatial sensory field—a paradigmatic structural feature of visual experience—possesses a distinctive epistemic role. Properly understood, this result undermines a widely assumed picture of how visual experience permits us to learn about the world. This paper defends an alternative picture in which visual experience provides at least two kinds of non-inferential justification for beliefs about the external world. Accommodating this justification in turn requires recognising a new way for visual experience to encode information about the world. Reflection upon the epistemic contribution of sensory experience’s structural features thus forces us to revise our understanding of how perception, cognition, and the world fit together.

Keywords: perceptual structure, perceptual justification, presupposition, sensory experience, perceptual presentation.

We must distinguish two factors that shape visual experience’s conscious character: facts about which particulars and properties a subject is visually aware of, and facts about her experience’s structural features. Structural features are relatively invariant aspects of our sensory phenomenology. For example, suppose a frolicsome puppy runs into view. In seeing the puppy, you’re visually aware of its shape, colour, and so on. You’re also aware of the region in view—your ‘spatial sensory field’—as having bounds (i.e., as bounded). This is a paradigmatic ‘structural’ feature because regardless of which objects you see, or what those objects look like, your experience involves being aware of the region as having bounds.

Recent interest in structural features has been driven by their potential to explain recalcitrant aspects of sensory experience’s conscious character (e.g. the experience of absence).¹ My interest lies instead with our paradigmatic structural feature’s *epistemic* role. A perceptual feature’s ‘epistemic role’ is its

¹ Martin (1992) is responsible for framing and launching the contemporary discussion of sensory experience’s structural features (see also Martin 1995). See Richardson (2010), Soteriou (2013), Mac Cumhaill (2015), Laasik (2019), and Textor (2019). See also Lande (2021) for a different—though related—notion of perceptual structure.

contribution to our ability to learn about the world. We learn about the world in part by forming beliefs in response to sensory experience. Some of these beliefs are reasonable to form. An object's looking red, for example, ordinarily makes it reasonable to believe that something red is before you. This is a case in which visual experience 'justifies' your belief.² It is also arguably a classic example of *non-inferential* perceptual justification. Perception provides 'non-inferential justification' when it justifies a belief without that justification deriving from perception-based justification for some other belief.³

Our paradigmatic structural feature—awareness of the spatial sensory field as bounded—possesses a distinctive epistemic role. Recognising this role causes trouble for a seemingly natural picture of how perception permits us to learn about the world. Philosophers frequently assume both that visual experience provides at most one fundamental kind of non-inferential justification for beliefs about the external world, and that non-inferential perceptual justification requires at most one way for visual experience to encode information.⁴ I use our paradigmatic feature's distinctive epistemic role to undermine these assumptions and the picture they frame. My view emerges from the ashes: Visual experience provides at least two fundamental kinds of non-inferential justification, and there are at least two corresponding ways for visual experience to encode information about the world. Reflection upon the epistemic contribution of sensory experience's structural features thus forces us to revise our understanding of how perception, cognition, and the world fit together.

The paper has five parts. Section I builds on Richardson (2010) and Soteriou (2013) to explicate our awareness of the spatial sensory field as bounded. Section II establishes this structural feature's distinctive epistemic role. Section III motivates the target picture. Section IV uses the epistemic role established in Section II to both undermine that picture and defend my alternative view. Section V traces a few consequences.

² I follow McGrath's (2018) liberating convention of treating a belief as 'justified' when it is reasonable to form, rather than only when it almost counts as knowledge. See also Pryor (2005), on whose view you have justification for a belief iff 'you're in a position where it would be epistemically appropriate for you' to have that belief.

³ Some reserve 'non-inferential perceptual justification' for the more demanding condition that a belief formed in response to sensory experience does not depend for its justification upon the justification for any other belief (cf. Pryor 2005; though see Smithies 2019: 74–5) for a version of my less demanding notion). However, those who permit non-inferential perceptual justification (in my sense) need not automatically take a stand on the disagreement between dogmatists like Pryor (2000) and anti-dogmatists such as Wright (2007).

⁴ This picture is arguably tacit in most frameworks that admit even relatively undemanding forms of non-inferential perceptual justification (e.g. Burge 2003; Huemer 2001; Pryor 2000; Silins 2011; Smithies 2019). I explore motivations for the picture in Section III.

I. BOUNDEDNESS AND THE REGION IN VIEW

I follow Richardson (2010) and Soteriou (2013: §5.1) in their core characterisation of our awareness of the spatial sensory field as bounded (i.e., as having bounds).⁵ They argue that this awareness has a special character: a subject is aware of the region as bounded *merely by her sensory limitations*. The contrast case is bodily awareness. Just as visual awareness is limited to the region in view, bodily awareness only extends so far. Yet the limits of bodily awareness do not strike us as merely sensory; they seem due to the boundaries of an object—the body. Now, suppose our awareness of the boundedness of the spatial sensory field were *agnostic* about whether the region's limits are merely sensory. Certain possibilities would remain open: it might be that the universe is so tiny that space does not extend beyond the region in view; or that the world is not tiny, but we who are large; or that the region in view is nothing but the interior of some object (as in the bodily case). But visual experience does not leave open these possibilities. We are thus plausibly aware of the region in view as bounded merely by our sensory limitations. Soteriou and Richardson exploit this result to defend an explanatory claim:

BOUNDEDNESS: Our awareness of the region in view as bounded merely by our sensory limitations in turn constitutes an awareness of that region as a sub-region of a larger space.⁶

There remains a question about exactly why visual experience possesses the relevant structural feature.⁷ But BOUNDEDNESS entails that our awareness of the region in view as a sub-region of a larger space is a feature of our visual experience rather than our cognitive 'take' on that experience.⁸ Here is Soteriou motivating BOUNDEDNESS:

⁵ Some will balk at the basic claim that visual experience makes us aware of the region in view as having bounds; in such disputes, I defer to Richardson and Soteriou. Martin (1992, 1995) shares something like this scepticism, since he distinguishes visual phenomenology from bodily and tactile phenomenology in part by denying that visual experience involves awareness of the boundaries of our spatial sensory field.

⁶ Soteriou (2013: 116–118) traces the *explanandum* in BOUNDEDNESS to Kant's (1998) transcendental aesthetic, following Melnick (1973: 11).

⁷ Cavedon-Taylor (2021) proposes that the feature reflects sensorimotor expectations. Of course, dependence upon sensorimotor expectations, or our capacity for action, does not prevent our awareness of the spatial sensory field's boundedness from counting as an aspect of visual phenomenology.

⁸ A central aim of this paper is to shed light on the nature of this awareness by examining its epistemic role. Yet the extensive literature on the perception of high-level properties has given rise to a plethora of fine distinctions concerning sensory awareness. Where does the awareness described by Soteriou and Richardson fit? Theirs is a claim about sensory phenomenology rather than epistemic looks or interpretive aspects of consciousness. Further distinctions are irrelevant to my main argument, though they may be fruitful for those who seek further elucidation of our awareness of the region in view as bounded.

That there are limits to what can now be sensed that are due to one's sensory limitations, rather than due to the limits of whatever it is that one is now sensing, brings with it the idea that there is more to be sensed beyond those sensory limits, hence the idea that your visual awareness of the region of space in front of you is in some sense an awareness of the region as a sub-region of a region of space that has that sub-region as part. (Soteriou 2013: 118; see also Richardson 2010: 234).

I simply take on board Richardson and Soteriou's arguments for BOUNDEDNESS. And that is because my primary concern is the epistemic role that acceptance of BOUNDEDNESS generates for our awareness of the region in view as bounded. However, I shall go beyond Richardson and Soteriou in elucidating what it is to be the region in view.

As I understand it, the region in view is the region that becomes *interactable* solely because of a subject's visual experience. Unpacking this 'modal criterion' requires some care. Being interactable is a modal property that we cannot assume reduces to non-modal features of a perceiving subject (e.g. their actual practical or cognitive mechanisms). For all the modal criterion says, then, a region might be interactable even if a given subject lacks the practical or cognitive capacity to interact with the region.⁹

This aspect of the modal criterion permits it to avoid the kind of counterexamples that threaten to undermine a purely dispositional criterion. On a purely dispositional criterion, the region in view is that region which a subject is disposed to interact with in response solely to her visual experience. As in my modal criterion, the 'solely' here is meant to rule out certain putative counterexamples, for example, those that involve a disposition to interact with (intuitively) unseen space partly in response to visual cues, and partly in response to an awareness of that region mediated by bodily awareness. Yet the purely dispositional criterion remains vulnerable to other counterexamples, such as agents who lack dispositions to interact with the outside world or whose dispositions are somehow deviant (e.g. they systematically ignore the margins of their visual field).¹⁰ By contrast, whether a region counts as interactable because of a visual experience does not depend upon idiosyncratic features of a perceiving subject's practical and cognitive capacities. Nevertheless, there remains an important connection between my modal criterion and the inadequate dispositional criterion. While a subject's dispositions are not always sufficient to fix which region becomes interactable because of her visual experience, in subjects that otherwise possess non-pathological dispositions, it is possible to use these dispositions as a guide when determining which region becomes interactable because of a visual experience.

⁹ Martin (2006: 364) makes a structurally parallel appeal to potentially primitive modal properties when elucidating his core notion of 'reflective indiscriminability'.

¹⁰ I thank referees for pressing the relevance of these kinds of counterexamples.

My ‘modal’ criterion has important consequences. These come into focus when the criterion is placed alongside an alternative. On this competing ‘presentational’ criterion, the region in view is that which a subject’s visual experience *presents*. As I use the term here, ‘presentation’ (or ‘perceptual presentation’) picks out the perceptual relationship characteristic of paradigm cases of object and property perception: seeing a red surface, or a round coin, and so on.

Perceptual presentation requires a characteristic kind of causal sensitivity. For example, if your perceptual system is operating properly, then your visual experience of a regularly shapeshifting object will evolve to reflect changes in the object’s shape. Were your visual experience not sufficiently sensitive to major changes in a perceived object’s shape, it would not present that shape. This connection between perceptual presentation and causal sensitivity explains why perceptual illusion and hallucination seem to involve failures of perceptual presentation: Standard illusion cases are those in which our visual system is insufficiently sensitive to a perceived object’s features, and hallucinations arise when there is no object to which our visual system is appropriately sensitive. Of course, this sensitivity cannot amount merely to brute reliability. Even if your experience tends to ‘match’ what an object is like, you avoid being victim of illusion or hallucination only if your experience matches what the object is like in virtue of what the object is like.¹¹

A sensitivity constraint on perceptual presentation places significant limits on when a subject’s visual experience counts as presenting a spatial region. In ordinary cases, the region that a visual experience presents, and thus which is causally responsible for the experience, will be the same region that becomes *interactable* solely because of a subject’s visual experience. The competing ‘modal’ and ‘presentational’ criteria come apart only in pathological cases.

Here is one such pathological case. Imagine you are fitted with prismatic glasses that you have not yet been entrained upon. There will then be objects that fall outside your ordinary field of view that you experience as being in front of you. While the region causally responsible for your visual experience is actually off to one side, you will remain disposed to act as if the presented objects are in front of you. On my modal criterion, this is a case where you are inclined to make a mistake about *which* region is in view (since your dispositions to act are not deviant, and are therefore a guide to which region becomes interactable because of your experience); the presentational criterion treats it as a case where you are inclined to mistake *the location* of the region in view.

Pure hallucination is another pathological case. In pure hallucination, your visual experience is not even a partial upshot of sensory uptake. On the

¹¹ I take this to be the central lesson of so-called ‘veridical hallucination’ and ‘veridical illusion’ cases. Compare Lewis (1980) and Johnston (2006).

presentational criterion, there will be no region that counts as the region in view; no region can satisfy the sensitivity constraint on perceptual presentation. On my modal criterion, by contrast, even in pure hallucination, a region may still count as the region in view. There may still be a region a subject is disposed to interact with in response to her visual experience, and thus a region that plausibly counts as becoming interactable because of her visual experience.

I initially introduced structural features as relatively invariant aspects of sensory phenomenology. For instance, regardless of which objects you see, or what those objects look like, you remain aware of the region in view as bounded. My modal criterion adds another dimension to this invariance: our paradigm structural feature characterises visual experience even in hallucination. Given BOUNDEDNESS, this result entails that even in cases of hallucination, visual experience makes us aware of the region in view as a sub-region of a larger space. And that is indeed my position.

My position permits a compelling intervention in a longstanding dispute over the commitments of visual experience, and this provides my view (and by extension my modal criterion) with some independent support. The dispute concerns whether the visual experience is *manifestly* a partial glimpse on a larger world.¹²

On the one side are ‘agnostics’ who allow that the spatial sensory field seems limited, but deny that this amounts to its seeming like a sub-region of a larger space. For example, sense datum theorists and their opponents often hold as common ground that visual experience presents a mosaic of sensible qualities.¹³ Sense datum theorists hold that the qualities are instantiated whenever presented, and sometimes by sense data (as opposed to particulars in the external world); opponents standardly insist that the qualities are not instantiated in cases of hallucination or illusion, and never by sense data. A sense datum theorist and her opponent thus disagree over *when* these qualities are instantiated and *what* instantiates them. However, they agree that the qualities—‘pure visibilia’—do not require for their instantiation that the region in view is a sub-region of a larger space.

To resist the agnostic, one must find a condition encoded by visual experience which is such that were it satisfied, visual experience would be a partial glimpse on a larger world. A mere array of sensible qualities does not provide such a condition since the qualities could be instantiated by sense data that do not themselves belong to (or constitute) a space that is a sub-region of a larger space. Likewise, it is not enough to appeal to the fact that we are aware of

¹² An important classic source for this dispute is the exchange between Strawson (1959: ch. 2; 1980) and Evans (1980, 1985). Recent discussions of the ‘objectivity’ of experience which pick up on this old debate include Campbell & Cassam (2014), Brewer (2020, 2021), and others. Siegel (2010) also discusses the nearby notion of the manifest ‘subject independence’ of certain objects of visual experience.

¹³ For an expression of this common ground, compare Strawson (1980: 277).

external particulars, since (by definition) hallucinatory visual experience does not include this awareness. One option is to say that visual experience makes it manifest that whatever particulars we see—if indeed we see—are denizens of a world that extends beyond what is in view. Yet this proposal has obvious problems: it adds implausible sophistication to our experience of particulars, leaves open the hard question of what exactly it is for a particular (on its own) to seem as if it belongs to a larger world, and makes mysterious how sensory experience can seem like a partial glimpse on a larger world even when there are no particulars in view.

I can provide a better response to the agnostic. My modal criterion permits our paradigm structural feature to characterise even hallucinatory visual experience. And BOUNDEDNESS entails that the structural feature encodes a condition that, if satisfied, guarantees that visual experience is a partial glimpse on a larger world (namely the condition that the region in view is a sub-region of a larger space). Apart from avoiding the pitfalls that plague other anti-agnostic views, my response has other virtues.¹⁴ It allows for cases where visual experience leaves a subject genuinely unsure about whether she sees an item in the external world (e.g. is a visibly discoloured patch an afterimage or a stain on my office wall?), and it explains how visual experience might not present particulars while still manifestly constituting a partial glimpse of a larger world.

A full defence of my response to the agnostic would require a developed argument for BOUNDEDNESS—for which I defer to Soteriou and Richardson. What I add with my modal criterion is an account of how BOUNDEDNESS interacts with hallucination and other pathological cases.

II. AN EPISTEMIC ASYMMETRY

Our awareness of the spatial sensory field's boundedness possesses a distinctive epistemic role. I defend this crucial claim in the present section by arguing for a specific epistemic asymmetry. The asymmetry obtains between our awareness of the boundedness of a spatial sensory field, on the one hand, and visual awareness of particulars and their properties, on the other. Beliefs that these features make it reasonable to form can be undermined in very different ways. Hence, ours is an asymmetry of *defeat*.

Absent special stage setting, if a subject's visual experience makes it seem that the world is one way, then it is reasonable for her to believe that things are that way. BOUNDEDNESS entails that the region in view seems to be a sub-region of a larger space. Hence, amongst the beliefs that visual experience permits

¹⁴ Anti-agnostic positions carry diverse commitments. Evans (1985), for instance, argues that visual space is a behavioural space and uses that claim as a foundation for resisting the agnosticism expressed by Strawson (1980). For an argument against Evans, see Campbell (2005).

you to reasonably form, there is a belief whose truth or falsity depends upon more than the layout of perceived objects and their visible features: the belief that the region in view is a sub-region of a larger space. This belief is reasonable to form in part because you are aware of the region in view as a sub-region of a larger space. According to **BOUNDEDNESS**, that awareness is constituted by awareness of the spatial sensory field's boundedness. Thus, when you believe that the region in view is a sub-region of a larger space—a belief formed in response to your visual experience—the belief counts as reasonable in virtue of the fact that your visual experience involves awareness of the spatial sensory field's boundedness. With this background in place, an example introduces my core asymmetry of defeat:

Sameen ingests what her (usually reliable) friends tell her is a psychotropic drug, and so takes herself to be victim of perceptual illusion. However, the drugs are duds, and so Sameen is not in fact a victim of illusion. Her visual experience includes awareness of a red chair as well as awareness of the boundedness of her spatial sensory field.

Sameen's mistaken belief about her perceptual situation acts as a defeater. Instead of believing that something red is in front of her—what it would otherwise have been reasonable to believe—Sameen should refrain from forming the belief (or abandon it if already formed). Yet defeat does not similarly threaten the belief that the region in view is a sub-region of a larger space. What does this felt asymmetry amount to?

II.1 Two doomed proposals

A first suggestion: Sameen's mistaken belief that she is a victim of illusion acts as a defeater for the first belief, but does not undermine the belief that the region in view is a sub-region of a larger space, because there cannot be illusions concerning whether space extends beyond what you see. But this flatfooted proposal does not withstand scrutiny. It tacitly relies on an arbitrarily narrow conception of perceptual illusion that restricts the phenomenon to familiar sensible qualities (colours, shapes, sizes, and the like). We ought instead recognise a broader—and perhaps more natural—conception on which perceptual illusions occur when something seems to possess one feature, and yet actually possesses an incompatible feature (e.g. when a round object looks oval). On this broader conception, Sameen's mistaken belief that she is a victim of illusion could—for all that has been said so far—manifest a concern that the region in view merely seems to be a sub-region of a larger region. And this concern would undermine the belief about the region in view. What form might the defeating concern take? Perhaps it is that for all Sameen knows, she lives in a 'tiny world': the universe contains no space except what is currently in view. The region in view would inherit its boundedness from the universe's limits,

and so not be a sub-region of a larger space; its bounds would instead reflect genuine ‘worldly’ limits.

Some may object that this sort of defeater is not coherent, and that the belief about the sub-region is distinctive simply in being immune to any defeater whatsoever. In particular, one might argue on broadly Kantian grounds that the relevant truth—namely that the region in view is a sub-region of a larger space—is epistemically necessary for those who possess spatial concepts (and thus that the truth is arguably *a priori*). If it were epistemically necessary, then a belief with the requisite content would not owe its epistemic status to a feature of visual experience. The proposal would be that the asymmetry of defeat is really an unremarkable asymmetry between a genuinely perception-based belief and a belief that does not owe its justification to perception.

Yet I contend that this Kantian approach fails to provide an adequate elucidation of the asymmetry of defeat illustrated by Sameen’s case. Even if it were epistemically necessary that every spatial region is a sub-region of a larger space—and I am far from sure that it is—this does not explain away the felt asymmetry of defeat. Instead of the belief that the space in view is a sub-region of a larger space, what would require epistemic support from perception is the belief that what happens to be in view is actually a region of space. This belief counts as reasonably formed only if visual experience provides grounds for taking what is in view to be a sub-region of a larger space. And yet the belief still seems immune to defeat from Sameen’s mistaken belief about her perceptual situation: an asymmetry of defeat remains unexplained.

Of course, these brief remarks are not likely to sway those deeply sympathetic to a Kantian treatment of space and perception.¹⁵ For committed Kantians, a region cannot seem to have the topological features characteristic of ‘manifest’ space without being a genuine spatial region. They thus refuse to recognise a gap between space as it shows up for us, and genuine or ‘real’ space. As a result, Kantians will simply deny that the epistemic good standing of the relevant belief—namely the belief that what happens to be in view is actually a region of space—requires additional support from the specific character of a subject’s visual experience. No region could be in view that is not a region of space (and so a sub-region of a larger space)! Or so says the Kantian.

I cannot refute this Kantian line here, though there is obviously much that could be (and has been) written about it. However, this silence does not undermine the interest of my project, or seriously threaten its viability. My project’s importance derives in part from how it accommodates an epistemic distinctiveness for structural features while rejecting (or, at the very least, not presupposing) a more thoroughgoing Kantian framework for thinking about perception. Furthermore, I expect that most readers are in fact hostile to the

¹⁵ For present purposes, I take this group to include Husserl, who (unlike Kant) saw space not as a form of intuition but sensuous objectivity.

Kantian's collapse of 'manifest space' and 'real space' (especially if they worry about the status of non-classical geometries). For these readers, an account of the felt asymmetry that presupposes the collapse is not viable, and so they must look elsewhere.

A second proposed explanation of the felt asymmetry exploits a distinction between illusion and hallucination. If Sameen took herself to be hallucinating, then this would also undercut her belief that there's something red in front of her, yet leave untouched the belief about the region in view. However, my modal criterion from Section I explains why this new asymmetry arises: were Sameen to hallucinate, there would be no red object that she sees; by contrast, there may still be a region with which she is disposed to interact in response to her experience, and which is therefore plausibly interactable because of her experience. The new asymmetry is thus not due to the nature of the justification for these beliefs but to what these beliefs are about. So it does not capture the original felt asymmetry.

The original asymmetry of defeat therefore traces to neither an *ad hoc* restriction on illusion nor a straightforward division between illusion and hallucination. I propose that the lesson to draw is that a correct account of the asymmetry must restrict Sameen's concern about illusion to a significant sub-class of perceptual illusions. Her concern could then undermine beliefs about the red chair she sees, yet leave undefeated the belief that the region in view is a sub-region of a larger space.

II.2 A sensitivity-involving solution

I take the original asymmetry to reflect a deep division in epistemic role. My challenge is therefore to isolate an appropriately *significant* sub-class of perceptual illusions. After all, some divisions amongst illusions generate merely trivial asymmetries of defeat. If Sameen were concerned only about being victim of some colour illusion, for example, then the belief about the region in view would remain undefeated simply in virtue of having nothing to do with colour. Avoiding this triviality requires typing illusions in a way that cross-cuts divisions amongst illusions that turn upon either the nature of the properties involved (e.g. colour illusions vs shape illusions) or the nature of what possesses those properties (e.g. illusions concerning the features of objects vs illusions concerning the features of spatial regions). We thus need an appropriate typing of illusions that permits us to distinguish the following: illusions that would undermine a belief that a fully perceived region—one whose boundaries fall wholly within the spatial sensory field—is a sub-region of a larger space; and illusions that would undermine a belief that the region in view is a sub-region of a larger space. Each belief treats a perceived region as a sub-region of a larger space, and so the distinction between the illusions cannot turn on the attributed properties or the nature of what is taken to possess these properties.

In fact, we *can* isolate an appropriately significant sub-class of perceptual illusions. Some illusions arise because our visual system misinterprets visual cues, either due to malfunction or something strange about the cues themselves. These might be used to defeat beliefs about both wholly visible sensible features, and those features that the visual system can detect despite occlusion (e.g. the straightness of an occluded edge). These illusions can certainly be used to defeat the belief that a fully perceived region—one whose boundaries fall wholly within the spatial sensory field—is a sub-region of a larger space. By contrast, the illusions that could be used to defeat the belief about the region in view are not of this general kind: They make no discernible difference to the sorts of cues that the visual system standardly exploits. Illusions in the first class exploit the sensitivity characteristic of perceptual presentation; illusions in the second class do not.

Having seen how to carve out an appropriately significant sub-class of perceptual illusion, let us consolidate the epistemic asymmetry illustrated by Sameen's case. A standard perception-based belief such as Sameen's *that is red* owes its epistemic status in part to a non-structural feature of visual experience. It is therefore susceptible to defeat that exploits the sensitivity characteristic of perceptual presentation. For example, if Sameen were to find out that her visual experience was not in fact sensitive to major changes in a perceived object's shape, and to that extent illusory, then this discovery would undermine her perception-based justification to form beliefs about its shape. By contrast, the belief about the region in view owes its epistemic status to a structural feature of Sameen's visual experience, and consequently is *not* susceptible to sensitivity-involving defeat. Sameen could take her visual system to be malfunctioning in a variety of ways, or the world to be presenting her visual system with misleading cues, and yet remain reasonable in holding fast to her belief.¹⁶ To reasonably abandon the belief, Sameen must possess a different type of defeater, such as an argument for a competing metaphysical picture of the world (e.g. solipsism).

This asymmetry of defeat is responsible for the distinctiveness of the epistemic role of our paradigmatic structural feature. Our awareness of the spatial sensory field's boundedness provides justification for the belief that the region in view is a sub-region of a larger space. And this justification is distinctive because it is immune to defeaters that exploit the perceptual sensitivity characteristic of perceptual presentation. These results matter because I use them

¹⁶ These would be cases where Sameen takes herself to be the victim of an illusion brought about by a failure of sensitivity. But not all failures of sensitivity generate illusions; others generate hallucinations. It may be that some such hallucinations, if taken by Sameen to obtain, would undermine her belief that the region in view is a sub-region of a larger space. Yet this poses no challenge. I require only a sensitivity-involving asymmetry of defeat with respect to illusion, not hallucination. For example, I can allow that modifying Sameen's case to make the drug a hallucinogen would not necessarily preserve the original asymmetry of defeat.

in Section IV to argue against a seemingly natural picture of how perception permits us to learn about the world. The alternative that emerges is important both for how it departs from the natural picture and for how much of that picture's broader motivation it preserves.

III. THE NATURAL PICTURE

I shall now introduce and motivate that natural picture. It has its roots in paradigm cases of non-inferential perceptual justification. For instance, the visual experience of a rolling ball can non-inferentially justify the belief that the ball is red because that experience presents the ball's colour. Theorists land on the natural picture when they treat this kind of case as a model for non-inferential perceptual justification. Two assumptions frame their picture.

EPISTEMIC PRESENTATIONALISM: there is at most one fundamental kind of non-inferential justification for those beliefs about the external world that are formed on the basis of visual experience, namely presentation-based non-inferential justification (or 'presentational justification').

EPISTEMIC PRESENTATIONALISM concerns only fundamental kinds of non-inferential justification. It therefore permits visual experience to support more fine-grained types of non-inferential justification. These types may be distinguished by variations in the presentation involved, or by what is presented (e.g. colours vs shapes).

PERCEPTUAL PRESENTATIONALISM: there is at most one way that information about the external world is encoded in visual experience which underwrites non-inferential perceptual-justification, namely the way characteristic of perceptual presentation.

PERCEPTUAL PRESENTATIONALISM allows that visual experience might encode information about the world in more than one way. It insists only that these other ways of encoding information are not required in order to accommodate non-inferential perceptual justification.¹⁷ Sensory experience 'encodes the information that p' if it would be reasonable for someone to believe that p in response to the experience. And the 'ways' for sensory experience to encode the information that p are individuated by how sensory experience encodes the information. Thus, all information that visual experience encodes in virtue of perceptual presentation counts as encoded in the same way, namely the way

¹⁷ **PERCEPTUAL PRESENTATIONALISM** is also intentionally coarse-grained due to the placeholder status of 'presentation'. A view in keeping with **PERCEPTUAL PRESENTATIONALISM** holds that accommodating non-inferential perceptual justification requires not only information encoded solely in virtue of a visual experience's phenomenology, but also information 'derivatively' encoded partly in virtue of the first level of information, and partly in virtue of other facts (e.g. about the causal antecedents of experiences with the relevant phenomenology). See Silins (2011) for a view of this type.

characteristic of perceptual presentation. An experience may even encode information in that way in the absence of successful perceptual presentation (e.g. in cases of hallucination). This last point is important. It explains why many (such as Pryor 2000 and Smithies 2019) would allow that ‘presentational justification’ may be present even when a subject hallucinates: So long as such a subject does not come to recognise that her perceptual situation does not involve actual perceptual presentation, she might retain presentational justification for her perception-based beliefs.

The basic distinction between ‘encoded information’ and ‘ways of encoding’ is most straightforwardly drawn within intentionalist frameworks that assign a fundamental explanatory role to perceptual content (though my arguments do not require the truth of intentionalism). For instance, intentionalists might identify some information encoded by visual experience with the content whose assignment to visual experience is necessary to characterise the experience’s conscious character (other information encoded by visual experience might then be identified with contents that play different explanatory roles). A way of encoding information thus becomes a relation to the encoded content that helps to distinguish, *inter alia*, the content of visual experience from content belonging to other types of sensory experience. By contrast, those who do not assign perceptual content a fundamental explanatory role within an account of visual experience are likely to understand perceptual information encoding in terms of the kinds of beliefs a sensory experience would make it reasonable to form.

EPISTEMIC PRESENTATIONALISM and PERCEPTUAL PRESENTATIONALISM are logically independent, but rarely prised apart.¹⁸ Beyond a tendency to generalise from paradigm cases of perceptual justification, one explanation for the natural picture’s persistent allure is non-inferential perceptual justification’s origins in a classic debate over the structure of justification. The traditional motivation for accepting non-inferential perceptual justification is a commitment to ‘foundationalism’ about justification. For old-fashioned foundationalists, usually a subject’s belief that *p* is justified only if she possesses an argument whose conclusion is *p*, and she already possesses justification to believe the argument’s premisses. Well-known regress arguments convince foundationalists that not all justification can work along these ‘inferentialist’ lines: Such justification must terminate in ‘basic beliefs’ whose justification does not depend upon whether a subject’s other beliefs are justified (and which foundationalists also traditionally regard as secure from doubt).¹⁹ To what do these basic beliefs owe their justification? Old-fashioned foundationalists argue that perception

¹⁸ One could read Munton (2016), for example, as accepting only PERCEPTUAL PRESENTATIONALISM: She argues that perceptual justification comes in degrees and, on that basis, motivates the view that perception always encodes both information about the world and confidence relations to that information.

¹⁹ Cf. Sosa (1980) on the classic debate over foundationalism.

can justify a belief without relying upon a subject's antecedent justification for some other belief, and without itself needing to be justified. In other words, perception can provide a demanding variety of 'non-inferential' justification.

Old-fashioned foundationalism is largely defunct. Many nevertheless share the foundationalist's conviction that perception possesses a distinctive epistemic role, namely a descendent of the foundationalist's regress-stopping role. These 'neo-foundationalists' therefore remain convinced that perception provides something akin to (though perhaps less demanding than) the foundationalist's non-inferential justification for beliefs about the external world.²⁰

EPISTEMIC PRESENTATIONALISM and PERCEPTUAL PRESENTATIONALISM are compelling in part because they naturally develop the basic neo-foundationalist framework. Neo-foundationalists introduce non-inferential perceptual justification to underwrite a core epistemic role. As a result, nothing in their framework forces us to posit more than one (fundamental) kind of non-inferential perceptual justification; indeed, pressure runs in the opposite direction. Similarly, perception is presumed to possess its distinctive epistemic role in part because of how it encodes information about the world. So, if you begin with the neo-foundationalist's single epistemic role, then non-inferential perceptual justification will seem to require at most one way for visual experience to encode information about the world. EPISTEMIC PRESENTATIONALISM and PERCEPTUAL PRESENTATIONALISM then appear to follow naturally, given perceptual presentation's acknowledged role in visual experience.

IV. AN ALTERNATIVE TO THE NATURAL PICTURE

The natural picture is a slippery target because the central notion of 'presentation' is schematic. For all I say here, visual experience might present properties in virtue of representing them directly, by representing conditions that they satisfy, or by placing a subject in a non-representational relation to them.²¹

I also take no immediate stand on what visual experience can present (i.e., only colours, shapes, and the like, or also kind properties, causal properties, and others equally strange).²² This non-committal understanding of 'perceptual presentation' generates a challenge. I must show that perceptual presentation, however exactly understood, cannot accommodate the epistemic significance of perceptual structure. This section answers that challenge. I use the epistemic asymmetry from Section II to undermine each component of the natural picture, and thereby establish my alternative.

²⁰ This group includes not only dogmatists like Pryor (2000) but also theorists such as Campbell (2002) and Smithies (2019) with otherwise radically divergent accounts of perception and perceptual justification.

²¹ For versions of these non-exhaustive alternatives, see Tye (1995), Chalmers (2004), and Campbell (2002).

²² Cf. Siegel (2010).

IV.1 Structural justification

The first target is

EPISTEMIC PRESENTATIONALISM: there is at most one fundamental kind of non-inferential justification for those beliefs about the external world that are formed on the basis of visual experience, namely presentation-based non-inferential justification (or ‘presentational justification’).

My core argument runs as follows. I argued in Section II that our awareness of the spatial sensory field’s boundedness can justify the belief that the region in view is a sub-region of a larger space, and that this justification is distinctive in part because it is immune to a certain kind of defeat. Presentational justification cannot be immune to this defeat. The justification provided by our awareness of the spatial sensory field’s boundedness is therefore not a variety of presentational justification. It is also non-inferential. Hence, visual experience permits a fundamental kind of non-inferential justification that is not presentational. So **EPISTEMIC PRESENTATIONALISM** is false. Instead, our awareness of the spatial sensory field’s boundedness provides what I dub ‘structural justification’:

Structural Justification: non-inferential perceptual justification that is distinct in kind from presentational justification due to immunity to sensitivity-involving defeaters.

Remember from Section II that these sensitivity-involving defeaters concern illusion, not hallucination. It is also possible to generate fine-grained varieties of structural justification by adjusting what counts as a sensitivity-involving defeater. For instance, if ‘perceptual presentation’ is understood narrowly, thereby restricting the class of sensitivity-involving defeaters, then structural justification’s scope may be comparatively wide.

The argument just outlined rests upon two new premisses: (i) presentational justification cannot be immune to the same kind of defeat as our perception-based justification for the belief that the region in view is a sub-region of a larger space; and (ii) the latter justification is non-inferential. I defend these premisses in reverse order.

In defending (ii), I do not seek to argue against global scepticism about non-inferential perceptual justification. My modest aim is to dislodge proponents of **EPISTEMIC PRESENTATIONALISM**. They already accept one variety of non-inferential perceptual justification, namely presentational justification. A standard reason for extending non-inferential perceptual justification to a given class of beliefs is that these beliefs capture how things seem to us visually. Soteriou and Richardson’s arguments for **BOUNDEDNESS** reinforce how visual experience makes the region in view seem to be a sub-region of a larger space. It is therefore hard to see how one might reasonably resist (ii). I suspect that resistance arises in part from confusing non-inferential justification with cognitively basic justification. It is perfectly consistent with the standard conception

of non-inferential justification that its presence (or at least its availability) depends in part upon sophisticated cognitive capacities on the part of agents. And it would certainly require relatively sophisticated cognitive capacities to form the belief that the region in view is a sub-region of a larger space.

I now turn to (i). Setting up the argument for (i) takes some work. Presentational justification's dependence upon perceptual presentation distinguishes it from other non-inferential justifications. What precise form this dependence takes will turn on our specific accounts of presentational justification and perceptual presentation. For instance, sometimes visual experience does not actually involve the presentation of an object (e.g. in hallucination). Can visual experience still provide presentational justification for an ordinary perception-based belief (e.g. *that is round*) in such cases? Opinions differ. By contrast, my argument requires only an uncontroversial claim about how presentational justification depends upon perceptual presentation. The claim concerns defeat. As a first pass, even if presentational justification might survive in the absence of perceptual presentation, it cannot survive if the subject discovers that her visual experience fails to involve perceptual presentation. More carefully:

(**): Visual experience cannot provide presentational justification for a belief if the subject discovers that her experience does not involve the relevant perceptual presentation.

An example brings out how this claim contrasts with the hard-line view that presentational justification requires actual perceptual presentation.

Naomi takes a bite from a brownie baked by a friend. She then notices what looks to her like a cowboy hat. A strange sight in an otherwise urban home! Yet unbeknownst to Naomi, the brownie contains a hallucinogen. When she later learns this from her friend, Naomi realises—with some relief—that she never actually saw a cowboy hat: she merely hallucinated.

Had Naomi come cross a straw hat prior to eating the brownie, her visual experience would plausibly have provided presentational justification for the belief *that's a straw hat*. What about when she seems to encounter a cowboy hat? Her experience does not involve actual presentation of an object, let alone a cowboy hat. Hardliners will insist that she lacks presentational justification for the belief *that's a cowboy hat*; moderates will argue that she retains presentational justification for the belief because it seems to her that she's seeing a cowboy hat.²³ Again, I take no side here. Instead, my claim is that once Naomi learns that she did not actually see a cowboy hat, she cannot be justified in believing *that's a cowboy hat*. Thus, even if she were justified originally, her later discovery that she was hallucinating acts as a defeater.

What would it take to know that your visual experience does not involve perceptual presentation? The sensitivity-involving condition on perceptual

²³ For the former view, see Campbell (2002) and Hellie (2011); for the latter, see Pryor (2000) and Smithies (2019).

presentation from Section I generates a sufficient condition for being in a position to know that your visual experience does not involve relevant perceptual presentation. For if perceptual presentation requires the sensitivity characteristic of perceptual presentation, then knowing that your visual experience is not appropriately sensitive to a given property will put you in a position to know that your visual experience does not present that property.

I can now argue for (i). Given (**), your visual experience provides presentational justification for a belief only if that justification would be undermined by discovery that the visual experience does not actually involve presentation of the sensible quality ascribed by the belief. A subject will be put in a position to know that her visual experience does not involve perceptual presentation if she knows that her visual experience does not involve the sensitivity characteristic of perceptual presentation. (i) then follows: presentational justification cannot be immune to the same kind of defeat as our perception-based justification for the belief that the region in view is a sub-region of a larger space. (i) follows because the defeat in question is precisely that which arises when the subject comes to know that her visual experience does not involve the sensitivity characteristic of perceptual presentation.

The difference in kind between structural and presentation justification, which I established using the asymmetry of defeat from Section II, provides resources for a deeper characterisation of structural features. Structural features are not unique in their relative invariance. Some facts about what visual experience presents remain unaffected by radical changes in which objects one sees or what those objects are like. This invariance is reflected in some of the beliefs for which sensory experience provides presentational justification. For example, presentational justification for certain general beliefs (e.g. ‘something is in view’) remains unaffected by a wide range of changes in what visual experience presents. Those who wish to cleave a sharp division between structural and non-structural features must therefore convert a mere difference of degree—greater or lesser invariance—into a genuine difference in kind. And that is precisely what structural justification permits us to do. If general beliefs are presentation-based, then they are vulnerable to sensitivity-involving defeat. For instance, while the general belief that something is in view is immune to defeat by many standard presentation-involving defeaters (e.g. those that involve changes in an object’s perceived colour, location, or identity), it is still vulnerable to sensitivity-involving defeaters that cannot undermine structural justification. The belief is vulnerable to sensitivity-involving defeat because it remains justified only if visual experience justifies some suitable non-general belief (e.g. the perceptual demonstrative belief that *that* is in view). Moving back to the perceptual level, what this vulnerability to defeat shows is that even relatively invariant aspects of perceptual presentation (e.g. that something or other seems to be in view) are grounded in the presentation of specific items. Structural features thus need not be more invariant than others. Rather, their

invariance needs only permit these features to structurally justify beliefs. Structural justification thus illuminates the invariance characteristic of perceptual structure.

IV.2 Perceptual presupposition

My target in this section is

PERCEPTUAL PRESENTATIONALISM: there is at most one way that information about the external world is encoded in visual experience which underwrites non-inferential perceptual-justification, namely the way characteristic of perceptual presentation.

I argue that PERCEPTUAL PRESENTATIONALISM is false because the asymmetry of defeat isolated in Section II—an asymmetry which also underwrites the distinction between presentational and structural justification—arises only if there is more than one way for visual experience to encode information about the world. Visual experience must not only encode information in the way characteristic of perceptual presentation but also sometimes encode information by ‘perceptually presupposing’ it.

Let’s revisit the asymmetry of defeat. A familiar type of defeat can undermine a subject’s perception-based belief that a fully perceived region (one whose boundaries fall wholly within her spatial sensory field) is a sub-region of a larger space. This defeat exploits a sensitivity characteristic of perceptual presentation: If a subject discovers that her perceptual relation to the world is defective either due to perceptual malfunction or misleading perceptual cues, then her belief is no longer reasonable. By contrast, a subject’s belief about the region in view is not susceptible to this type of defeat. Only a quite different type of defeater should force her to abandon that belief.

It is not just that only one of the subject’s beliefs remains reasonable. Rather, this asymmetry is clear even *from her perspective*. Upon acquiring knowledge of her defective perceptual situation, the subject is in a position to recognise that only one of her beliefs is still reasonable to maintain. This sensitivity to the asymmetry of defeat requires explanation. It cannot be explained by a difference in what her beliefs are about. Each belief treats a perceived region as a sub-region of a larger space. Instead, the subject’s sensitivity to the asymmetry of defeat must trace to an apparent difference within how perception permits her access to the outside world. And this appearance must in turn be constituted by a difference within perception’s ‘encoding’ of information about the world. The subject formed one belief in response to perceptual presentation of a sub-region’s relation to the space around it; she formed another in response to her awareness of the boundedness of her spatial sensory field. It is this distinction between two aspects of her visual experience—perceptual presentation and awareness of her spatial sensory field’s boundedness—that reflects the division

between two ways for the subject's visual experience to encode information about the outside world.

Not just any difference between ways of encoding information undermines PERCEPTUAL PRESENTATIONALISM. PERCEPTUAL PRESENTATIONALISM permits distinctions between how visual experience and non-visual experience encode information, as well as divisions within the way of encoding information characteristic of perceptual presentation. Yet our subject's beliefs were formed in response to visual experience. And, more importantly, the asymmetry of defeat concerns the invulnerability of a perception-based belief to the defeat characteristic of perceptual presentation. So the requisite distinction between ways for perception to encode information about the external world cannot be understood as a division within the way of encoding information characteristic of perceptual presentation. Hence, the subject's sensitivity to the asymmetry of defeat—and thus our sensitivity to it—must be a response to a way for visual experience to encode information that is genuinely distinct from the way of encoding information characteristic of perceptual presentation.

This argument against PERCEPTUAL PRESENTATIONALISM also doubles as a case for my positive claim that visual experience encodes information about the world in a way that is distinct from the way characteristic of perceptual presentation. In particular, visual experience can distinctively encode that the region in view is a sub-region of a larger space. I dub this way of encoding 'perceptual presupposition':²⁴

Perceptual Presupposition: A perceiver's sensory experience counts as 'perceptually presupposing' p iff the perceiver would be non-inferentially justified in believing that p in response to her experience, and this justification does not depend upon what her experience presents (or appears to present).

Visual experience perceptually presupposes that the region in view is a sub-region of a larger space, and does so because it involves awareness of the spatial sensory field as bounded by our sensory limitations. However, my argument leaves open what else sensory experience might perceptually presuppose. Perceptual presupposition's scope depends in part on perceptual presentation. If sensory experience presents a wide range of properties, then a broad class of beliefs potentially possess non-inferential presentation-based justification; if sensory experience presents only a narrow range of properties, then few beliefs potentially possess non-inferential presentation-based justification. The former option permits a more restricted scope for perceptual presupposition, while the latter potentially allows a broad swath of information to count as perceptually presupposed.

²⁴ My use of 'perceptual presupposition' recognises a parallel to a neglected strand in Strawson's seminal discussion of linguistic presupposition (cf. Strawson 1950: 330).

Perceptual presupposition's scope also depends upon the range of beliefs that admit non-inferential justification by sensory experience. In each new case, a central challenge is to establish that a given belief can be non-inferentially justified by perception. Soteriou and Richardson's arguments for BOUNDEDNESS are important in part because they motivate an extension of visual phenomenology to include our awareness of the region in view as a sub-region of a larger space. Adding this complexity to visual phenomenology generates a corresponding extension of non-inferential justification's scope—or so I have argued—since our primary guide to the latter is our sensory phenomenology.

In addition to these scope questions, perceptual presupposition and structural justification generate hard questions about cases that suddenly threaten to fall in the middle. For instance, suppose you glimpse a carpeted room through a door and believe on that basis that the as-yet-unseen parts of the room are similarly carpeted. Maybe your belief is merely inferentially justified. However, if it is not, then your belief's content is plausibly neither wholly perceptually presupposed nor encoded wholly by perceptual presentation. Any non-inferential perceptual justification for the belief would depend both on what is perceptually presented (e.g. the visible carpeting) and on what is perceptually presupposed (i.e., that the region in view is a sub-region of a larger space). We then face a choice between a few positions:

1. Perceptual presentation and perceptual presupposition do not exhaust the ways for perception to encode information about the external world (and so presentational and structural justification are plausibly not exhaustive). Instead, there is a form of perceptual encoding that incorporates at least both what is perceptually presented and what is perceptually presupposed.
2. Perceptual presentation and perceptual presupposition are exhaustive, but either (i) presentational justification can depend upon perceptual presupposition or (ii) visual experience supports a form of non-inferential justification that is neither structural nor presentational.

This choice illustrates the kinds of nice but hard questions that arise once we introduce perceptual presupposition and structural justification.

V. CONCLUSION

Much work has already been done on structural features' implications for sensory phenomenology. Less has been written about their epistemic import. I have shown that reflection upon the epistemic significance of visual experience's structural features expands our conception of perception's contribution to our ability to know about the world. I shall close by highlighting some consequences of this expansion.

When a subject believes that the region in view is a sub-region of a larger space—a belief reasonably formed in response to her visual experience—her belief counts as justified because her visual experience involves awareness of the spatial sensory field’s boundedness. This belief is immune to defeaters that exploit the sensitivity characteristic of perceptual presentation. There is thus a deep asymmetry of defeat between ordinary perception-based beliefs and the perception-based belief that the region in view is a sub-region of a larger space. This asymmetry forces us to abandon a pair of widely held assumptions that together frame an otherwise compelling neo-foundationalist picture about how perception permits us to learn about the external world: **EPISTEMIC PRESENTATIONALISM** and **PERCEPTUAL PRESENTATIONALISM**. Rejecting these assumptions, and the picture they frame, is a major advance.

Another is the introduction of a new type of non-inferential perceptual justification to accommodate the asymmetry of defeat: *structural justification*. On the one hand, it permits us to expand the range of beliefs that count as basic (in the neo-foundationalist’s sense) without widening the scope of perceptual presentation. And that is significant not only for those who defend some descendant of foundationalism about justification but also for those who wish to retain a robust epistemic role for visual experience while restricting perceptual presentation to a narrow range of properties. On the other hand, the introduction of structural justification poses a challenge for extant neo-foundationalist accounts of non-inferential perceptual justification. These accounts are designed to accommodate what I dubbed ‘presentational justification’. They therefore risk falling short not only as general accounts of non-inferential perceptual justification but also as accounts of what it is about perception that permits perceptual justification to play its seemingly distinctive epistemic role.

The asymmetry of defeat from Section II also introduces a new way for visual experience to encode information: *perceptual presupposition*. There are nice questions about how perceptual presupposition can be made to fit with accounts of sensory experience that are built around perceptual presentation. Perceptual presupposition’s introduction also reshapes the longstanding debate over which properties visual experience puts us into perceptual contact with. Objections to the inclusion of specific properties—e.g. kind properties—often undermine only the view that these properties are perceptually presented by visual experience. And, on the other side, arguments that purport to show that visual experience presents non-traditional properties risk establishing no more than that we perceive these properties, since perceptual presentation and perceiving no longer necessarily coincide.²⁵

²⁵ For written feedback or discussion that improved the paper, thanks to Fatema Amijee, Mike Martin, Imogen Dickie, Alex Moran, Kevin Lande, Tashi Alford-Duguid, Henry Clarke, Clare Mac Cumhail, Bill Brewer, Umut Baysan, Rachel Fraser, Nick Hughes, Lea Salje, Peter Epstein, Ori Beck, Maarten Steenhagen, Dave Suarez, Anthony Bruno, Ivan Ivanov, Nick Shea,

REFERENCES

- Brewer, B. (2020) 'Perception of Continued Existence Unperceived', *Analytic Philosophy*, 30/1: 24–38.
- . (2021) 'The Objectivity of Perception', *Proceedings of the Aristotelian Society*, 121/1: 1–20.
- Burge, T. (2003) 'Perceptual Entitlement', *Philosophy and Phenomenological Research*, 67/3: 503–48.
- Campbell, J. (2002) *Reference and Consciousness*. Oxford: Oxford University Press.
- . (2005) 'Information-Processing, Phenomenal Consciousness and Molyneux's Question', in J. Bermudez (ed.) *Thought, Reference and Experience: Themes from the Philosophy of Gareth Evans*. Oxford: Oxford University Press.
- Campbell, J. and Cassam, Q. (2014) *Berkeley's Puzzle*. Oxford: Oxford University Press.
- Cavedon-Taylor, D. (2021) 'Sensorimotor Expectations and the Visual Field', *Synthese*, 198: 3991–4006.
- Chalmers, D. (2004) 'The Representational Character of Experience', in B. Leiter (ed.) *The Future for Philosophy*. Oxford: Oxford University Press.
- Evans, G. (1980) 'Things Without the Mind', in Z. Van Straaten (ed.) *Philosophical Subjects*. Oxford: Clarendon Press.
- . (1985) 'Molyneux's Question', in A. Brodersen (ed.) *Collected Papers*. Oxford: Oxford University Press.
- Hellie, B. (2011) 'There It Is', *Philosophical Issues*, 21: 110–64.
- Huemer, M. (2001) *Skepticism and the Veil of Perception*. Lanham: Rowman & Littlefield
- Johnston, M. (2006) 'Better than Mere Knowledge? The Function of Sensory Awareness', in J. Hawthorne and T. Gendler (eds) *Perceptual Experience*. Oxford: Oxford University Press.
- Kant, I. (1998) *Critique of Pure Reason*. translated by P. Guyer and A. W. Wood (eds) Cambridge: Cambridge University Press.
- Laasik, K. (2019) 'Visual Field and Empty Space', *European Journal of Philosophy*, 27/2: 403–11.
- Lande, K. (2021) 'Mental Structures', *Nóus*, 55/2: 649–77.
- Lewis, D. (1980) 'Veridical Hallucination and Prosthetic Vision', *Australasian Journal of Philosophy*, 58/3: 239–49.
- McGrath, M. (2018) 'Looks and Perceptual Justification', *Philosophy and Phenomenological Research*, 96/1: 110–33.
- Mac Cumhail, C. (2015) 'Perceiving Immaterial Paths', *Philosophy and Phenomenological Research*, 90/3: 687–715.
- Martin, M. G. F. (1992) 'Sight and Touch', in T. Crane (ed.) *The Contents of Experience*. New York, NY: Cambridge University Press.
- . (1995) 'Bodily Awareness: A Sense of Ownership', in J. Bermudez, A. Marcel and N. Eilan (eds) *The Body and the Self*. Cambridge, MA: MIT Press.
- . (2006) 'On Being Alienated', in T. Gendler, and J. Hawthorne (eds) *Perceptual Experience*. Oxford: Oxford University Press.
- Melnick, A. (1973) *Kant's Analogies of Experience*. Chicago, IL: University of Chicago Press.
- Munton, J. (2016) 'Visual Confidences and Direct Perceptual Justification', *Philosophical Topics*, 44/2: 301–26.
- Pryor, J. (2000) 'The Skeptic and the Dogmatist', *Nous*, 34: 517–49.
- . (2005) 'There Is Immediate Justification', in M. Steup and E. Sosa (eds) *Contemporary Debates in Epistemology*. London: Blackwell.
- Richardson, L. (2010) 'Seeing Empty Space', *European Journal of Philosophy*, 18/2: 227–43.
- Siegel, S. (2010) *The Contents of Visual Experience*. Oxford: Oxford University Press
- Silins, N. (2011) 'Seeing through the "Veil of Perception"', *Mind*, 120/478: 329–67.

Matt Parrott, and Anil Gomes. Thanks also to audiences at Oxford, Cambridge, King's College London, the New College of the Humanities, the 2018 CPA Congress, the 2018 Joint Session, and the 2019 Pacific APA Meeting. I also owe a special debt to Matt Soteriou for his invaluable feedback and encouragement, as well as for the book that inspired this paper. Finally, most of the work on this paper was completed during my stint as a Leverhulme Early Career Fellow at Oxford, and so I gratefully acknowledge the support of the Leverhulme Trust (ECF-040-2018).

- Smithies, D. (2019) *The Epistemic Role of Consciousness*. Oxford: Oxford University Press.
- Soteriou, M. (2013) *The Mind's Construction*. Oxford: Oxford University Press.
- Sosa, E. (1980) 'The Raft and the Pyramid: Coherence versus Foundations in the Theory of Knowledge', *Midwest Studies in Philosophy*, 5/1: 3–25.
- Strawson, P. F. (1950) 'On Referring', *Mind*, 59/235: 320–44.
- . (1959) *Individuals*. London: Methuen.
- . (1980) 'Reply to Evans', in Z. Van Straaten (ed.) *Philosophical Subjects*. Oxford: Clarendon Press.
- Textor, M. (2019) 'Perceptual Objectivity and the Limits of Perception', *Phenomenology and the Cognitive Sciences*, 18/5: 879–92.
- Tye, T. (1995) *Ten Problems of Consciousness*. Cambridge, MA: MIT Press.
- Wright, C. (2007) 'The Perils of Dogmatism', in S. Nuccetelli and G. Seay (eds) *Themes from G. E. Moore*. Oxford: Oxford University Press.

University of British Columbia, Canada