

Naïve realism with many fundamental kinds

Neil Mehta, Yale-NUS College

This is a manuscript. Please cite the final version published in Acta Analytica.

Naïve realism is a theory of perception with great explanatory ambitions. According to its advocates, this theory can explain the distinctively presentational phenomenology of perception.¹ It can also explain why perceptions put us in a position to make singular reference to, and to acquire singular knowledge about, what we perceive.² It is primarily on these grounds that so many philosophers have been attracted to naïve realism (§1).

Michael Martin has argued that, in order to make good on these explanatory ambitions, the naïve realist should adopt an additional commitment about what *fundamental kind of thing* any perception is. More than a dozen naïve realist works have joined him in taking this commitment on board (§2).³ I think, however, that adopting this commitment does not particularly help the naïve realist to realize her explanatory ambitions, and so is not warranted (§4–§9). This result is significant because by relinquishing this commitment about fundamental kinds, we make room for the development of some new and surprising forms of naïve realism (§3).

1. Naïve realism and its explanatory ambitions

To perceive is to take in the world with your senses, so to speak. Perceptions are certainly to be distinguished from total hallucinations, while the precise relationship between perceptions and illusions is a matter of some controversy.

Perceptions have a number of distinctive features. Here are three:

1. *Presentational character.* Suppose that I perceive a mango. Later, I form a singular belief about it; later still, I visually imagine the mango. In all three cases, I am aware of the mango. But there is an obvious difference between my perception, on the one hand, and my singular belief and visual imagining, on the other hand: my perception introspectively seems to just *present* the mango to me – it seems to make me aware of the mango in some peculiarly direct way. Indeed, quite generally, perceptions have *presentational character*: they introspectively seem to put us in a peculiarly direct relation, *presentation*, to mind-independent objects.⁴

¹ See Alston (1999); Martin (2002); Crane (2006); Hellie (2007); Kennedy (2009) and (2013); Nudds (2009); Fish (2009); Brewer (2011, p. 2); Genone (2016, §4.3).

² The seminal work here is Campbell (2002). Similar ideas appear in Fish (2009, pp. 26–8); Brewer (2011, ch. 6); Genone (2016, §4.1).

³ See Martin (2004) and (2006); Soteriou (2005, p. 178); Crane (2006, p. 139); Neta (2008, pp. 311–312); Nudds (2009, p. 337); Brewer (2011, p. 94); Logue (2012a, p. 211), (2012b, p. 174), and (2013, p. 109); French (2013, p. 1735); Genone (2016, p. 7); Gomes (2017, p. 534); Beck (2018, p. 2).

⁴ For naïve realists interested in explaining this datum, see fn. 1. Note that hallucinations also have presentational character.

2. *Reference-enabling character*. When I perceive a novel mind-independent object or property-instance – a novel fruit, animal, scent, or texture – I acquire a new ability: I acquire the ability to form singular thoughts about *it*.⁵ Similarly –

3. *Knowledge-enabling character*. When I perceive a novel mind-independent object or property-instance, I typically acquire a new ability: absent defeaters, I acquire the ability to gain singular knowledge about *it*.⁶

Now consider:

Core posit of naïve realism: Any perception involves a relation of *acquaintance* – i.e., a relation of direct, non-representational awareness to mind-independent objects and their property-instances.⁷

Sometimes the core posit of naïve realism is formulated more narrowly as a claim about the *phenomenal character* of perception. You are welcome to use such a formulation of the posit, and of all of this paper’s further claims about naïve realism, if you like.⁸

One crucial feature of naïve realism is that it treats perceptions as *relations*. This distinguishes naïve realism at least from pure versions of representationalism, which treat perceptions purely as *representations*. A representation is something that can be correct/accurate or incorrect/inaccurate. Thus, it is possible to represent something that does not exist. By contrast, a relation is something that cannot go wrong, so to speak: you can stand in a relation to something only if it exists.

With this distinction in hand, it is clear enough, at least in outline, how naïve realists can explain the above data about perception. In particular, the naïve realist can say that perceptions have presentational character, while singular thoughts or sensory imaginings do not, because perceptions are *relational*, while singular thoughts and visual imaginings are not – they are merely representational.⁹ The naïve realist can add that because perception is relational, it can ground our capacities to form singular thoughts about, and to acquire singular knowledge of, the mind-independent objects and property-instances that we perceive.¹⁰

It is primarily on the basis of these explanatory powers that naïve realism has become such an influential theory of perception.

2. Monistic naïve realism introduced

⁵ For naïve realists interested in explaining this datum, see fn. 2.

⁶ For naïve realists interested in explaining this datum, see fn. 2.

⁷ For advocates of naïve realism, see fns. 1-3.

⁸ I explain why I prefer to avoid this narrower formulation in [Author’s Work B], but the issue will not affect the arguments of this paper.

⁹ Naïve realists still must explain why hallucinations have presentational character, but there are many standard strategies for doing this. I mention a new strategy later in this paper.

¹⁰ Many would argue that naïve realism is not *uniquely* positioned to explain these data – see for example [Author’s Work A]; Schellenberg (2018). A few would argue that naïve realism is not even particularly well-positioned to explain some of these data – see for example Millar (2014). Here I bracket such concerns.

There is something noteworthy about the naïve realist account that I have sketched so far: it does not say anything about fundamental kinds. It makes claims about explanations, relations of direct awareness, and mind-independent objects and property-instances, but the words “fundamental” and “kind” have not made an appearance.

Nevertheless, it has become utterly standard for naïve realists to commit to the following theory:

Monistic naïve realism: Every perception belongs to *exactly one fundamental kind*. This fundamental kind includes only entities that involve relations of acquaintance.

This theory, or something very close to it, is endorsed in more than a dozen naïve realist works.¹¹ But why should a naïve realist accept it? Indeed, what is a fundamental kind in the first place?

The most influential answers to these questions come from Michael Martin. As we will learn, Martin understands a fundamental kind as a *maximally specific essential kind*. He argues that naïve realists should accept monistic naïve realism both to avoid *trivializing* their theories and to escape a certain *screening-off argument* against naïve realism.¹² I will begin my critical discussion of monistic naïve realism by scrutinizing these ideas and related ones (§4-§7).

Many other naïve realists adopt monism without any explanation (besides, perhaps, a citation to Martin), so I presume that they accept the thrust of Martin’s reasoning.¹³ One notable fact, however, is that unlike Martin, some of these same naïve realists have understood fundamental kinds as maximally specific kinds, *whether or not these kinds are essential*.¹⁴ These naïve realists have not explicitly justified their deviation from Martin’s preferred notion, but they do cite Martin, so perhaps they think that his arguments better support their version of monistic naïve realism. I will therefore weave my criticisms of this idea into my discussion of Martin.

There is another approach, due to Heather Logue, that deserves separate consideration. Logue takes inspiration from Martin in many ways – for example, she is sympathetic towards monistic naïve realism. However, unlike Martin, she understands fundamental kinds as *basic psychological kinds*,¹⁵ and she offers new reasons for preferring monistic naïve realism. So I will end my critical discussion by examining Logue’s ideas, along with a few others (§8-§9).

Before I get to any of that, however, let me explain what is at stake in the discussion.

¹¹ See the theorists cited in fn. 3.

¹² See Martin (2004, pp. 61-62) and (2006, pp. 360 and 369-372).

¹³ See especially French (2013, p. 1735); Genone (2016, p. 7); Beck (2018, p. 2).

¹⁴ See Crane (2006, p. 139); Nudds (2009, p. 337); Brewer (2011, p. 3).

In addition, other theorists – in particular, Soteriou (2005, p. 178); Gomes (2017, p. 534) – understand fundamental kinds as maximally specific *phenomenal* kinds. These theorists also do not explain why naïve realists should be monistic naïve realists, but I think that it is easy enough to fill in some good reasoning here. So I will set this view aside.

¹⁵ See Logue (2012b, p. 174); (2013, p. 109).

3. Naïve realism without monism

The thesis of this paper is that naïve realists have no good reason to accept monistic naïve realism – again, the view that every perception belongs to exactly one fundamental kind, where this kind includes only entities that are constituted by acquaintance relations. The most that naïve realists have good reason to accept is *moderate naïve realism*, the view that any perception essentially belongs to a metaphysically robust kind that includes only entities that involve acquaintance relations. What is the significance of this point?

My most basic answer is that it is *very hard* to develop a good theory of perception. For example, naïve realism has some deeply appealing features, but it is very hard to couple this theory with a satisfying account of hallucinations. Perhaps the problem is that naïve realists have been working in a straitjacket, by accepting unnecessary commitments about fundamental kinds.

What unnecessary commitments, precisely? Here is one possibility. Like many other theorists of perception, the monistic naïve realist accepts the following assumption:

Monism: Every perception belongs to exactly one fundamental kind.

Once we make this assumption, we must make a choice between the following two theories:

The common kind theory: For any perception and any matching hallucination, there is a fundamental kind that they both belong to.

Disjunctivism: For any perception and any matching hallucination, there is a fundamental kind that the perception belongs to but the hallucination does not.

Unsurprisingly, once forced to choose, naïve realists have almost all ended up on the side of disjunctivism.

To my mind, however, the choice between disjunctivism and the common kind theory is a bitter one. After all, perceptions and matching hallucinations seem to be importantly similar in some respects, but also importantly different in other respects. (Note well: you can accept this point even if you disagree with my first answer about what these particular similarities and differences are!) So it seems worthwhile to at least *explore the possibility* of explaining the common features via a common fundamental kind, while explaining the distinctive features of perception via a fundamental kind that is unique to that case. But we must reject monism to make room for such a theory.

To bring this point into focus, I would like to sketch a specific version of non-monistic naïve realism, which I will call *pluralistic naïve realism*. As it happens, I believe that pluralistic naïve realism is the best version of naïve realism. But I will not make that case here. I will, however, use this theory later as a concrete illustration for my ideas.

Pluralistic naïve realism posits two very different kinds of awareness. The first is *deep awareness*. This is a non-representational form of awareness whose

targets¹⁶ are always universals,¹⁷ and it reveals a substantial part – not necessarily all! – of their essences. The second kind of awareness is *acquaintance*. Like deep awareness, acquaintance is non-representational. But its targets are always particulars, and it does not reveal any part of their essences.

Consider a perception of a red apple and a mere hallucination of one. According to pluralistic naïve realism, the perception of the apple fundamentally involves deep awareness of universals such as redness, *together with* acquaintance with an ordinary, mind-independent apple. By contrast, the hallucination fundamentally involves only deep awareness of universals such as redness. Acquaintance is not involved at all. Thus, the pluralistic naïve realist accepts the common kind theory, as defined above: she thinks that any perception and any matching hallucination fundamentally involve deep awareness of exactly the same universals. But she also accepts disjunctivism, as defined above: she thinks that any perception fundamentally involves acquaintance with particulars, while any matching hallucination does not involve acquaintance with anything.

Elsewhere, I lay out many motivations for pluralistic naïve realism in detail.¹⁸ For now, I will mention just one cluster of them.

Recall the familiar example of the brilliant color scientist, Mary, who has never experienced anything red.¹⁹ Consider two different ways that she might first have such an experience. In the first and more familiar version of the story, one day Mary actually *sees* a red apple. In the second version of the story, Mary merely *hallucinates* a red apple: her brain is directly stimulated, via implanted electrodes, so as to put her in precisely the same internal condition she would have occupied upon actually seeing a red apple.

Pluralistic naïve realism lets us capture many plausible claims about these two experiences.

To begin with, it is plausible that seeing the red apple positions Mary to know a great deal about *what it is* to be red – e.g., that it lies in the nature of red to be more similar to orange than to green. It is plausible that hallucinating the apple positions Mary to gain the very same knowledge. The pluralistic naïve realist honors this intuition: she says that hallucinations and perceptions both involve deep awareness of redness, a form of awareness that is partly essence-revealing. There are, by the way, many serious objections against the bolder claim that seeing red lays bare the *complete* essence of redness.²⁰ By saying that deep awareness need not reveal the *entire* essence of universals, the pluralistic naïve realist avoids these objections.

Now consider how Mary's perception of the apple positions her with respect to the apple itself. On the one hand, her experience positions her to make

¹⁶ I say “targets” where most would say “objects,” since the latter term is also used to refer to

¹⁷ See [Author's Work A]; [Author's Work C]. I take from Johnston (2004), and subsequently Conduct (2012), the idea that perceptions and hallucinations both involve awareness of universals, while only perceptions involve awareness of particulars. Where I disagree with Johnston is about his view that perception involves only a single kind of awareness of universals *and* particulars. For reasons that will be made clear very shortly, I think that these are two quite different kinds of awareness.

¹⁸ See [Author's Work C].

¹⁹ Of course, the original story was introduced in Jackson (1982).

²⁰ Sympathy for the bold claim is expressed in Johnston (1992). Objections to the bold claim appear in McLaughlin (2003, p. 98); Allen (2011, pp. 163-167); Broi (2020).

singular reference to *that particular object*. It also positions her to know many *accidental* facts about that particular apple – e.g., that it is red. At the same time, seeing an apple surely does not position Mary to know anything in particular about *what it is* to be an apple. For all that Mary’s perception reveals, what it is to be an apple is to be an ostrich egg. The pluralistic naïve realist explains these facts by saying that Mary is merely *acquainted* with an apple, where acquaintance is a relation that does not reveal any part of the essence of its targets. Acquaintance *does* permit singular reference and singular knowledge, however, which explains why Mary’s perception positions her to refer to, and come to know about, *that apple* in particular.

Finally, Mary’s hallucinatory experience does not position her to make singular reference to, or acquire singular knowledge about, any particular apple. Why not? The pluralistic naïve realist answers that merely hallucinating an apple does not involve acquaintance with any particular apple. It merely seems to do so. (Why that misleading appearance? I will suggest an answer later.)

At this point, you might worry that pluralistic naïve realism does not compare well to a certain monistic rival, which I will call *sensing naïve realism*. The sensing naïve realist posits just one fundamental kind of awareness, call it *sensing*. She says that we can sense particulars and universals. However – and this is the crucial twist – she suggests that there is something in the nature of properties such that when we sense *them*, parts of their very essences are revealed to us.²¹ How does sensing naïve realism stack up against pluralistic naïve realism?

Sensing naïve realism delivers many of the same predictions as its rival. For instance, it predicts that seeing or hallucinating a red apple will reveal part of the essence of redness, but not of any particular apple; that seeing an apple will position the subject to make singular reference to, and gain singular knowledge about, that particular apple; and that hallucinating an apple will not position the subject to make singular reference to, and gain singular knowledge about, any particular apple. At the same time, sensing naïve realism is clearly the simpler view. It relies on just one fundamental posit; the pluralistic naïve realist relies on two. I acknowledge that simplicity is a great theoretical advantage.

But the view has some potential disadvantages, too. Here is one. I find it phenomenologically plausible that when I see a pine tree, I *perceive* it as a pine tree.²² The most natural way for the sensing naïve realist to accommodate this claim is to say that I sense the property of being a pine tree. Thus she should predict that my perception reveals to me part of the very nature of the property of being a pine tree. This is not the case: for all I can know just on the basis of seeing a pine tree, what it is to be a pine tree is to be an alien spaceship. The pluralistic naïve realist can do better: she can say that seeing a pine tree merely *acquaints* me with the property of being a pine tree. Thus, even though I perceive that property, nothing of its essence is revealed to me.

This is the beginning of the debate, not the end of it.²³ Still, I submit that pluralistic naïve realism is at least worth *considering*. But this view, though compatible with moderate naïve realism, is incompatible with monistic naïve

²¹ Sensing naïve realism is similar to the view of Johnston (2004).

²² For some arguments to this effect, see Siegel (2006); Bayne (2009); Masrour (2011).

²³ I continue the discussion in [Author’s Work C].

realism. That is another reason that it is worthwhile to compare the latter two brands of naïve realism.

Looking forward. The focal question of this paper is whether naïve realists should accept monism. If not, then there is a flowering of possibilities for the naïve realist – for instance, she can combine disjunctivism with the common kind theory, and she can endorse pluralistic naïve realism.

Now that we understand why the debate matters, it is time to engage in it. The next step is to consider the most popular way of understanding what a fundamental kind is.

4. Martin: Fundamental kinds as maximally specific essential kinds

Here is how Martin introduces the notion of a fundamental kind:

... I will assume the following: for all [objects and events] there is a most specific answer to the question, ‘What is it?’ In relation to the mental, and to perception in particular, I will assume that for mental episodes or states there is a unique answer to this question which gives *its most specific kind*; it tells us what *essentially* the event or episode is.²⁴

There are two importantly different interpretations of this passage. According to the first interpretation, what Martin is assuming is that (i) for any entity *E* – at least if *E* is a mental entity – there is a single maximally specific kind *K* to which *E* belongs, and (ii) whatever *K* is, *E*’s belonging to *K* is essential to *E*. *K* is then the *fundamental kind* to which *E* belongs.²⁵

This assumption can be decomposed into three claims. First, that every entity belongs to *at least one* maximally specific kind. Second, that every entity belongs to *at most one* maximally specific kind. Third and finally, that for any maximally specific kind to which an entity belongs, belonging to that kind is essential to it.

Others have expressed skepticism about the first and second of these claims.²⁶ But, to my mind, it is the third claim that is most problematic. For consider the famous philosopher Hypatia. Perhaps she belongs to the kind *woman*, understood as a biological kind corresponding to the property of *being a human female adult*.²⁷ Perhaps we should regard the property *woman* as corresponding to a kind because there is some important explanatory work done by the entire property of being a human female adult. This important work might even go beyond the conjunction of the important work done by the properties of being human, being female, and being an adult. And perhaps *woman* is the maximally specific kind to which Hypatia belongs. But belonging to that kind was not

²⁴ Martin (2006, p. 361, emphases mine). Martin acknowledges a debt to Wiggins (1980) and (2001). On the notion of essence, see Fine (1994).

²⁵ This interpretation is suggested in Logue (2013, p. 109).

²⁶ Regarding the first claim, Pautz (2007, p. 528) worries that particles of dust do not belong to at least one maximally specific kind. Regarding the second claim, Mehta (2014, p. 327) worries that a single perception might belong to several maximally specific but “cross-cutting” kinds. This worry relies on the idea that the conjunction of two kinds does not always yield a third kind.

²⁷ Here the property of being female is to be understood as a sex property, not as a gender property.

essential to her: Hypatia was still Hypatia when she was a child and so did not belong to the kind *woman*.

Thus, if this is what Martin means to assume, then I believe that his assumption is simply false.

But there is another way to interpret the passage from Martin. Let us say that an *essential kind* is any kind that an entity belongs to essentially. Then we might instead read Martin as assuming that every entity belongs to at least one essential kind, with one of these essential kinds being more specific than the rest. This maximally specific essential kind is the entity's *fundamental kind*.

When read in this way, Martin's assumption has the virtue of not mentioning non-essential kinds at all. For this reason, the assumption does not entail the first or the third of the potentially problematic claims just mentioned.²⁸ It does not entail (first) that every entity belongs to *at most one* maximally specific kind, full stop. For it allows that there might be a maximally specific essential kind and a non-essential kind, neither of which is more specific than the other. More importantly, this assumption does not entail (third) that for any maximally specific kind to which an entity belongs, belonging to that kind is essential to it. For this reason, I shall henceforth adopt this interpretation of Martin's assumption.

The assumption *is* still substantial. For even when we focus just on the kinds that an entity belongs to essentially, it is not immediately obvious why one of these kinds must be more specific than all the rest. Why can't an entity belong to several maximally specific essential kinds? This question becomes pressing once we realize that the intersection of two kinds need not itself be a kind. Just think of the properties *being rational* and *being made of matter* (as opposed to anti-matter): each of these properties corresponds to a kind, but presumably there is not a further kind associated with the conjunctive property of *being a rational being made of matter*. The naïve realist who makes Martin's assumption is committed to thinking that things *must* be different when it comes to essential kinds. That is a real cost!

But Martin and others have given several arguments for bearing this cost. Let us see whether any of these arguments is sound.

5. The anti-triviality argument for monistic naïve realism

The core idea of Martin's first argument is that naïve realists should accept monistic naïve realism *in order to ensure that their preferred explanations end up being non-trivial*. As Martin puts it:

[1] Naïve realism can be preserved only at the expense of denying the Common Kind Assumption.... [2] There are ways of construing the Common Kind Assumption on which it comes out as trivially false. [3] If we relax our conception of a kind of event sufficiently then any description of an event mirrors a kind of event.... [4] For the Common Kind Assumption to be a non-trivial falsehood, therefore, we need some conception of the privileged descriptions of experiences. [5] For it to be a substantive matter that perceptions fail to be the

²⁸ Technically, it does not even entail the second problematic claim that every entity belongs to *at least one* maximally specific kind, full stop. For it allows that an entity might belong to an infinite series of ever-more specific maximally specific kinds. However, this does not help with Pautz's worry, as described in fn. 26.

same kind of mental episode as illusions or hallucinations, we need some characterisations of events which reflect their nature or what is most fundamentally true of them (2006, p. 360).

I have numbered Martin's remarks for ease of reference.

I would unpack this argument – which I will call the *anti-triviality argument* – as follows. Start by observing that the naïve realist is trying to explain various facts about perception by giving a non-trivial theory about the metaphysical character of perceptions – by giving a non-trivial theory about what kind of entity a perception is ([2] and [4]). For example, naïve realism stands in contrast to the *common kind theory*, according to which a perception and a matching hallucination are the same kind of thing ([1]).

Now suppose, for the sake of *reductio*, that there is a kind corresponding to every description, no matter how arbitrary. Then there will be infinitely many arbitrary descriptions, and hence infinitely many kinds, bundling perceptions together with hallucinations. There will also be infinitely many arbitrary descriptions, and hence infinitely many kinds, splitting perceptions apart from hallucinations. So any theory about what kind of entity a perception is will be trivial ([3]).

This contradicts the premise that naïve realism is a non-trivial theory on precisely this topic, so – the argument runs – the naïve realist must jettison her assumption: she should conclude that there are not kinds corresponding to every description ([4]); she should conclude that the debate is over fundamental kinds, i.e., maximally specific essential kinds ([5]). For purposes of the discussion below, I especially want to emphasize that Martin's argument moves directly from [4] to [5].

When I consider this argument, I notice two superficial problems and one deep problem.

The first superficial problem is this. In order to realize her explanatory ambitions, the naïve realist must say that the relation of acquaintance explains the presentational character, reference-enabling character, and knowledge-enabling character of perceptions. But she can say this without ever speaking about *kinds*. So it is not yet clear why the naïve realist should commit to any notion of a kind, let alone a metaphysically robust notion of a kind.

In fact, however, I think that there is a compelling (though not irresistible) reason for the naïve realist to commit to a metaphysically robust notion of kinds. In particular, reflect on the contrast between the following pairs of properties:

- (i) The property of *having presentational character* vs. the conjunctive property of *having presentational character and being distinct from the planet Mars*.
- (ii) The property of *being reference-enabling* vs. the disjunctive property of *being reference-enabling or being identical to a piranha*.
- (iii) The property of *being knowledge-enabling* vs. the trivial property of *being such that $2 + 2 = 4$* .

Perceptions have all of the properties on this list. But I hope that you will find it obvious that, for each of these pairs of properties, the first property is *metaphysically robust*, while the second property is *metaphysically arbitrary*. If you are a

metaphysician, then you might call the first property in each pair *natural*, *structural*, or *joint-carving*, and you might call the second property in each pair *gruesome*.²⁹

So, if the naïve realist accepts this distinction, then she is committed to thinking that the property of involving a relation of acquaintance explains not just why perceptions have *many* properties, but why perceptions have many *metaphysically robust* properties. And I find it very reasonable to think that if a property *F* explains the existence of many other metaphysically robust properties, then property *F* is itself metaphysically robust. From there, I find it reasonable to think that if certain entities share a metaphysically robust property *F*, then they thereby belong to a metaphysically robust kind – i.e., a natural, structural, or joint-carving kind – that is characterized by property *F*.

Thus I agree with Martin that naïve realists have a good reason to accept that every perception belongs to a metaphysically robust kind that includes only entities that involve relations of acquaintance.

But there is a second problem that has been raised for the anti-triviality argument: why should the naïve realist say anything about what metaphysically robust kind(s) a perception belongs to *essentially*? What difference would it make if a perception were to belong to one or more of these metaphysically robust kinds *accidentally*?³⁰

Again, however, I think that there is a reasonable response to this problem. This is that the features of perception that naïve realists want to explain – such as its presentational, reference-enabling, and knowledge-enabling character – seem to be *necessary* features of perception. For example, it seems that, *necessarily*, any perception must be such that it puts the subject in a position to form singular concepts of whatever object is perceived. Of course, pigs and dogs might lack the cognitive capacities required to actually form such concepts, but that seems to be a *later* failure, so to speak. It seems that the *perception itself* offers the materials necessary for singular reference, even if many perceiving creatures cannot avail themselves of that opportunity. And if we want to explain why perceptions are *necessarily* presentational, reference-enabling, and knowledge-enabling, then it strikes me as at least reasonable to hypothesize that it is *essential* to any perception to involve an acquaintance relation.

Thus I think that naïve realists might well have a good reason to accept a theory that I will call:

Moderate naïve realism: Any perception essentially belongs to a metaphysically robust kind that includes only entities that involve relations of acquaintance.

But this brings me to a third, and to my mind very serious, problem for the anti-triviality argument. According to Martin's claim [4], naïve realism characterizes a metaphysically robust kind, and I have just granted the even stronger claim that naïve realism characterizes a *metaphysically robust essential kind*.

²⁹ On the notion of naturalness, see Lewis (1983); on the notion of structure, see Sider (2011); on the notion of gruesomeness, see Goodman (1983); on the notion of being joint-carving, see just about any realist metaphysician you please.

³⁰ As Byrne and Logue put it: "Let the essences look after themselves – we can still perceive the world as 'directly' as we would wish" (2008, p. 82).

From this Martin arrives directly at [5], the claim that naïve realism characterizes a *maximally specific essential kind*. Meanwhile, perhaps some of Martin's followers would prefer to arrive at the claim that naïve realism characterizes a *maximally specific* (but perhaps not essential) *kind*.

Either way, I think that this inference is a howler. Even if we have a metaphysically robust essential kind, that kind does not have to be a maximally specific essential kind, nor does it have to be a maximally specific kind *simpliciter*. This point should be obvious, but let us have an example to drive it home. Grant that Hypatia belonged to the metaphysically robust essential kind *material object*. Still, the property of being a material object does not characterize her maximally specific essential kind, much less her maximally specific kind *simpliciter*. For presumably she belonged to other, more specific essential kinds, such as the kinds *animal* and *human being*.

Thus, the anti-triviality argument is not a good one.

Summary. Martin does not articulate why naïve realists would benefit by using the notion of a metaphysically robust kind or the notion of an essential kind. Still, I think that he might be right on these points, and I have attempted to fill in the missing reasoning.

But the anti-triviality argument concludes that naïve realists should also make claims about the *maximally specific* essential kind to which any perception belongs, and other naïve realists have thought it worth making similar claims about maximally specific (but perhaps not essential) kinds. So far, we have found no support for these further conclusions.

But let us continue to look.

6. The screening-off argument

Elsewhere, Martin considers the following *screening-off argument*³¹:

- (SO1) Any perception and any matching hallucination will cause similar beliefs, actions, etc.³²
 - (SO2) If any perception and any matching hallucination will cause similar beliefs, actions, etc., then any perception and any matching hallucination belong to a common kind.
 - (SO3) If any perception and any matching hallucination belong to a common kind, then the acquaintance relation does no explanatory work.
 - (SO4) If the acquaintance relation does no explanatory work, then naïve realism is false.
-
- (SO5) Naïve realism is false.

Martin thinks that if we accept monistic naïve realism, then we can reject premise (SO3). (Why Martin thinks this will not matter for our purposes.) Martin focuses his energies on giving a thorough examination of (SO3) itself. Still, this suggests

³¹ See Martin (2004, pp. 61-62) and (2006, pp. 369-372).

³² *Similar*, but not *identical*. For example, a perception might result in a genuinely singular belief, whereas a matching hallucination will instead result only in a putatively singular belief.

en passant a possible reason for the naïve realist to prefer the monist version of the view: perhaps this is the only good way for her to resist the screening-off argument.

To this, I respond that the naïve realist can easily reject (SO3) regardless of whether she accepts monistic naïve realism. For it is perfectly clear what explanatory work is supposed to be done by the acquaintance relation: this property is supposed to help explain why perceptions (but not hallucinations) are presentational, reference-enabling, and knowledge-enabling with respect to ordinary, mind-independent particulars. For instance, the naïve realist can say that when a subject sees an apple, she is acquainted with that very apple, and that her being acquainted with the apple at least partly constitutes her being presented with that very apple.

Martin anticipates this response, but argues that it only gets the naïve realist into further trouble. His concern is that the naïve realist wishes to explain what constitutes the subject's *conscious perspective* (or, if you like, the *phenomenal character* of the subject's experience): it is supposed to be constituted by what she actually perceives. But the hallucinating subject has a conscious perspective as well! So if a perception of an apple and a matching hallucination of an apple belong to a common kind, then presumably it is this common kind that fully constitutes the subject's conscious perspective in both cases.³³

I submit that the naïve realist can reject the last step of this argument. I will illustrate the point by focusing on pluralistic naïve realism.

Compare once more a subject who sees an apple with a subject who merely hallucinates one. According to the pluralistic naïve realist, the conscious perspectives of these subjects share *one* constituent: deep awareness of certain universals, such as redness. This common constituent in conscious perspective helps to explain why both subjects can know part of the very *essences* of those universals.³⁴ At the same time, the pluralistic naïve realist identifies *another* constituent that is distinctive to the conscious perspective of the perceiver: the property of being a relation of acquaintance – a direct, non-representational relation that is not even partly essence-revealing – to the apple and its particular instances of redness. Nothing like this constitutes the conscious perspective of the hallucinator. This explains why the perceiver is genuinely conscious *of the apple and its particular instances of redness*, and why the hallucinator is not. In short, the idea is that the constituents of the perceiver's conscious perspective *include but outstrip* the constituents of the hallucinator's conscious perspective.³⁵

But there is a further concern that Martin might press against this approach. It is clear that the hallucinating subject, just like the perceiving subject, will (normally) believe on the basis of introspection that she is conscious of mind-independent objects and their property-instances. So there must be some common property – call it *C* – that is found in hallucination, and that suffices to explain the

³³ See Martin (2004, p. 64).

³⁴ Again, this idea is anticipated by Johnston (2004) and developed carefully in [Author's Work C]. But there are other possibilities as to what the commonality might be: perhaps these subjects are both employing conscious representational capacities of the same type, or perhaps they are both exercising the same conscious capacities to single out and discriminate particulars (see Schellenberg (2018)).

³⁵ On this last point, I am in perfect agreement with Johnston (2004).

subject's belief that she is conscious of mind-independent objects and their property-instances. Now, *C* is presumably also a property of perceptions. Thus, *even in cases of perception*, it turns out to be *C*, and not the acquaintance relation, that explains the subject's belief that she is aware of mind-independent objects and their property-instances. But then the acquaintance relation does not explain what it should.

I believe that, in the end, this argument from Martin does not withstand scrutiny. However, it will take me some time to explain my reasoning. I begin by laying out the pluralistic naïve realist's account of presentational character.

According to pluralistic naïve realism, involuntariness³⁶ is necessary and sufficient for the introspective appearance of being presented with *something*. Once this is in place, there is a further *mind-independent gestalt* that causes the introspective appearance of being presented with *mind-independent objects* – i.e., presentational character. The idea is that perceptions typically involve deep awareness that is vivid, stable, and involuntary. In addition, the universals that are the objects of this deep awareness are typically integrated with one another, they typically change in somewhat predictable ways as the subject moves around her environment, and so on. Of the remaining features, no single one is sufficient for presentational character. Presentational character instead occurs when there is the right overall gestalt of these features. Think for example of the difference between seeing redness on a tomato and experiencing a floating red phosphene: in the latter case the redness is not integrated with the environment, nor does the phosphene move in the expected ways as the subject moves. That is why the subject of the phosphene does not introspectively seem to be presented with a *mind-independent* red object.

In short, the pluralistic naïve realist explains presentational character in terms of involuntary deep awareness that has a mind-independent gestalt. This is Martin's common property *C*. Call this approach the *gestalt account*.

I will briefly mention one piece of independent evidence for the gestalt account. Suppose that, unbeknownst to me, you set up some lights within my field of vision. You design them so that a perception of the lights will be introspectively indiscriminable from the phosphene experience that I get when I vigorously rub my eyes – for example, you set up the lights to be very faint and to flash only briefly before disappearing. Now I vigorously rub my eyes and open them. I simultaneously experience phosphenes together with the real lights that you have set up. I will surely take them to be phosphenes. It will introspectively seem to me that I am not presented with real, mind-independent lights. Why does this occur? With the gestalt account in hand, we can say that although my visual perception *in fact* presents mind-independent lights, my perception does not have a mind-independent gestalt and thus does not *introspectively seem* to present them.

This makes it possible to address the present concern. For, contrary to the concern, any (typical) perception will have this common property *C*, the gestalt property, *because the perception is partly constituted by an acquaintance relation* – it is because the subject is acquainted with (say) a red tomato that she has a vivid, stable, and involuntary deep awareness of redness. That is true even if matching hallucinations must have *C for some other reason*, say because of the deliberate

³⁶ Or the appearance thereof, if that is different.

manipulations of a neuroscientist. Thus, contrary to the concern, the acquaintance relation is still crucial to the explanation of why the *perceiver* introspectively believes that she is conscious of mind-independent objects and their property-instances.

An analogy will drive the point home. A first window was smashed because a baseball struck it; the baseball struck it because Grace threw it; and Grace's having thrown it partly constitutes her having done something wrong (it being wrong to smash people's windows). A second window was smashed because a baseball struck it; this baseball struck it because it was swept up in a hurricane and then discharged in the general direction of the window; and this does not partly constitute the hurricane's having done anything wrong. The fact that Grace threw the baseball still helps to explain the smashing of the *first* window – it is not as though this explanation is screened off by the fact that *any* way of getting a baseball to strike a window at a certain velocity will result in smashing.

The explanatory structure given by the pluralistic naïve realist is perfectly parallel. An ordinary perception of a tomato has presentational character because it has a mind-independent gestalt; it has a mind-independent gestalt because it is a relation of acquaintance to a mind-independent tomato; and its being a relation of acquaintance to a mind-independent tomato constitutes the subject's actually being presented with a mind-independent tomato. A matching hallucination has presentational character because it has a mind-independent gestalt; it has a mind-independent gestalt because (say) a devious neuroscientist has stimulated the subject's brain in a certain way; and the activity of the devious neuroscientist does *not* constitute the subject's actually being presented with anything. The fact that the perception is a relation of acquaintance to a mind-independent tomato still helps to explain why the *first* experience has presentational character – it is not as though this explanation is screened off by the fact that *any* way of getting an experience to have a mind-independent gestalt will result in presentational character.³⁷

But wait: according to the gestalt account, acquaintance with an apple does not even partly *constitute* the presentational character of a perception of an apple. It merely causes it. How, then, can acquaintance partly constitute the subject's conscious perspective, as the pluralistic naïve realist says it does? I reply that presentational character – the property of *introspectively seeming* to present mind-independent objects – is merely *one* aspect of the subject's conscious perspective. Even if *this* aspect is not even partly constituted by acquaintance, other aspects of the subject's conscious perspective can be. For instance, the property of *actually* presenting mind-independent objects can be partly constituted by acquaintance, and the latter can in turn *reliably cause* (rather than constituting) presentational character. That is the pluralistic naïve realist's view.

I will mention in passing that I take the gestalt account to have a further virtue whose importance is hard to overstate. One of the greatest challenges for naïve realism is to explain why matching hallucinations have presentational character – why they *introspectively seem* to present mind-independent objects. The gestalt account lets the pluralistic naïve realist meet this challenge. The idea is

³⁷ Martin may be right to think that screening-off worries apply in cases where we are comparing an explanation of something in terms of a *determinate* to an explanation in terms of a *determinable*. My point is that the present explanation does not have that structure.

that matching hallucinations have presentational character because they are constituted by deep awareness of universals with a mind-independent gestalt. A mind-independent gestalt is *by and large* an excellent indicator that the subject is genuinely presented with mind-independent objects. But it is not an *infallible* indicator, and in cases of matching hallucination it fails. I develop this idea elsewhere.³⁸

This might trigger one final concern. In favorable circumstances, the perceiver seems to have an introspective *guarantee* that the object is present. Martin can say that perceivers do have such a guarantee; the pluralistic naïve realist has just denied this. This might seem to be an important disadvantage of the latter view.

Perhaps. But Martin's view encounters its own difficulties: Martin must explain how introspection could guarantee the object's presence for *some* perceptions but not others. Now, this asymmetry makes perfect sense if we are comparing subjects who differ in *non-experiential* respects: say, in their rationality, alertness, or attentiveness to their experiences. But the subjects that I discuss – one who sees an apple in ordinary circumstances and another who sees a faint light – are not like this. We may suppose that they are both highly rational, alert, etc.

Perhaps the asymmetry also makes sense if the experiences themselves differ in some relevant respect. That is just what happens in our examples: the apple-experience is vivid, the light-experience faint. But we can reconstruct the cases to remove this difference. Think of what happens when you close your eyes and then briefly crack them open to see an apple in front of you: your experience, though faint, will obviously be presentational. Or consider cases in which a boxer gets hit in the side of the head and experiences a bright, vivid flash of light. At the same time, a photographer takes a flash photograph, producing another bright, vivid flash of light. The boxer might easily take *both* of his flash-experiences to be sensations – experiences of mind-*dependent* objects. This gives us an example of a vivid perception that does not seem to be presentational.

On Martin's view, all of these perceptions are instances of acquaintance with mind-independent objects. Why, then, doesn't introspection on these perceptions uniformly provide an introspective guarantee that they have real, mind-independent objects? Martin cannot explain the difference by appealing to non-experiential factors alone, and it is hard to see what experiential factors could do the work. So it is not easy to see how his account can be filled out in a satisfying way. The pluralistic naïve realist avoids this problem by saying that perceptions do *not* provide introspective guarantees, but only reliable yet fallible indicators, that a mind-independent object is present.

Summary. Even if we do not accept monistic naïve realism, we can reject premise (SO3) of the screening-off argument. We can do this while saying that the acquaintance relation helps to explain what is distinctive about the perceiver, and indeed partly constitutes her conscious perspective. At the same time, we can allow that this perspective shares a different constituent – deep awareness of universals, perhaps – with the conscious perspective of the hallucinator.

So far, then, we have found little support for the view that naïve realists should accept monistic naïve realism, where fundamental kinds are understood

³⁸ See [Author's Work C].

either as maximally specific essential kinds or simply as maximally specific kinds. But let us consider one last argument for this view.

7. The argument from explanatory priority

Heather Logue writes: “[S]omething needs to be said about why philosophers of perception have spilled so much ink in arguing over claims about an experience’s most specific kind. What’s so special about the most *specific* characterization [S] we can give of a perceptual experience? ... Here’s something special about *S*: my experience satisfies other psychological characterizations ultimately *in virtue of* being *S*.”³⁹

This might seem to suggest the following argument:

The *argument from explanatory priority*:

- (EP1) If K_1 is the maximally specific kind to which a perception belongs, then for any other kind K_2 to which that perception belongs, it belongs to K_2 in virtue of belonging to K_1 .
- (EP2) If (EP1) is true, then naïve realists should accept that every perception belongs to exactly one maximally specific kind, where this kind includes only entities that involve acquaintance relations.
- (EP3) Naïve realists should accept that every perception belongs to exactly one maximally specific kind, where this kind includes only entities that involve acquaintance relations.⁴⁰

I see several problems with this argument.

To begin with, I think that (EP1) is not well-supported. To be sure, membership in a more specific kind *sometimes* explains membership in a less specific one. But the explanation can easily run in the other direction. For example, my water bottle is not silver and cylindrical because it is a silver cylinder. Rather, it is a silver cylinder partly because it is silver and partly because it is cylindrical. The pluralist thinks that this is just what happens in perception. The thought is that when I see a silver water bottle, my perception involves deep awareness of the universals of silverness and cylindricality, as well as acquaintance with the particular water bottle. Only in virtue of belonging to these less specific kinds does my experience belong to the kind *experience of a water bottle as silver and cylindrical*.

There is also a problem with (EP2). We can evaluate this premise by granting its antecedent just for the sake of argument – that is, by assuming that (EP1) is true. Should naïve realists then accept that every perception belongs to exactly one maximally specific kind, where this kind that includes only entities that involve acquaintance relations? No! The reason is that *for all (EP1) tells us, there might still be perceptions that belong to two or more maximally specific kinds*. After

³⁹ See Logue (2013, pp. 108-109). All emphases are hers. Notice that in the first quotation, Logue moves smoothly from talk about *kinds* to talk about *characterizations*.

⁴⁰ To be clear, I am not certain that Logue would endorse (EP2). All that is clear is that she endorses (EP1).

all, (EP1) is just a conditional claim about what is the case *if* a perception belongs to exactly one maximally specific kind. If we had a good reason to think that every perception belonged to exactly one maximally specific kind, then we might find (EP2) very plausible. But we have not yet been given such a reason.

Thus, both premises of the argument from explanatory priority are insufficiently supported.

Intermission. I have examined several arguments for the claim that the naïve realist should accept monistic naïve realism, where fundamental kinds are understood either as maximally specific essential kinds or simply as maximally specific kinds. As far as I can tell, none of these arguments are successful.

But, independently of the argument from explanatory priority, Heather Logue has laid out a second way for the fundamental kind theorist to elaborate on the notion of a fundamental kind. Logue offers some further reasons for preferring this alternative version of monistic naïve realism. Let us consider her ideas.

8. Logue: Fundamental kinds as unique basic psychological kinds

Logue's proposal is to understand fundamental kinds as *unique basic psychological kinds*:

What an experience fundamentally consists in are the features of it that provide the ultimate *personal-level psychological* explanations of certain phenomenal, epistemological, and behavioral facts.... (2012b, p. 174, emphasis hers)

The idea is this. For any perception *P*, there will be certain crucial phenomenal, epistemological, and behavioral facts about *P*. For a perception to have *K* as a *basic psychological kind* is (i) for these facts to be explained by the fact that *P* belongs to *K*, and (ii) for there to be no other psychological kind *K** such that the perception's belonging to *K** metaphysically explains its belonging to *K*.

With this definition in hand, we can then express Logue's theory as follows:

Logue's monistic naïve realism: Every perception belongs to exactly one basic psychological kind, namely, the *naïve kind*.⁴¹

Here, the *naïve kind* is the kind that includes all and only entities that involve acquaintance relations.

However, I find Logue's monistic naïve realism to be over-committal. For I find it over-committal to claim – as her theory does – that while a perception might belong to many psychological kinds that explain various phenomenal, epistemological, and behavioral facts, it must nevertheless belong to *just one* psychological kind that explains all of the others. For any given perception, why must all of these chains of psychological kinds bottom out in exactly the same anchor? Why can't different chains bottom out in different anchors instead, as on pluralistic naïve realism?

⁴¹ See Logue (2012a, p. 211); (2012b, p. 174); (2013, p. 109).

Logue makes several remarks that might seem to help answer this question. First, she explains why it is worth discussing basic psychological kinds in the first place: “Of course, there are further *subpersonal* psychological facts concerning the information processing that generates experience, and further *non-psychological* facts concerning the biological and chemical underpinnings of such processing. Such facts are of course explanatorily relevant, but it’s not the job of a *philosopher* to identify them” (2012b, p. 174). In short, Logue’s idea is that by identifying the basic psychological kind(s) to which any perception belongs, we are offering the deepest kind of characterization of perception that a philosopher could possibly offer.

For the sake of argument, grant that Logue is right about this. Still, nothing here suggests that every perception belongs to *exactly one* basic psychological kind. The pluralistic naïve realist thinks that every perception belongs to exactly two basic psychological kinds. She has thereby offered the deepest kind of characterization of perception that is possible from the armchair, just like a monistic naïve realist.

Elsewhere, however, Logue offers an argument that is supposed to support her version of monistic naïve realism. In brief, the argument is that perceptions put us in a position to know what objects and property-instances are like independently of our perceiving them, and we can explain how perceptions do this only if we say that perceptions involve acquaintance relations.⁴²

An initial problem is that, at best, this argument supports only the conclusion that every perception involves an acquaintance relation. The argument does not support the further conclusion that acquaintance is a basic psychological kind. But perhaps we can get defeasible support for this further conclusion by observing that no one has yet offered a good explanation, at the level of personal psychology, for why perception is an acquaintance relation. There remains a deeper problem: nothing in this argument suggests that every perception belongs to *exactly one* basic psychological kind characterized by acquaintance. At best, the argument supports only the conclusion that every perception belongs to *at least* this basic psychological kind. The pluralistic naïve realist accepts this conclusion while thinking that any perception belongs to an additional basic psychological kind, the kind whose members all involve deep awareness.

Thus Logue’s arguments do not carry us all the way to her version of monistic naïve realism.

9. Other purported reasons to commit to monistic naïve realism

There is an assortment of further reasons that the naïve realist might give for committing to some relevant claim about fundamental kinds, whether these are understood as maximally specific essential kinds, as maximally specific (but perhaps not essential) kinds, or as basic psychological kinds.

For one thing, the naïve realist might think that without some such notion, she cannot distinguish her theory from rival theories. To see the problem, return to:

⁴² See Logue (2012a).

Moderate naïve realism: Any perception essentially belongs to a metaphysically robust kind that includes only entities that involve acquaintance relations.

This theory is perfectly compatible with both of the following theories:

The *moderate common kind theory:* For any perception and any matching hallucination, there is a metaphysically robust kind K such that the perception and the hallucination both essentially belong to K .

Moderate representationalism: For any perception and any matching hallucination, there is a metaphysically robust *representational* kind K such that the perception and the hallucination both essentially belong to K .

But surely the naïve realist should be able to distinguish her theory from these competitors!

My response is to deny this last claim. There are many desiderata for a theory of perception: for example, to explain the presentational, reference-enabling, and knowledge-enabling character of perception. But it is just not a desideratum on a theory of perception to be incompatible with various common kind or representationalist theories. Thus, the ecumenical character of moderate naïve realism is not a problem. It may even be a virtue, insofar as it lets the moderate naïve realist borrow various useful theoretical posits from these other theories.

Still, there is a very different kind of reason, namely a sociological one, for not counting what I have called “moderate naïve realism” and “pluralistic naïve realism” as genuine forms of naïve realism. In recent memory, many (though certainly not all) philosophers on various sides of the table – including self-avowed “naïve realists,” “common kind theorists,” and “representationalists” – have presupposed that any perception belongs to exactly one privileged kind. Their idea is that this privileged kind that will do just about all of the explanatory work that there is to do. Thus they have tended to define terms like “naïve realism,” “the common kind theory,” and “representationalism” in a way that makes the first theory incompatible with the second and the third. I have been arguing that we have not yet heard a good reason to accept this presupposition, but given the historical facts, you might think that it is useful to keep defining things in a way that preserves these incompatibilities.

Fair enough! Then you can read me as saying two things. First, that there is not much reason to accept naïve realism itself. Second, that there might well be good reason to accept a related but weaker view – roughly, the view that you get by starting with standard-issue naïve realism and then paring away any commitments about fundamental kinds. You are welcome to call this new view whatever you like.

10. Conclusion

It is utterly standard for naïve realists to accept:

Monistic naïve realism: Every perception belongs to exactly one fundamental kind that includes only entities that involve acquaintance relations.

I have argued that if there is a good reason for naïve realists to commit to monistic naïve realism, then no one has yet identified it. At best, naïve realists have only been given good reasons to commit to:

Moderate naïve realism: Any perception essentially belongs to a metaphysically robust kind that includes only entities that involve acquaintance relations.

I believe that undoing the shackles of monism lets the naïve realist solve the most serious problems for her view. I have hinted at how that solution might go without arguing for it in detail. What I *have* argued in detail is that there is nothing to stop the naïve realist from making the attempt.

REFERENCES

- [Author's Work A]
 [Author's Work B]
 [Author's Work C]
- Allen, K. (2011). "Revelation and the nature of colour." *Dialectica* 65 (2): 153-176.
- Alston, W. (1999). "Back to the theory of appearing." *Philosophical Perspectives* 13: 181-203.
- Bayne, T. (2009). "Perception and the reach of phenomenal content." *Philosophical Quarterly* 59 (236): 385-404.
- Beck, O. (2018). "Rethinking naïve realism." *Philosophical Studies* (online): 1-27.
- Beck, O. (forthcoming). "Two conceptions of naïve realism." *Philosophers' Imprint*.
- Brewer, B. (2011). *Perception and Its Objects*. Oxford: Oxford University Press.
- Broi, A. (2020). "Revelation and phenomenal relations." *Philosophical Quarterly* 70 (278): 22-42.
- Byrne, A. and Logue, H. (2008). "Either/or." In *Disjunctivism: Perception, Action, Knowledge*, eds. Adrian Haddock and Fiona Macpherson (New York: Oxford University Press), pp. 57-94.
- Campbell, J. (2002). *Reference and Consciousness*. Oxford: Oxford University Press.
- Crane, T. (2006). "Is there a perceptual relation?" In *Perceptual Experience*, eds. T. Gendler and J. Hawthorne (Oxford University Press), pp. 126-146.
- Fine, K. (1994). "Essence and modality." *Philosophical Perspectives* 8: 1-16.
- Fish, W. (2009). *Perception, Hallucination, and Illusion*. Oxford: Oxford University Press.
- French, C. (2013). "Perceptual experience and seeing that *p*." *Synthese* 190: 1735-1751.
- Genone, J. (2016). "Recent work on naïve realism." *American Philosophical Quarterly* 53 (1): 1-24.
- Gomes, A. (2017). "Naïve realism in Kantian phrase." *Mind* 126: 529-578.
- Goodman, N. (1983). *Fact, Fiction, and Forecast* (fourth edition). Cambridge, MA: Harvard University Press.
- Hellie, B. (2007). "Factive phenomenal characters." *Philosophical Perspectives* 21: 259-306.
- Jackson, F. (1982). "Epiphenomenal qualia." *Philosophical Quarterly* 32: 127-36.
- Johnston, M. (1992). "How to speak of the colors." *Philosophical Studies* 68: 221-263.
- Johnston, M. (2004). "The obscure object of hallucination." *Philosophical Studies* 120: 113-83.
- Kennedy, M. (2009). "Heirs of nothing: the implications of transparency." *Philosophy and Phenomenological Research* 79 (3):574-604.
- Kennedy, M. (2013). "Explanation in good and bad experiential cases." In *Hallucination: Philosophy and Psychology*, eds. F. Macpherson and D. Platchias (Cambridge, MA: MIT Press), pp. 221-254.
- Lewis, D. (1983). "New work for a theory of universals." *Australasian Journal of Philosophy* 61 (4): 343-377.

- Logue, H. (2012a). "Why naïve realism?" *Proceedings of the Aristotelian Society* 112 (2): 211-237.
- Logue, H. (2012b). "What should the naïve realist say about total hallucinations?" *Philosophical Perspectives* 26: 173-199.
- Logue, H. (2013). "Good news for the disjunctivist about the bad cases." *Philosophy and Phenomenological Research* 86 (1): 105-133.
- Lyons, J. (2005). "Perceptual belief and nonexperiential looks." *Philosophical Perspectives* 19: 237-56.
- Martin, M. (2002). "The transparency of experience." *Mind and Language* 17 (4): 376-425.
- Martin, M. (2004). "The limits of self-awareness." *Philosophical Studies* 120: 37-89.
- Martin, M. (2006). "On being alienated." In *Perceptual Experience*, eds. T. Gendler and J. Hawthorne (Oxford University Press), pp. 354-410.
- Masrour, F. (2011). "Is perceptual phenomenology thin?" *Philosophy and Phenomenological Research* 83 (2): 366-397.
- Masrour, F. (forthcoming). "On the possibility of hallucinations." *Mind*.
- McLaughlin, B. (2003). "The place of color in nature." In *Colour Perception: Mind and the Physical World*, eds. R. Mausfeld and D. Heyer (New York: Oxford University Press), pp. 475-502.
- Millar, B. (2014). "The phenomenological directness of perceptual experience." *Philosophical Studies* 170: 235-253.
- Neta, R. (2008). "In defence of disjunctivism." In *Disjunctivism: Perception, Action, Knowledge*, eds. A. Haddock and F. Macpherson (Oxford: Oxford University Press), pp. 311-329.
- Nudds, M. (2009). "Recent work in perception: Naïve realism and its opponents." *Analysis Reviews* 69 (2): 334-346.
- Pautz, A. (2007). "Intentionalism and perceptual presence." *Philosophical Perspectives* 21 (1): 495-541.
- Schellenberg, S. (2018). *The Unity of Perception: Content, Consciousness, Evidence*. Oxford: Oxford University Press.
- Sider, T. (2011). *Writing the Book of the World*. Oxford: Oxford University Press.
- Siegel, S. (2006). "Which properties are represented in perception?" In *Perceptual Experience*, eds. T. Gendler & J. Hawthorne (New York: Oxford University Press), pp. 481-503.
- Soteriou, M. (2005). "The subjective view of experience and its objective commitments."
- Wiggins, D. 1980. *Sameness and Substance*. Oxford: Blackwell.
- Wiggins, D. 2001. *Sameness and Substance Renewed*. Cambridge: Cambridge University Press.
- Williamson, T. (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press.