



Perceptual variation and ignorance

John Morrison¹

Received: 18 June 2019 / Accepted: 29 December 2020 / Published online: 24 January 2021
© The Author(s), under exclusive licence to Springer Nature B.V. part of Springer Nature 2021

Abstract

There is variation in how people perceive colors and other secondary qualities. The challenge of perceptual variation is to say whose perceptions are accurate. A natural and influential response is that, whenever there's variation in two people's perceptions, at most one of their perceptions is accurate. I will argue that this leads to an unacceptable kind of ignorance.

Keywords Philosophy of mind · Epistemology · Perception · Color · Perceptual variation

1 Introduction

There's a lot we don't know, and shouldn't expect to know: We don't know the current number of stars in distant clusters, because light from those clusters won't reach our telescopes for billions of years. We don't know the aggregate weight of all the chocolate in existence, because it's created and consumed too quickly. We don't know Socrates's exact height when he drank hemlock, because his corpse decomposed long ago.

In most cases, ignorance is philosophically uninteresting. But in some cases, it forces us to rethink our basic assumptions. Consider velocity. At some point in the Earth's orbit, it moves 50 km/s relative to Mars, and 30 km/s relative to the sun. Many have assumed that, in addition to these relative velocities, the Earth also has an absolute velocity—that is, a velocity that isn't relative to any other bodies or frames of reference. For Newtonians, this follows from their view of space as a container-like entity. For others, it's a natural pre-theoretical way of thinking about motion. However, if the Earth has an absolute velocity, we can't know it, and shouldn't expect to know it, given Maxwell's equations (see Bell 2004, Ch. 9). According to many philosophers and physicists, this should lead us to give up the assumption that the Earth has an

✉ John Morrison
jmorrison@barnard.edu

¹ Barnard College, Columbia University, 3009 Broadway, 326b Milbank Hall, New York, NY 10027, USA

absolute velocity. They claim that, while our ignorance of the current number of stars in distant clusters shouldn't lead us to deny that there are stars in distant clusters, our ignorance of absolute velocities should lead us to deny that bodies have absolute velocities.¹

I'll develop a related objection to a natural and influential response to the puzzle of perceptual variation. In particular, I'll argue that it should be rejected, because, like the assumption that the Earth has an absolute velocity, it leads to an unacceptable kind of ignorance. As part of my argument, I will explain why this kind of ignorance should lead us to give up otherwise attractive assumptions and responses, while other kinds of ignorance, such as our ignorance of the number of stars in distant clusters, shouldn't.

What is the challenge of perceptual variation? Suppose that Aaron and Miriam are looking at the same lemon, and their perceptions differ. To convey how their perceptions differ, let's use your perceptions as a reference point. Let *phenomenal-greenish-yellow* be the phenomenal character of your perception when you report that a surface looks greenish yellow, and *phenomenal-pure-yellow* be the phenomenal character of your perception when you report that a surface looks unmixed, pure yellow in the same context. Suppose that Aaron's perception of the lemon is phenomenal-greenish-yellow, and Miriam's perception of the lemon is phenomenal-pure-yellow. The challenge of perceptual variation is to say whose perception is accurate.

This challenge is not a mere hypothetical, because there's compelling evidence that perceptual variation is widespread. Elsewhere, I review the physiological evidence, as well as the behavioral evidence from experiments in which subjects are asked to identify the pure colors (see Morrison 2020b). Here, I'll just review the behavioral evidence from matching experiments (Wyszecki and Stiles 1982, Ch. 5; Webster and MacLeod 1988). Suppose we select subjects who pass all the standard tests of color acuity, such as the Farnsworth-Munsell Hue Test. We then show them two panels, one illuminated by a single monochromatic light and the other illuminated by a mixture of three monochromatic lights. We then ask subjects to adjust the mixture of three monochromatic lights until both panels look the same. The same subject will reliably produce the same mixture. But different subjects will reliably produce different mixtures. Moreover, if we show subjects the mixtures produced by other subjects, they'll often say that the panels look slightly different. If these subjects perceived the panels in the same way, they would reliably produce the same mixture, and they wouldn't disagree about which panels match. Thus, even subjects with normal color acuity perceive the panels differently. This is a short and hopefully uncontroversial step away from the conclusion that even these subjects have perceptions with different phenomenal characters. Aaron and Miriam are supposed to be subjects like that. So, whose perception is accurate?

According to the response I have in mind ("one-ism"), at most *one* of their perceptions is accurate. Its proponents include Stroud (2000); Tye (2000); Byrne and Hilbert (2003); Allen (2016). According to these philosophers, when the perceptions of normal human subjects have different phenomenal characters, one of those phenomenal characters is special, in that all of the accurate perceptions have that phenomenal char-

¹ For arguments of this kind, see Ismael and van Fraassen (2003), Maudlin (2012, Ch. 3), and Dasgupta (2016). Note that they disagree about *why* our ignorance of absolute velocities should lead us to deny that bodies have absolute velocities. I'll return to this issue in Sect. 3.

acter. It follows that, if the perceptions of two normal human subjects have different phenomenal characters, at most one of their perceptions is accurate. In the next section (Sect. 2), I'll argue that one-ism leads to ignorance about whose perception is accurate. In the following section (Sect. 3), I'll explain why this ignorance is just as unacceptable as our ignorance about absolute velocity. I'll conclude that one-ism is an unacceptable response to the puzzle. I'll end (Sect. 4) with a brief discussion of the more extreme kind of variation found in so-called "phenomenal inversions". I will argue that this kind of variation is not as effective against one-ism.

I'm not the first to object that one-ism leads to ignorance. See, in particular, Jackson and Pargetter (1987, p. 133); Hardin (1988, p. 89; 2003, pp. 199–201); Block (1999, pp. 46, 54); Clark (2000, pp. 215–217); Kalderon (2007, p. 566); Cohen (2009, pp. 45–64); Brogaard (2010, p. 146) and Egan (2010, p. 70). But their objections fall short. First, they don't establish that one-ism really does lead to ignorance, because they consider only a small subset of the evidence that we'd expect a one-ist to use to learn whose perception is accurate. Second, they don't explain why ignorance of this kind is unacceptable, and thus why they've identified a reason to reject one-ism.² They often spend only a few paragraphs developing their objections, perhaps because they believe that the reasons for rejecting one-ism are obvious and self-contained. As we'll see, however, the reasons for rejecting one-ism are subtle and draw on seemingly unrelated principles across metaphysics and epistemology.

Why does one-ism deserve so much attention? Perhaps the most popular argument for one-ism starts by appealing to a weak version of representationalism about color perception:

REPRESENTATIONALISM

If two perceptions have the same phenomenal character, they represent the same color.

Recall that, by stipulation, phenomenal-greenish-yellow is the character of your perception when you report that a surface looks greenish yellow, and phenomenal-pure-yellow is the phenomenal character of your perception when you report that a surface looks pure yellow in the same context. If you were looking at those surfaces

² Cohen (2009, pp. 51–52) identifies a precedent for thinking that this would be an unacceptable kind of ignorance. In particular, he points out that if two people disagree about whether a joke is funny, few would insist that only one of them is right. He then asks, in essence: Why think that perceptual variation is any different? But this is too open-ended, in that he merely says that we don't have any reason to think that colors *aren't* like humor, inviting the one-ist to respond that we also don't have any reason to think that colors aren't like sizes, shapes, and distances. See Tye (2012, p. 299).

Brogaard (2010, p. 146) argues that this would be an unacceptable kind of ignorance, because if we don't know which objects are pure yellow, we can't learn the meanings of color terms such as 'pure yellow'. But a one-ist could say that we learn the meaning of 'pure yellow' by perceiving a lemon as pure yellow, even if we don't know our perception is accurate.

McLaughlin (2003, p. 122) argues that one-ism implies that our perceptions of fine-grained colors are usually inaccurate. But it's not clear to me why that's objectionable, given that our perceptions of coarse-grained colors might still be mostly accurate, and our perceptions of fine-grained colors would still allow us to discriminate and re-identify objects.

at the same time, and your perceptions had the same respective phenomenal characters, then, in most cases, you would perceive the first surface as greener than the second surface. And for that to be true, you must be perceiving the surfaces as having *incompatible* colors. Consider that, if your perceptions left open the possibility that the surfaces have the same color, you wouldn't perceive one as greener than the other. Analogously, if a friend told you that his brother is older than his sister, he would have told you that his brother and sister have incompatible ages. If what he said left open the possibility that his siblings are the same age, he wouldn't have told you that his brother is *older*.

So far, this is just a claim about your perceptions. But given REPRESENTATIONALISM, if *your* phenomenal-greenish-yellow perception and *your* phenomenal-pure-yellow perception represent incompatible colors, then *all* phenomenal-greenish-yellow perceptions and *all* phenomenal-pure-yellow perceptions represent incompatible colors. Thus, because Aaron's perception has a phenomenal-greenish-yellow character and Miriam's perception has a phenomenal-pure-yellow character, their perceptions represent incompatible colors, and at most one of their perceptions is accurate.

I'll say more about this argument later. But I hope this is enough to explain why so many are drawn to one-ism.

Importantly, one-ism doesn't take a stand on the metaphysics of colors. A reflectance physicalist might think that Aaron and Miriam are perceiving different ways of reflecting light, and that the lemon has only one of these reflectances. A realist primitivist might think that they're perceiving different non-physical properties, and that the lemon can instantiate only one of these properties because there's something about them that prevents their co-instantiation. A dispositionalist might think that they're perceiving dispositions to cause different kinds of perceptions in a certain kind of subject in a certain kind of context, and that the lemon causes only one of those perceptions in the relevant kind of subject in the relevant kind of context.

2 Ignorance

There's an important disanalogy between our ignorance of absolute velocity and our ignorance of whose perception is accurate. Given Maxwell's equations, we can *prove* that there's no way to set up a causal link between absolute velocities and observable effects. As a result, we can prove that there is no evidence that might allow us to choose between the hypothesis that the Earth has an absolute velocity of 0 km/s and the hypothesis that it has an absolute velocity of 20 km/s (see again Bell 2004, Ch. 9). With respect to color perception, there are no proofs, because there are no physical laws that entail that there is no evidence about whose perception is accurate. We instead need to canvass the nine kinds of evidence that might seem relevant, and argue that each of them is insufficient. While we won't end up with the same degree of certainty in our conclusion, we can still be highly confident, provided that we are sufficiently thorough.

To make our discussion more concrete, let's assume reflectance physicalism, the view that colors are ways of reflecting light. After we've considered all nine kinds of evidence, we can drop that assumption and generalize to other views about the

metaphysics of colors, such as realist primitivism and dispositionalism. Let's also assume (incredibly!) that these are the only two possible perceptions of the lemon among normal human subjects. Given this assumption, one-ism implies that there are only two hypotheses: that Miriam's perception is accurate and Aaron's perception is inaccurate, and that Miriam's perception is inaccurate and Aaron's perception is accurate. This assumption is helpful to us, because it will simplify our discussion of each kind of evidence. This assumption is also helpful to the one-ist, because, as a result, she just needs to identify evidence that decides between these two hypotheses. Without this assumption, she would need to identify evidence that decides between a much larger number of hypotheses, one for each kind of perception.

(1) The first kind of evidence is about how the lemon reflects light. If colors are reflectances, this evidence might seem especially helpful, because knowledge of the lemon's reflectance would be knowledge of the lemon's color. However, even in that case, it's unclear how this evidence could help justify our belief that Miriam's perception is more accurate. Even if we know the reflectance of the lemon, we're left wondering: Who is perceiving *that reflectance*? Is it Aaron, in virtue of his phenomenal-greenish-yellow perception, or is it Miriam, in virtue of her phenomenal-pure-yellow perception? Some of the evidence we'll consider below might seem to help us answer these questions. For now, my point is just that evidence about the reflectance of the lemon is not, by itself, enough.

(2) The second kind of evidence is about how other objects reflect light. If colors are reflectances, perhaps we could thereby learn how the lemon's color relates to the colors of other objects. If we learned that the lemon's reflectance falls midway between reddish yellow and greenish yellow, we could justifiably believe that the lemon is pure yellow, and thus that Miriam is accurately perceiving it. But it's not that simple. Reflectances vary along infinitely many different dimensions, one for each wavelength at which light is reflected, and thus for each real value between 400 and 700nm. We can order reflectances using any weighted combination of these dimensions. If we order reflectances in some ways, the lemon's reflectance will be in the middle. But if we order reflectances in other ways, it won't be. We're thus left wondering: Which is the ordering of reflectances from reddish yellow to greenish yellow? Is it an ordering with the lemon's reflectance in the middle, or an ordering with the lemon's reflectance closer to one of the endpoints?

(3) The third kind of evidence is about their reports. Aaron might report that the lemon looks greenish yellow, while Miriam reports that the lemon looks pure yellow. But these reports don't help us choose the accurate perception. If the lemon's color is its reflectance, we're still left wondering: Who is describing *that reflectance*?

There's a helpful contrast with reports about shape. Suppose that Aaron and Miriam both look at an ellipse. If Aaron reports that the ellipse looks perfectly circular, while Miriam reports that it looks slightly longer along its horizontal axis, we can settle the matter with a ruler. The direct approach is to measure the ellipse's major and minor axes. The indirect approach is to measure the major and minor axes of other ellipses, order them by the degree to which they depart from perfect circularity, i.e., by their eccentricity, and see where the ellipse fits in. Either way, we can verify their reports. Aaron and Miriam's reports about the lemon's color aren't similarly verifiable. While

we can use a spectrometer to measure how the lemon and other objects reflect light, we can't thereby settle whose report is accurate.³

(4) The fourth kind of evidence is about the physiological differences between Aaron's and Miriam's eyes. Miriam's eyes might have more of the detectors primarily responsible for perceiving colors (S-cones, M-cones, and L-cones). Or, Miriam's eyes might have proportionally more of a certain kind of detector (e.g., proportionally more S-cones). Or Miriam's detectors might be more sensitive to light than Aaron's detectors. Or Miriam's detectors might be more sensitive to lower wavelengths than Aaron's detectors. And so on. However, it's unclear how this evidence alone could justify our belief that Miriam's perception is more accurate. We'd still need to determine that people with eyes like Miriam's perceive colors more accurately than people with eyes like Aaron's. Future scientific advances are unlikely to help, because we already know many of the physiological causes of variation, and they don't give us any justification for believing that Miriam's perception is accurate. We'll just continue to learn more about the underlying mechanisms, without any principled way to label one as defective.

In other domains, this is easier to determine. If two microscopes produce different images, and one of them has a scratched lens, then we know which microscope is defective. But it's not similarly obvious why Aaron's eyes, rather than Miriam's eyes, are defective. The differences between their eyes are not like the differences between scratched and unscratched lenses.

This doesn't mean that it's impossible for us to discover that Aaron's eyes are defective. If we had independent evidence about whose perception is accurate, we could work backwards and figure out whose eyes are defective and in what respect. But evidence about their eyes is not, by itself, enough to reach any conclusions.

(5) The fifth kind of evidence is about the phenomenal characters of their perceptions. To see why some philosophers might think that this evidence is helpful, let's again consider shape perception.

There are two ways in which philosophers might think that evidence about their phenomenal characters can help us determine whose shape perception is accurate. According to some philosophers, we can deduce the degrees of eccentricity that Aaron and Miriam are perceiving from the phenomenal characters of their perceptions. How? According to these philosophers, the relevant phenomenal characters are picture-like in that they *exemplify* the degrees of eccentricity they represent. If these philosophers are right, and if Aaron and Miriam are perceiving different degrees of elongation, we might be able to use their phenomenal characters to figure out whose perception is accurate. One approach would be to first use their phenomenal characters to deduce what degree of eccentricity each is perceiving, and then use a ruler to determine the eccentricity of the ellipse.

³ There is an exception. If the meaning of their reports depends on the reports of other subjects, we might be able to verify Miriam's report, on the basis of what others report. For example, if Miriam reports 'The Footnote 3 continued

lemon is pure yellow', and that report is true just in case most other subjects would report that it's pure yellow, we might be able to verify her report by asking other subjects for their reports. But if the meaning of Miriam's report depends on what others report, her reports aren't a guide to what she's perceiving. I'll return to these issues later, in my discussion of the sixth kind of evidence.

According to other philosophers, we can “read off” the degrees of eccentricity that Aaron and Miriam are perceiving from the phenomenal characters of their perceptions, even though those phenomenal characters don’t exemplify a degree of eccentricity. This is a hard view to summarize, but the basic thought is that there’s a necessary connection between certain phenomenal characters and certain shapes, and those necessary connections are immediately obvious to us in perception (Pautz 2019, pp. 392–396). If these philosophers are right, we once again might be able to use Miriam’s and Aaron’s phenomenal characters to figure out whose perception is accurate.

Of course, it might be hard, and perhaps even impossible, to gather evidence about the phenomenal characters of their perceptions. Nonetheless, this kind of evidence might be enough to figure out whose perception is accurate, at least according to the philosophers I just mentioned.

Returning to color perception, it’s unclear how we could figure out whose perception is accurate, even if we knew the phenomenal characters of their perceptions. If pure yellowness is a certain kind of reflectance, we can’t deduce which reflectances Aaron and Miriam are perceiving from their phenomenal characters. In particular, their phenomenal characters do not exemplify the reflectances they perceive, and there isn’t an immediately obvious connection between their phenomenal characters and those reflectances, at least if those reflectances are described in a way that would allow us to match them with the lemon’s reflectance (unlike, e.g., “the reflectance I’m perceiving right now”).

Just to be clear, I’m not sure that the aforementioned views of shape perception are right. I just think that they help us appreciate why this kind of evidence doesn’t help answer our questions; even if one of these views of shape perception is right, the analogous view of color perception is untenable. The phenomenal characters of our color perceptions do not exemplify or necessitate colors in a way that allows us to figure out whose perception is accurate.

(6) The sixth kind of evidence is about the brain activity that underlies their perceptions. Let’s again consider shape perception to see why this kind of evidence might seem helpful. In principle, the level of neural activity that underlies a subject’s perception of an ellipse’s eccentricity could vary along one dimension, with a minimum and maximum, so that each level corresponds to a number between 0 and 1. Also in principle, increases in the eccentricity of an ellipse might roughly correspond to increases in the level of that neural activity. We might then reasonably conclude that the level of neural activity corresponds to the eccentricity that a subject is perceiving, so that a subject perceives an eccentricity of 0 when the level of neural activity is 0, and an eccentricity of 1 when the level of neural activity is 1. In that case, if Miriam and Aaron perceive the same ellipse, and the level of neural activity is 1 in Miriam and .9 in Aaron, we could use a ruler to discover whose perception is accurate.

Of course, this presupposes substantial empirical assumptions about the structure of the relevant neural activity and the way it correlates with the actual eccentricities of ellipses. I don’t know of any empirical evidence in support of these assumptions. My point is just that, given these assumptions, it might be possible to learn who is accurately perceiving the eccentricity of an ellipse.

Color perception is different, because it’s not similarly possible to learn who is accurately perceiving the reflectance of a surface, even given equivalent assumptions.

To appreciate why, assume that the level of neural activity that underlies a subject's perception of yellow varies along one dimension, with a minimum and maximum, so that each level corresponds to a number between 0 and 1. The crucial difference is that, whereas there's only one way of ordering ellipses by their eccentricities, there are infinitely many ways of ordering objects by their reflectances. Some orderings will better correspond to increases in the level of neural activity in Miriam, and other orderings will better correspond to increases in the level of neural activity in Aaron. Because there's no principled way of choosing which ordering corresponds to increases and decreases in the relevant neural activity, we don't know which reflectance the maximum level of neural activity represents. Thus, if Miriam and Aaron perceive the same lemon, and the level of neural activity is 1 in Miriam and .9 in Aaron, we have no way of discovering which reflectances they're representing, and thus no way of discovering who is accurately representing the lemon's reflectance. In some orderings, the lemon is at the maximum, so that Miriam's perception would be accurate, and in other orderings, the lemon isn't at the maximum, so that Aaron's perception might be accurate. This issue didn't arise for shape perception, because there was only one ordering of ellipses by eccentricity.

Of course, if we already knew that Miriam's perception of the lemon's reflectance was accurate, we might be able to work backwards, to determine which reflectance the maximum level of neural activity represents. But we can't go in the opposite direction: from the activity in Miriam's and Aaron's brains alone, we can't deduce which reflectance the maximum level of neural activity represents or, more generally, which reflectances different levels of neural activity represent. Thus, evidence about the brain activity that underlies Miriam's and Aaron's perceptions does not, by itself, allow us to determine whose perception is accurate.

Note that this isn't a problem if, contrary to one-ism, Miriam's and Aaron's perceptions might both be accurate. In that case, even if the level of neural activity is 1 in Miriam and .9 in Aaron, they might be representing the same reflectance, because their neural activity might be tracking different orderings of reflectances, such that the lemon's reflectance is at the maximum of the ordering tracked by Miriam's neural activity, and below the maximum of the ordering tracked by Aaron's neural activity (e.g., see Morrison 2020b).

(7) The seventh kind of evidence is about evolution. Why might this kind of evidence seem helpful? Evolution seems to explain a lot about color perception. For example, it seems to explain why color perceptions are sensitive to electromagnetic radiation between 400 and 700nm (i.e., light), rather than electromagnetic radiation at significantly lower or higher wavelengths. Briefly: in our ancestral environments, the way that objects reflect light between 400 and 700nm is the most useful for avoiding predators, hunting prey, picking fruits, and performing all of the other activities that contributed to our survival. As a result, we evolved so that our color perceptions are sensitive to electromagnetic radiation between 400 and 700nm. Deep-sea animals, in contrast, evolved to have color perceptions (or something equivalent) that are sensitive to shorter wavelengths, because, at those depths, the way surfaces reflect electromagnetic radiation at shorter wavelengths is more useful for survival (for an overview, see Lythgoe and Partridge 1989).

It's natural to think that evolution also explains *what* we perceive. In particular, if, from an evolutionary point of view, the function of color perception is to detect the reflectances of predators, prey, fruits, and so on, it's natural to think that's what we perceive.

There are a number of well-known objections to using evolution to explain what we perceive (for an overview, see Neander 2012). I'm sympathetic with many of these objections, but I won't rehearse them here. I'll instead argue for a weaker conclusion that's easier to establish, namely that, even if evolution explains why Aaron and Miriam are perceiving reflectances, it's unlikely that evolution explains *which* reflectances they are perceiving, and thus unlikely that this kind of evidence can help justify our belief that only Miriam's perception is accurate.

Aaron's and Miriam's perceptions differ, in that Aaron's perception is phenomenal-greenish-yellow while Miriam's perception is phenomenal-pure-yellow. If evolution explains why only Miriam accurately perceives the reflectance of the lemon, evolution must explain why only phenomenal-pure-yellow perceptions have the function of detecting that reflectance, and thus why, from an evolutionary point of view, phenomenal-pure-yellow perceptions are the ideal reaction to the lemon. But that's unlikely. From an evolutionary point of view, the differences between Miriam's and Aaron's perceptions are probably just as uninteresting as the slight differences in their heights, skin pigments, and hair densities, and nobody should think that, from an evolutionary point of view, there's an *exact* height, skin pigment, or hair density that's ideal. It's far more likely that, from an evolutionary point of view, there is a range of acceptable variation within each of these traits. In that case, as long as Miriam's and Aaron's perceptions fall within the range of acceptable variation, evolution doesn't explain why only phenomenal-pure-yellow perceptions have the function of detecting the lemon's reflectance (Tye 2006, pp. 175–176 and Gert 2006, p. 579, make similar points).

Why, from an evolutionary point of view, is it far more likely that there is an acceptable range of variation? To start, from an evolutionary point of view, what matters is *behavior*, and Miriam's and Aaron's behavior isn't a simple function of their phenomenal characters. Even if they have slightly different phenomenal characters, they might behave in ways that are equally conducive to survival. Keep in mind that, because Aaron and Miriam are ordinary subjects who pass all the standard tests of color acuity, they are equally adept at color discrimination. It is therefore unlikely that Miriam is better at detecting predators, prey, fruits, and so on.

In addition, even if there is a slight difference in Miriam's and Aaron's behavior, and it benefits Miriam, that benefit is likely to depend on transient features of her environment. It might depend on the amount of rainfall and cloud cover, for example. Thus, even if, in a given environment, Miriam's behavior is more beneficial, that benefit is likely to be highly environment-sensitive, because in a slightly different environment Aaron's eyes would be more advantageous.

Building on this last point, it's unlikely that phenomenal-pure-yellow perceptions were *selected for* the detection of a specific reflectance and thus have the function of detecting that reflectance. Our ancestors presumably inhabited a number of different environments, with different amounts of rainfall, cloud cover, etc., so that a trait that was the most beneficial to one ancestor in their environment might not have been

the most beneficial to her grandchildren or to her cousin, and a trait that was most beneficial to ancestors on one side of her family might not have been the most beneficial to ancestors on the other side. In that case, it's unlikely that phenomenal-pure-yellow perceptions were selected for, regardless of one's preferred account of selection.

There's another reason to think it's unlikely that phenomenal-pure-yellow perceptions were selected for the detection of a specific reflectance. The genetic material at the locus of a genotype typically affects many different phenotypic traits. For a large number of phenotypic traits, many different loci are involved. As a result, in order to estimate the effect of any given trait on survival and reproduction, one would have to consider all of the other traits associated with the same underlying genotypes. Thus, even if Miriam's behavior was beneficial, that wouldn't establish that any of her distinctive phenotypic traits were selected for. One would need to show that it conferred a benefit greater than the benefits conferred by all the other phenotypic traits effected by the same genotypes. Given the extremely minor behavioral differences under discussion, that seems unlikely. Thus, it's unlikely that phenomenal-pure-yellow perceptions, rather than phenomenal-greenish-yellow perceptions, were selected for detecting the lemon's reflectance and thus have the function of detecting that reflectance.

These considerations also establish a stronger conclusion: that it's unlikely that *any* of our perceptions were selected for the detection of a specific reflectance. We thus shouldn't hold out hope that some other kind of perception was selected for its detection of a specific reflectance, thereby calibrating all of our perceptions, giving each the function of detecting a specific reflectance (contrary to Byrne and Hilbert 2007, fn. 6).

(8) The eighth kind of evidence is about other subjects. This approach is inspired by the popular view that the referents of terms such as 'arthritis' and 'cold' depend on how other people use them. According to one version of this view, if doctors use 'arthritis' to refer to a certain kind of ailment, then 'arthritis' refers to that kind of ailment, even if some non-experts misuse the term (see, e.g., Burge 1979). According to another version, if a majority of the people in your linguistic community use 'cold' to refer to temperatures below a certain threshold, then 'cold' refers to temperatures below that threshold, even if a minority misuse the term. These might seem like promising models for color perception. We could then learn whether Miriam is accurately perceiving the lemon by learning how others perceive it.

Let's first consider the version that appeals to experts. There's no question that people can improve at various perceptual tasks. Chefs can learn to identify the ingredients in a dish by tasting it. Photographers can learn to identify the tint of a film by looking at it. In these cases, it might make sense to defer to the experts. But we're interested in variation with an immutable, physiological cause. While everyone in the kitchen might defer to the chef's perception of a dish, if the difference is ultimately due to different proportions of the relevant detectors on the tongue, or some other physiological difference, this is mere social deference, without any clear implications as to whose perceptions are accurate. Likewise, if Aaron and Miriam defer to a painter, a photographer, or some other expert about which objects are pure yellow, this is mere social deference, given that it's ultimately due to immutable physiological causes, and not the result of expertise. Of course, if we knew that the lemon was pure yellow, we could determine whom we should defer to, rather than just whom we happen to defer

to for social reasons. But without a prior reason to think that the lemon is pure yellow, we can only guess.

Let's next consider the version that appeals to an average across the population. One might think that if most people's perceptions of the lemon are phenomenal-pure-yellow, then phenomenal-pure-yellow perceptions represent the reflectance of the lemon, because that's what fixes the referent of that kind of perception. However, the statistics aren't that neat, especially if we temporarily set aside our simplifying assumption that there are only two possible perceptions of the lemon. In particular, perceptual variation isn't between two alternatives, and there isn't a simple majority who have the same perception. Thus, one would have to rely on a more abstract way of averaging perceptions, and it's unclear how one could choose between the alternatives. To start, it's unclear how one should assign numbers to all of the perceptions so that one could then calculate their average. It's also unclear what kind of average one should calculate. A simple average? A truncated average? A weighted average? And why not something different, such as an interpolated median? For these reasons, appealing to an average across the population merely pushes the question back, because in order to justify our belief that only Miriam's perception is accurate, we would first need to justify the decision to assign numbers in one way rather than another way, as well as to use one statistical measure rather than another statistical measure.

There's also a general problem with any view of perception that's modeled after the popular view about the referents of terms such as 'arthritis': what we perceive doesn't seem to depend on contingent facts about what others perceive. If everyone with eyes and brains unlike Miriam's died tomorrow, that wouldn't change what property she perceives, and thus wouldn't change whether her perception is accurate. It also wouldn't make a difference if they died shortly before her birth, or even never existed. While the meaning of 'arthritis' would change if doctors started using 'arthritis' differently, the color that Miriam is perceiving wouldn't change if people in her community started perceiving differently. I think this reflects a fundamental difference in the kinds of intentionality involved in perception and linguistic communication. When we speak, we intend to use words as others use them, but when we perceive, our perception doesn't include an intention to see what others see. Perception involves a less sophisticated kind of intentionality, at least in this respect. Thus, while the referents of our terms might depend on how others use them, what we perceive doesn't seem to depend on what others perceive.

(9) The ninth and final kind of evidence is about *one's own* perception of the lemon. According to some philosophers, if *you* perceive the lemon as pure yellow, you have defeasible justification for believing it's pure yellow, rather than greenish yellow. In that case, you might have defeasible justification for believing that anyone else who perceives the lemon as pure yellow is perceiving it accurately. As a result, you might have defeasible justification for believing that Miriam's perception is accurate.

But once you find out that others perceive the lemon differently, you cannot justifiably believe that the lemon is pure yellow on the basis of your own perception. In the jargon: your justification is *defeated*. Why? You now have evidence that the lemon is greenish yellow, because you have evidence that others perceive it as greenish yellow. You also have evidence that you're unreliable at perceiving fine-grained colors, because you have evidence that others perceive them differently, and no evidence that

you're the lucky person whose perceptions are most reliable. You thus have evidence that your perception is inaccurate.

There's a helpful analogy with lottery tickets. If you discover that your lottery ticket is one among many, you cannot justifiably believe that your ticket will win. Likewise, once you find out that others perceive the lemon differently, and that these differences are as arbitrary as the differences between lottery tickets, your evidence assigns a low probability to your perception being accurate, and you cannot justifiably believe that the lemon is pure yellow based on your perception.

To better understand why your evidence defeats the justification provided by your own perception of the lemon, let's consider your perception of motion. It's at least plausible that you *perceive* absolute motion. If you watch a bird fly past you on the beach, the bird doesn't just look like it is moving relative to you and the beach, so that you and the beach also look like you're moving away from it. The bird looks like it is moving *absolutely*. Likewise, when you look down at the ground, it doesn't just look motionless relative to you and your shoes, so that it might be moving relative to something else. The Earth looks like it is motionless *absolutely*. It's a surprise to learn that all motions are relative precisely because we usually perceive motion as absolute. At least, so one might argue. Set aside any reservations, and assume that this argument establishes that you perceive absolute velocities. Does it follow that you can know the absolute velocities of bodies on the basis of your perception? No. For example, suppose you perceive the Earth as absolutely motionless, and thus as having an absolute velocity of 0 km/s. Even if the Earth had an absolute velocity of 0 km/s, you couldn't know that based on your perception, because your evidence would assign a low probability to your perception being accurate. According to that evidence, if you were standing on another planet, you would perceive the Earth as moving, and there isn't a greater probability of the perception from your current location being accurate. Thus, believing that the Earth is absolutely motionless based on your perception would be like believing that your lottery ticket will win. Even if your ticket will in fact win, this isn't something you can know, given that your evidence assigns such a low probability to that possibility. Likewise, even if the Earth in fact has an absolute velocity of 0 km/s, and even if you perceive absolute velocities, this isn't something you could know on the basis of your perception.

Here's the underlying principle:

PERCEPTUAL DEFEASIBILITY

If your evidence assigns a low probability to your perception that p being accurate, you cannot justifiably believe that p based on your perception.

Two points might help clarify the appeal of this principle: First, it would otherwise be extremely difficult for you to lose justification for believing that p , since evidence against your reliability wouldn't be enough, even if that evidence assigned a low probability to your perception being accurate. For example, you could justifiably believe that the lemon is pure yellow even after you learned that you're probably wearing yellow-tinted contacts. Second, we otherwise couldn't reject absolute velocities on the grounds that they'd be unknowable, even though this seems like the right reason to

reject them.⁴ Without PERCEPTUAL DEFEASIBILITY, you could know that the Earth is absolutely motionless on the basis of your perception, even if your evidence assigned a low probability to your perception being accurate.

For these reasons, PERCEPTUAL DEFEASIBILITY should be uncontroversial. If this principle seems controversial, it's because it sounds like principles that have been rejected by Pryor (2000); Wright (2002); Goldman (2008); Williamson (2014); Lasonen-Aarnio (2014); Kelly (2010); and Wedgwood (2007). In a four-page appendix (Morrison 2020a), I distinguish PERCEPTUAL DEFEASIBILITY from those other principles, and explain why the philosophers who reject those principles would still accept PERCEPTUAL DEFEASIBILITY.

There is a way around PERCEPTUAL DEFEASIBILITY. I just argued that *after* you learn about perceptual variation, you can't use your own perceptions to know that Miriam's perception is accurate. This leaves open the possibility that you can use your own perceptions to know that Miriam's perception is accurate *before* you learn about perceptual variation. If you don't know about perceptual variation, your evidence might assign a high probability to your perception being accurate. Thus, under these special circumstances, you might be able to use your own perceptions to know that Miriam's perception is accurate. I'm going to set this possibility to the side, because it will be easier to explain what's unsatisfying about it once I've laid out the rest of my argument.

We just reviewed the nine kinds of evidence that might seem to help us justifiably believe that only Miriam's perception (or only Aaron's perception) is accurate, and thus know whose perception is accurate.

In other domains, there are other kinds of evidence. Consider the metaphysical debate about which objects combine to form wholes. Participants in that debate appeal to the intuition that there are objects that combine to form you and me (e.g., van Inwagen 1990, p. 73), the argument that for every n it's not vague whether n objects exist (e.g., Lewis 1986, pp. 212–213), and the ontological parsimony of denying that wholes exist (e.g., Dorr and Rosen 2002). We didn't review similar kinds of evidence, because the hypothesis that only Miriam's perception is accurate isn't supported or undermined by a priori philosophical arguments or intuitions, and neither hypothesis is more parsimonious.

Consider also the epistemological debate about whether we have justification for believing that material objects exist. Participants in that debate appeal to the explanatory power of the hypothesis that material objects exist (e.g., Russell 1912, Ch. 2), and the coherence of this hypothesis with our other beliefs (e.g., Bonjour 1976). We didn't review similar kinds of evidence, because both of our hypotheses are equally explanatory, and cohere equally well with our other beliefs.

Some philosophers have expressed the hope that our ultimate psychosemantic theory will settle whose perception is accurate (Byrne and Hilbert 2003, p. 17). But this underestimates the significance of the challenge of perceptual variation. In particular, we won't be able to settle on the correct psychosemantic theory *until* we settle on a response to the challenge of perceptual variation. Consider two psychosemantic theo-

⁴ Some reject absolute velocities on the grounds that they'd be *redundant*. But Dasgupta (2016, pp. 844–850) persuasively argues that this is a good reason to reject absolute velocities only if 'redundant' is taken to mean something like 'undetectable' or 'unknowable'.

ries that are otherwise equivalent, except that one implies that Miriam's perception is accurate and Aaron's perception is inaccurate, while the other implies that Aaron's perception is accurate and Miriam's perception is inaccurate. Perhaps one theory implies that Miriam is an ideal perceiver while the other theory implies that Aaron is an ideal perceiver. Given what we just established, we shouldn't expect any of our evidence to support one theory over the other. For example, we shouldn't expect one to be simpler or more explanatory. It is blind optimism to hope that we'll eventually find evidence that allows us to justifiably choose between these theories.

I conclude that our current and future evidence is unlikely to justify our belief that only Miriam's perception is accurate; the most promising kinds of evidence don't support this hypothesis over the alternative. We can generalize this conclusion in a number of ways.

Once we drop the simplifying assumption that there are only two possible perceptions of the lemon, our arguments establish that we can't know whose perception is *more accurate*. Suppose that another kind of perception is accurate, and Miriam's and Aaron's perceptions are both inaccurate. We can still ask which of their perception is more accurate. Analogously, 'Beirut is exactly 100km from Damascus' and 'Beirut is exactly 200km from Damascus' are both inaccurate descriptions of Beirut's location, because Beirut is really 106.19km from Damascus, but we can still ask which description is more accurate. If the evidence we reviewed doesn't justify the belief that Miriam's perception is accurate, it also doesn't justify the belief that Miriam's perception is more accurate, because it doesn't give us any reason to prefer her perception to Aaron's perception. Thus, we're similarly ignorant of whose perception is *more accurate*.

Next, our ignorance extends beyond Miriam's and Aaron's perceptions of individual colors to their perceptions of color relations, such as yellower, bluer, and greener. For example, because Miriam perceives the lemon as pure yellow, she doesn't perceive it as greener than any other objects, but because Aaron perceives the lemon as greenish yellow, he perceives it as greener than at least some other objects. Whose perceptions of those relations are accurate? None of the evidence justifies one choice or the other, and for the same reasons it didn't justify a choice between their perceptions of the lemon's individual color. Thus, given one-ism, our ignorance extends to their perceptions of color relations.

Third, in our review of the evidence, we assumed reflectance physicalism, the view that colors are reflectances. Our conclusion straightforwardly generalizes to the other views. If colors are primitive properties, we might be able to know which colors Aaron and Miriam are perceiving, by relying on evidence about their phenomenal characters. In particular, we might be able to rely on a necessary and immediately obvious connection between their phenomenal characters and the colors that they are perceiving (e.g., Campbell 1993). But we couldn't know which of these colors supervenes on the lemon's surface, assuming with the one-ist that it has only one color. Evidence about their eyes, their reports, the lemon's reflectance, and so on, wouldn't help, because this evidence still wouldn't help us know whether the lemon has the color Miriam is perceiving or the color Aaron is perceiving, and thus wouldn't help us know whose perception is accurate.

If colors are dispositions, we could know even less, assuming with the one-ist that only one kind of subject and only one kind of context is relevant. None of the evidence we reviewed would help us choose between the many different normal subjects who perceive the lemon differently, or between the many different normal contexts in which the same subject will perceive the lemon differently. Thus, we couldn't know which colors Aaron and Miriam are perceiving, because they would be perceiving dispositions involving an unknown kind of subject and an unknown kind of context. We also couldn't know whose perception is accurate, because we couldn't know how the lemon would affect the unknown kind of subject in the unknown kind of context.

3 Unacceptable Ignorance

Suppose I'm right, and we shouldn't expect to know whose perception is accurate. So what? There's a lot we shouldn't expect to know, including the current number of stars in distant clusters, the aggregate weight of all the chocolate in existence, and Socrates's exact height when he drank hemlock. And nobody should respond by denying that there is still a fact of the matter in all these cases. So why would it be a problem if, given one-ism, we're ignorant about whose perception of the lemon is accurate?

This, in essence, is how color one-ists such as Stroud (2000, pp. 173–176); Tye (2000, p. 108); Byrne and Hilbert (2003, pp. 16–17; 2004, pp. 37–39; 2007) and Allen (2016, pp. 58–65) respond when asked to identify the person whose perception is most accurate. Here are two representative passages:

From the fact that we have no good reason to believe, of any chip, that it is unique green, it does not follow that we have no good reason to believe that there are any unique green chips. That would be like arguing that we have no good reason to believe that Professor Plum has been murdered, on the ground that there is no particular person who is clearly the culprit. (Byrne and Hilbert 2003, p. 17)

God knows precisely which hue chip 527 has, but we may very well never know. Our only access to the colours of things is via a single sense and the colour detectors nature has endowed us with are limited. We do not suppose that objects do not have precise lengths because of the limitations of our measuring equipment. Why suppose that the situation is fundamentally any different for the case of colour? (Tye 2006, pp. 177–178)

According to these philosophers, we shouldn't reject one-ism even if it implies we can't know whose perception is accurate.

But some kinds of ignorance are more acceptable than others. In all of the examples of ignorance that I listed, there are identifiable causal processes preventing us from collecting the relevant evidence. We can't know the number of stars in distant clusters because, given the distance, the relevant evidence (e.g., propagating light from those stars) isn't fast enough to have reached us. We can't know the aggregate weight of all the chocolate in existence because, given the number of locations where there's chocolate, and the speed at which that chocolate is created and consumed, we can't keep track of all of the relevant evidence. We can't know Socrates's exact height when

he drank hemlock because, due to decomposition, the relevant evidence was lost. There are also identifiable causal processes in the examples from Byrne and Hilbert and Tye. The murderer of Professor Plum might have left the room and disposed of the weapon, thereby destroying the only evidence that would have allowed us to know who did it (see Cohen 2003). Due to oscillations in the internal temperature of our equipment, that equipment might vibrate, expand, and contract, and as a result provide only probabilistic evidence about an object's exact length. In all of these examples, there is (or was) observable evidence, and that evidence is (or was) in principle collectable, but we're unable to collect it due to an identifiable causal process. As a result, there's no pressure to reject the natural assumptions that the stars in distant clusters are countable, that the aggregate of all chocolate in existence has a weight, that Socrates had an exact height at the time of his death, that someone murdered Professor Plum, or that objects have exact lengths.

Assuming one-ism, our ignorance about the accuracy of Aaron's and Miriam's perceptions goes deeper. We would still be ignorant even if we had perfect evidence about Miriam's and Aaron's eyes, brains, phenomenal characters, shared evolutionary history, etc. Thus, unlike our ignorance of the past, etc., this ignorance isn't the result of an identifiable causal process. There also aren't any a priori philosophical arguments or intuitions, and appeals to explanatory power, ontological parsimony, and coherence don't favor one hypothesis over the other. In all these respects, our ignorance seems more like our ignorance of absolute velocity (see Dasgupta 2015, p. 610). Thus, to the extent that we're inclined to reject the assumption that there are absolute velocities, we should also be inclined to reject the assumption that only one person's perception is accurate. In both cases, our ignorance would be so unacceptable that we should instead reject whatever assumption implies that there's something we don't know.

The underlying principle is hard to state. But it goes something like this:

UNACCEPTABLE IGNORANCE

Suppose that an assumption implies that a list of hypotheses is exhaustive. If we don't expect to justifiably believe the first hypothesis, the second hypothesis, etc., even given unlimited time and resources, and our ignorance doesn't have the right kind of causal explanation, then we should reject that assumption.

This principle is about what we should believe; it's epistemological rather than meta-physical. There's thus no guarantee that it will lead to a true belief. But it's still our best bet. In this respect, it's similar to induction and inference to the best explanation

I'm going to argue that this principle is incompatible with one-ism. In particular, given our simplifying assumption that there are only two possible perceptions of the lemon, one-ism implies that there are only two hypotheses about Miriam and Aaron: that Miriam's perception is accurate and Aaron's perception is inaccurate, and that Miriam's perception is inaccurate and Aaron's perception is accurate. In the last section, I argued that we shouldn't expect to justifiably believe either of these hypotheses, even given unlimited time and resources. I'll later argue that there isn't the right kind of causal explanation of our ignorance. Assuming these arguments are successful, UNACCEPTABLE IGNORANCE implies that we should reject one-ism.

Proponents of one-ism sometimes accuse their critics of logical positivism (e.g., Byrne and Hilbert 2007, p. 88). It's therefore worth mentioning that this principle is much weaker than the principles logical positivists use to deny that metaphysical debates are substantial (for an overview, see Creath 2017, Sect. 4.1). For example, it doesn't place any constraints on the kinds of evidence that we might use to justifiably believe a hypothesis. Consider again the metaphysical debate about which objects combine to form wholes. Participants in that debate take themselves to have evidence for their hypotheses, enough to justify believing those hypotheses. UNACCEPTABLE IGNORANCE thus doesn't imply that there can't be a substantial debate about which objects combine to form wholes. It likewise doesn't imply that there can't be substantial debates about the ontological status of numbers, propositions, properties, and possible worlds. In contrast, the principles used by logical positivists are supposed to imply that these debates are unintelligible. In addition, UNACCEPTABLE IGNORANCE doesn't require the relevant hypotheses to be *verifiable*, at least in the sense required by many logical positivists. It just requires the hypotheses to be justifiably believable. In this respect too, it is much weaker than the principles used by the logical positivists.

I already described the role UNACCEPTABLE IGNORANCE plays in our thinking about motion. In what follows, I'll describe the role it plays in our thinking about time, vagueness, and ethics. These examples will serve a number of purposes. First, they will motivate UNACCEPTABLE IGNORANCE by showing that it unifies the way we think about diverse phenomena. Second, they will clarify what counts as the "right kind of causal explanation." Third, they will clarify the narrow range of circumstances in which some philosophers might be willing to accept ignorance despite UNACCEPTABLE IGNORANCE, helping us later establish that ignorance about whose perception is accurate would be less acceptable than the kinds of ignorance some are willing to accept in other domains.

After presenting the three examples, I will restate the problem for one-ism, taking into account what we learned.

(1) Let's start with time. In particular, let's start with absolute simultaneity, i.e., simultaneity that isn't relative to any particular frame of reference. If two stars erupt, forming supernovas, it's natural to assume that their eruptions are either absolutely simultaneous or absolutely sequential. However, that's not something we can know, or expect to know, given the laws of the Special Theory of Relativity. More generally, given the laws of the Special Theory of Relativity, if two events are absolutely simultaneous, we can't know it. From some reference frames, the eruptions will appear simultaneous; from other reference frames, the eruptions will appear sequential; and we can't know which appearance is accurate, assuming that only one of them is.

Most philosophers and physicists conclude that we should reject the assumption that the eruptions are absolutely simultaneous or absolutely sequential. They instead relativize simultaneity to reference frames, so that events can be simultaneous relative to one reference frame, and sequential relative to another reference frame. These philosophers and physicists thereby avoid committing us to ignorance about the "true" frame of reference. According to them, we can know that both appearances are accurate. Their view of absolute simultaneity thus mirrors their view of absolute velocity: just as they reject the assumption that bodies are either absolutely moving or resting, they reject the assumption that events are either absolutely simultaneous or sequential.

While I'm not aware of any philosophers or physicists who hold onto the assumption about absolute velocity, there are philosophers and physicists who hold onto the assumption about absolute simultaneity. These philosophers and physicists recognize that they owe us a reason. To satisfy this burden, Zimmerman appeals to presentism, the view that only the present moment exists. He thinks that this is a sufficient reason to hold onto the assumption, because he regards presentism as close to a truism (see Zimmerman 2007, p. 221f; Prior 1972, p. 323; Crisp 2003, pp. 232–235). Bell takes a different approach. He argues that quantum mechanics might require a “true” reference frame, and thus absolute simultaneity. In particular, he says it might be necessary to explain quantum non-locality (Bell 2004, pp. 179–180; see also Bohm 1952; Maudlin 1994; Lucas 1998, p. 55). If it does, that would be a reason to hold onto the assumption that all events are absolutely simultaneous or sequential, even if we can't always justifiably choose between competing hypotheses about their temporal ordering.

For our purposes, it's unimportant whether Zimmerman and Bell are right about presentism or quantum mechanics. What's important is that they recognize that they owe us a reason to hold onto this assumption despite our ignorance, and that they regard truisms and other scientific theories as sufficient.

(2) Let's now turn to vagueness. It is natural to assume that all declarative sentences are either true or false.⁵ But, due to vagueness, it seems that we can't know whether some declarative sentences are true or false. For example, if Boaz is a borderline case of baldness, it seems that we can't know that 'Boaz is bald' is true and we can't know that 'Boaz is bald' is false.

Many philosophers infer that 'Boaz is bald' is not true and 'Boaz is bald' is not false, thereby rejecting the assumption that all declarative sentences are either true or false. These philosophers evaluate 'Boaz is bald' differently, perhaps as neither-true-nor-false. According to these philosophers, we can know that 'Boaz is bald' is neither-true-nor-false (or whatever other status they assign to borderline claims) because we can know that Boaz is a borderline case of baldness. These philosophers thereby avoid committing us to ignorance about whether 'Boaz is bald' is true by *expanding* the list of hypotheses to include another evaluation. Other philosophers *contract* the list of hypotheses by removing an evaluation. They claim that all sentences containing vague terms are false, and thus we can know that 'Boaz is bald' is false (Braun and Sider 2007). Still other philosophers avoid committing us to ignorance by claiming that it's indeterminate whether we know that it's true or know that it's false (Barnett 2011; Dorr 2003). According to these philosophers, we aren't ignorant, because it's not determinate that we lack knowledge.

There are also some philosophers (“epistemicists”) who hold onto the natural assumption that all declarative sentences are either true or false. They insist that 'Boaz is bald' is true or 'Boaz is bald' is false, despite our ignorance of which. They recognize that they owe us a reason to hold onto this assumption, given that it leads to ignorance. To satisfy this burden, some attribute our ignorance to a causal process. For example, according to Williamson (1994, p. 230f), the meaning of 'bald' depends

⁵ Possible counterexamples include ambiguous sentences (“Your friends are at the closest bank”), sentences about the future (“There will be a sea battle tomorrow”), and epistemic modals (“It might be raining outside”). Even if these are genuine counterexamples to bivalence, it is natural to assume that all other declarative sentences, including those involving vague predicates, are either true or false.

on how this term is currently used throughout our linguistic community, including the current dispositions of other members of our community to describe certain people as bald. On this basis, Williamson argues that we can't know the exact meaning of 'bald' at any given time, because its current use changes too quickly. Thus, according to Williamson, our ignorance of the truth-value of 'Boaz is bald' is like our ignorance of the exact weight of all the chocolate in the world, in that it depends on a large amount of evidence that is in constant flux. This evidence is in principle collectable. It's just that we're unable to collect all of it. According to Williamson, we can therefore believe that that this sentence is true or false, despite our ignorance about which.

Other epistemicists appeal to principles they regard as close to truisms. For example, Horwich (1990, pp. 81–87) appeals to classical logic and the truth schema. According to Horwich, if these principles are true, then 'Boaz is bald' is true or 'Boaz is bald' is false, even though we can't know which.

For our purposes, it's irrelevant whether Williamson and Horwich are right. What's important is that they recognize that they owe us either a causal explanation or a reason to hold onto our assumption, and that in this case a sufficient reason might be a bundle of principles that are close to truisms.

One of the reasons why this is a helpful example is that it shows that UNACCEPTABLE IGNORANCE applies to semantic assumptions, in this case the assumption that 'Boaz is bald' is *true* or 'Boaz is bald' is *false*. This is helpful because we're interested in a similar kind of assumption, namely the assumption that there's only one kind of perception of the lemon that's *accurate*. Ignorance can be unacceptable even when it's about semantics rather than physics or metaphysics.

(3) Let's finally turn to ethics, in particular to ethical permissibility. It is natural to assume that ethical permissibility is absolute, in that it isn't relative to a culture, community, individual, time of assertion, etc. But some disagreements about ethical permissibility seem unresolvable, at least initially. Let's focus on an example mentioned earlier, the disagreement about whether it's ethically permissible to eat meat.

Many philosophers claim that this disagreement is resolvable. Among other sources of evidence, they think we can rely on intuitions about cases (e.g., Norcross 2004), rights theory (e.g., Regan 1983), or a thorough accounting of all the effects of eating meat (e.g., Singer 1975). According to these philosophers, we can know whether eating meat is absolutely permissible.

Philosophers who insist that we can't know often infer that ethical permissibility is relative to a culture, community, individual, etc. These philosophers thereby avoid committing us to ignorance about whether eating meat is ethically permissible, because we just need to consult the relevant culture, community, individual, etc.

I'm not aware of any philosopher who actually says that eating meat is absolutely permissible or absolutely impermissible, even though we can't know which. However I can think of at least two groups of philosophers who could say that, at least in principle.

The first are utilitarians. Utilitarians might think that we can't know whether eating meat is ethically permissible, because we can't predict all of its effects. In that case, our ignorance might be due to the number of effects of eating meat, as well as uncertainty about those effects.

The second are philosophers who think that ethical knowledge requires “imaginative acquaintance.” According to Smith, Lewis, and Johnston (1989), ignorance is sometimes due to our own psychological limitations, in particular our inability to become imaginatively acquainted with all the relevant facts. They think that these psychological limitations have a causal explanation: we don’t have the right kind of brain to become imaginatively acquainted with all the relevant facts, just as we don’t have the right kind of brain to become imaginatively acquainted with what it’s like to be a pig. Thus, they could give the right kind of causal explanation of our ignorance. For example, we might not know whether eating pork is morally permissible because we don’t know what it’s like to be a pig.

For our purposes, it’s again irrelevant whether any of these positions are ultimately defensible. What’s important is that everyone seems to agree that ethical ignorance is acceptable only if there’s the right kind of causal explanation.⁶

These examples motivate UNACCEPTABLE IGNORANCE by showing that it underlies our thinking about diverse phenomena. Giving up UNACCEPTABLE IGNORANCE would thus come at a high cost. For example, it would be hard to reject absolute motion. It would also be hard to deny that there is a “true” reference frame, even if you don’t regard presentism as a truism, or think that quantum mechanics requires it. It would be similarly hard to deny that there are unknowable semantic facts, including not just the truth of borderline claims, but also the referents of our terms and the validity of our arguments. And it would be hard to rule out ethical views that imply we can’t know which actions are ethically permissible. Rejecting UNACCEPTABLE IGNORANCE has the potential to disrupt our thinking about parts of philosophy that wouldn’t have otherwise seemed connected.

These examples also give us a better sense of what counts as a causal explanation of the “right kind,” and which reasons might be sufficient to nonetheless believe that one of the relevant hypotheses is true.

What counts as a causal explanation of the “right kind”? A causal explanation of the right kind presupposes that there are observable effects, and explains why we can’t observe them. Such an explanation might identify the process that destroyed the evidence, the amount of time it will take for the evidence to reach us, the rate at

⁶ Philosophers who accept the principle that “ought implies can” might have another reason to avoid ignorance about moral permissibility. Understood in a certain way, this principle might imply that, if eating meat is morally impermissible, we must be able to know that it is morally impermissible, because otherwise we can’t fulfill our obligation not to eat meat. But even if this principle helps explain why some philosophers regard ethical ignorance as unacceptable, it doesn’t explain why others do. First, while “ought implies can” is popular, it is controversial. As far as I’m aware, the philosophers who deny that “ought implies can” still reject ethical ignorance (e.g., Sinnott-Armstrong 1984; Graham 2011; King 2014). They must have other reasons for rejecting it. Second, those who accept the principle usually focus on obligations that we can’t fulfill because of physical limitations (e.g., a donor who can’t give his kidney to two friends), psychological limitations (e.g., an addict who can’t refuse alcohol), or practical limitations (e.g., a homeowner who can’t repay his loan) (see King 2019, Ch. 1). That is, they focus on the act itself, rather than our knowledge about the act. It’s therefore unclear whether the principle, as understood by its proponents, implies that we must be able to know about our moral obligations. Third, regardless of how the principle is understood, it’s unclear that it implies that we must be able to know about *all* of our moral obligations. For example, if we’re just ignorant about the moral permissibility of a small range of actions (e.g., eating meat, hiring prostitutes, committing suicide), we can make sure that we fulfill our obligations by avoiding all of those actions, at the potential cost of avoiding some actions that might not really be morally impermissible.

which that evidence is changing, the distorting influence of our measuring instruments, indeterminacies in the processes responsible for future events, limitations on our ability to imagine unfamiliar kinds of consciousness, and so on. A causal explanation of this kind doesn't just establish that we can't collect the relevant evidence. After all, that's true for absolute velocities, because the physical laws entail that there's no way to set up a causal link between absolute velocities and observable effects.

Which reasons are sufficient to believe one of the relevant hypotheses is true, despite our ignorance of which? It would be enough that it's indispensable to our best scientific understanding of the world that one of these hypotheses is true. It would be enough if there wasn't a coherent alternative. More controversially, it might be enough to preserve principles that are so central to the way we think and talk that they approach truisms.

Let's now restate the problem, taking into account what we just learned. Given one-ism, there are only two hypotheses: that Miriam's perception is accurate and Aaron's perception is inaccurate, and that Miriam's perception is inaccurate and Aaron's perception is accurate. But this conflicts with UNACCEPTABLE IGNORANCE, the principle that:

UNACCEPTABLE IGNORANCE

Suppose that an assumption implies that a list of hypotheses is exhaustive. If we don't expect to justifiably believe the first hypothesis, the second hypothesis, etc., even given unlimited time and resources, and our ignorance doesn't have the right kind of causal explanation, then we should reject that assumption.

This principle lists two conditions. Let's consider them separately:

First, we shouldn't expect to justifiably believe any of these hypotheses, because we shouldn't expect our evidence to support either of them (see Sect. 2). Recall that, for simplicity, we're assuming that Miriam's and Aaron's perceptions of the lemon are the only two kinds. If we drop this assumption, and take into account all of the other perceptions of the lemon, then there will be many more hypotheses. But in each hypothesis, one kind of perception will be accurate, and we should shouldn't expect our evidence to support any of them, given our arguments in the last section.

Second, our ignorance doesn't have the right kind of causal explanation, because our ignorance isn't due to factors such as information loss, physical distance, or instrument error. Even if we had complete information about Aaron's and Miriam's eyes, brains, phenomenal characters, and evolutionary history, as well as complete information about the lemon's reflectance, we'd still be ignorant. We *can* explain why the lemon affects Aaron and Miriam differently. But that's not the right kind of causal explanation. To see why, consider that we can explain why the Earth appears motionless to someone standing on it, and in motion to someone on another planet; why two supernovas appear simultaneous from one reference frame, and sequential from another reference frame; why eating meat seems ethically permissible in one culture, and ethically impermissible in another culture. But nobody should think that these explanations are enough to establish that ignorance of absolute motion, simultaneity, or ethical permissibility would be unproblematic. Our ignorance would still call into question the assumptions that the Earth is absolutely at rest or in motion, that the supernovas are

absolutely simultaneous or sequential, and that eating meat is absolutely permissible or impermissible.

Is there a sufficient reason to believe that one of their perceptions might be accurate, even though we can't know whose? It would be enough if our scientific understanding of the world required it. But our scientific understanding of the world doesn't imply anything about the accuracy of our color perceptions. It doesn't even imply that objects are colored. Scientists can explain every stage in the causal progress leading up to Miriam's and Aaron's perceptions without describing the lemon as colored. It's notable that, whereas most chemists presumably believe that chemicals are real, many perceptual psychologists deny that anything is really colored (see Byrne and Hilbert 2003, pp. 3–4). It's our natural ways of thinking and talking, not our scientific theories, that push us to attribute colors to objects. Thus, our scientific theories are compatible with at least one of the alternatives to one-ism, namely the view that all perceptions of the lemon's color are inaccurate, because the lemon isn't really colored.

It might be enough if there were a principle approaching a truism that implied that one of Aaron's and Miriam's perceptions is accurate. But the motivations for one-ism are much weaker. Let's consider them one-by-one. First, our natural ways of thinking and talking motivate one-ism. For example, if Miriam reports that the lemon is pure yellow, Aaron might disagree, insisting that it is greenish yellow. But these ways of thinking and talking shouldn't be given too much weight. Consider that we also naturally think and talk about motion as absolute, but that isn't a reason to think that motion is absolute, despite arguments to the contrary. Our natural ways of thinking and talking are often ill-informed and unreflective, and should be revised as the result of new evidence and philosophical reflection (a point that Hawthorne and Kovakovich 2006, pp. 180–181, also emphasize). If our natural ways of thinking and talking motivate one-ism, it's because we naturally assume that objects appear only one way to normal subjects under normal conditions. Once we learn about perceptual variation, we should be prepared to start talking and thinking in new ways, or else stop treating our old way of talking and thinking as a guide to reality.

Second, as mentioned in the introduction, a weak version of representationalism about color perception motivates one-ism:

REPRESENTATIONALISM

If two perceptions have the same phenomenal character, they represent the same color.

Given that *your* phenomenal-pure-yellow perceptions and *your* phenomenal-greenish-yellow perceptions represent incompatible colors, this version of representationalism implies that Miriam's phenomenal-pure-yellow and Aaron's phenomenal-greenish-yellow perception represent incompatible colors, and thus at most one of their perceptions can be accurate.

Representationalism is a popular view, with a number of attractive features (see Lycan 2014). But it's nowhere near a truism. As with our natural ways of thinking and talking, we should therefore be prepared to give it up as the result of new evidence and further philosophical scrutiny. One option is to weaken it, perhaps restricting it to a person. For example:

INTRAPERSONAL REPRESENTATIONALISM

For any person *S*: If two perceptions *in S* have the same phenomenal character, they represent the same color.

In that case, we can't assume that *your* phenomenal-greenish-yellow perceptions and *Aaron's* phenomenal-greenish-yellow perceptions represent the same color, or that *your* phenomenal-pure-yellow perceptions and *Miriam's* phenomenal-pure-yellow perceptions represent the same color, thereby blocking the conclusion that *Miriam's* phenomenal-pure-yellow perceptions and *Aaron's* phenomenal-greenish-yellow perceptions represent incompatible colors. This opens up the possibility that Miriam and Aaron are perceiving the same color, as well as the possibility that they are perceiving different but compatible colors, and thus are both accurately perceiving the lemon, despite their phenomenal differences (see Block 1999). This weakening of REPRESENTATIONALISM might not preserve all of the attractive features of the original, but it might preserve enough of them. For example, it might still preserve what's called "transparency", because transparency is about how things seem to a subject, and is thus preserved by a principle restricted to a subject. In any case, REPRESENTATIONALISM isn't a truism, and thus doesn't give us a sufficient reason to think that one of their perceptions is accurate.

In addition, REPRESENTATIONALISM isn't sufficient for one-ism. REPRESENTATIONALISM is consistent with the hypothesis that all of our color perceptions are illusory, and thus that neither Miriam nor Aaron is accurately perceiving the lemon (see, e.g., Pautz 2010, pp. 58–60; Chalmers 2006 on Edenic colors). REPRESENTATIONALISM is sufficient for one-ism only when combined with the additional assumption that at least some of our color perceptions are accurate, and this isn't anywhere near a truism, for the reasons mentioned above. As a result, even if REPRESENTATIONALISM were a truism, that wouldn't be enough for one-ism, because it would just push us to reject the assumption that at least some of our color perceptions are accurate.

Third, one-ists sometimes report having a *direct intuition* that all phenomenal-greenish-yellow perceptions and all phenomenal-pure-yellow perceptions represent incompatible colors, without explicitly mentioning a principle like REPRESENTATIONALISM (see, e.g., Tye 2012, pp. 300–301; Byrne and Hilbert 2007, pp. 87–89). I suspect that these philosophers are implicitly relying on a principle like REPRESENTATIONALISM. After all, their intuition is the result of reflection on their own perceptions, so to infer anything about other people's perceptions, they must be relying on a principle like REPRESENTATIONALISM. In that case, this motivation collapses into the motivation we just considered. Moreover, even if there were a direct intuition for the conclusion that all phenomenal-greenish-yellow perceptions and all phenomenal-pure-yellow perceptions represent incompatible colors, we should reject that intuition, perhaps in favor of the weaker intuition that *one's own* phenomenal-greenish-yellow perceptions and phenomenal-pure-yellow perceptions represent incompatible colors. Philosophical intuitions are important, but they aren't sacred, and we should be prepared to give them up once they're shown to lead to unacceptable ignorance.

We're considering whether any of the motivations for one-ism are powerful enough to justify holding it despite its conflict with UNACCEPTABLE IGNORANCE. There's a final motivation worth considering. One thing that Aaron and Miriam have in common

is that they both perceive the lemon *as yellow*. Plausibly, both of their perceptions are accurate. According to some philosophers, the simplest explanation of why they're both accurately perceiving the lemon as yellow is that they're both perceiving shades of yellow, such as pure yellow and greenish yellow. Assuming that the lemon is only one shade of yellow, it seems to follow that only one of their perceptions of the lemon's shade is accurate. Thus, one-ism might be a consequence of the simplest explanation of why our perceptions of more coarse-grained colors like yellow are accurate (see Allen 2009, p. 202; Allen 2016; it might also be motivating Tye 2006, see p. 177). I'm not convinced that this is the simplest explanation. I think that an equally simple explanation, mentioned above in our discussion of REPRESENTATIONALISM, is that Aaron and Miriam are perceiving the same shade of yellow despite the phenomenal differences between their perceptions (for proposals along these lines, see Jackson and Pargetter 1987; McLaughlin 2003; Matthen 2009; Morrison 2020b). But that would take a long time to establish. For now, I just want to point out that, even if this were the simplest explanation, it wouldn't be enough to believe that only one of their perceptions is accurate. Consider motion. If absolute motion is fundamental, the planets have just eight fundamental motions. But if relative motion is fundamental, the planets have twenty-eight fundamental motions. Thus, the simplest explanation of planetary motion is that the planets have absolute motions. But, even if that is the simplest explanation, we should still deny that the planets have absolute motions, because we can't know which absolute motions. Also consider vagueness. The simplest explanation of why the disjunction 'Boaz is bald or not bald' is true might be that one of its disjuncts is true. But, even if that is therefore the simplest explanation, as long as there isn't a causal explanation of our ignorance, etc., we should still deny that one of its disjuncts is true, because we can't know which. Simplicity is a theoretical virtue, and it gives us a reason to prefer one hypothesis to another. But it does not give us a reason to accept assumptions that lead to unacceptable ignorance.

I conclude that there isn't a sufficient reason to hold onto one-ism, despite its conflict with UNACCEPTABLE IGNORANCE.

The one-ist has at least one remaining option, mentioned earlier: deny that her view gives rise to ignorance, on the grounds that you can know that Miriam's perception is accurate on the basis of your own perception *before* you learn that others (e.g., Aaron) perceive the lemon differently. One problem with this option is that it seems incompatible with three plausible conditions on knowledge. First, your belief might then be too lucky to count as knowledge, because only the lucky few who accurately perceive the lemon's color could know that Miriam's perception is accurate. Second, your belief might not count as knowledge because it's not shareable. As soon as you starting telling people, you'd learn about perceptual variation, thereby undermining your knowledge. Third, your belief might not count as knowledge because you can't rule out all the relevant alternatives. The possibility that others perceive the lemon differently might be so obviously relevant that you might need evidence to rule it out even before you learn that perceptual variation is actual.

More would need to be said in support of these points. But, for our purposes, a less controversial point should suffice: Otherwise someone could know that the Earth is absolutely motionless, because if they are lucky enough to accurately perceive the Earth as absolutely motionless, they could know that it is absolutely motionless

before they learn that it would appear to be moving from another location. Similarly, someone could know that two eruptions are absolutely simultaneous before they learn that they appear sequential from another reference frame. And so on, for all of our other examples. When deciding which fundamental assumptions to accept, what's relevant is what *we* can justifiably believe, given our collective evidence, not what an individual can justifiably believe, given her evidence at a particular time.

In Sect. 2, I argued that one-ism leads to ignorance. In Sect. 3, I argued that this is an unacceptable kind of ignorance. I conclude that one-ism should be rejected because it commits us to an unacceptable kind of ignorance—a kind of ignorance far less acceptable than what people are willing to accept in other domains. We should therefore prefer another response to the challenge of perceptual variation. An alternative, mentioned above, is that neither of their perceptions is accurate, because all of our color perceptions are illusory. In that case, there's only one hypothesis (viz., Miriam's and Aaron's perceptions are both inaccurate) and we can justifiably believe it. Another alternative, also mentioned above, is that both of their perceptions are accurate, despite their phenomenal differences. In that case, there's another hypothesis (viz., Miriam's and Aaron's perceptions are both accurate) and we might be able to justifiably believe it. Thus, unlike one-ism, these other responses don't lead to unacceptable ignorance.

4 Conclusion

We focused on actual variation in color perception, which is moderate. More extreme variations are conceivable. For example, we can conceive of a creature, Adam, whose perception of the lemon is phenomenal-pure-red, and whose perception of an apple is phenomenal-pure-yellow. More generally, we can suppose that the phenomenal characters of his perceptions are the inverse of the phenomenal characters of Miriam's perceptions. Adam is what's called a “phenomenal invert”. This gives rise to the parallel challenge of saying whose perception is accurate, Adam's or Miriam's.

Why did we focus on a challenge involving actual variation, rather than merely possible variation (unlike Pautz 2006; Chalmers 2006)? Suppose that, in the actual world, all and only phenomenal-pure-yellow perceptions represent the color of the lemon (e.g., its reflectance). It would then be tempting to *identify* having that phenomenal character with representing that color. As a result, it would be tempting to deny the metaphysical possibility of a perception that has a different phenomenal character but nonetheless represents that color. Even if such a perception is conceivable, it would be tempting to insist that any perception with a different phenomenal character must represent a different color, and thus cannot be an accurate perception of the lemon's color. The resulting debate would hinge on subtle questions about the relation between conceivability and possibility, just like the debate about mind-body identity (see Byrne 2016, Sec 2.4; Pautz 2014, Sec 5; Tye 2015, p. 200). Focusing on actual variation allowed us to sidestep that debate.

But that doesn't mean the challenge involving merely possible variation isn't worth considering. In fact, given what we established about actual variation, we're now in a better position to address merely possible variation. We established that, if Miriam's phenomenal-pure-yellow perception of the lemon's color is accurate, then Aaron's

phenomenal-greenish-yellow perception of the lemon's color might also be accurate. We thus established that, if all phenomenal-pure-yellow perceptions represent the color of the lemon, then perceptions of other kinds also represent the color of the lemon, undermining the temptation to identify having that phenomenal character with representing that color. We thereby clarified and strengthened the challenge involving merely possible variation.⁷

With that in mind, how should we respond to this other challenge? Is Adam's or Miriam's perception accurate? I don't think that we should say that only one of their perceptions is accurate, and for the same reasons we shouldn't say that only Aaron's or only Miriam's perception is accurate. In particular, as long as Adam is just as successful at avoiding traffic, selecting fruit, matching socks, and performing all of the other tasks that our color perceptions contribute to, I don't think there's any evidence that could justify our belief that Miriam, rather than Adam, is accurately perceiving the lemon. Thus, insisting that only one of their perceptions is accurate would give rise to a kind of ignorance that's just as unacceptable as our ignorance with respect to actual variation.

This sharpens the alternatives to one-ism: either Adam's and Miriam's perceptions are both inaccurate, and color perception is necessarily illusory, or Adam's and Miriam's perceptions can both be accurate, and color perceptions can be accurate regardless of their phenomenal characters. While I prefer the second option, I won't try to motivate that choice here. For now, I just want to point out that we're left with fundamental questions about the role of phenomenal characters in perception. The challenge of perceptual variation thereby shines a spotlight on some of the most fundamental questions we can ask about the mind.⁸

⁷ There's another reason why it's better to start with actual variation. Suppose that, in the actual world, everyone who accurately perceives the color of the lemon has some property P , and that this property Footnote 7 continued

explains the accuracy of their perceptions. P might have something to do with the configuration of their eyes and brains, the kinds of reports they make, their phenomenal characters, their evolutionary history, or their statistical relation to the rest of the population, among other options. We might not even know much about P . Whether spectrum inversion is metaphysically possible would then depend on whether it's possible for a creature to have a different phenomenal character, but nonetheless have P , and thus to accurately perceive the lemon's color. To argue for that metaphysical possibility, we would need to consider as many of the candidates for P as we can. This would invite the one-ist to respond that we hadn't considered the relevant P , or to modify their initial proposal about P . As an example, see Byrne and Tye's (2006, pp. 252–254) response to Pautz (2006, pp. 220–227). By first establishing that perceptions with *moderately* different phenomenal characters can be accurate, we establish that perceptions with moderately different phenomenal characters can all have P , without our needing to individually consider all of the candidates for P . This makes it easier to then establish that perceptions with *extremely* different phenomenal characters can have P . We would just need to argue that, because perceptions with moderately different phenomenal characters can all have P , whether a perception has P has to do with relations to its environment that are independent of that perception's particular phenomenal character. It would follow that perceptions with extremely different phenomenal characters can all have P , and thus be accurate, without our needing to individually consider all of the candidates for P . This is another way in which the challenge involving actual variation strengthens the puzzle involving merely possible variation.

⁸ This paper began as a section of my "Perceptual Variation and Structuralism," and then expanded into a separate paper. Thanks to Alex Byrne for several helpful conversations, including the conversation that convinced me that a separate paper was needed. Thanks to Justin Clarke-Doane and Shamik Dasgupta for reading both an early draft and a late draft. Thanks to Philip Kitcher for improving my discussion of genotypes and phenotypes. And thanks especially to David J. Barnett for patiently guiding me through the literatures on perceptual justification and peer disagreement.

References

- Allen, K. (2009). Inter-species variation in colour perception. *Philosophical Studies*, 142(2), 197–220.
- Allen, K. (2016). *A Naïve realist theory of colour*. Oxford: Oxford University Press.
- Barnett, D. (2011). Does vagueness exclude knowledge? *Philosophy and Phenomenological Research*, 82(1), 22–45.
- Bell, J. S. (2004). *Speakable and unspeakable in quantum mechanics: Collected papers on quantum philosophy* (p. 1987). Cambridge: Cambridge University Press.
- Block, N. (1999). Sexism, racism, ageism, and the nature of consciousness. *Philosophical Topics*, 26(1/2), 39–70.
- Bohm, D. (1952). A suggested interpretation of the quantum theory in terms of “hidden” variables. *Physical Review*, 85, 166–193.
- Bonjour, L. (1976). The coherence theory of empirical knowledge. *Philosophical Studies*, 30(5), 281–312.
- Braun, D., & Sider, T. (2007). Vague, so untrue. *Noûs*, 41(2), 133–156.
- Brogaard, B. (2010). Color in the theory of colors? or: Are philosophers’ colors all white? In G. Yancy (Ed.), *The center must not hold: White women philosophers on the whiteness of philosophy* (pp. 131–152). Washington: Lexington Books.
- Burge, T. (1979). Individualism and the mental. *Midwest Studies in Philosophy*, 4(1), 73–121.
- Byrne, A. (2016). Inverted qualia. In E. N. Zalta (Ed.), *The stanford encyclopedia of philosophy*. Stanford: Metaphysics Research Lab, Stanford University.
- Byrne, A., & Hilbert, D. R. (2003). Color realism and color science. *Behavioral and Brain Sciences*, 26(1), 3–21.
- Byrne, A., & Hilbert, D. R. (2004). Hardin, Tye, and color physicalism. *The Journal of Philosophy*, 101(1), 37–43.
- Byrne, A., & Hilbert, D. R. (2007). Truest blue. *Analysis*, 67(1), 87–92.
- Byrne, A., & Tye, M. (2006). Qualia ain’t in the head. *Noûs*, 40(2), 241–255.
- Campbell, J. (1993). A simple view of colour. In J. Haldane & C. Wright (Eds.), *Reality, representation, and projection* (pp. 257–269). Oxford: Oxford University Press.
- Chalmers, D. (2006). Perception and the fall from Eden. In T. S. Gendler & J. Hawthorne (Eds.), *Perceptual experience* (pp. 49–125). Oxford: Oxford University Press.
- Clark, A. (2000). *A theory of sentience*. Oxford: Oxford University Press.
- Cohen, J. (2003). Perceptual variation, realism, and relativization, or: How I learned to stop worrying and love variations in color vision. *Behavioral and Brain Sciences*, 26(1), 25–26.
- Cohen, J. (2009). *The red and the real: An essay on color ontology*. Oxford: Oxford University Press.
- Creath, R. (2017). Logical empiricism. In E. N. Zalta (Ed.), *The stanford encyclopedia of philosophy*. Stanford: Metaphysics Research Lab, Stanford University.
- Crisp, T. M. (2003). Presentism. In M. J. Loux & D. W. Zimmerman (Eds.), *The oxford handbook of metaphysics* (pp. 211–246). Oxford: Oxford University Press.
- Dasgupta, S. (2015). Substantivalism vs relationalism about space in classical physics. *Philosophy Compass*, 10(9), 601–624.
- Dasgupta, S. (2016). Symmetry as an epistemic notion. *The British Journal for the Philosophy of Science*, 67(3), 837–878.
- Dorr, C. (2003). Vagueness without ignorance. *Philosophical Perspectives*, 17(1), 83–113.
- Dorr, C., & Rosen, G. (2002). Composition as a fiction. In R. Gale (Ed.), *The blackwell guide to metaphysics* (pp. 151–174). Oxford: Blackwell.
- Egan, A. (2010). Projectivism without error. In B. Nanay (Ed.), *Perceiving the world* (pp. 68–96). Oxford: Oxford University Press.
- Gert, J. (2006). A realistic colour realism. *Australasian Journal of Philosophy*, 84(4), 565–589.
- Goldman, A. I. (2008). Immediate justification and process reliabilism. In Q. Smith (Ed.), *Epistemology: New essays* (pp. 63–82). Oxford: Oxford University Press.
- Graham, P. A. (2011). ‘Ought’ and ability. *The Philosophical Review*, 120(3), 337–382.
- Hardin, C. L. (1988). *Color for philosophers: Unweaving the rainbow* (expanded). London: Hackett.
- Hardin, C. L. (2003). A spectral reflectance doth not a color make. *The Journal of Philosophy*, 100(4), 191–202.
- Hawthorne, J., & Kovakovich, K. (2006). Disjunctivism. *Aristotelian Society Supplementary*, 80(1), 145–183.
- Horwich, P. (1990). *Truth*. Oxford: Basil Blackwell.

- Ismael, J., & van Fraassen, B. C. (2003). Symmetry as a guide to superfluous theoretical structure. In K. Brading & E. Castellani (Eds.), *Symmetries in physics: Philosophical reflections* (pp. 371–392). Cambridge: Cambridge University Press.
- Jackson, F., & Pargetter, R. (1987). An objectivist's guide to subjectivism about color. *Revue Internationale de Philosophie*, 41(160), 127–241.
- Kalderon, M. E. (2007). Color pluralism. *The Philosophical Review*, 116(4), 563–601.
- Kelly, T. (2010). Peer disagreement and higher-order evidence. In R. Feldman & T. A. Warfield (Eds.), *Disagreement* (pp. 111–174). Oxford: Oxford University Press.
- King, A. (2014). Actions that we ought, but can't. *Ratio*, 27(3), 316–327.
- King, A. (2019). *What we ought and what we can*. Abingdon: Routledge.
- Lasonen-Aarnio, M. (2014). Higher-order evidence and the limits of defeat. *Philosophy and Phenomenological Research*, 88(2), 314–345.
- Lewis, D. K. (1986). *On the plurality of worlds*. Oxford: Blackwell.
- Lucas, J. R. (1998). Transcendental tense ii. *Aristotelian Society Supplementary*, 72, 45–56.
- Lycan, W. (2014). Representational theories of consciousness. In E. N. Zalta (Ed.), *The stanford encyclopedia of philosophy*. Stanford: Metaphysics Research Lab, Stanford University.
- Lythgoe, J. N., & Partridge, J. C. (1989). Visual pigments and the acquisition of visual information. *Journal of Experimental Biology*, 146(1), 1–20.
- Matthen, M. (2009). Truly blue: An adverbial aspect of perceptual representation. *Analysis*, 69(1), 48–54.
- Maudlin, T. (1994). *Quantum non-locality and relativity: metaphysical intimations of modern physics*. Oxford: Blackwell.
- Maudlin, T. (2012). *Philosophy of physics: Space and time*. Princeton: Princeton University Press.
- McLaughlin, B. (2003). Colour, consciousness, and colour consciousness. In Q. Smith & A. Jokic (Eds.), *Consciousness: New philosophical perspectives* (pp. 97–154). Oxford: Oxford University Press.
- Morrison, J. (2020a). Appendix to “perceptual variation and ignorance”. Columbia University Academic Commons [digital repository]. <https://doi.org/10.7916/d8-d20y-1g83>.
- Morrison, J. (2020b). Perceptual variation and structuralism. *Noûs*, 54(2), 290–326.
- Neander, K. (2012). Teleological theories of mental content. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy*. Stanford: Metaphysics Research Lab, Stanford University.
- Norcross, A. (2004). Puppies, pigs, and people: Eating meat and marginal cases. *Philosophical Perspectives*, 18(1), 229–245.
- Pautz, A. (2006). Sensory awareness is not a wide physical relation: An empirical argument against externalist intentionalism. *Noûs*, 40(2), 205–240.
- Pautz, A. (2010). A simple view of consciousness. In R. C. Koons & G. Bealer (Eds.), *The waning of materialism* (pp. 25–66). Oxford: Oxford University Press.
- Pautz, A. (2014). The real trouble with armchair arguments against phenomenal externalism. In M. Sprevak & J. Kallestrup (Eds.), *New waves in philosophy of mind* (pp. 153–181). London: Palgrave Macmillan.
- Pautz, A. (2019). How can brains in vats experience a spatial world?: A puzzle for internalists. In A. Pautz & D. Stoljar (Eds.), *Blockheads!: Essays on Ned Block's philosophy of mind and consciousness* (pp. 379–420). Cambridge: The MIT Press.
- Prior, A. N. (1972). The notion of the present. In J. T. Fraser, F. C. Haber, & G. H. Müller (Eds.), *The study of time* (pp. 320–323). Berlin: Springer.
- Pryor, J. (2000). The skeptic and the dogmatist. *Noûs*, 34(4), 517–549.
- Regan, T. (1983). *The case for animal rights*. California: University of California Press.
- Russell, B. (1912). *The problems of philosophy*. London: Home University Library.
- Singer, P. (1975). *Animal liberation: A new ethics for our treatment of animals*. New York: HarperCollins.
- Sinnott-Armstrong, W. (1984). ought' conversationally implies 'can. *Philosophical Review*, 93(2), 249–261.
- Smith, M., Lewis, D., & Johnston, M. (1989). Dispositional theories of value. *Proceedings of the Aristotelian Society*, 63(1), 89–174.
- Stroud, B. (2000). *The quest for reality: Subjectivism and the metaphysics of colour*. Oxford: Oxford University Press.
- Tye, M. (2000). *Consciousness, color, and content*. Cambridge: The MIT Press.
- Tye, M. (2006). The puzzle of true blue. *Analysis*, 66(3), 173–178.
- Tye, M. (2012). Cohen on color relationism. *Analytic Philosophy*, 53(3), 297–305.
- Tye, M. (2015). Phenomenal externalism, Lolita, and the planet Xenon. In T. Horgan, M. Sabatés, & D. Sosa (Eds.), *Qualia and mental causation in a physical world: themes from the philosophy of Jaegwon Kim* (pp. 190–208). Cambridge: Cambridge University Press.

- van Inwagen, P. (1990). *Material beings*. New York: Cornell University Press.
- Webster, M. A., & MacLeod, D. I. A. (1988). Factors underlying individual differences in the color matches of normal observers. *Journal of the Optical Society of America A*, 5(10), 1722–1735.
- Wedgwood, R. (2007). *The nature of normativity*. Oxford: Oxford University Press.
- Williamson, T. (1994). *Vagueness*. Abingdon: Routledge.
- Williamson, T. (2014). Very improbable knowing. *Erkenntnis*, 79(5), 971–999.
- Wright, C. (2002). (Anti-) sceptics simple and subtle: GE Moore and John McDowell. *Philosophy and Phenomenological Research*, 65(2), 330–348.
- Wyszecki, G., & Stiles, W. S. (1982). *Color science: Concepts and methods, quantitative data and formulae* (p. 2). Hoboken: Wiley.
- Zimmerman, D. (2007). The privileged present: Defending an ‘A-theory’ of time. In T. Sider, J. Hawthorne, & D. W. Zimmerman (Eds.), *Contemporary debates in metaphysics* (pp. 211–225). Oxford: Blackwell.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.