

Perceptual Variation and Ignorance

draft

John Morrison
October 9, 2017

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.

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. It's natural to assume 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. However, if it does, we can't know it, and shouldn't expect to know it, given the Special Theory of Relativity (see Bell 2004, p.75–76). According to many philosophers and physicists, this is an unacceptable kind of ignorance. 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 (see Ismael and van Fraassen 2003, Maudlin 2012, Dasgupta 2016).

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 it leads to ignorance that is just as unacceptable as our ignorance of absolute velocity.

What is the challenge of perceptual variation? Suppose that Aaron and Miriam are looking at the same lemon, and their perceptions differ. I think the best way to convey how their perceptions differ is to 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. 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 forthcoming). Here, I'll just review the behavioral evidence from matching experiments (Wyszecki and Stiles 1982, Ch 5, and 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 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 ordinary observers 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 ordinary observers perceive the panels differently. This is a short and hopefully uncontroversial step away from the conclusion that even ordinary observers have different phenomenal characters. Aaron and Miriam are supposed to be observers 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 (2002), Byrne and Hilbert (2003), and Allen (2016). In the next section (Section 2),

I'll argue that one-ism leads to ignorance about whose perception is accurate. In the following section (Section 3), I'll argue that 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 (Section 4) with a brief discussion of the more extreme kind of variation found in so-called "phenomenal inversions."

I'm not the first to object to one-ism along these lines. In particular, Jackson and Pargetter (1987, p.133), Hardin (1988, p.89; 2003, p.199–201), Block (1999, p.46, 54), Clark (2000, p.215–217), Gert (2006, p.567), Kalderon (2007, p.566), Cohen (2009, p.45–64), Brogaard (2010, p.146), and Egan (2010, p.70) also object that one-ism leads to ignorance about whose perception is accurate. But their objections are too quick, in that they consider only a small subset of the evidence that we'd expect a one-ist to use to learn whose perception is accurate. They also don't explain why ignorance of this kind is unacceptable. As a result, they haven't yet explained what's wrong with one-ism.¹

Why does one-ism deserve so much attention? Here's what I regard as its strongest motivation: 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. If you're looking at those surfaces at the same time, and your perceptions have the same phenomenal characters, you'll perceive the first surface as *greener* than the second surface. In that case, you're perceiving the surfaces as having incompatible colors. If your perceptions left open the possibility that the surfaces have the same color, you wouldn't perceive one as greener than the

¹Cohen (2009, p.51–52) identifies a precedent for thinking that this would be an unacceptable kind of ignorance. In particular, he points out that there's variation in which jokes we think are funny, and when people disagree about whether a joke is funny, few insist that at most 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.

other. Analogously, if a friend tells you that his brother is older than his sister, he has told you that his brother and sister have incompatible ages. If what he said has left open the possibility that his siblings are the same age, he hasn't told you that his brother is *older*.

A one-ist might then appeal to a weak version of representationalism about color perception:

REPRESENTATIONALISM

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

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, returning to Aaron and Miriam, at most one of their perceptions of the lemon is accurate, because Aaron's perception has a phenomenal-greenish-yellow character and Miriam's perception has a phenomenal-pure-yellow character.

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 reflectances, and that the lemon has at most one of these reflectances. A realist primitivist might think that they're perceiving different non-physical properties, and that the lemon can instantiate at most one of them because there's something about these properties 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 observer in a certain kind of context, and that the lemon causes at most one of those perceptions in the relevant kind of observer in the relevant kind of context.

2 Ignorance

There are seven 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 to know whose perception is accurate. To make our discussion more concrete, let's assume reflectance physicalism. After we've considered all seven kinds of evidence, we can drop that assumption and generalize.

(1) The first 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. 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.

It's also unlikely that future scientific advances will help. We already know many of the physical causes of variation in our perceptions of color, and they don't give us any justification for believing that Miriam's perception is accurate. Further details are unlikely to make a difference. It's unlikely we'll find anything as obviously defective as a microscope's scratched lens. We'll just discover more details about the underlying mechanisms, without any principled way to label one as defective.

It's worth keeping in mind that Aaron and Miriam are both healthy and actual. Because they're both healthy, it doesn't matter whether a one-ist can convince us that there's a link between accuracy and health, perhaps by convincing us that there's a link between the accuracy of our perceptions and the healthy functioning of our perceptual organs. For example, it doesn't matter whether a one-ist can convince us that those with unhealthy eyes (e.g., those suffering from glaucoma) misperceive objects because their eyes are malfunctioning. Similarly, because they're both actual, it doesn't matter whether a one-ist can convince us that there's a link between accuracy and actuality, perhaps by convincing us that there's a link between the accuracy of our perceptions and the evolution of our perceptual organs. For example, it doesn't matter whether a one-ist can convince us that non-actual people with the same eyes, but different phenomenal characters, must have developed those eyes randomly, rather than through evolution, and are thus misperceiving (see Byrne and Tye 2006, p.252–253).

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.

(2) The second kind of evidence is about how the lemon reflects light. This evidence might seem especially helpful if colors are reflectances, 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? We'll next consider kinds evidence that might seem to help us answer this question.

(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 answer our question. If the lemon's color is its reflectance, we're still left wondering: Who is describing that reflectance? There's a helpful contrast with shape perception. 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. If they don't have the vocabulary to describe the ellipse, we might ask them to estimate how much they'd have to stretch the ellipse along each of its axes to form it into a perfect circle. We might also ask them to grab the ellipse as we measure the aperture of their grips. Regardless, we could still settle the matter with a ruler. Their reports about the lemon's color aren't similarly verifiable.

(4) The fourth kind of evidence is about the phenomenal characters of Aaron's and Miriam's perceptions, and the brain activity that underlies those characters. To see why some philosophers might think that this evidence is helpful, let's again consider shape perception. 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 can 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 by measuring its major and minor axes.

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 forthcoming). If these philosophers are right, we can once again use Miriam’s and Aaron’s phenomenal characters to figure out whose perception is accurate.

Still other philosophers might hope that we can deduce the degrees of eccentricity that Aaron and Miriam are perceiving from the brain activity underlying their perceptions of the ellipse. At least in principle, the brain could encode information about the ellipse’s eccentricity in a way that we could easily decode. For example, it could encode that information using a pattern of brain activity that itself forms an ellipse with the relevant eccentricity, so that the ellipse would be visible in an fMRI scan. Or, more realistically, it could use a pattern of brain activity that’s a relatively simple transformation of the relevant eccentricity, such as a logarithmic transformation, allowing us to easily decode it. We could discover that transformation by taking into account the effects of that brain activity, such as whether it leads people to report, “That’s a perfect circle.” If the brain encodes information in one of these ways, we can use Miriam’s and Aaron’s brain activity to figure out whose perception is accurate.

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, and the brain activity underlying those characters. 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").

The brain also doesn't encode information about reflectances in a way that we can easily decode. If pure yellowness is a certain kind of reflectance, and we already knew that Miriam's perception of the lemon's reflectance was accurate, then we might be able to work backwards, to determine how her brain encodes information about that reflectance. But we can't go in the opposite direction: from the activity in her brain alone, we can't deduce which reflectance she's perceiving. One of the underlying problems is that, whereas eccentricity varies along only one dimension, reflectances vary along infinitely many dimensions, and we can't deduce which dimensions a brain is encoding just by observing its activity.

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, and the brain activity underlying those perceptions, do not exemplify or encode colors in a way that allows us to figure out whose perception is accurate.

(5) The fifth 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 400nm 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 400nm 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 400nm 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 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, and 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, p.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 observers 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 of 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.

In fact, these considerations establish an even 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).

(6) The sixth kind of evidence is about other observers. 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. Photographers can learn to identify the tint of a film. 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, 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 perceive the lemon as pure yellow rather than greenish yellow, then Miriam is accurately perceiving the lemon correctly, because that’s what fixes the referent of that kind of perception. However, the statistics aren’t that neat. 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 it would make a difference if doctors started using 'arthritis' differently, it wouldn't make a difference to our perceptions if people in our 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.

(7) The seventh 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. You might hope that this gives you justification for believing that Miriam's perception is accurate. If it did, you could justifiably believe that Miriam's perception is accurate on the basis of your own perception.

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 debate about whether your perception provides evidence for its own accuracy. If it does, set that evidence aside. Your remaining evidence assigns a low probability to your perception being accurate. That is, your *independent* evidence assigns a low probability to your perception of the lemon being accurate. This defeats your justification for believing that the lemon is pure yellow, and thus for believing that Miriam's perception is accurate.

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. There's a debate about whether you *can* know that your ticket *won't* win. But everyone should agree that you *cannot* justifiably believe, let alone know, that your ticket *will* win, and that's the analogue here. Once you find out that others perceive the lemon differently, and that these differences are as arbitrary as the differences between lottery tickets, your independent 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 independent 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. Assuming this argument succeeds, would 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 independent 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. For this reason, even if the Earth had 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 independent evidence assigns a low probability to your perception that p being accurate, you cannot justifiably believe that p based on your perception.

This principle is widely accepted. Two points might help clarify its appeal: First, it would otherwise be extremely difficult for you to lose justification for believing that p , since evidence for not- p or 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 no matter what you learned from others, including that it's probably a tangerine, and 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.²

If PERCEPTUAL DEFEASIBILITY seems controversial, it's perhaps because it's easily confused with a number of similar-sounding principles. To reassure ourselves that, unlike those other principles, PERCEPTUAL DEFEASIBILITY is widely accepted, let's consider some of the philosophers who reject the similar-sounding principles, and what they would say about PERCEPTUAL DEFEASIBILITY.

Pryor (2000) (a "liberal"), Wright (2002) (a "conservative"), and Goldman (2008) (a "reliabilist") all claim that you can justifiably believe that p based on your perception, even if you lack independent evidence for the accuracy or the reliability of your perceptions. For example, you can justifiably believe that you have two hands, even if you lack independent evidence that you're not hallucinating hands. Indeed, on Pryor's view, your perceptual

²Some reject absolute velocities on the grounds that they'd be *redundant*. But Dasgupta (2016, p.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'.

evidence can itself justify the belief that you're not hallucinating hands. But these philosophers also say that, once you get independent evidence assigning a low probability to your perception being accurate, you can no longer justifiably believe that you have two hands.

Williamson (2014) and Lasonen-Aarnio (2014) claim that in some cases your evidence assigns a high probability to your belief that p being accurate, and a low probability to your belief that p being justified. They conclude that you can justifiably believe that p , even if your evidence assigns a low probability to your belief being justified.³ The cases they describe aren't counterexamples to PERCEPTUAL DEFEASIBILITY, because PERCEPTUAL DEFEASIBILITY is about cases in which your evidence assigns a low probability to your *perception* being *accurate*, not about cases in which your evidence assigns a low probability to your *belief* being *justified*.

It is thus widely agreed that, if your independent evidence assigns a low probability to your perception that p being accurate, you cannot justifiably believe that p based on your perception. Let's now reassure ourselves that it is also widely agreed that, once you learn about perceptual variation, your independent evidence assigns a low probability to your perception being accurate. In particular, let's consider two of the philosophers who might seem least likely to agree (as "anti-conciliationists").

Kelly (2010) claims that if you justifiably believe that p , and discover a disagreement about p with epistemic peers who have the same evidence, you can still justifiably believe that p . For example, if you justifiably believe

³To better understand how that might happen, it might be helpful to consider Williamson's famous clock example (p.979–980). While the original example is about knowledge, he later (p.981–982) says that it works just as well for justified belief. Suppose you're looking at the second hand on an unmarked clock, and that the most exact proposition you can justifiably believe is that the second hand is between 8 and 12, where 8 and 12 are the locations that would be marked "8" and "12" if every second were marked. Because your belief is justified, your evidence must assign a high probability to the second hand being in that range. As a result, your evidence must assign a high probability to your belief being *accurate*. Nonetheless, Williamson argues, your evidence assigns a low probability to your belief being *justified*, because your belief is justified only if the second hand is at 10 exactly. According to Williamson, if the second hand is at 11, the most exact proposition you can justifiably believe is that the second hand is between 9 and 13. More generally, if the second hand is at x , the most exact proposition you can justifiably believe is that the second hand is between $x - 2$ and $x + 2$. Thus, even though your evidence assigns a high probability to the second hand being between 8 and 12, it assigns a low probability to your belief being justified, because it assigns a low probability to the second hand being at 10 exactly.

that it will rain tomorrow on the basis of the meteorological data, you can still justifiably believe that it will rain, even if an epistemic peer believes that it won't based on the same data. Kelly says that this is because the evidence for your belief includes more than just the psychological facts about what you believe and what your epistemic peer believes. Your evidence also includes the meteorological data itself, and that might still be enough to justify your belief. Thus, even when you discover the disagreement, your evidence assigns a high probability to your belief being accurate. But when your disagreement is based entirely on different perceptions, there is no other evidence. Your evidence is exhausted by the psychological facts about what you perceive and what the other person perceives, and the reliability of these perceptions. Thus, according to Kelly, when you learn that someone else perceives things differently, your independent evidence assigns a low probability to your perception being accurate, and you should suspend judgment (p.150–152).

Wedgwood (2007, p.257–266) claims that if you have a normative intuition that p , and discover that others have different intuitions, you can still justifiably believe that p . For example, if you have the normative intuition that eating meat is ethically permissible, and discover that another person has the intuition that it's impermissible, you can still justifiably believe that it is permissible. Even so, Wedgwood insists that your justification can be defeated, either by incoherence with the rest of your beliefs, or by evidence that the other person's normative intuitions are at least as reliable. With respect to your normative intuitions, this second kind of evidence is rare. Wedgwood's only example is when the other person has previously brought about significant changes in your own normative intuitions. In most cases, you don't have any evidence about the reliability of your own normative intuitions or the other's person's normative intuitions, and thus your independent evidence doesn't assign any probability to the accuracy of your intuition. When that happens, Wedgwood thinks that you're justified in relying more on your own intuition, and thus that your intuition can still justify your belief. Wedgwood doesn't say as much about perceptions, but he does say the justification for beliefs based on perception can be defeated in the same ways, namely either by incoherence with the rest of your beliefs, or by evidence that the other person's perceptions are at least as reliable. Unlike normative intuitions, this second kind of evidence is plentiful. You can, for example, confirm that others perform equally well on standard perceptual tasks, such as the Farnsworth-Munsell hue test. You also know that per-

ceptions of fine-grained colors are the result of haphazard genetic mutations and environmental interactions, much like your exact height, skin pigment, and hair density. Thus, with respect to perception, the situation is quite different; you have a lot of independent evidence about the reliability of your own perceptions and other people's perceptions, and that evidence assigns a low probability to your own perception being accurate.

We just reviewed the seven 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 many different kinds of evidence, including the intuition that there are objects that combine to form you and me (e.g., van Inwagen 1990, p.73), the argument that for any n it's not vague whether n objects exist (e.g., Lewis 1986, p.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 hypotheses 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.

It might be tempting to hope that our ultimate psychosemantic theory about how our perceptions manage to represent properties will settle the matter (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. Analogously, consider a philosopher of language who says that she needn't worry about Frege cases, because she can just wait until we have our final theory of mental content. Such a philosopher underestimates the significance of the challenge created by Frege cases, because we won't be able to settle on the final theory of mental content *until* we settle on a response to Frege cases. One of the reasons why the challenge

of perceptual variation is so important is that settling on a response is an important step towards a developing a satisfying psychosemantic theory of perception.

I conclude that our current and future evidence is unlikely to justify our belief that Miriam's perception is accurate; the most promising kinds of evidence don't support this hypothesis over the alternative. This isn't to deny the possibility of such evidence. If God were to reveal to us that Miriam's perception is accurate, we'd certainly gain justification. If a new and well-verified scientific theory were to entail that Miriam's perception is accurate, we'd also gain justification. But we have to reason to expect that anything like that will happen.

Significantly, this ignorance extends much further than ignorance about whose perception *is accurate*. Even if both Miriam's and Aaron's perceptions are inaccurate, we can still ask whose 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 also ignorant of whose perception is *more accurate*.

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 whose perception is accurate, and thus wouldn't help us know whether the lemon has the color Miriam is perceiving or the color Aaron is perceiving.

If colors are dispositions, we could know even less, assuming with the one-ist that only one kind of observer and only one kind of context is relevant. None of the evidence we reviewed would help us choose between the many

different normal observers who perceive the lemon differently, or between the many different normal contexts in which the same observer 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 observer 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 observer 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 deny that there is still a fact of the matter in all these cases. So why would it be a problem if we're ignorant about whose perception of the lemon is accurate?

This, in essence, is how color one-ists such as Stroud (2000, p.173–176), Tye (2002, p.108), Byrne and Hilbert (2003, p.16–17; 2004, p.37–39; 2007), and Allen (2016) 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, p.177–178)

According to these philosophers, there's a fact of the matter, even though we shouldn't expect to know it. For example, even though we can't know that Miriam's perception is accurate and Aaron's perception is inaccurate, it might still be a fact that Miriam's perception is accurate and Aaron's perception is inaccurate.

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. 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 precise 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 deny that there are facts of the matter.

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. Our ignorance seems more like our ignorance of absolute velocity. Recall from the introduction that, given the laws of the Special Theory of Relativity, we can't know the absolute velocities of bodies. This can be established rigorously, because the laws themselves entail that there's no way to set up a causal link between absolute velocities and observable effects. Thus, the laws themselves entail that there is no evidence that we could in principle collect that would allow us to choose between the hypothesis that 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, p.75–76). With respect to color perception, there are no rigorous demonstrations, because there are no physical laws that entail that we can't collect evidence about whose perception is accurate. But I think that our informal canvassing of evidence in the last section is still enough to establish that this isn't something we could learn, assuming one-ism. As with the hypotheses about the Earth's absolute velocity (see Dasgupta 2015, p.610), 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. Thus, to the extent that we're inclined to deny that there are absolute velocities, we should also be inclined to deny that only one person's perception is accurate. In both cases, our ignorance would be so unacceptable that we should instead believe that there isn't a fact of the matter.

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

UNACCEPTABLE IGNORANCE

For a given list of hypotheses, if we don't justifiably believe the first hypothesis, we don't justifiably believe the second hypothesis, etc., and we don't expect to justifiably believe any of these hypotheses, and our ignorance doesn't have the right kind of causal explanation, then we should believe that none of these hypotheses is a fact.

For example, because we don't expect to justifiably believe any of the hypotheses about the absolute velocities of the planets (e.g., that the Earth's absolute velocity is 10 k/s), and our ignorance doesn't have the right kind of causal explanation, we should believe that none of these hypotheses is a fact, and thus that there are no facts about the absolute velocities of the planets.

I'm going to argue that this principle is incompatible with 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 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 are no facts about which objects combine to form wholes. It likewise doesn't imply that there

are no facts about the ontological status of numbers, propositions, properties, and possible worlds.

This principle is about what we should believe; it's epistemological, rather than metaphysical. Similar to induction and inference to the best explanation, there's no guarantee that it will lead to a true belief, but it's still worth following.

I already described the role this principle plays in our thinking about motion. In what follows, I'll describe the role it plays in our thinking about time, vagueness, and ethics. This will motivate UNACCEPTABLE IGNORANCE by showing that it unifies the way we think about diverse phenomena. These examples will also clarify what counts as the "right kind of causal explanation," and help 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. Afterward, I will restate the problem for one-ism, taking into account what we learned from these examples.

These examples will serve another purpose as well. As we'll see, some philosophers think that, despite our ignorance, we should still believe that one of the hypotheses under consideration is a fact, provided there's what I'll call a "sufficiently weighty reason." These philosophers regard UNACCEPTABLE IGNORANCE as a defeasible principle. I'm not convinced that these philosophers are right. I'm inclined to think that ignorance of this kind should always lead us to believe that none of the hypotheses is a fact, except when our other scientific theories require at least one of the hypotheses to be a fact. Nonetheless, it will be helpful to consider what these philosophers regard as sufficiently weighty reasons, so that we can determine whether one-ists can offer similar reasons.

(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 think that their explosions 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 explosions will appear simultaneous, from other reference frames, the explosions will appear sequential, and we can't know which appearance is accurate, assuming that at most one of

them is.

Most philosophers and physicists conclude that we should 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: they deny that there are unknowable facts about either absolute motion or absolute simultaneity. They countenance only knowable facts about relative motion and relative simultaneity.

While I’m not aware of any philosophers or physicists who believe that there are still facts about absolute velocity, there are philosophers and physicists who believe that there are still facts about absolute simultaneity. These philosophers and physicists recognize that they owe us a weighty reason. To satisfy this burden, Zimmerman appeals to presentism, the view that only the present moment exists. He thinks that this is a weighty reason, because he regards presentism as close to a truism (see Zimmerman 2007, p.221f; Prior 1972, p.323; Crisp 2003, p.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, p.77; see also Bohm 1952; Maudlin 1994; Lucas 1998, p.55). If it does, that would be a reason to think that there are facts of the matter about absolute simultaneity, even if we can’t justifiably believe hypotheses about it.

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 weighty reason to accept ignorance, such as a truism or another scientific theory.

(2) Let’s now turn to vagueness. If Boaz is a borderline case of baldness, we can’t know that ‘Boaz is bald’ is true and we can’t know that ‘Boaz is bald’ is false. This might even be constitutive of what it is for Boaz to be a *borderline* case of baldness.

Most philosophers infer that ‘Boaz is bald’ is not true and ‘Boaz is bald’ is not false. These philosophers evaluate ‘Boaz is bald’ differently, perhaps as neither-true-nor-false. These philosophers thereby avoid committing us to ignorance about whether ‘Boaz is bald’ is true. Because we can know

that Boaz is a borderline case of baldness, we can know that ‘Boaz is bald’ is neither-true-nor-false, or whatever other status is assigned to borderline claims. Other philosophers avoid ignorance by claiming that all sentences containing vague terms are false, in which case we can know ‘Boaz is bald’ is false (Braun and Sider 2007). Still others avoid ignorance by claiming that it’s indeterminate whether we know that it’s true or know that it’s false (Barnett 2011, Dorr 2003). All of these philosophers avoid ignorance by expanding or contracting the list of hypotheses about how to evaluate vague sentences.

But there are some philosophers (“epistemicists”) who insist that ‘Boaz is bald’ is true or ‘Boaz is bald’ is false, despite our ignorance of which. These philosophers recognize that they owe us a weighty reason. 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 on how this term is currently used throughout our linguistic community, including the current disposition 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 number of facts that are in constant flux. The evidence is in principle collectable. It’s just that we’re unable to collect it. We can therefore believe that it’s either a fact that this sentence is true or a fact that this sentence is false, despite our ignorance.

Other epistemicists appeal to principles they regard as close to truisms. For example, Horwich (1990, p.81–87) appeals to classical logic and the truth schema. According to Horwich, if these principles are true, then it is either a fact that ‘Boaz is bald’ is true, or a fact that ‘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 weighty reason, and that in this case a weighty reason might be a principle that’s close to a truism.

One of the reasons why this is a helpful example is that it shows that UNACCEPTABLE IGNORANCE applies to semantic facts, in this case facts about which sentences are *true*. This is helpful because we’re interested in a similar kind of fact, namely a fact about which perceptions are *accurate*.

(3) Let's finally turn to ethics, in particular to ethical permissibility. It is natural to think 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, if we're ignorant about whether eating meat is ethically permissible, we needn't expect to remain ignorant.

Philosophers who insist that this disagreement is unresolvable often conclude that ethical permissibility is not absolute. According to the philosophers I have in mind, 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. They deny that there are unknowable facts about absolute ethical permissibility. They countenance only knowable facts about relative ethical permissibility.

I'm not aware of any philosopher who actually says that there's a fact about whether eating meat is absolutely ethically permissible but we can't know it. 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, 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 bat. Thus, they can 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.

These examples motivate UNACCEPTABLE IGNORANCE by showing that it underlies our thinking about diverse phenomena. They also give us a better sense of what counts as a causal explanation of the "right kind," and which reasons might be "sufficiently weighty" to nonetheless believe that one of the relevant hypotheses is a fact.

What counts as a causal explanation of the "right kind"? 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 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. A causal explanation of the right kind presupposes that there are observable effects, and explains why we can't observe them. Of course, it would be preferable to have a more precise definition of the "right kind of causal explanation," but for our purposes this is enough.

Which reasons are "sufficiently weighty" to believe one of the relevant hypotheses is a fact, 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 a fact. It might also be enough to preserve principles that are so central to the way we think and talk that they approach truisms. As mentioned earlier, I'm not sure that would be enough, because I'm inclined to think that we should accept this kind of ignorance only when it's indispensable to our best scientific understanding of the world, even if that forces us to give up apparent truisms. Nonetheless, these examples help give us a sense of what might be enough, at least according to some philosophers.

Let's now restate the problem, taking into account what we just learned. Given one-ism, there are only three hypotheses: that Miriam's perception is accurate and Aaron's perception is inaccurate, that Miriam's perception is inaccurate and Aaron's perception is accurate, and that both of their percep-

tions are inaccurate. But this conflicts with UNACCEPTABLE IGNORANCE, the principle that:

For a given list of hypotheses, if we don't justifiably believe the first hypothesis, we don't justifiably believe the second hypothesis, etc., and we don't expect to justifiably believe any of these hypotheses, and our ignorance doesn't have the right kind of causal explanation, then we should believe that none of these hypotheses is a fact.

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

First, we don't justifiably believe any of these hypotheses. We also shouldn't expect to justifiably believe any of these hypotheses, because we shouldn't expect our evidence to support any of them (see Section 2).

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 whether there's a fact of the matter.

Is there a sufficiently weighty 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, p.3–4). It's our natural ways of thinking and talking, not our scientific theories, that pushes us to attribute colors to objects.

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. Our natural ways of thinking and talking motivate one-ism, but they 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, p.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 observers under normal conditions. Once we reject that assumption, 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.

In the introduction I described what I regard as one-ism's strongest motivation. It began by appealing to the apparent incompatibility between the colors represented by *your* phenomenal-pure-yellow perceptions and *your* phenomenal-greenish-yellow perceptions. It then appealed to a weak version of representationalism about color perception:

REPRESENTATIONALISM

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

Representationalism is a popular view, with a number of attractive features (see Lycan 2014). But it's nowhere near a truism, as suggested by the fact that so many think it's false. 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 REPRESENTATIONALISM, 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-greenish-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 sufficiently weighty 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 2009, p.58–60, Chalmers 2006 on Edenic colors). For this reason, even if REPRESENTATIONALISM were a truism, that wouldn't be enough for one-ism.

There's a final motivation worth considering. Suppose that Aaron and Miriam are both accurately perceiving the lemon as yellow. The simplest explanation might be that each is perceiving a different shade of yellow. If the lemon has at most one shade of yellow, it would follow that at most one of their perceptions is accurate. Thus, one-ism might be a consequence of the simplest explanation of why our perceptions of more course-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 is that REPRESENTATIONALISM is false, and Aaron and Miriam are perceiving the same shade despite the differences in their phenomenal characters (for proposals along these lines, see Jackson and Pargetter 1987, McLaughlin 2003, Matthen 2009, and Morrison forthcoming). 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 at most one of their perceptions is accurate. Consider motion. The simplest explanation of the relative motions of the planets might be that they're consequences

of the planets' absolute motions, because we could then derive all of their twenty-eight relative motions from their eight 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 'Boaz is bald or not bald' is true might be that either 'Boaz is bald' is true or 'Boaz is bald' is false. But, even if that is the simplest explanation, as long as there isn't a causal explanation of our ignorance, etc., we should still deny that either 'Boaz is bald or not bald' is true or 'Boaz is bald or not bald' is false, 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 think there are facts that we cannot know.

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. In that case, there's again only one hypothesis (viz., Miriam's and Aaron's perceptions are both accurate) and we can again justifiably believe it. Thus, unlike one-ism, these other responses don't lead to unacceptable ignorance.

As I said before, the underlying principle is hard to state, and there's room for disagreement about whether UNACCEPTABLE IGNORANCE is the best formulation. But I doubt reformulating the principle will help the one-ist, because there are too many similarities between our ignorance of the lemon's color and the unacceptable kinds of ignorance we surveyed.

I tried to motivate UNACCEPTABLE IGNORANCE by showing how it underlies our thinking about phenomena as diverse as motion, time, vagueness, and ethics. If I'm right, giving up UNACCEPTABLE IGNORANCE would 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. As these examples suggest, rejecting UNACCEPTABLE IGNORANCE has the potential to disrupt our thinking about parts of philosophy that wouldn't have otherwise seemed connected.

There are other motivations for UNACCEPTABLE IGNORANCE, though they'll appeal to smaller audiences. For example, it might follow from certain versions of scientific naturalism, or a version of the principle of sufficient reason that says we can eventually know the sufficient reason for every truth. Some might even regard UNACCEPTABLE IGNORANCE as bedrock. For our purposes, we don't need to choose between these motivations. It's enough that, if we're forced to choose between UNACCEPTABLE IGNORANCE and one-ism, we should prefer 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 we call 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 and 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.

It also let us preemptively undermine the tempting response just mentioned. In particular, 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 is also 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.⁴

With that in mind, whose perception is accurate, Adam's or Miriam's? 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. This gives rise to a kind of ignorance that's just as unacceptable as our ignorance with respect to actual variation.

We're thus left with an even starker choice: conclude that color perception is necessarily illusory, or conclude that color perceptions can be accurate regardless of their phenomenal characters. I favor the second option, because I think it better preserves our natural ways of thinking and talking. I also think it better explains our success at avoiding traffic, selecting fruit, and matching socks. But that's a topic for another paper. For now, I just want

⁴There'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 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, p.252–254) response to Pautz (2006, p.220–227). By first establishing that perceptions with *moderately* different phenomenal characters have P , it is easier to then establish that it's possible for perceptions with *extremely* different phenomenal characters to have P , and thus to be accurate.

to point out that this leaves us grappling with fundamental questions about the role of phenomenal characters in perception. The challenge of perceptual variation thus shines a spotlight on some of the most fundamental questions we can ask about the mind.⁵

References

- Allen, Keith (2009). Inter-species variation in colour perception. *Philosophical Studies*, 142(2):197–220.
- Allen, Keith (2016). *A Naive Realist Theory of Colour*. Oxford University Press UK.
- Barnett, David (2011). Does vagueness exclude knowledge? *Philosophy and Phenomenological Research*, 82(1):22–45.
- Bell, J. S. (2004). *Speakable and Unsayable in Quantum Mechanics: Collected Papers on Quantum Philosophy*. Cambridge University Press. First published 1987.
- Block, Ned (1999). Sexism, ageism, racism, and the nature of consciousness. *Philosophical Topics*, 26(1):39–70.
- Bohm, David (1952). A suggested interpretation of the quantum theory in terms of "hidden" variables. *Physical Review*, 85:166–193.
- Bonjour, Laurence (1976). The coherence theory of empirical knowledge. *Philosophical Studies*, 30(5):281–312.
- Braun, David and Sider, Theodore (2007). Vague, so untrue. *Noûs*, 41(2):133–156.

⁵This 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.

- Brogaard, Berit (2010). Color in the theory of colors? or: Are philosophers' colors all white? In Yancy, George, editor, *The Center Must Not Hold: White Women Philosophers on the Whiteness of Philosophy*, pages 131–152. Lexington Books.
- Burge, Tyler (1979). Individualism and the mental. *Midwest Studies in Philosophy*, 4(1):73–121.
- Byrne, Alex (2016). Inverted qualia. In Zalta, Edward N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, winter 2016 edition.
- Byrne, Alex and Hilbert, David R. (2003). Color realism and color science. *Behavioral and Brain Sciences*, 26:3–21.
- Byrne, Alex and Hilbert, David R. (2004). Hardin, Tye, and color physicalism. *The Journal of Philosophy*, 101:37–43.
- Byrne, Alex and Hilbert, David R. (2007). Truest blue. *Analysis*, 67(293):87–92.
- Byrne, Alex and Tye, Michael (2006). Qualia ain't in the head. *Noûs*, 40(2):241–255.
- Campbell, John (1993). A simple view of colour. In Haldane, John and Wright, Crispin, editors, *Reality, Representation, and Projection*, pages 257–269. Oxford University Press.
- Chalmers, David (2006). Perception and the fall from Eden. In Gendler, Tamar Szabó and Hawthorne, John, editors, *Perceptual Experience*, pages 49–125. Oxford University Press.
- Clark, Austen (2000). *A Theory of Sentience*. Oxford University Press.
- Cohen, Jonathan (2009). *The Red and the Real: An Essay on Color Ontology*. Oxford University Press.
- Crisp, Thomas M. (2003). Presentism. In Loux, Michael J. and Zimmerman, Dean W., editors, *The Oxford Handbook of Metaphysics*, pages 211–246. Oxford University Press.

- Dasgupta, Shamik (2015). Substantivalism vs relationalism about space in classical physics. *Philosophy Compass*, 10(9):601–624.
- Dasgupta, Shamik (2016). Symmetry as an epistemic notion. *British Journal for the Philosophy of Science*, 67(3):837–878.
- Dorr, Cian (2003). Vagueness without ignorance. *Philosophical Perspectives*, 17(1):83–113.
- Dorr, Cian and Rosen, Gideon (2002). Composition as a fiction. In Gale, Richard, editor, *The Blackwell Guide to Metaphysics*, pages 151–174. Blackwell.
- Egan, Andy (2010). Projectivism without error. In Nanay, Bence, editor, *Perceiving the World*, pages 68–96. Oxford University Press.
- Gert, Joshua (2006). A realistic colour realism. *Australasian Journal of Philosophy*, 84(4):565–589.
- Goldman, Alvin I. (2008). Immediate justification and process reliabilism. In Smith, Quentin, editor, *Epistemology: New Essays*, pages 63–82. Oxford University Press.
- Hardin, C. L. (1988). *Color for Philosophers: Unweaving the Rainbow* (expanded edition). Hackett.
- Hardin, C. L. (2003). A spectral reflectance doth not a color make. *The Journal of Philosophy*, 100:191–202.
- Hawthorne, John and Kovakovich, Karson (2006). Disjunctivism. *Aristotelian Society Supplementary Volume*, 80(1):145–183.
- Horwich, Paul (1990). *Truth*. Basil Blackwell.
- Ismael, Jenann and van Fraassen, Bas C. (2003). Symmetry as a guide to superfluous theoretical structure. In Brading, Katherine and Castellani, Elena, editors, *Symmetries in Physics: Philosophical Reflections*, pages 371–392. Cambridge University Press.
- Jackson, Frank and Pargetter, Robert (1987). An objectivist’s guide to subjectivism about color. *Revue Internationale de Philosophie*, 41:127–241.

- Kalderon, Mark Eli (2007). Color pluralism. *The Philosophical Review*, 116:563–601.
- Kelly, Thomas (2010). Peer disagreement and higher-order evidence. In Feldman, Richard and Warfield, Ted A., editors, *Disagreement*, pages 111–174. Oxford University Press.
- Lasonen-Aarnio, Maria (2014). Higherorder evidence and the limits of defeat. *Philosophy and Phenomenological Research*, 88(2):314–345.
- Lewis, David K. (1986). *On the Plurality of Worlds*. Blackwell.
- Lucas, J. R. (1998). Transcendental tense ii. *Aristotelian Society Supplementary Volume*, 72:45–56.
- Lycan, William (2014). Representational theories of consciousness. In Zalta, Edward N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, winter 2014 edition.
- Lythgoe, J. N. and Partridge, J. C. (1989). Visual pigments and the acquisition of visual information. *Journal of Experimental Biology*, 146(1):1–20.
- Matthen, Mohan (2009). Truly blue: An adverbial aspect of perceptual representation. *Analysis*, 69(1):48–54.
- Maudlin, Tim (1994). *Quantum Non-Locality and Relativity: Metaphysical Intimations of Modern Physics*. Blackwell.
- Maudlin, Tim (2012). *Philosophy of Physics: Space and Time*. Princeton University Press.
- McLaughlin, Brian (2003). Colour, consciousness, and colour consciousness. In Smith, Quentin and Jokic, Aleksander, editors, *Consciousness: New Philosophical Perspectives*, pages 97–154. Oxford University Press.
- Morrison, John (forthcoming). Perceptual variation and structuralism. *Noûs*.
- Neander, Karen (2012). Teleological theories of mental content. In Zalta, Edward N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, spring 2012 edition.

- Norcross, Alastair (2004). Puppies, pigs, and people: Eating meat and marginal cases. *Philosophical Perspectives*, 18(1):229–245.
- Pautz, Adam (2006). Sensory awareness is not a wide physical relation: An empirical argument against externalist intentionalism. *Noûs*, 40(2):205–240.
- Pautz, Adam (2009). A simple view of consciousness. In Koons, Robert C. and Bealer, George, editors, *The Waning of Materialism: New Essays*, pages 25–66. Oxford University Press.
- Pautz, Adam (2014). The real trouble with armchair arguments against phenomenal externalism. In Sprevak, Mark and Kallestrup, Jesper, editors, *New Waves in Philosophy of Mind*, pages 153–181. Palgrave Macmillan UK.
- Pautz, Adam (forthcoming). How can brains in vats experience a spatial world?: A puzzle for internalists. In Pautz, Adam and Stoljar, Daniel, editors, *Themes from Block*. The MIT Press.
- Prior, A. N. (1972). The notion of the present. In Fraser, J. T., Haber, F. C., and Müller, G. H., editors, *The Study of Time*, pages 320–323. Springer Berlin Heidelberg.
- Pryor, James (2000). The skeptic and the dogmatist. *Noûs*, 34(4):517–549.
- Regan, Tom (1983). *The Case for Animal Rights*. University of California Press.
- Russell, Bertrand (1912). *The Problems of Philosophy*. Home University Library.
- Singer, Peter (1975). *Animal Liberation: A New Ethics for Our Treatment of Animals*. HarperCollins.
- Smith, Michael, Lewis, David, and Johnston, Mark (1989). Dispositional theories of value. *Proceedings of the Aristotelian Society*, 63:89–174.
- Stroud, Barry (2000). *The Quest for Reality: Subjectivism and the Metaphysics of Colour*. Oxford University Press.
- Tye, Michael (2002). *Consciousness, Color, and Content*. The MIT Press.

- Tye, Michael (2006). The puzzle of true blue. *Analysis*, 66:173–178.
- Tye, Michael (2012). Cohen on color relationism. *Analytic Philosophy*, 53(3):297–305.
- Tye, Michael (2015). Phenomenal externalism, Lolita, and the planet Xenon. In Horgan, Terence, Sabatés, Marcelo, and Sosa, David, editors, *Qualia and Mental Causation in a Physical World: Themes from the Philosophy of Jaegwon Kim*, pages 190–208. Cambridge University Press.
- van Inwagen, Peter (1990). *Material Beings*. Cornell University Press.
- Webster, Michael A. and MacLeod, Donald 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, Ralph (2007). *The Nature of Normativity*. Oxford University Press.
- Williamson, Timothy (1994). *Vagueness*. Routledge.
- Williamson, Timothy (2014). Very improbable knowing. *Erkenntnis*, 79(5):971–999.
- Wright, Crispin (2002). (Anti-)Sceptics Simple and Subtle: G. E. Moore and John McDowell. *Philosophy and Phenomenological Research*, 65(2):330–348.
- Wyszecki, G and Stiles, W.S. (1982). *Color Science: Concepts and Methods, Quantitative Data and Formulae*. Wiley, 2nd edition.
- Zimmerman, Dean (2007). The privileged present: Defending an ‘a-theory’ of time. In Sider, Theodore, Hawthorne, John, and Zimmerman, Dean W., editors, *Contemporary Debates in Metaphysics*, pages 211–225. Blackwell.