Perceptual Variation and Relativism

for *Epistemology After Sextus Empiricus*, Vogt and Vlastis (eds.)

John Morrison
September 7, 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. According to Sextus, Protagoras's response is that all of our perceptions might be accurate. As this response is traditionally developed, it is difficult to explain color illusion and color constancy. I will argue that this difficulty is due to a widespread assumption that I call perceptual atomism. I will conclude that, if we want to develop Protagoras's response, we need to give up perceptual atomism. I will end with a brief sketch of my preferred alternative, perceptual structuralism.

1 Introduction

According to Sextus, children perceive brighter colors than their grandparents (*Outlines of Scepticism*, I:105). Suppose that Sextus is right, and that Jacob perceives brighter colors than Sarah, his grandmother. In particular, suppose that Jacob and Sarah are looking at the same stone tile, and that while Sarah has the kind of perception you have when you look at the square on the left, Jacob has the kind of perception you have when you look at the square on the right:

Sarah

Jacob

Whose perception of the tile is accurate, Jacob's or Sarah's?

Sextus rejects the standard responses. One response is that *neither* of their perceptions is accurate, because the tile isn't really colored. Sextus

attributes this response to Democritus (I:213; II:63; see also *Against the Logicians* I:135–140). Sextus's objection is that we can't know *anything* about the tile, not even whether it is or is not colored. This reflects his more general skepticism about what we can know (I:15). But one needn't be a skeptic like Sextus to find this first response objectionable. One might find it objectionable even if one thinks we can know a lot about the tile, including what our best scientific theories say about it. For example, an influential non-skeptical objection is that it's central to the way we think and talk that objects are really colored, and that, by default, we should try to preserve these ways of thinking and talking (see, e.g., Lewis 1997, p.325–328; Johnston 1992, p.221–222; Cohen 2009, p.15, 65). There's obviously a lot more to say about this non-skeptical objection, but I hope it's familiar enough.

Another response is that *one* of their perceptions is accurate. Sextus's objection is that we cant know whose perception is accurate. For example, suppose that only Jacob's perception is accurate. According to Sextus, the mere fact that Jacob is younger does not give us any justification for believing that his perception is accurate (I:105–112). Neither does the mere fact that Jacob is healthier (I:102-103) or in the majority (I:89). More generally, none of the facts about Jacob and Sarah give us justification for believing that one of them, rather than the other, is accurately perceiving the tile. Sextus also rejects a proposal, available to the Dogmatists, that one should rely on one's own perception, so that if one has the same kind of perception as Jacob, one has justification for believing that Jacob's perception is accurate, and if one has the same kind of perception as Sarah, one has justification for believing that Sarah's perception is accurate. Sextus objects that relying on one's own perception would be question-begging, given that one also doesn't have justification for believing that one's own perception is accurate (I:90, 112–113). Finally, Sextus denies that we can rely on facts about the relevant object, in this case the tile, because he denies that we can know anything about external objects (e.g., I.99). Because there don't seem to be any other relevant facts, Sextus concludes that, if only one of their perceptions is accurate, we can't know whose (I:112; II:51–53). In other work (Morrison manuscript), I argue that one should accept this conclusion, even if, unlike Sextus, one thinks that we can know a lot about external objects, including what our best scientific theories tell us. I also explain why this conclusion is unacceptable, even if, unlike Sextus, one is willing to say similar things in other domains — for example, that only one estimate about the size of the universe can be true, even if we can't identify it. Once again, there's obviously a lot more to

say about this non-skeptical objection, but I hope it's familiar enough (see, e.g., Block 1999, p.46, 54).

A third response is that *both* of their perceptions might be accurate. Let's call this "perceptual relativism," because it implies that, in some sense, what Jacob is perceiving is relative to him, and what Sarah is perceiving is relative to her. As we'll see, there are many senses in which their perceptions might be relative. In one sense, the colors themselves are relations to perceivers (e.g., medium-gray-for-Jacob, dark-gray-for-Jacob). But in another sense, it's the descriptions under which we perceive the colors that are relative to perceivers. I elsewhere call this response "both-ism" to emphasize that it subsumes any view that implies that both perceptions might be accurate (Morrison forthcoming).

Sextus attributes relativism to Protagoras:

> He [Protagoras] says that ... our senses are rearranged and altered depending on age and the other constituents of the body. He also says that the reasons for all apparent things are present in matter, so that matter can, as far as itself is concerned, be all the things it appears to anyone to be. Men grasp different things at different times, depending on their different conditions: someone in a natural state apprehends those things in the matter which can appear to those in a natural state, someone in an unnatural state apprehends what can appear to those in an unnatural state. And further, depending on age, and according to whether we are sleeping or waking, and in virtue of each sort of condition, the same account holds. According to him, then, man is the standard for what is; for all things apparent to men actually are, and what is apparent to nobody is not (I:217–218, trans. Annas and Barnes 2000; see also I:79–87, *Against the Logicians* I:60–64, and Plato's *Theaetetus* 153d–154b).

I think that Sextus's description captures what's right about perceptual relativism. Suppose we fill one hundred unmarked tubes with different amounts of mercury and then place them in a warm room. Their mercury will rise to different levels. It would be silly to then ask which tube out of all the tubes is *accurately* measuring the room's temperature. Each tube just registers the temperature in a different way. Likewise, nature has filled our eyes with different amounts of the relevant kinds of detectors, distributed those detectors

in different patterns along our retinas, wired our brains to respond differently to the signals sent from our eyes, and so on. Consequently, people's eyes and brains respond differently to the same object, producing different perceptions. As with the unmarked tubes of mercury, it seems misguided to ask which perceptual system out of all perceptual systems is accurately measuring the tile's color. Different perceptual systems just register the tile's surface in different ways. For this reason, I think there's something right about perceptual relativism.

Sextus disagrees. The motivation for perceptual relativism depends on a characterization of matter as in flux, and our senses as rearranged and altered depending on our age and other factors. Sextus objects that we can't know *anything* about matter, not even whether it is in flux, nor about our senses, not even how they're constituted (I:219). He concludes that perceptual relativism is unmotivated. However, few contemporary philosophers would endorse this objection. Most contemporary philosophers believe that we can know a lot about matter and our senses, including what our best scientific theories tell us. Perhaps as a result, perceptual relativism continues to have many proponents, including Cohen (2009), Egan (2010), Gert (2013), Jackson and Pargetter (1987), Kalderon (2007), Matthen (2009), and McLaughlin (2003).[1]

But one needn't be a skeptic to object to perceptual relativism, as least when it is developed in the way Sextus describes. In particular, if relativism is developed so that it implies that "all things apparent to men actually are," then it has trouble explaining two phenomena that are central to color perception: color inaccuracy and color constancy. I will go into more detail later, so, for now, here is a quick sketch. Let's start with color inaccuracy. In some cases, it seems that our perceptions are inaccurate, as when we're looking at an object under a misleading illuminant, or against a misleading background. But it's unclear how any of our perceptions can be inaccurate if "all things apparent to men actually are," because this seems to imply that our perceptions are always accurate. Likewise, in many cases, it seems that we perceive the same colors across different contexts, as when we view an object under different illuminants, or against different backgrounds. But it's unclear how that's possible if "all things apparent to men actually are," be-

[1]Some of these authors use 'perceptual relativism' more restrictively, and therefore don't describe themselves as relativists (see, e.g., Cohen 2009, p.10; Kalderon 2007, p.584). But they're all perceptual relativists in our sense, because they'd all say that Jacob's and Sarah's perceptions might both be accurate.

cause this seems to imply that, with each new appearance, we're accurately perceiving a new color, and thus that there is no color that we perceive as constant. I will also argue that varieties of relativism with weaker implications (e.g., that "all things apparent to Sarah in normal conditions actually are") are similarly unable to explain color inaccuracy and color constancy.

I would like to identify a way of developing perceptual relativism that can explain these phenomena. With this goal in mind, I will argue that the difficulty in explaining color inaccuracy and color constancy is the result of combining perceptual relativism with a plausible and widespread assumption about perception, which I call "perceptual atomism." I will conclude that, if we want to preserve what's right about perceptual relativism, we should develop it in a way that doesn't rely on perceptual atomism. Toward the end of the paper I'll briefly sketch an alternative to perceptual atomism that I think perceptual relativists should develop, and which I call "perceptual structuralism." The key difference between atomistic relativism and structural relativism is that, whereas atomistic relativists say that color perception is relative to a kind of perception, structural relativists say that color perception is relative to a "structure" built out of comparisons to other objects. I elsewhere develop perceptual structuralism in more detail.

Of course, Sextus would reject perceptual relativism even when combined with perceptual structuralism rather than perceptual atomism. His objection applies to all variants of perceptual relativism, including variants that can explain color inaccuracy and color constancy. Thus, my focus in this paper is on two themes within Sextus's philosophy, namely the puzzle of perceptual variation and the feasibility of perceptual relativism, rather than on his philosophy as a systematic whole. My focus is also limited to our perception of secondary qualities, with color as the central example. Sextus, in contrast, extends the debate to *all* of our judgments about the external world, including our judgments about shapes, causes, times, places, bodies, morality, and God (see Vogt 2016, Sec 4.4, 5.4).

Sextus gives two other examples of interpersonal variation in color perception: that those with jaundice perceive yellower colors, and that those with blood-suffusion perceive redder colors (I:44, 101; II:51). While the difference between children and their grandparents is genuine, and attributable to the hardening of their grandparents' cataracts, I'm not aware of any evidence that corroborates these other examples. Like other ancient authors, Sextus mistakenly assumes that, because jaundice yellows the eye, it yellows one's perceptions, and because blood-suffusion reddens the eye, it reddens one's

perceptions (see Burnyeat 2012, p.288). But nothing important depends on these examples, because we now have compelling empirical evidence of widespread interpersonal variation in color perception, even among those with perfectly healthy eyes. This is partly due to variation in the density, placement, sensitivity, and tuning of the detectors responsible for color perception. I review the empirical evidence elsewhere (Morrison manuscript and forthcoming; and see Cohen 2009, p.29–30; Hardin 1997, p.76–80, 163–164). We could focus on any of these other examples. Nonetheless, given that this is a volume on Sextus, let's continue to focus on the variation between children and their grandparents, specifically the variation between Jacob's and Sarah's perceptions of the stone tile.

2 Perceptual Atomism

Perceptual atomism is a view about how we manage to perceive secondary qualities, and thus how we manage to perceive them accurately or inaccurately. Let's introduce perceptual atomism by elaborating on our initial example. If Sarah is perceiving the stone tile alongside a porcelain tile, she doesn't just perceive the stone tile as dark gray. She also perceives the porcelain tile as bright white, and the stone tile as darker than the porcelain tile. Let's separate these three elements:

1. Sarah perceives the stone tile as dark gray.

2. Sarah perceives the porcelain tile as bright white.

3. Sarah perceives the stone tile as darker than the porcelain tile.

What explains (1), (2), and (3)?

We're looking for a kind of explanation that's familiar from the philosophy of language. Philosophers of language want to explain how 'Napoleon' manages to refer to Napoleon. Some appeal to causal relations between 'Napoleon' and Napoleon. Others appeal to descriptions associated with 'Napoleon' that single out Napoleon (e.g., "the French general who lost at Waterloo in 1815"). These aren't the only explanations of how 'Napoleon' manages to refer to Napoleon, but I hope they're enough to get a grip on the relevant kind of explanation. Explanations of this kind are called "metasemantic," because they don't just indicate the objects and properties that our

terms refer to (that's semantics); they explain how our terms manage to refer to those objects and properties.

We're searching for a similar kind of explanation. We want to explain how Sarah manages to perceive the stone tile as dark gray, the porcelain tile as bright white, and the stone tile as darker than the porcelain tile. Explanations of this kind are often called "psychosemantic." They don't just indicate the objects and properties that our psychological states refer to; they explain how our psychological states manage to refer to those objects and properties. (For this reason, 'psycho*meta*semantics' is a better description.)

Perceptual atomists start by explaining facts like (1) and (2), and they explain them by appealing to relations involving perceptions of the same kind. We might think of atomists as first filling in schema like:

1.′ Sarah perceives the stone tile as dark gray because _____ .

2.′ Sarah perceives the porcelain tile as bright white because _____ .

Atomists often appeal to *causal* relations. For example, an atomist might say that Sarah perceives the stone tile as dark gray because dark gray objects cause that kind of perception in ideal perceivers under ideal conditions (Tye 2002, p.138). Or an atomist might say that Sarah perceives the stone tile as dark gray because dark gray objects cause that kind of perception in normal perceivers under normal conditions (Peacocke 1984, p.373). Or an atomist might say that Sarah perceives the stone tile as dark gray because dark gray objects cause that kind of perception in Sarah under normal conditions (Egan 2010, p.88–92, and Chalmers 2006, p.59, on what he calls "imperfect colors"). Or an atomist might say that Sarah perceives the stone tile as dark gray because that's the color shared by all the objects that caused that kind of perception in Sarah during her learning period (Dretske 1981, Ch 8). Or, an atomist might say that Sarah perceives the stone tile as dark gray because that's the color shared by objects that most frequently caused that kind of perception in Sarah in the past (extrapolating Rupert 1999). Etc.

Atomists don't always appeal to causal relations. An atomist might instead appeal to *teleological* relations. For example, an atomist might say that Sarah perceives the stone tile as dark gray because that kind of perception was selected for its ability to indicate dark gray objects (extrapolating Millikan 1989). Or an atomist might say that Sarah perceives the stone tile as dark gray because that's part of the best explanation of how Sarah and

members of her species use that kind of perception to avoid predators, avoid prey, and perform other, similarly significant zoological actions (extrapolating Burge 2010, Ch 8).

Still other atomists appeal to phenomenal relations, or relations that obtain solely because of the phenomenal character of one's perception. An atomist might say that Sarah perceives the stone tile as dark gray because the phenomenal character of her perception perfectly matches dark grayness (Pautz 2010b, p.58–60, and Chalmers 2006, p.66, on what he calls "perfect colors"). Or an atomist might say that Sarah perceives the stone tile as dark gray because the phenomenal character of her perception was in some sense inherited from the dark grayness of the stone tile (Campbell 1993, p.268).

As I hope this list makes clear, atomism is a big tent, covering a diverse group of philosophers. What all these philosophers have in common is that they start by explaining facts like (1) and (2), and explain each kind of perception by appealing only to relations involving that kind of perception. For example, their explanation of (1) doesn't mention the kind of perception involved when we perceive a porcelain tile as bright white, and their explanation of (2) doesn't mention the kind of perception involved when we perceive a stone tile as dark gray. In this sense, their explanations of facts like (1) and (2) are independent of each other.

After explaining facts like (1) and (2), atomists explain facts like (3). Atomists claim that Sarah perceives the stone tile as darker than the porcelain tile *because* she perceives the stone tile as dark gray and the porcelain tile as bright white. They think of color perception like a by-the-numbers painting, where what's explanatorily basic are assignments of colors to individual objects, such as dark gray to the stone tile and bright white to the porcelain tile. This assignment of colors to individual objects gives rise to color relations between the objects.

Explaining (3) in this way is less straightforward than it might first appear. At least in principle, it seems possible to *think* that the stone tile is dark gray and *think* that the porcelain tile is bright white, without *thinking* that the stone tile is darker, perhaps because it never crossed one's mind to compare them. In some cases, it might even be possible to *perceive* that the stone tile is dark gray and *perceive* that the porcelain tile is bright white, without *perceiving* that the stone tile is darker, perhaps because the stone tile and porcelain tile are sufficiently far apart. Thus, the atomist must explain what's different about this case, so that Sarah perceives the stone tile as darker than the porcelain tile. That is, they must explain what it is about

Sarah and her perceptions such that she perceives the stone tile as darker than the porcelain tile because she perceives the stone tile as dark gray and the porcelain tile as bright white. We might think of atomists as filling in schema like:

3.′ Sarah perceives the stone tile as darker than the porcelain tile because Sarah perceives the stone tile as dark gray, Sarah perceives the porcelain tile as bright white, and _____.

Importantly, an atomist needn't say that Sarah *first* perceives the stone tile as dark gray and the porcelain tile as bright white, and *then* perceives the stone tile as whiter, perhaps as the result of something analogous to an inference. That is, an atomist needn't postulate a delay between facts like (1) and (2) and facts like (3). The connection between these facts could be much tighter, so that (3) is simultaneous with (1) and (2), and no further psychological activity is necessary.

I don't think this issue has received the attention it deserves.[2] But we don't need to worry about it here, because, for our purposes, it's enough that, for the atomist, facts like (1) and (2) are supposed to be part of the explanation of facts like (3).

Why do I call this "perceptual atomism"? I think there's a helpful parallel with metaphysical atomism. Metaphysical atomists claim that stone and porcelain tiles result from combining more fundamental objects that are metaphysically independent of one another. At a minimum, the most fundamental objects are independent of each other in that we can say what each is without mentioning the others. They might also be independent of each other in that it's possible for each to exist without the others. Likewise, perceptual atomists claim that Sarah's total perception of the porcelain and stone tiles results from combining more basic perceptions, namely her perception of the stone tile as dark gray and her perception of the porcelain tile as bright white, and that these perceptions are explanatorily independent of one another. At a minimum, the atomist's explanations are independent of each other in that the atomist's explanation of each doesn't mention the others. They might also be independent in that it's possible for her to perceive

[2]The notable exceptions are Byrne (2003) and Matthen (2010). Clark (2000, p.232–235), Hilbert and Kalderon (2000, p.200–201), and Pautz (2011, p.397) briefly address it. Thanks to Nemira Gasiunas for many helpful conversations on this topic.

the stone tile as dark gray even if she's unable to perceive anything as bright white.

I also call this "perceptual atomism" because there's a helpful parallel with conceptual atomism. Conceptual atomists explain why Sarah has the concept of dark gray without appealing to relations involving other concepts, such as the concept of bright white (see Jackman 2017). But there are also dissimilarities. An obvious dissimilarity is that conceptual atomism is about how we manage to *think* about dark gray, whereas perceptual atomism is about how we manage to *perceive* dark gray, and these capacities might have different explanations. Another dissimilarity is that conceptual atomism implies that all our concepts are explanatorily independent of each other, including our concepts of relations, such as our concept of darker than. Perceptual atomism, in contrast, just implies that our perceptions of individual colors, such as gray and white, are explanatorily independent on each other. According to perceptual atomists, our perceptions of relations, such as darker, are explanatorily dependent on our perceptions of individual colors. A final dissimilarity is that conceptual atomism is about what's required for us to *possess* a concept, not about what's required for us to then *employ* that concept. Even if conceptual atomism is wrong, and we can possess the concept of gray only if we possess many other concepts, including the concept of white, we might still be able to think about an object as gray without thinking about anything as white. Perceptual atomism, in contrast, is about what's required to perceive a color relation between two objects, and thus about what's required to *employ* a perceptual representation, rather than merely about what's required to possess that representation. According to perceptual atomists, we perceive an object as darker than another only if we perceive the individual colors of those objects and our perceptions satisfy whatever further condition is slotted into (3′).

Perceptual atomism is widespread and plausible. Nonetheless, I will argue that, when perceptual atomism is combined with perceptual relativism, it becomes hard to explain color inaccuracy and color constancy. At root, the problem is that, given atomism, the relativist is pushed to say that perception is relative to contexts as well as perceivers, and thus to accept that "all things apparent to men actually are." In the next two sections I'll explain why.

Figure 1: Kitaoka's illusion

3 Relativism and Context

So far, we've been considering variation across two perceivers, namely Jacob and Sarah. But there's also variation within the same perceiver. Sextus provides a number of examples: if you look at the sun, and then look at a book, the letters will look golden (I:45); if you smear your lamp-wick with bronze rust, everything will look bronzed (I:46); if you smear your lamp-wick with cuttlefish ink, everything will look black (I:46); if you look at a lamp under sunlight, it will look dimmer than when you view it in the dark (I:118); and if you look at a dove's neck as it turns its head, you will see a range of colors (I:120).

These are vivid examples, but hard to replicate. Let's therefore focus on Kitaoka's (2006) lightness illusion (Figure 1). Perhaps surprisingly, the left and right squares are intrinsically alike. More precisely, they reflect light in the same way, and, as a result, could be interchanged without affecting the illusion. Which square, if either, are you accurately perceiving?

As we defined 'relativism', it is a view about variation across different perceivers, not a view about variation within the same perceiver. It thus doesn't imply any particular answer to this question. Nonetheless, a relativist who accepts atomism should answer that you're accurately perceiving *both* squares. To see why, let's consider why they shouldn't give a different answer.

To start, an atomistic relativist shouldn't answer that you're accurately perceiving only *one* of the squares. For the sake of argument, suppose that you're accurately perceiving only the left square. The problem is that, from

an atomistic point of view, nothing seems to justify the belief that you're accurately perceiving only the left square. Let's consider the available evidence: Your perception of the left square is phenomenal-dark-gray, while your perception of the right square is phenomenal-medium-gray. But there's no reason to believe that the accurate perception is phenomenal-dark-gray rather than phenomenal-medium-gray. Similarly, your perception of the left square is the result of viewing it against a relatively light background, while your perception of the right square is the result of viewing it against a relatively dark background. But there's no reason to believe that the accurate perception results from viewing the left square against a lighter background rather than a darker background. Finally, underlying each perception is a different pattern of activity in your eye and brain. But, once again, there's no reason to believe that the accurate perception resulted from one causal process rather than the other. More generally, none of the intrinsic features of these perceptions, including their phenomenal characters, and none of their relations to the squares, including their causal relations, seem to justify the belief that you're accurately perceiving only the left square. From an atomistic point of view, it's unclear what else could justify that belief, because atomists say that only your perception's intrinsic features and relations to objects are relevant. Thus, an atomist who claims that you're accurately perceiving only one of the squares is apparently committed to ignorance about which square it is. But perceptual relativism is motivated, in part, by our unwillingness to say that only one perceiver is accurate, because that would lead us to this very kind of ignorance. Thus, an atomistic relativist should be similarly unwilling to say that only one of your perceptions is accurate.

An atomistic relativist also shouldn't answer that *neither* perception is accurate. According to relativism, you, Sarah, Jacob, and all other normal observers *can* accurately perceive the colors of objects. Thus, an atomist who denies that you're accurately perceiving the left square's color in its current context must think that you accurately perceive its color in another context. But, similarly to before, there's no reason to believe that you accurately perceive the square's color in another context rather than in its current context. More generally, there's no reason to believe that you accurately perceive the square's color in any particular context rather than all the other contexts.

I conclude that an atomistic relativist should answer that you're accurately perceiving *both* squares.

Importantly, this line of reasoning doesn't presuppose that there's no way to distinguish between normal and abnormal contexts. How might atomistic

relativists distinguish them? As noted in the introduction, our natural ways of thinking and talking are what push us to maintain that things are colored. Thus, it would be natural for anyone who maintains that things are colored to also rely on these ways of thinking and talking to distinguish between normal and abnormal contexts. Alternatively, one might use statistical regularities. But, even granting the atomist this distinction, the problem remains, because there will still be enough variation across normal contexts. For example, the atomistic relativist will be pushed to classify the left square's current context as normal, because it's so similar to paradigmatic examples of normal contexts. We're surprised to learn that the squares are intrinsically alike precisely because there doesn't seem to be anything abnormal about their current contexts. We often view objects against dark and light backgrounds, and we don't regard either background as abnormal. More generally, our eyes and brains aren't perfect at tracking lightness constancies, and as a result intrinsically alike objects sometimes look different, even in paradigmatic examples of normal contexts, such as classrooms and playgrounds. Kitaoka's illusion is just a vivid example of this widespread phenomenon (see, e.g., Gilchrist 2006, Ch 10). Thus, even if the atomistic relativist can motivate a distinction between normal and abnormal contexts, there's still no reason to believe that you accurately perceive the square's color in any particular normal context, rather than any of the other normal contexts in which they cause perceptions with different phenomenal characters.

For these reasons, I think that atomistic relativists should respond that you're accurately perceiving *both* squares. But how is that possible? After all, there are obvious differences between your perceptions. In the next section, I'll introduce what I think are the two most promising proposals. After explaining why these proposals make it hard to explain color inaccuracy and constancy, I will consider other proposals, and explain why they either make it just as hard to explain color inaccuracy and constancy, or have other, more serious drawbacks. I will conclude that perceptual relativists should give up perceptual atomism.

4 Two Proposals

According to perceptual atomists, you perceive the left square's color because of a certain relation involving that kind of perception. Which relation? To what kind of perception? And in what context? I think that an atom-

istic relativist should say that it's a *causal* relation to perceptions with *that phenomenal character* in *that context*. In particular, I think an atomistic relativist should say that you perceive the left square as dark gray because the left square causes phenomenal-dark-gray perceptions in you in that context. The appeal of this suggestion will become clear when we consider the alternatives. But briefly: it doesn't lead to ignorance about who can accurately perceive the squares; it doesn't lead to ignorance about the contexts in which they can accurately perceive the squares; it doesn't attribute any properties to the squares that aren't already attributed by our best scientific theories; and it doesn't appeal to any relations between you and the squares that aren't already mentioned by our best scientific theories.

There are two proposals along these lines. The first is about the colors themselves. The second is about the descriptions that pick out the colors. More specifically:

Proposal 1 The first proposal is that you are perceiving the left square as instantiating *causes phenomenal-dark-gray perceptions in me in this context*. According to this proposal, this is among the many colors of the left square. Thus, you're accurately perceiving its color. Likewise, you are perceiving the right square as instantiating *causes phenomenal-medium-gray perceptions in me in this context*. Thus, you're accurately perceiving its color as well. This first proposal likewise explains why Sarah and Jacob are both accurately perceiving the stone tile. In particular, Sarah perceives the stone tile as instantiating *causes phenomenal-dark-gray perceptions in me [Sarah] in this context* and Jacob perceives the stone tile as instantiating *causes phenomenal-medium-gray perceptions in me [Jacob] in this context*. Because these are among the stone tile's many colors, they are both accurately perceiving a color of the stone tile. Cohen (2004, 2009) endorses and develops this proposal. It's the proposal about the colors themselves.

Proposal 2 The second proposal is that you are perceiving the left square as instantiating whatever property satisfies a description like: the property that disposes the left square to cause phenomenal-dark-gray perceptions in me in this context. Likewise, you are perceiving the right square as instantiating whatever property satisfies the description: the property that disposes the right square to cause phenomenal-medium-gray perceptions in me in this context. Each description picks out a property of the relevant square. Thus,

you're accurately perceiving both. Moreover, you're accurately perceiving *the colors* of both squares, because to be a color is just to satisfy a description of this kind. This second proposal also explains why Sarah and Jacob are both accurately perceiving the stone tile. In particular, Sarah perceives the stone tile as instantiating whatever property satisfies a description like: the property that disposes the stone tile to cause phenomenal-dark-gray perceptions in me [Sarah] in this context. And Jacob perceives the stone tile as instantiating whatever property satisfies a description like: the property that disposes the stone tile to cause phenomenal-medium-gray perceptions in me [Jacob] in this context. Thus, Sarah and Jacob are both accurately perceiving its color. Jackson and Pargetter (1987) and McLaughlin (2003) endorse and develop this proposal.

Some further observations about this second proposal might help clarify it. To start, it is natural to think of the relevant descriptions as *modes of presentation*, so that you're perceiving the left square's color under a different mode of presentation than the right square's color, and Sarah is perceiving the stone tile's color under a different mode of presentation than Jacob is perceiving the stone tile's color (see Jackson 2000, p.158; 2007, p.175–178). But one needn't think this. One might just think of these descriptions as singling out the properties you, Sarah, and Jacob are perceiving, without corresponding to anything like a mode of presentation.

Moreover, the second proposal doesn't imply that the relevant properties belong to a specific metaphysical category. For example, they might be the percentage of light that the relevant objects reflect at each wavelength, i.e., their "reflectances" (as Byrne and Hilbert (1997, 2003) claim). Or they might belong to a metaphysical category that isn't mentioned by our best scientific theories (as Campbell (1993) claims). In other possible worlds, they might even be shapes, because in those words an object's shape is what disposes it to cause perceptions with the relevant phenomenal character. Thus, whereas the first proposal tells us exactly which properties you, Sarah, and Jacob are perceiving, the second proposal just implies that the relevant properties dispose objects to cause certain perceptions, leaving the metaphysical category of those properties as a question for further investigation.

Finally, the second proposal leaves open how these properties relate to one another. For concreteness, suppose that the colors we perceive are reflectances. One possibility is that when you perceive the left square, you're perceiving its reflectance, and when you're perceiving the right square, you're perceiving its reflectance. In that case, because both squares have the same

reflectance, you're perceiving each square as having that reflectance. Another possibility is that you're perceiving each square as having a different determinable of that reflectance. Perhaps (massively oversimplifying) the left square is disposed to cause phenomenal-dark-gray perceptions because of the percentage of light it reflects around 450nm, and the right square is disposed to cause phenomenal-medium-gray perceptions because of the percentage of light it reflects around 650nm. In that case, when you perceive the left square in that context, you're perceiving a determinable of its reflectance that has to do with how it reflects light around 450nm, and when you perceive the right square in that other context, you're perceiving a determinable of its reflectance that has to do with how it reflects light around 650nm. One might think of these determinables as different colors, or, following Kalderon, one might think of them as different "aspects" of the same color (Kalderon 2011, p.248–249). A third possibility is that you're perceiving completely unrelated properties. While this is unlikely in the actual world, given the systematic correlation between our perceptions and how objects reflect light, it's not hard to imagine a counterfactual in which phenomenal-dark-gray and phenomenal-medium-gray perceptions have radically different causes. In one such counterfactual, phenomenal-dark-gray perceptions are caused by elliptical objects, and phenomenal-medium-gray perceptions are caused by objects with sufficient density, and thus in that world you're perceiving completely unrelated properties, namely a shape and a density.

The difference between these proposals is analogous to the difference between role functionalism and filler functionalism about mental states (see Cohen 2009, p.76, 176;s Pautz 2010a, fn 1). Functionalists of both kinds agree that there's a close connection between a mental state, such as the belief that it is raining, and the function of that mental state, such as to prompt you to carry an umbrella. According to role functionalists, the belief just is that function (more carefully: the property of being in that functional state). According to filler functionalists, the belief is whatever state in an organism carries out ("fills") that function, regardless of whether it's a neural state, silicon state, or state of some other kind.

I will now argue that these proposals have difficulty explaining color inaccuracy and color constancy.

4.1 Inaccuracy

The problem with both proposals is that almost all color perceptions would be accurate, even though at least some of our perceptions seem inaccurate. Consider the first proposal. If you have a phenomenal-dark-gray perception in a context, and your perception is caused by an external object, that object instantiates *causes phenomenally-dark-gray perceptions in you in that context*. Moreover, if you have a phenomenal-dark-gray perception in a context, you perceive an object as instantiating *causes phenomenally-dark-gray perceptions in me in that context*. This is because, whatever your perception's phenomenal character, you perceive something as the cause of a perception with that phenomenal character. Thus, unless you're misidentifying the object causing your perception, or the causal chain is deviant, your perception is accurate. The second proposal has the same consequence, with the unimportant difference that your perception represents the property that *disposes* the object to cause a phenomenally-dark-gray perception in you in that context. Thus, at least with respect to color perception, both proposals imply that "all things apparent to men actually are."

This is a problem, because at least some of our perceptions seem inaccurate. In Kitaoka's illusion, we perceive the left square as darker, even though the squares are intrinsically the same, and thus we seem to be misperceiving at least one of the squares. This illusion is due to variation in the background. In Munker-White's illusion (Figure 2), we perceive the gray rectangles on the left as darker, even though they are intrinsically the same, and thus we seem to be misperceiving at least some of the rectangles. Our misperception is due to the differences in the occlusion and illumination cues. In Cornsweet's illusion (Figure 3), we perceive the left square as darker than the right square, even though they reflect light in exactly the same way, and thus we seem to be misperceiving at least one of the squares. Our misperception is due to a misleading edge cue, in particular the line between the two squares. Our misperceptions have many other causes. For example, even if two objects reflect light in the same way, we might perceive one as darker, because of an undetected asymmetry in how the objects are illuminated, or because of misleading motion and depth cues. These other causes of misperception are important, but harder to illustrate in print.

It would be hard for anyone to deny that in these illusions we're misperceiving the individual color of at least one of the relevant objects. But it would be especially hard for an atomist. According to the atomist, we per-

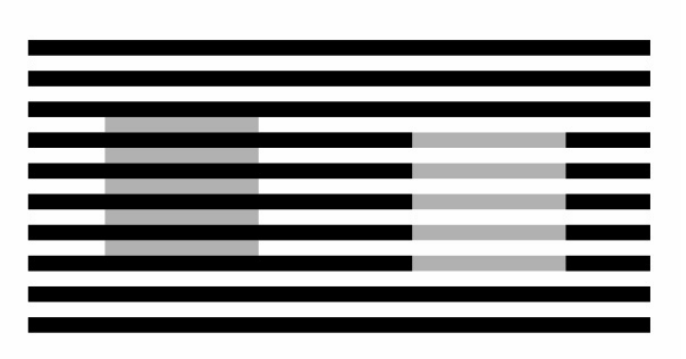

Figure 2: Munker-White's illusion

Figure 3: Cornsweet's illusion

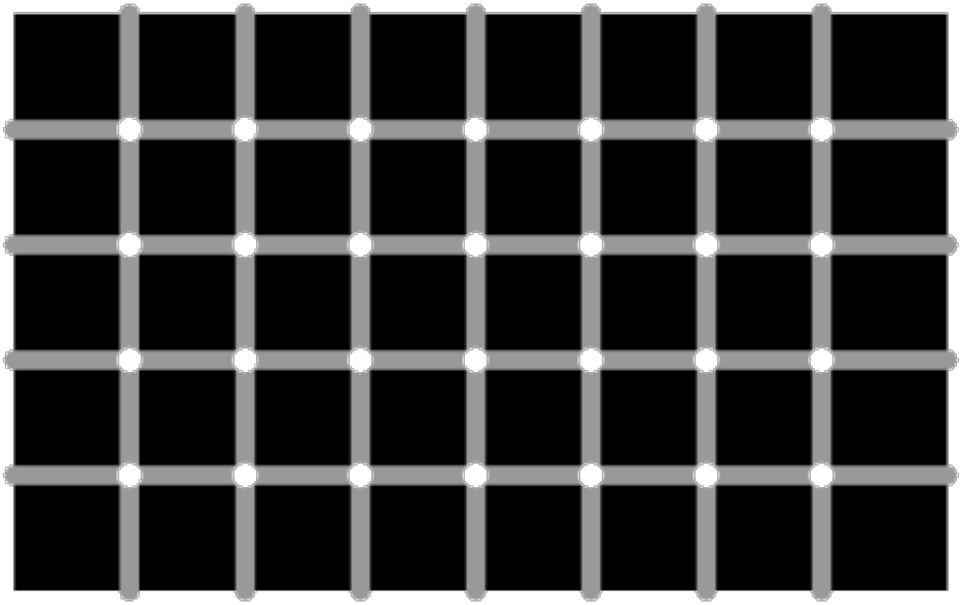

Figure 4: Bergen's illusion

ceive a relation between two objects because of how we perceive each of their individual colors. In all of the examples above, we misperceive a relation between two objects. Thus, the atomist seems forced to concede that we're misperceiving the color of at least one of the relevant objects.

Other illusions provide equally compelling examples of misperception, even though they don't involve a misperception of a relation between two objects. Consider Bergen's grid illusion, in which black dots appear and disappear inside the white dots as you look around (Figure 4). Whenever a black dot appears inside a white dot, you seem to be misperceiving the white dot.

Next, consider the "filling in" illusion that results when your visual system "fills in" the region of the visual field obstructed by the optic nerve (Figure 5). Close your left eye and center your right eye on the cross. When the your eye is about six centimeters away from the page, move the page a little closer or further. At some point, the dot will disappear. When it does, you seem to be misperceiving the page where the dot used to be.

Finally, consider the illusion that results when you focus your gaze on the center of a hazy circle (Figure 6). After a few seconds, the haze will seem to

Figure 5: "Filling in" illusion

Figure 6: Haze illusion

shrink. When it does, you seem to be misperceiving the page where the haze used to be.

In all of these examples, you seem to misperceive the colors of objects. But both proposals imply that you're perceiving the colors of all the objects accurately.

Cohen acknowledges this difficulty (2007, p.338–339; see also Clark 2000, p.226–227).[3] He tries to mitigate it by claiming that standard examples of color illusion involve inaccurate predictions at the level of thought about how an object will look in other contexts. With respect to Kitaoka's illusion, you

[3]Other than hallucinations, McLaughlin mentions just one way for perceptions to be inaccurate, namely that the relevant property isn't shared by everything that disposes your current phenomenal character in your current context (2003, p.123). I'm setting this aside, because it doesn't seem to explain the inaccuracy of your perceptions of these illusions.

might inaccurately predict that the squares would look different if viewed by certain perceivers in certain contexts. Which perceivers? Which contexts? Cohen suggests that the relevant context is determined by our linguistic community, because illusions involve the misapplication of color predicates, and our the linguistic community uses certain perceivers and contexts to establish the correct application of predicates such as 'gray'. Cohen thereby relocates color inaccuracy to cognition and language. For example, in Kitaoka's illusion, your false predictions might be about how the relevant objects would look if viewed against a uniform white background, rather than a graded or black background, because that's the context our linguistic community uses to establish the correct application of 'gray'. Likewise for the other illusions.

This is unsatisfying, because, as Cohen (2007, p.346–347) acknowledges, it implies that animals are incapable of color inaccuracy, unless they have a non-linguistic way of singling out certain perceivers and contexts as normal, and also represent those perceivers and contexts at the level of thought. It also implies that whether a perception is illusory can change as a result of a change in our language, or perhaps even just a change in which language we're inclined to use at that moment, as long as users of those languages single out different perceivers and contexts as normal. But most fundamentally, this proposal is unsatisfying because our errors are sometimes entirely perceptual. In Kitaoka's illusion, we're *misperceiving* the left square as darker, and thus we're *misperceiving* at least one of the squares. In Munker-White's illusion, we're *misperceiving* the rectangle on the left as darker, and thus we're *misperceiving* at least some of the rectangles. And so on. Our errors are perceptual, not cognitive or linguistic. There is disagreement about where to draw the border between perception and cognition. But I think anyone who draws the border so that these errors are non-perceptual is drawing the border in the wrong place.[4]

This difficulty is ultimately due to perceptual atomism. From an atomistic point of view, there's no way to justify the belief that you're accurately

[4]Tye (2012, p.300–302) also objects to Cohen's account of color inaccuracy. Tye considers two squares in Adelson's checker-shadow illusion. As Tye understands Cohen's account: You perceive the two squares as different because of the phenomenal characters those squares cause in their respective contexts, and this perception is inaccurate because of the phenomenal characters those squares would cause if they were cut out and viewed in the same context. Tye's objection (which I won't reproduce here) is effective. But it depends on a misunderstanding of Cohen's account, because Cohen denies that your perception is itself inaccurate.

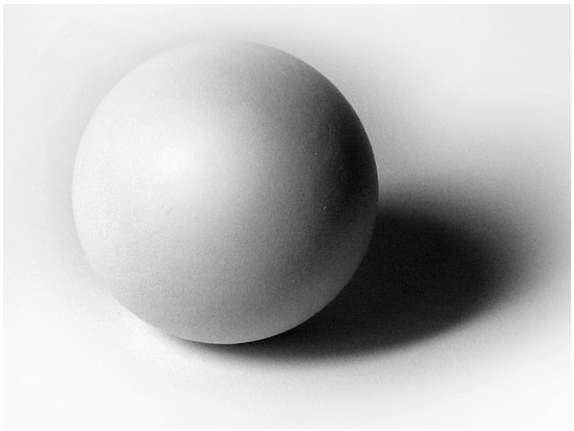

Figure 7: Color constancy

perceiving in one normal context rather than the other. This pushes atomists to say that you're accurately perceiving in both contexts. But, as we just saw, the best proposals imply that almost all of your perceptions are accurate, because what you perceive is relative to your current context, and in your current context the relevant object really is causing a perception with that phenomenal character.

4.2 Constancy

Both proposals also have trouble explaining color constancy. Consider a sphere illuminated from only one side (Figure 7). Different regions on the sphere's surface cause perceptions with different phenomenal characters. Even so, there's *a sense* in which you perceive the entire sphere as the same color. Both proposals have trouble capturing this aspect of color constancy.

Proposal 1 Let's start with Cohen's (2008) explanation. Consider a region of the sphere's surface that's causing phenomenal-dark-gray perceptions in you. According to Cohen, you perceive that region as *causes phenomenal-dark-gray perceptions in me in that shadow context*. While other regions on the sphere are causing perceptions with different phenomenal characters, they would cause phenomenal-dark-gray perceptions if rotated into the same shadow context. Thus, they also have the dispositional property *causes*

phenomenal-dark-gray perceptions in me in that shadow context. According to Cohen, this explains the sense in which the sphere looks the same color: you perceive each region as having all these dispositional properties, and you also perceive this uniformity, even though each region manifests only one of those dispositions.[5]

I don't find Cohen's explanation satisfying, because I think it misdescribes color constancy. We don't just perceive all regions of the sphere's surface as having the same bundle of dispositional properties. We perceive something above and beyond those dispositions. We think we can pick out that extra something using a demonstrative like 'that color'. There's a helpful parallel with shape constancy. Rotate the page while focusing your eyes on the shape in Figure 8. At each angle of rotation, the figure is disposed to cause a perception with a different phenomenal character. But you don't just perceive the figure as having a constant bundle of dispositions. Assuming you perceive these dispositions at all, you also perceive a property that's above and beyond them — namely, the figure's shape. Your perceptions of colors and shapes differ in many other respects. But in this respect I think they're the same.

Other philosophers have made a similar point about constancy in connection with intentionalism, the view that for every change in the phenomenal character of our perceptions, there is a change in what color we perceive (most famously, Peacocke 1983, p.12–13). These philosophers point out that there also seems to be a color that we perceive as constant. I think they're

[5]I'm setting aside a complication: Cohen (2008) is noncommittal about whether you *perceive* each region as instantiating all these dispositional properties, and thus about whether you *perceive* constancy. He wants to leave open the possibility that this occurs later, at the level of cognition, because he doesn't want to commit to any specific way of distinguishing perception from cognition. He thus merely describes constancy as *driven by* our visual system (p.80). In a more recent paper (2015, p.632–634) he says that the empirical evidence strongly suggests that color constancy isn't "exclusively perceptual." But this doesn't mean he's now committed to saying that constancy occurs only later, at the level of cognition. To start, he never makes this stronger claim. In addition, the only evidence he cites is that constancy is enhanced for familiar objects. Even granting that this is evidence of a top-down influence on perception, constancy might still be exclusively perceptual in the sense that it might be represented by a state that's exclusively perceptual.

Despite Cohen's official neutrality, I'm going to present him as offering a proposal about the sense in which you *perceive* the sphere as the same color, because I think that's the best version of the proposal he offers. I'll explain why later. I'll also explain why the problem with this proposal doesn't depend on whether the relevant representation is perceptual or cognitive.

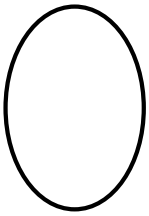

Figure 8: Shape constancy

right, and that, regardless of whether this is really a problem for intentionalism, it seems to undermine Cohen's attempt to explain color constancy. In particular, I think that, even if Cohen has explained *a sense* in which we perceive all regions on the cube's surface as the same, he hasn't explained color constancy, because color constancy involves something over and above the relevant dispositions.

As a way of reinforcing this point, consider the *number* of dispositions that you'd have to perceive in order to perceive the entire sphere as the same color (Gert 2010, p.678). Of course, it's hard to count the number of dispositions, because there's at least one disposition for each phenomenal character, and it's hard to count your phenomenal characters. Nonetheless, because you presumably have more than fifty phenomenal characters, there are presumably more than fifty dispositions. Thus, according to Cohen, you're perceiving at least fifty dispositions *at each location* on the sphere's surface, you're perceiving that uniformity, and *this* is the sense in which you perceive the sphere's color as constant. I'm doubtful that you perceive so many dispositions at each location, at least if your perception of the sphere is like mine, because I don't find anything like that. In addition, even if you do perceive all of these colors, this doesn't seem to be the sense in which you perceive the sphere's color as constant. There seems to be a special color that you perceive as constant, and which you can pick out using a demonstrative like 'that color'. If anything, it's your perception of this special color that seems to underlie your predictions about how each region on the sphere would look if you rotated it to the left or right, moving it farther into or farther out of the shadow. Likewise, it's your perception of the ellipse as an ellipse that seems to underlie your predictions about how it would look if you moved it closer or farther, to the left or to the right.

There's another way of bringing out this same problem. Your perception

of the sphere is an example of synchronic color constancy, or constancy at a given time. There's also diachronic color constancy, or constancy across times. Suppose at t_1 under a bright light you perceive an object as *causes phenomenal-dark-gray perceptions in me in context*$_1$, and then a few seconds later at t_2 someone dims the light so that you perceive the same object as *causes phenomenal-light-gray perceptions in me in context*$_2$. According to Cohen, if you perceive its color as constant, it's because at t_2 you still perceive it as *causes phenomenal-dark-gray perceptions in me in context*$_1$. If we reversed the order in which you viewed the object, so that you first viewed it under the dimmer light, and then under the brighter light, you would have perceived its color as constant because at the later time you would still perceive it as *causes phenomenal-light-gray perceptions in me in context*$_2$. Thus, if we reversed the order in which you viewed the object, you would perceive its color as constant because of your perception of a different disposition. This loses an important feature of color constancy, which is that, in many cases, the constant color is the same regardless of the order of our perceptions, such as whether you perceive the object under a dim light that's getting brighter, or under a bright light that's getting dimmer. In this respect, color constancy is again like shape constancy, because the constant shape you perceive doesn't seem to depend on the order of your perceptions, such as whether you start rotating it from its current position, with its major axis up-down, or from another position, such as with its major axis left-right.[6] As you rotate an object, or move it closer to the illuminant, the shifts in your phenomenal character might allow you to become aware of additional dispositions. But this doesn't affect your perception of it as a constant color.

I conclude that Cohen's proposal fails to explain color constancy.

It might be helpful to distinguish this objection from existing objections. Armstrong (1987, p.4–5) and many others argue that colors don't look like relations, and thus don't look like dispositions to cause certain perceptions in us (see Cohen 2010, fn 5, for additional references). Unlike Armstrong, I'm hesitant to assume that the metaphysical nature of the colors is revealed

[6]As E.J. Green pointed out to me, this isn't *always* true for shape constancy. If you perceive a diamond, and it is then rotated forty-five degrees, you might perceive it as a diamond on its side, whereas if you perceive a square, and it is then rotated forty-five degrees, you might perceive it as a square on its side, even though it's the same geometrical shape in both cases. However, color constancy isn't like that. It's like most cases of shape constancy, in which the order of your perceptions doesn't matter.

to us by our phenomenology (for more discussion, see Cohen 2010, p.15–22, and Johnston 1992, p.226–228, 254–259). I think we should be open to the possibility that colors are relations, even if that's not how they seem. We didn't make that assumption. We argued, on phenomenological grounds, that a perception of the sphere's constant color is not just a perception of each region as instantiating *causes phenomenal-medium-gray perceptions in me in that shadow context, causes phenomenal-dark-gray perceptions in me in that other shadow context*, and so on. That is, we argued, on phenomenological grounds, for the *existence* of a color above and beyond these dispositions, without making any assumptions about its metaphysical nature.

Johnston (1992, p.226–227) observes that we can perceive a series of rapidly changing dispositions without perceiving a constant color, as when we look at the back of a compact disk. Building on this observation, Kalderon (2008, p.950–953) argues that there's no principled explanation for why perceiving the compact disk's dispositions isn't sufficient for color constancy, but (to use our example) perceiving the sphere's dispositions is sufficient for color constancy. He concludes that color constancy involves more than just perceiving dispositions. While I agree with the conclusion, I'm not convinced by the argument. Perhaps the compact disk's dispositions are too unstable, in that even the slightest shift can change which dispositions are manifest. Or perhaps, as Gert (2013, p.194–197) suggests, the compact disk's dispositions just haven't been unified under the kind of concept that's necessary for color constancy. As these alternative explanations suggest, it might be possible to explain why perceiving dispositions is sometimes, but not always, sufficient for perceiving a constant color. Our objection is different. We argued that, even if the relevant dispositions are stable and unified, the color we perceive as constant is above and beyond these dispositions.

It's unclear how else Cohen could explain color constancy, at least while respecting the motivations for atomism. Atomism is motivated, in part, by the thought that there's a tight connection between our perceptions of individual colors and our perceptions of their color relations, and that this connection is intelligible through introspection. For example, when we perceive a stone tile as gray, a porcelain tile as white, and the stone tile as darker than the porcelain tile, our perception of the tiles' individual colors doesn't merely seem *associated* with our perception of their relation, as a perception of a banana might be associated with a perception of yellowness. The connection is much tighter. Within Cohen's framework, it's unclear how the individual colors of the sphere's regions could have a sufficiently tight

connection to anything except disjunctions and conjunctions of those colors.

Second Proposal McLaughlin, Jackson, and Pargetter don't offer explanations of color constancy. But, given their account, a natural explanation is that we perceive the sphere as the same color because the same property satisfies all of the relevant descriptions (e.g., 'the property that disposes the sphere to cause phenomenal-dark-gray perceptions in me in that shadow context').

However, given McLaughlin, Jackson, and Pargetter's account of color perception, perceiving the same color isn't sufficient for color constancy. In Kitaoka's illusion, the squares are intrinsic duplicates, so the same property might satisfy both of the relevant descriptions. Yet we don't perceive the squares as the same color; the left square looks darker. It also wouldn't be sufficient to perceive the squares as different determinables of the same determinate, because then your perception would leave open the possibility that the squares are the same as well as the possibility that they're different, and thus you wouldn't perceive the left square as darker. Analogously, if I tell you that my son Ian is between 5 and 10 years old, and my son Isaac is between 3 and 6 years old, I haven't told you that Ian is older. I've left open the possibility that they're the same age as well as the possibility they're different ages.

In addition, given McLaughlin, Jackson, and Pargetter's account of color perception, perceiving the same color also isn't necessary for color constancy, because there are illusions of color constancy. We can perceive objects as the same color even though they're intrinsically quite different. For example, if we modify Kitaoka's illusion by lightening the left square until the two squares look exactly the same, we'll perceive their color as constant even though the squares will have different reflectances. Thus, we can perceive objects as the same, even though different properties satisfy the relevant descriptions.

It's unclear how else they might explain color constancy. As far as I can tell, their only remaining option is to appeal to the descriptions themselves. Perhaps you perceive the sphere as a uniform color because of a *similarity* between the descriptions under which you perceive each region. But the relevant descriptions just specify contexts and phenomenal characters, and there's little similarity between the contexts and phenomenal characters of your perceptions of the sphere. The contexts involve a range of different

illuminations, and the phenomenal characters are as different as phenomenal-dark-gray and phenomenal-medium-gray. Perhaps there's some other relation between the relevant descriptions, besides similarity, that explains why you perceive the sphere as a constant color. But I can't think of one.

Cohen, McLaughlin, Jackson, and Pargetter could always deny that we *perceive* the sphere as a constant color (see fn 4; Chalmers 2006, p.85). But I don't think they should. From a first-personal perspective, your representation of the sphere as a constant color seems paradigmatically perceptual. Among its hallmarks, it's conscious, automatic, accessible, dissociable from what you believe, directed toward nearby objects and properties, and fast enough that we can't detect any delay. From a third-personal perspective, it also seems paradigmatically perceptual, given that color constancy is the product of activity in the visual cortex (Gegenfurtner 2003, p.566–567). It also influences perceptual grouping, among other paradigmatically perceptual phenomena (Palmer 2002, p.103–104).

But more fundamentally, denying that constancy is perceptual wouldn't help, because the problems we identified don't depend on classifying constancy as perceptual. In particular, even if color constancy involves a post-perceptual judgment, that judgment still seems to represent something over and above the dispositions shared by all the regions, and the relevant descriptions still won't be necessary or sufficient for representing constancy.

I conclude that the relativist has difficulty explaining color constancy. Like the relativist's difficulty explaining color inaccuracy, this difficulty is ultimately due to perceptual atomism. From an atomistic point of view, there's no way to justify the belief that you're accurately perceiving one square but not the other. This pushes the relativist to say that you're accurately perceiving both squares. But, as we just saw, the best proposals imply that you're perceiving a different color in every context, making it hard to explain color constancy.

4.3 Stepping Back

We just considered two proposals for combining perceptual relativism and perceptual atomism. I argued that both proposals have difficulty explaining color inaccuracy and color constancy. Importantly, these difficulties aren't peripheral to perceptual atomism. Recall that perceptual atomists make two claims. The first is that our explanations of how we perceive individual col-

ors, such as dark gray and bright white, appeal to a relation involving that kind of perception, and don't appeal to relations involving perceptions of other kinds. Thus, this first claim is about our perceptions of individual colors. But I argued that atomistic relativists don't have a good explanation of our perceptions of individual colors, because they're pushed to say that *all* our perceptions of individual colors are accurate. The atomist's second claim is that we perceive color relations, including the relations involved in color constancy, because of how we perceive individual colors. For them, color perception is like a by-the-numbers painting. But I argued that atomistic relativists don't have a good explanation of our perceptions of color relations, because they're pushed to say that the colors we perceive are relative to the context, and thus they have trouble explaining how we can perceive an object as the same color across contexts. For example, they have trouble explaining how we can perceive an object as the same color across different lighting conditions within the same scene, as when an object is partially covered by a shadow, or across sequential changes in its background or illumination. Thus, these explanatory shortcomings call into question perceptual atomism's definitive commitments.

5 Other Proposals

We just considered two proposals for combining perceptual relativism and perceptual atomism. Given that these proposals make it difficult, if not impossible, to explain color inaccuracy and color constancy, it's natural to wonder about the alternatives. In this section I'll consider the most plausible alternatives. They differ from the proposals we considered in that they appeal to other kinds of perceptions (7.1), other kinds of perceivers (7.2), other kinds of contexts (7.3), or other relations (7.4). I will argue that either these alternatives don't help explain color inaccuracy and color constancy, or they have other, more serious drawbacks. I'll conclude that perceptual relativists should reject perceptual atomism. I'll then sketch the view that I think perceptual relativists should prefer.

5.1 Other Perceptions

According to both of the proposals we considered, you perceive the color of one of the sphere's shaded regions (hereafter just: the shaded region)

because of its causal relation involving *phenomenal-dark-gray* perceptions. Specifically:

> Proposal 1: You perceive the shaded region as instantiating *causes phenomenal-dark-gray perceptions in me in this context.*
>
> Proposal 2: You perceive the shaded region as instantiating whatever property satisfies the description: the property that disposes the shaded region to cause phenomenal-dark-gray perceptions in me in this context.

What other kind of perception might we appeal to? It will be helpful to state this question more schematically. We're trying to fill in the schema:

> Proposal 1b: You perceive the shaded region on the sphere as instantiating *causes* ____ *perceptions in me in this context*
>
> Proposal 2b: You perceive the shaded region as instantiating whatever property satisfies the description: the property that disposes that region to cause ____ perceptions in me in this context.

How should we fill these schema?

The first suggestion, due to Chalmers (2006, p.86–91), is to identify a respect in which perceptions with different phenomenal characters can be the same.[7] For example, one might claim that, while your perceptions of different regions on the sphere's surface have different phenomenal characters, their phenomenal characters are the same in some respect. In particular, they are the same in that they all involve the phenomenal character of perceiving a specific shade of white (see also Hilbert 2005, p.151). Let's call this *phenomenal-underlying-whiteness*$_7$ (the subscript is to indicate that it's a specific shade). Chalmers's suggestion is that this is the relevant kind of perception.

What is the relation between phenomenal-underlying-whiteness$_7$ and the other phenomenal characters of your perceptions, such as phenomenal-dark-gray and phenomenal-medium-gray? There's room for a number of different views. One possibility is that phenomenal-underlying-whiteness$_7$ is a *part* of these other phenomenal characters. Another possibility is that phenomenal-underlying-whiteness$_7$ is a *determinable* of these other phenomenal characters. Yet another possibility is that phenomenal-underlying-whiteness$_7$ is an

[7] Chalmers merely says that he is *inclined* toward this suggestion (p.87).

additional character, above and beyond these other phenomenal characters, in that it can exist without them, and vice versa. There are other possibilities, but, for our purposes, these possibilities will be sufficiently representative.

Following Chalmers's suggestion, one might say that whenever two regions cause phenomenal-underlying-whiteness$_7$ perceptions in you, there is a shade of white such that you perceive both regions as that shade of white. By design, this would explain color constancy, because all of the regions on the sphere cause the relevant kind of perception, namely a phenomenal-underlying-whiteness$_7$ perception, and thus would be perceived as that shade of white. It's therefore a promising suggestion. Nonetheless, I don't think it withstands scrutiny.

For phenomenal-underlying-whiteness$_7$ to help the atomist, it must be included in the phenomenal character of your perception of each region, whether as a part, determinable, or addition. The atomist wants to explain your perception of the sphere's constant color by appealing to your perceptions of each region's color, and without appealing to perceptions of other kinds. Thus, the atomist must appeal to a phenomenal character that's included in your perception of each region. This is perhaps easier to appreciate in cases when you're perceiving several different objects. Suppose you perceive the tiles in a bathroom as a constant color, despite their uneven illumination. The atomist wants to explain your perception of the tiles' constant color by appealing to your perceptions of each tile's color, and without appealing to perceptions of other kinds. An atomist must therefore appeal to a phenomenal character that's included in your perception of each tile. More generally, an atomist can't appeal to global, Gestalt-like phenomenal characters. She can appeal only to what we might call "local" phenomenal characters.[8]

While I agree with Chalmers that there's more to your phenomenology than characters like phenomenal-dark-gray and phenomenal-medium-gray, I don't think there's a local phenomenal character that helps the atomist. To start, when I introspect, I don't find what Chalmers describes. Consider again the squares from the introduction:

[8]Chalmers makes exactly this point. He considers the view that the color you perceive at each region depends on your global phenomenology, and says that "this view requires a certain anti-atomism about perceptual content: the veridicality conditions of an experience of a color at a location are not determined just by the local phenomenology associated with the location, but by the phenomenology of the entire visual experience"(2006, p.86).

I perceive these squares as different colors; this is not an instance of color constancy. I also don't perceive either square as white. Thus, if there is such a thing as phenomenal-underlying-whiteness$_7$, it isn't included in my perceptions of these squares. Yet when I then look at the sphere, there are regions such that my perceptions of those regions have the same local phenomenal characters as my perceptions of these squares. This implies that phenomenal-underlying-whiteness$_7$ isn't included in the local phenomenal characters of my perceptions of these regions of the sphere's surface, as either a part, a determinable, or an additional element. For this reason, Chalmers's suggestion doesn't seem true to the phenomenology. While my perception of the sphere has additional phenomenal characters, they are more global and Gestalt-like, and thus unhelpful to the atomist.[9,10]

[9]This is similar in spirit to Cohen's (2008, p.67–69) objection. Cohen refers to experiments in which subjects create a match between unevenly illuminated surfaces by adjusting the hue of one of the surfaces. Cohen infers that subjects don't perceive the surface's color and illumination as separate elements. He concludes that there aren't separate phenomenal characters for color and illumination. One advantage of our objection is that we don't need to worry that subjects are merely misperceiving the change in hue as a change in illumination, and are thus still perceiving the surface's hue and illumination as separate elements. Another advantage of our objection is that we don't need to worry that subjects are reporting a match, despite phenomenal differences (see Hilbert 2005, p.151–152; 2012, p.207). A third advantage of our objection is that we don't need to worry that these experiments don't generalize from two-dimensional displays to natural scenes (see Gert 2010, p.276–278). More generally, by relying on first-personal introspection of the phenomenal characters, rather than on third-personal data about when subjects report a match between surfaces, we can be more confident that the local phenomenal characters are the same.

[10]I think there's another way of bringing out this problem, internal to Chalmers's system. Chalmers says that when you look at a white object covered by a shadow, the Edenic content of our perception includes Edenic white and what he calls an "Edenic shadow"(p.87). But, while I have a grip on what he calls Edenic white, I can't get a grip on Edenic shadows. I think this shows that the phenomenal character of our perception can't be divided into a shadow component and a color component. Moreover, the alternative suggestion that you are perceiving a "certain mode" of Edenic white (see again p.87) seems to merely label the problem without solving it. We're left wondering: How can Edenic colors have "modes"?

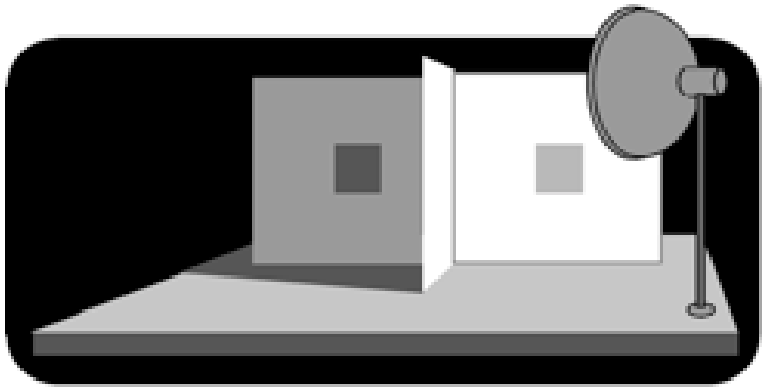

Figure 9: Light and shadow matching experiment, from Gilchrist 2006, p.289

There's another, related problem. In an experiment that's standard among color researchers, subjects view surfaces under two different illuminants, and are asked whether they perceive the same color (Figure 9). For example, on the left side of a partition you might view a surface under a dim illuminant, and on the right side of the partition you might view a surface under a bright illuminant. Suppose you perceive the surface on the left, surface-L, and the surface on the right, surface-R, as the same color. If Chalmers's suggestion were true, your perceptions of both surfaces would share a phenomenal character. Without loss of generality, let's suppose it's phenomenal-underlying-whiteness$_7$.

For almost any pair of illuminants, it's possible to gradually transition from one to the other while you continue to perceive a surface as a constant color. This can be done synchronically, as in Figure 7, by placing a surface under an uneven illuminant. It can also be done diachronically, by gradually adjusting the illumination's location, intensity, and hue. Suppose we gradually increase the illumination of surface-L. As the illumination increases, you will perceive surface-L as the same color. Thus, if Chalmers's suggestion were true, your perception of surface-L would still include phenomenal-underlying-whiteness$_7$.

Recall that, according to Chalmers's suggestion, if your perceptions of two surfaces both include phenomenal-underlying-whiteness$_7$, you perceive them as the same shade of white, and thus as the same color. Chalmers's suggestion thereby generates a prediction: that when the illumination of

surface-L increases enough to equal the right illuminant, you will perceive surface-L and surface-R as the same color. But that's not what happens: you will perceive surface-L as brighter than surface-R. One of the basic results of lightness research is that, in experiments like this, subjects perceive a target as the same color as more reflective surfaces under dimmer illumination, and as less reflective surfaces under brighter illumination, and thus as the same color as surfaces that would look lighter or darker under equal illumination (see Gilchrist 2006, Ch 10).[11]

Chalmers says that his suggestion merely requires the mechanisms of color constancy to work "reasonably well"(2006, p.89). But it actually requires those mechanisms to be optimal, in that it requires them to be able to identify the same reflectance under all normal lighting conditions, because that's the only way they can avoid producing phenomenal characters like phenomenal-underlying-whiteness$_7$ when viewing less reflective surfaces under brighter illuminants, while also producing phenomenal characters like phenomenal-underlying-whiteness$_7$ when viewing a surface under an illuminant that's gradually increasing or decreasing.

I conclude that color constancy isn't explained by a phenomenal character like phenomenal-underlying-whiteness$_7$. This isn't to say that your perception of the sphere doesn't involve *any* additional phenomenal characters. Here's what I find:

First, I find an additional phenomenal character that accompanies my perception of the entire sphere *as the same color*, and which is over and above the phenomenal characters of my perceptions of each region (for a related claim, see James 1950, Volume 1, p.495). I find the same phenomenal character in my perception of shaded apples and lemons, and thus it doesn't seem specific to my perception of white objects. It's a more general feeling of uniformity. For this reason, it won't help the atomist explain why I perceive the ball as white rather than red or yellow, because it isn't caused only by white objects. It also won't help the atomist because it's too global and Gestalt-like; it's over and above the phenomenal characters of my perceptions of each region.

[11]One might try to defend Chalmers's suggestion by insisting that the local phenomenal character of your perception of surface-R changes as the illumination of surface-L increases. But that doesn't seem true in my own case; the local phenomenal characters of my perception of surface-R seem to remain exactly the same. This is due, in part, to the fact that the surfaces are separated by a partition, and are illuminated by different lights, and thus seem completely independent of each other.

Second, I find a phenomenal character that accompanies my appreciation that the sphere is the same color *as objects I've seen before* (for more, see my forthcoming, Sec 7.3). That is, there's a feeling of recognition; I'm perceiving another instance of a familiar kind of object. As with the phenomenal character I just mentioned, this character is over and above the phenomenal characters of my perceptions of each region. I also think it's too coarse-grained for the perceptual atomist, because I'm appreciating a similarity to a relatively diverse group of objects; I'm just appreciating that the sphere is the same as objects with the same general shade of white, rather than objects with the more specific shade of white that I perceive as coating the ball's surface. Therefore, this phenomenal character doesn't help the atomist. The atomist wants to explain my perception of the sphere's more specific shade, and therefore needs to identify a phenomenal character that's unique to my perception of that shade.

It's also worth mentioning a significant shortcoming of Chalmers's suggestion: it wouldn't help explain color inaccuracy. Consider again Kitaoka's illusion. Atomists are pushed to say that your perceptions of both squares are accurate, because there doesn't seem to be any justification for believing that one perception, rather than the other, is accurate. Chalmers's suggestion doesn't help in this regard, because you perceive the left square as darker than the right square, and thus the phenomenal characters of your perceptions can't be the same in a respect like phenomenal-underlying-whiteness$_7$. Thus, even if there were such a thing as phenomenal-underlying-whiteness$_7$, atomists are still pushed to say that both of these perceptions are accurate, making it hard to explain color inaccuracy.

Chalmers's suggestion isn't the only option. Recall that we're trying to fill in the schema:

> Proposal 1b: You perceive the shaded region on the sphere as instantiating *causes* ____ *perceptions in me in this context.*
>
> Proposal 2b: You perceive the shaded region as instantiating whatever property satisfies the description: the property that disposes that region to cause ____ perceptions in me in this context.

I think that Chalmers's suggestion is the most promising way of dividing perceptions according to their *phenomenal characters*. But this isn't the only way of dividing perceptions. We can also divide them according to their

functional roles. Let's say your perception of a region on the left side of the sphere has functional kind *left-function* and your perception of a region on the right side of the sphere has functional kind *right-function*. Let's also suppose that there's a respect in which these functional kinds are the same, and call that respect *left&right-function*. Another suggestion is that you perceive a region as white whenever your perception has functional type left&right-function.

One problem with this suggestion is that, whereas we can use introspection to determine whether there's a respect in which the phenomenal characters of our perceptions are the same, we don't have the same kind of access to the functional roles of our perceptions. For this reason, supposing that there is a respect in which the functions are the same is an empirical conjecture, and thus a risky bet.

But, even setting that aside, this suggestion wouldn't explain color inaccuracy, and for the same reason that Chalmers's suggestion doesn't explain color inaccurcy. In particular, your perceptions of the squares in Kitaoka's illusion must not have functions that are the same in the relevant respect, because you perceive the left square as darker than the right square. Thus, even if there is such a thing as left&right-function, atomists are still pushed to say that both perceptions are accurate, and the most natural explanation is that your perception of a region on the left represents whatever color is shared by objects that cause perceptions with left-function in that context, and your perception of a region on the right represents whatever color is shared by objects that cause perceptions with right-function in that context. And we're back into familiar territory, because that quickly leads to the conclusion that all our perceptions are accurate.

Another possibility is that there's a functional state, *global-function*, that's over and above our perception of each region, and that suffices for your perceiving the entire sphere as white. But this won't help an atomist, and for the same reason that a general feeling of uniformity won't help the atomist. In particular, an atomist wants to explain color constancy by first explaining how we perceive the color of each individual region, and thus needs to identify a functional role that belongs to each of those perceptions.

So far, we considered dividing perceptions by their phenomenal characters and functional roles. There are other ways of dividing them — for example, by the underlying neural activity. But we've identified a general problem with appeals to other kinds of perception: If it's a kind shared by your perceptions of all the regions, it's unclear there is such a kind, and it won't help explain

color inaccuracy. And if it's not a kind shared by your perceptions of all the regions, it doesn't help the atomist.

5.2 Other Perceivers

According to both of the proposals we considered, you perceive the color of the shaded region because of a causal relation involving *you*. What other kind of perceiver might one appeal to? Once again, it will be helpful to restate this question more schematically. We're trying to fill in the schema:

> Proposal 1c: You perceive the shaded region as instantiating *causes phenomenal-dark-gray perceptions in ____ in this context*
>
> Proposal 2c: You perceive the shaded region as instantiating whatever property satisfies the description: the property that disposes that region to cause phenomenal-dark-gray perceptions in ____ in this context.

How should we fill in these schema?

One suggestion is with the phrase: *all normal perceivers*. But this is a non-starter, because the same objects cause perceptions with different phenomenal characters in different normal perceivers. In our example, the same stone tile causes perceptions with different phenomenal characters in Sarah and Jacob.

Another suggestion is: *at least one normal perceiver*. An advantage of this suggestion is that color inaccuracy would be possible. If you are not a normal perceiver, then you can misperceive the colors of objects. Perhaps when your perception has a phenomenal-dark-gray character, you perceive the object as causing phenomenal-dark-gray perceptions in normal perceivers. If you're not normal, then the relevant object might instead cause phenomenal-medium-gray perceptions in normal perceivers, in which case you're misperceiving its color.

One problem with this suggestion is that you won't be capable of misperceiving the color of an object if you're a normal perceiver. And that seems wrong. When looking at Kitaoka's illusion, normal perceivers will perceive the left square as darker, and thus seem to misperceive at least one square. Likewise for the other illusions.

Another problem with this suggestion is that it doesn't explain our perceptions of color relations. At bottom, the problem is that we'd perceive colors that are too coarse-grained. Consider again the two squares:

Let's suppose the phenomenal character of your perception of the left square is phenomenal-dark-gray, and the phenomenal character of your perception of the right square is phenomenal-medium-gray. As just noted, the same objects cause perceptions with different phenomenal characters in different perceivers. Thus, leaving some distracting qualifications in a footnote, the right square causes phenomenal-dark-gray perceptions in another normal perceiver.[12] According to the suggestion we're considering, you would perceive the left square as having a color that's shared by both squares. Moreover, the right square causes phenomenal-medium-gray perceptions in another normal perceiver. According to the suggestion we're considering, you would perceive the right square as having a color that's shared by both squares. Putting these conclusions together, you would perceive colors such that it's possible for both squares to be exactly the same color. For example, you might perceive determinable colors with some of the same determinants (i.e., overlapping determinables). In that case, you wouldn't perceive the left square as *darker*. When we perceive one object as darker than another, we're perceiving them as having incompatible colors, one darker than the other. Returning to our earlier analogy, suppose I tell you that my son Ian is between 5 and 10 years old, and that my son Isaac is between 3 and 6 years old. I haven't thereby told you that Ian is *older* than Isaac, because what I've said leaves open the possibility that they're the same age. To tell you that Ian is older, I have to specify non-overlapping, and thus incompatible, ranges for each child.

A third suggestion is: *a certain perceiver*, perhaps a certain normal perceiver, or perhaps an idealized perceiver. The problem is that this is inconsistent with perceptual relativism. If Sarah perceives the stone tile as causing phenomenal-dark-gray perceptions in that perceiver, and Jacob perceives the stone tile as causing phenomenal-medium-gray perceptions in that same per-

[12]Two qualifications: First, this might not follow if you're on one of the extremes of variation within normal perceivers. In that case, we'd just need to refocus the argument on someone closer to the mean. Second, this might not follow if the variation within normal perceivers isn't wide enough. In that case, we'd just need to refocus the argument on squares that are more similar.

ceiver, then at most one of their perceptions is accurate, because the stone tile can cause a perception with at most one of these phenomenal characters in the relevant perceiver. Perceptual relativism, however, implies that Sarah and Jacob can both accurately perceive the stone tile.

5.3 Other Contexts

According to both of the proposals we considered, you perceive the shaded region because of a causal relation in *that context*. What other kind of context might one appeal to? Once again, it will be helpful to restate this question more schematically. We're trying to fill in the following schema

> Proposal 1d: You perceive the shaded region as instantiating *causes phenomenal-dark-gray perceptions in you in* ___.
>
> Proposal 2d: You perceive the shaded region as instantiating whatever property satisfies the description: the property that disposes that region to cause phenomenal-dark-gray perceptions in you in ___.

How should we fill in these schema?

One suggestion is with the phrase: *all normal contexts*. But this is a non-starter, because the same object can cause perceptions with different phenomenal characters in different normal contexts. Kitaoka's illusion is just one example. Thus, nothing causes phenomenal-dark-gray perceptions in *all* normal contexts.

Another suggestion is: *at least one normal context*. But then it's not possible to misperceive an object's color in a normal context, because, if an object causes a phenomenal-dark-gray perception in a normal context, you'll accurately perceive that object as causing phenomenal-dark-gray perceptions in at least one normal context. However, the illusions we considered in Section 3 seem to involve misperception. Moreover, at least some of them seem to occur in normal contexts. Kitaoka's illusion is perhaps the most compelling example, because we normally view objects against white and black backgrounds. It's possible to generate similar illusions in other ways. For example, we can make one object look darker than another, even if they are intrinsically the same, by varying the lighting, distance, viewing angle, and objects previously seen at that location. And we normally view objects under dim lighting and bright lighting, from a few inches away and several feet

away, directly ahead of us and slightly to the side, and after viewing dark objects and bright objects.

This suggestion also has difficulty explaining our perceptions of color relations. As the illusions in Section 3 demonstrate, almost all objects along the white-black continuum cause phenomenal-dark-gray perceptions in at least one normal context. We just need to vary the background or lighting. With this in mind, consider the squares yet again:

The left square is causing a phenomenal-dark-gray perception. Assuming we make sufficient adjustments to the background and lighting, the right square will also cause a phenomenal-dark-gray perception. Thus, according to the suggestion we're considering, whenever you have a phenomenal-dark-gray perception, you're perceiving a color shared by both of these squares. Likewise, the right square is causing a phenomenal-medium-gray perception. Assuming we make sufficient adjustments to the background and lighting, the right square will also cause a phenomenal-medium-gray perception. Thus, according to the suggestion we're considering, whenever you have a phenomenal-medium-gray perception, you're perceiving a color shared by both of these squares. But that makes it hard for the atomist to explain why you perceive the left square as darker.[13]

A third suggestion is: *a certain context*, perhaps a certain normal context, or perhaps an idealized context. In that case, your phenomenal-dark-gray perception of an object is accurate only if that object would cause a phenomenal-dark-gray perception in you in the specified context. The problem is that this doesn't cohere with the motivations for relativism, at least given atomism. Relativism is motivated by the thought that, because we don't have any justification for believing that one *perceiver* rather than another *perceiver* is accurately perceiving the color of an object, we should say that both perceivers are accurately perceiving it (for more discussion, see my

[13]Chalmers says that it wouldn't be "too bad" if our phenomenal-medium-gray perceptions and phenomenal-dark-gray perceptions represented compatible properties (2006, p.91). But, for the reason mentioned above, I think it would be incompatible with atomism.

Morrison (manuscript)). Likewise, given atomism, we don't seem to have any justification for believing that one *context* rather than another *context* is the place where our color perceptions are accurate. It would thus be strange for an atomistic relativist to insist that whenever there's variation across contexts, we're nonetheless accurately perceiving an object's color in only one of those contexts. If ignorance is unacceptable with respect to perceivers, it is also unacceptable with respect to contexts.

A fourth suggestion is to specify a phenomenal character for *each* context. For example, if we develop this proposal along the lines of (2d), you perceive the shaded region as instantiating whatever property satisfies the description: causes a phenomenal-dark-gray perception in this context, causes a phenomenal-medium-gray perception in that other context, and so on, for each of the different combinations of illuminants, backgrounds, foregrounds, distances, and viewing angles. An obvious problem with this suggestion is that there are *so many* different combinations of these factors, and thus *so many* different contexts, that it's unclear what grounds all the specifications. In other words, it's unclear why you perceive the property that satisfies *this* description, rather than another description, such as a description that specifies a different phenomenal character for just one of the contexts. The proposals we considered in Section 4 don't give rise to the same problem, because they require the specification of only one phenomenal character, namely the phenomenal character of your current perception, and it's clear what grounds that specification. This fourth suggestion also makes it hard for you to accurately perceive an object's color, because the object's color would need to satisfy a description that specifies a phenomenal character for each context, and it's hard for the object's color to satisfy such an elaborate description. As noted in our discussion of Chalmers's suggestion, our predictions about how an object will look under other illuminants are systematically inaccurate. Thus, if what grounds the specifications is also what's responsible for these predictions, an object's color would rarely satisfy the relevant description, and thus we wouldn't perceive it. Both of these problems also arise if this suggestion is developed along the lines of (1d), so that the specifications are included in the color itself.

5.4 Other Relations

Finally, according to both of the proposals we considered, you perceive the shaded region because of a *causal* relation. What other kind of relation might

one appeal to?

One suggestion is: a *teleological* relation. In particular, perhaps phenomenal-dark-gray perceptions were naturally selected for their ability to indicate a certain reflectance, and therefore represent that reflectance. The problem is that natural selection is rarely so specific. Color perceptions as a whole were plausibly selected for their ability to help us discriminate and identify objects. But it's hard to believe that color perceptions of a specific kind, such as phenomenal-dark-gray perceptions, were selected for their ability to help us identify a specific kind of object, such as those with a specific reflectance. This is a point others have made (e.g., Tye 2006, 2007; Gert 2006, p.579), and I go into more detail in another paper (manuscript). For now, I just want to point out that natural selection is especially unhelpful for the *relativist*. Natural selection provides explanations that are species-wide, so that if phenomenal-dark-gray perceptions were naturally selected for their ability to indicate a certain reflectance, they would represent that reflectance in everyone, and if phenomenal-medium-gray perceptions were naturally selected for their ability to indicate a certain other reflectance, they would represent that other reflectance in everyone. According to the relativist, Sarah's phenomenal-dark-gray perception and Jacob's phenomenal-medium-gray perception are both accurate perceptions of the stone tile, and thus represent compatible properties. Because natural selection is species-wide, it would follow that *your* phenomenal-dark-gray perceptions and *your* phenomenal-medium-gray perceptions represent compatible properties. But that makes it hard to explain why, in our example involving two squares, you perceive the left square as darker than the right square.

Another suggestion is: a *phenomenal* relation to a color, i.e., a relation that obtains solely because of the phenomenal character of your perception (Campbell 1993, Kalderon 2011, and Pautz 2007; see also Chalmers 2006 on Edenic colors). This suggestion is hard to reconcile with relativism. First, Sarah's perception and Jacob's perception have different, incompatible phenomenal characters. If what they perceive is determined solely by the phenomenal characters of their perceptions, it presumably follows that they're perceiving different, incompatible colors. But then at most one of their perceptions can be accurate, contrary to relativism. Moreover, if they are perceiving compatible colors, despite the difference in their phenomenal characters (as Kalderon 2007 suggests), then any pair of perceptions with the same phenomenal characters would also be perceptions of compatible colors, including your perception of the squares from the last subsection.

But that makes it hard to explain why you perceive the left square as darker. Second, it would be hard to explain color constancy, such as when you perceive the sphere as the same color despite the phenomenal differences due to the shadow. If what we perceive is determined solely by the phenomenal characters of our perceptions, it presumably follows that you're perceiving a constant color only if there is a constant phenomenal character included in all your perceptions of the sphere's regions. And, as I argued earlier, I don't think there is any such phenomenal character.

We're now in a position to appreciate a general problem with appealing to other relations: Whereas causal relations can vary from person to person, these other relations are usually unvarying across the entire population. As a result, appealing to these other relations can make it hard for the relativist to explain how Sarah's phenomenal-dark-gray perception of the stone tile and Jacob's phenomenal-medium-gray perception can both be accurate, without thereby implying that your phenomenal-dark-gray perception of the left square and your phenomenal-medium-gray perception of the right square are insufficient for perceiving the left square as darker. If Sarah's and Jacob's perceptions are both accurate, they're perceiving compatible colors, and if your perceptions are accurate, you're perceiving incompatible colors.

Perceptual atomism is a big tent, and there are a staggering number of variants. It would be impossible to consider all of them. But I can't find a variant that's compatible with perceptual relativism and that avoids the problems just mentioned. For this reason, I think that a perceptual relativist should give up perceptual atomism and look for an alternative. In the final section, I'll sketch the alternative that I prefer.

6 Perceptual Structuralism

Let's return to Kitaoka's illusion. We asked: Which square, if either, are you accurately perceiving? Because you're perceiving one of the squares as darker than the other, I think the natural answer is that you're accurately perceiving at most one of the squares. This is the kind of answer we're inclined to give in other domains: if you perceive one person as taller than another, and they're really the same height, the natural conclusion is that you're accurately perceiving at most one of their heights. But an atomistic relativist is pushed to give another answer. From an atomistic point of view,

there's nothing that justifies the belief that you're inaccurately perceiving one square rather than the other, because none of the intrinsic features of your perceptions, and none of their relations to the squares, seem to indicate which perception is inaccurate. From a relativist point of view, this seems like a good reason to conclude that both perceptions are accurate. Thus, the atomistic relativist is pushed to say that you're accurately perceiving both squares, leading to conclusions like "all things apparent to men actually are," and "all things apparent to Sarah in normal contexts actually are." But that makes it hard to explain color inaccuracy and constancy. What's the alternative?

One alternative is that there is a *contradiction* between your perception of the individual colors of the squares and your perception of their relation. In particular, there is a color c such that you perceive the left square as c and the right square as c, but nonetheless perceive the left square as darker. This alternative would allow one to say that both perceptions are accurate without relativizing color perception to contexts. However, this would leave it mysterious why we rarely, if ever, notice contradictions between our perceptions of colors and our perceptions of color relations. It would also still be hard to explain color inaccuracy. If you accurately perceive the color of both squares in their respective contexts, despite the dramatic differences between those contexts, it is hard to resist concluding that you accurately perceive the color of that kind of square in all contexts, leaving us a short step away from the conclusion that you accurately perceive the color of every object in every context, even if you sometimes misperceive their relations. That's not a satisfying explanation of color inaccuracy.

According to the alternative I prefer ("perceptual structuralism"), what's basic to color perception are comparisons between objects, including that the left square looks darker than the right square. Your perceptions of the individual colors are explained by these comparisons. There's a helpful analogy with maps. Consider:

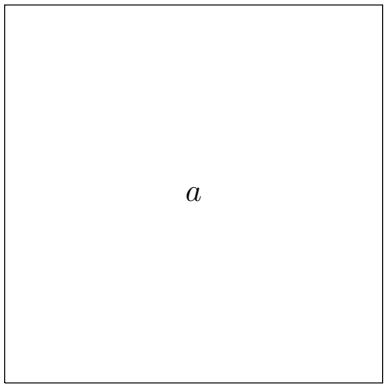

This map doesn't give us any information about a's location; it could be anywhere. Now consider:

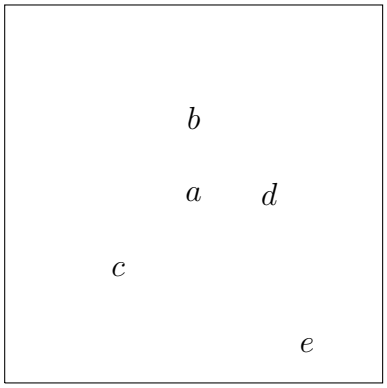

This map describes a's location by describing its relations to other objects. Similarly, a map of landmarks in New York City might describe Penn Station's location by describing its relations to Grand Central Station, the Empire State Building, Central Park, and the Statue of Liberty.

Suppose that a is in fact in the same location as b. In that case, the map is accurately describing at most one of their locations. But is it accurately describing a's location or b's location? It depends on what the map says about their relations to the other objects. If the map accurately describes a's relations to the other objects (namely: c, d, e), and it inaccurately describes b's relations to these other objects, then its description of b's location is inaccurate, while its description of a's location might be accurate. Similarly, if a map of landmarks in New York City misdescribes Penn Station as a mile south of Madison Square Garden, even though they are in fact in the

same location, we can ask whether it is inaccurately describing Penn Station's location or Madison Square Garden's location. If the map accurately describes all of Penn Station's relations to the other landmarks (the Empire State Building, Central Park, and the Statue of Liberty), while it inaccurately describes Madison Square Garden's relations to the other landmarks, then its description of Madison Square Garden's location is inaccurate, while its description of Penn Station's location might be accurate.

Structuralists say something similar about Kitaoka's illusion. According to structuralists, you are accurately perceiving the color of at least one of the squares, because of the relations you are perceiving to other objects. Which objects? Which relations? In other work (Morrison forthcoming), I argue that you are perceiving relations to other objects you're currently perceiving, objects you saw recently, and groups of objects you perceived in the more distant past. The details are of course important, but this isn't the right place to introduce them. For now, I just want to point out that structuralists can explain why you're inaccurately perceiving at least one of the squares, and thus aren't pushed to say that you're accurately perceiving both. More generally, they can give a straightforward account of color inaccuracy.

They can also give a straightforward explanation of color constancy. For the structuralist, sameness in color is one of the relations out of which the structure is built. Thus, unlike the atomist, they don't need to first explain your perceptions of individual colors, and then, on that basis, explain your perception of constancy. They can take it for granted that you perceive constancy; it's not something that needs to be explained in terms of a more basic kind of perception (for more details, again see Morrison forthcoming). This isn't to say that color constancy is *computationally* straightforward. Structuralism is a psychosemantic theory, not a computational theory. Analogously, the causal theory of reference provides a straightforward explanation of how 'Napoleon' manages to refer to Napoleon, without giving any explanation at all of the computational processes that underlie our acquisition and use of that term.

One final point in favor of structuralism: it coheres nicely with perceptual relativism. Once again, the analogy with maps is helpful. Consider:

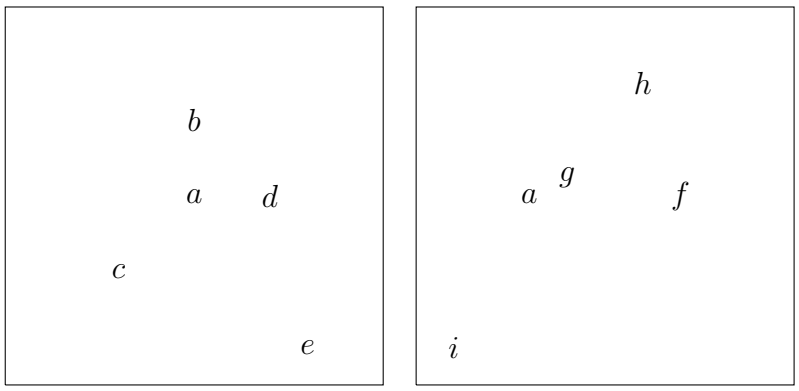

Each map describes a's relations to a different set of objects. Nonetheless, both maps might accurately describe a's location, by accurately describing its relations to those objects. Structuralist say something similar about Sarah's and Jacob's perceptions, namely that they might both be accurately perceiving the color of the stone tile, despite perceiving different relations to different objects. In particular, they might both be perceiving relations that describe the same color (or at least compatible colors). Structuralists can thereby explain why both Sarah and Jacob are accurately perceiving the stone tile, while also explaining why you're inaccurately perceiving at least one of the squares. Both are consequences of the fact that the relevant "structure" of comparisons is specific to each individual. What sets structuralists apart from atomists is that, whereas atomists claim that color perception is relative to a kind of perception (e.g., a phenomenal-dark-gray perception), structuralists claim that color perception is relative to a "structure" built out of relations between objects.

Of course, there's a lot more to say about perceptual structuralism. But, for present purposes, what's most important is that perceptual relativism doesn't lead inextricably to the conclusion that "all things apparent to men actually are," or to an equally problematic conclusion, such as "all things apparent to Sarah in normal contexts actually are." It leads to these conclusions only when combined with perceptual atomism. Alternative views about how we manage to perceive the colors of objects, such as perceptual structuralism, give us the resources to explain why you're misperceiving one of the squares, and thus to avoid relativizing color perception to contexts in a way that makes it difficult to explain color inaccuracy and color constancy. Perceptual structuralists are therefore well-advised to rethink their

basic assumptions about the nature of perception.[14]

References

Annas, Julia and Barnes, Jonathan, editors (2000). *Sextus Empiricus: Outlines of Scepticism*. Cambridge University Press.

Armstrong, D. M. (1987). Smart and the secondary qualities. In Pettit, Philip, Sylvan, Richard, and Norman, Jean, editors, *Metaphysics and Morality: Essays in Honour of J. J. C. Smart*, pages 1–15. Basil Blackwell.

Block, Ned (1999). Sexism, ageism, racism, and the nature of consciousness. *Philosophical Topics*, 26(1):39–70.

Burge, Tyler (2010). *Origins of Objectivity*. Oxford University Press.

Burnyeat, M. F. (2012). Conflicting appearances. In *Explorations in Ancient and Modern Philosophy*, pages 276–315. Cambridge University Press. Originally published 1979.

Byrne, Alex (2003). Color and similarity. *Philosophy and Phenomenological Research*, 66:641–665.

Byrne, Alex and Hilbert, David R. (1997). Colors and reflectances. In Byrne, Alex and Hilbert, David R., editors, *Readings on Color: The Philosophy of Color*, pages 263–288. The MIT Press.

Byrne, Alex and Hilbert, David R. (2003). Color realism and color science. *Behavioral and Brain Sciences*, 26:3–21.

Campbell, John (1993). A simple view of colour. In Haldane, John and Wright, Crispen, 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.

[14]This paper grew out of a section of my "Perceptual Variation and Structuralism." Thanks to all those who read that paper over the years. Thanks also to EJ Green and Adam Pautz for reading one of the final drafts of this paper.

Clark, Austen (2000). *A Theory of Sentience*. Oxford University Press.

Cohen, Jonathan (2004). Color properties and color ascriptions: A relationalist manifesto. *The Philosophical Review*, 113:451–506.

Cohen, Jonathan (2007). A relationalist's guide to error about color perception. *Noûs*, 41(2):335–353.

Cohen, Jonathan (2008). Colour constancy as counterfactual. *Australasian Journal of Philosophy*, 86(1):61–92.

Cohen, Jonathan (2009). *The Red and the Real: An Essay on Color Ontology*. Oxford University Press.

Cohen, Jonathan (2010). Color relationalism and color phenomenology. In Nanay, Bence, editor, *Perceiving the World*, pages 13–32. Oxford University Press.

Cohen, Jonathan (2015). Perceptual constancy. In Matthen, Mohan, editor, *The Oxford Handbook of Philosophy of Perception*, pages 621–639. Oxford University Press.

Dretske, Fred I. (1981). *Knowledge and the Flow of Information*. The MIT Press.

Egan, Andy (2010). Projectivism without error. In Nanay, Bence, editor, *Perceiving the World*, pages 68–96. Oxford University Press.

Gegenfurtner, Karl R. (2003). Cortical mechanisms of colour vision. *Nature Reviews Neuroscience*, 4:563–572.

Gert, Joshua (2006). A realistic colour realism. *Australasian Journal of Philosophy*, 84(4):565–589.

Gert, Joshua (2010). Color constancy, complexity, and counterfactual. *Noûs*, 44(4):669–690.

Gert, Joshua (2013). Color constancy and dispositionalism. *Philosophical Studies*, 162(2):183–200.

Gilchrist, Alan (2006). *Seeing Black and White*. Oxford University Press.

Hardin, C. L. (1997). Reinverting the spectrum. In Byrne, Alex and Hilbert, David R., editors, *Readings on Color: The Philosophy of Color*, volume 1, pages 289–301. Oxford University Press.

Hilbert, David R. (2005). Color constancy and the complexity of color. *Philosophical Topics*, 33(1):141–158.

Hilbert, David R. (2012). Constancy, content, and inference. In Hatfield, Gary and Allred, Sarah, editors, *Visual Experience: Sensation, Cognition, and Constancy*, pages 199–211. Oxford University Press.

Hilbert, David R. and Kalderon, Mark Eli (2000). Color and the inverted spectrum. In Davis, Steven, editor, *Color Perception: Philosophical, Psychological, Artistic, and Computational Perspectives*, volume 9 of *Vancouver Studies in Cognitive Science*, pages 187–214. Oxford University Press.

Jackman, Henry (2017). Meaning holism. In Zalta, Edward N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, spring 2017 edition.

Jackson, Frank (2000). Philosophizing about color. In Davis, Steven, editor, *Color Perception: Philosophical, Psychological, Artistic, and Computational Perspectives*, pages 152–162. Oxford University Press.

Jackson, Frank (2007). Colour for representationalists. *Erkenntnis*, 66(1):169–185.

Jackson, Frank and Pargetter, Robert (1987). An objectivist's guide to subjectivism about color. *Revue Internationale de Philosophie*, 41:127–141.

James, William (1950). *The Principles of Psychology*. Dover. First published in 1890.

Johnston, Mark (1992). How to speak of the colors. *Philosophical Studies*, 68:221–263.

Kalderon, Mark Eli (2007). Color pluralism. *The Philosophical Review*, 116:563–601.

Kalderon, Mark Eli (2008). Metamerism, constancy, and knowing which. *Mind*, 117(468):935–971.

Kalderon, Mark Eli (2011). The multiply qualitative. *Mind*, 120(478):239–262.

Kitaoka, Akiyoshi (2006). Luminance-gradient-dependent lightness illusion. `http://www.psy.ritsumei.ac.jp/~akitaoka/gilchrist2006mytalke.html#LGLI`.

Lewis, David (1997). Naming the colours. *Australasian Journal of Philosophy*, 75:325–342.

Matthen, Mohan (2009). Truly blue: An adverbial aspect of perceptual representation. *Analysis*, 69(1):48–54.

Matthen, Mohan (2010). Color experience: A semantic theory. In Cohen, Jonathan and Matthen, Mohan, editors, *Color Ontology and Color Science*, pages 67–90. The MIT 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.

Millikan, Ruth Garrett (1989). Biosemantics. *The Journal of Philosophy*, 86(6):281–297.

Morrison, John (forthcoming). Perceptual variation and structuralism. *Noûs*.

Morrison, John (manuscript). Perceptual variation and ignorance.

Palmer, Stephen E. (2002). Perceptual grouping: It's later than you think. *Current Directions in Psychological Science*, 11(3):101–106.

Pautz, Adam (2007). Intentionalism and perceptual presence. *Philosophical Perspectives*, 21:495–541.

Pautz, Adam (2010a). Review of Jonathan Cohen, The Red and the Real. *Notre Dame Philosophical Reviews*. `http://ndpr.nd.edu/news/the-red-and-the-real-an-essay-on-color-ontology/`.

Pautz, Adam (2010b). A simple view of consciousness. In Koons, Robert C. and Bealer, George, editors, *The Waning of Materialism*, pages 25–66. Oxford University Press.

Pautz, Adam (2011). Can disjunctivists explain our access to the sensible world? *Philosophical Issues*, 21:384–433.

Peacocke, Christopher (1983). *Sense and Content: Experience, Thought, and Their Relations*. Oxford University Press.

Peacocke, Christopher (1984). Colour concepts and colour experience. *Synthese*, 58:365–382.

Rupert, Robert D. (1999). The best test theory of extension: First principle(s). *Mind & Language*, 14(3):321–355.

Tye, Michael (2002). *Consciousness, Color, and Content*. The MIT Press.

Tye, Michael (2006). The puzzle of true blue. *Analysis*, 66:173–178.

Tye, Michael (2007). True blue redux. *Analysis*, 67:92–93.

Tye, Michael (2012). Cohen on color relationism. *Analytic Philosophy*, 53(3):297–305.

Vogt, Katja (2016). Ancient skepticism. In Zalta, Edward N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, winter 2016 edition.