# Appendix for "Perceptual Variation and Relativism"

John Morrison

## 1 Introduction

In "Perceptual Variation and Relativism" (forthcoming b) I criticized two proposals for combining perceptual relativism with perceptual atomism. Both proposals are atomistic in that they explain, in atomistic fashion, how we manage to perceive the colors of objects. Both proposals are relativistic in that they imply, in accordance with relativism, that two people can perceive an object's color even if their perceptions have different phenomenal characters. For concreteness, I focused on a particular perception, namely your perception of a shaded region on a certain sphere's surface (see Figure 1). Both proposals agree that you perceive that region's color because of a causal relation between the shaded region and the phenomenal character of your perception, namely phenomenal-dark-grayness. But they disagree about which properties of the shaded region are its colors, and thus about which property you're perceiving.

> Proposal 1a: While the shaded region causes a phenomenal-dark-gray perception in you, you perceive the shaded region as instantiating *causes phenomenal-dark-gray perceptions in me in this context*.
>
> Proposal 2a: While the shaded region causes a phenomenal-dark-gray perception in you, 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.

In "Perceptual Variation and Relativism" I argued that these proposals make it difficult, if not impossible, to explain color inaccuracy and color constancy. It's therefore natural to wonder about the alternatives. In this appendix I'll consider the most plausible alternatives. They differ from the proposals above in that they appeal to other phenomenal characters (Sect 2), other properties of perceptions (Sect 3), other perceivers (Sect 4), other contexts (Sect 5), or other relations (Sect 6). I will argue that either these alternatives don't help explain color inaccuracy and color constancy, or they have other,

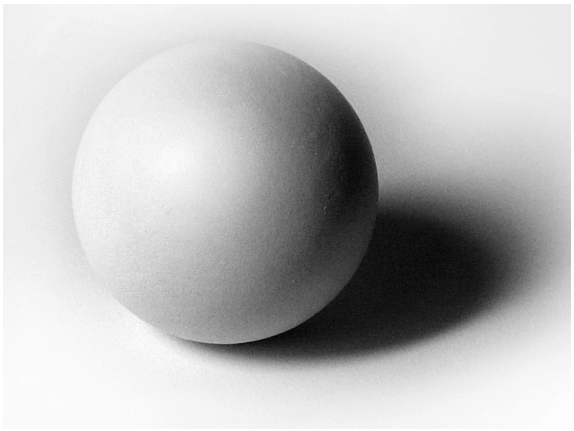

Figure 1: Color constancy

more serious drawbacks. I'll conclude that perceptual relativism shouldn't be combined with perceptual atomism.

## 2 Other Phenomenal Characters

Your perceptions of different regions on the sphere's surface have different phenomenal characters, such as phenomenal-dark-gray and phenomenal-medium-gray. Chalmers (2006, p.86–91) suggests that there's a respect in which all of these phenomenal characters are the same. In particular, he suggests that, because you perceive each region as the same shade of white (let's call it 'white$_7$'), all of these phenomenal characters involve the phenomenal character associated with perceiving that shade of white (let's call it 'phenomenal-underlying-whiteness$_7$'). More generally, Chalmers suggests that, if you perceive an object as a certain color, your perception has a certain phenomenal character, independent of how you perceive that object's illumination (see also Hilbert 2005, p.151).

What is the relation between phenomenal-underlying-whiteness$_7$ and the other phenomenal characters of your perceptions, such as phenomenal-dark-grayness and phenomenal-medium-grayness? 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 phenome-

nal characters. Yet another possibility is that phenomenal-underlying-whiteness$_7$ is an *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.

Chalmers suggests that phenomenal-underlying-whiteness$_7$, rather than phenomenal-dark-gray, is the relevant phenomenal character.[1] This yields two alternative proposals:

> Proposal 1b: While the shaded region causes a phenomenal-underlying-white$_7$ perception in you, you perceive the shaded region as instantiating *causes phenomenal-underlying-white$_7$ perceptions in me in this context*.

> Proposal 2b: While the shaded region causes a phenomenal-underlying-white$_7$ perception in you, you perceive the shaded region as instantiating whatever property satisfies the description: the property that disposes the shaded region to cause phenomenal-underlying-white$_7$ perceptions in me in this context.

Both proposals seem to explain color constancy. Given (1b), you might perceive the sphere's color as constant because you perceive all of the regions on the sphere's surface as instantiating *causes phenomenal-underlying-whiteness$_7$ perceptions in me in this context*. Given (2b), you might perceive the sphere's color as constant because your perceptions all have the same phenomenal character, namely phenomenal-underlying-whiteness$_7$, rather than because of which colors you're perceiving. If it seems strange to explain color constancy in this second way, it might help to note that proponents of proposals like (2b), such as McLaughlin (2003, p.111–116), offer parallel explanations of color similarity. They claim that you perceive objects as similar in color because your perceptions have similar phenomenal characters, rather than because you're perceiving similar colors.

Given that these proposals seem to explain color constancy, and that Chalmers's suggestion is essential to them, it's a suggestion worth taking seriously. 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,

[1] Chalmers never endorses this suggestion. He merely says that he is *inclined* toward it (2006, p.87).

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.[2]

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:

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 (Figure 1), there are regions such that my perceptions of those regions have the same local phenomenal characters as my perceptions of these squares. That is, there is a region on the sphere's surface that seems to have exactly the same phenomenal character as my perception of the square, and likewise for the

[2]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).

right square. 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.[3]

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, and thus that there isn't a phenomenal character, such as phenomenal-underlying-whiteness$_7$, that's constant across changes in 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.676–678). 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.

There's another, related problem with Chalmers's suggestion. In an experiment that's standard among color researchers, subjects view surfaces under two different illuminants, and are asked whether they perceive the same

[3]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"?

color (Figure 2). 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 as the same color as the surface on the right. 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-white$_7$. Suppose we then gradually decrease the illumination of the right surface. As the illumination decreases, you will probably perceive the right surface as remaining the same color, despite the change in its illumination, just as you ordinarily perceive surfaces as remaining the same color despite changes in their illumination, such as when you carry them outdoors. Thus, if Chalmers's suggestion were true, your perception of the right surface would still include phenomenal-underlying-whiteness$_7$, despite the change in illumination. Your perception of the left surface, meanwhile, would presumably also still include phenomenal-underlying-whiteness$_7$, given that its illumination hasn't changed. Summarizing: your perceptions of both surfaces are initially phenomenal-underlying-white$_7$, and your perceptions of both surfaces would still be phenomenal-underlying-white$_7$ after we decrease the illumination of the right surface until it equals the illumination of the left surface. Recall that, according to Chalmers's suggestion, if your perceptions of two surfaces are phenomenal-underlying-white$_7$, 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 the right surface decreases enough to equal the illumination of the left surface, you will still perceive the two surfaces as the same color. But that's not what happens: you will perceive the left surface as darker than the right surface. 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 thus as the same color as surfaces that would look lighter under equal illumination (see Gilchrist 2006, Ch 10).

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. Otherwise, there will be counterexamples like the one I just described, given that we ordinarily perceive surfaces as remaining the same color despite changes in their illumination.

I conclude that color constancy isn't explained by a local phenomenal

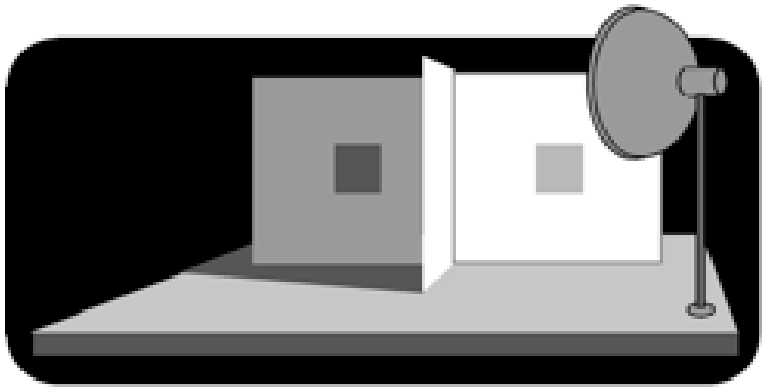

Figure 2: Light and shadow matching experiment, from Gilchrist 2006, p.289

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.

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 a, 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

Figure 3: Kitaoka's illusion

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 another shortcoming of Chalmers's suggestion: it wouldn't help explain color inaccuracy. Consider again Kitaoka's illusion (Figure 3). 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$_7$. Perhaps your perception of the left square is phenomenal-underlying-gray$_5$ and your perception of the right square is phenomenal-underlying-gray$_{15}$. 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.

# 3 Other Properties

So far, we focused on ways 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 other properties, such as their functional roles and neural realizers, for example.

Let's start with 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*. Perhaps these are the relevant properties of your perceptions. This would yield two proposals:

> Proposal 1c: While the shaded region causes a left-function perception in you, you perceive the shaded region as instantiating *causes left-function perceptions in me in this context*
>
> Proposal 2c: While the shaded region causes a left-function perception in you, you perceive the shaded region as instantiating whatever property satisfies the description: the property that disposes that region to cause left-function perceptions in me in this context.

These proposals make it difficult to explain color constancy and color inaccuracy, and for exactly the same reasons as (1a) and (2a). For example, given (1c), you would perceive each region as a different color. Likewise, given (2c), there doesn't seem to be any common element that might explain the sense in which you perceive the sphere as a constant color. While the arguments about (1a) and (2a) from Section 4 of "Perceptual Variation and Relativism" are about phenomenal characters, they apply equally to functional roles. Thus, we should reject (1c) and (2c) for the same reasons we rejected (1a) and (2a).

Inspired by Chalmers's suggestion, one might try to identify a functional role that's shared by all of our perceptions of the sphere's surface. Let's call it *left&right-function*. That would yield:

> Proposal 1d: While the shaded region causes a left&right-function perception in you, you perceive the shaded region as instantiating *causes left&right-function perceptions in me in this context*
>
> Proposal 2d: While the shaded region causes a left&right-function perception in you, you perceive the shaded region as instantiating whatever property satisfies the description: the property that disposes that region to cause left&right-function perceptions in me in this context.

One problem with these proposals 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. For example, it might turn out that the only kind of functional role shared by these perceptions is also shared by our perceptions of green, in which case we would perceive white and green objects as the same color, at least according to these proposals.

But, even setting this problem aside, these proposals wouldn't explain color inaccuracy, and for the same reason that Chalmers's suggestion doesn't explain color inaccuracy. 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 their underlying neural activity. But we've identified a general problem with appeals to other kinds of perception: If it's a kind specific to each perception (e.g., phenomenal-dark-gray, left-function), it's hard to explain color constancy and color inaccuracy. If it's a kind shared by your perceptions of all the regions (e.g., phenomenal-underlying-whiteness$_7$, left&right-function), 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 (e.g.,

Gestalt phenomenal characters, global-function), it doesn't help the atomist.

# 4 Other Perceivers

According to all of the proposals we considered, you perceive the color of the shaded region because of a causal relation involving *you*. What other perceiver might one appeal to? It might be helpful to restate this question more schematically. We're looking for an alternative way of filling in the schema:

> Proposal 1e: While the shaded region causes a phenomenal-dark-gray perception in you, you perceive the shaded region as instantiating *causes phenomenal-dark-gray perceptions in ____ in this context*
>
> Proposal 2e: While the shaded region causes a phenomenal-dark-gray perception in you, 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.

What are our options?

One option is: *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 initial example from "Perceptual Variation and Relativism," the same stone tile causes perceptions with different phenomenal characters in Sarah and Jacob. Thus, nothing has the property *causes phenomenal-dark-gray perceptions in all normal perceivers in this context*, and nothing satisfies the description: the property that disposes that region to cause phenomenal-dark-gray perceptions in all normal perceivers in this context. Thus, nobody's perception of the stone tile would be accurate. But, according to perceptual relativism, Sarah's and Jacob's perceptions might both be accurate.

Another option is: *at least one normal perceiver*. An advantage of this option is that color inaccuracy would be possible. If you are not a normal perceiver, then you can misperceive the colors of objects. Suppose that when your perception has a phenomenal-dark-gray character, you perceive the object as instantiating *causes phenomenal-dark-gray perceptions in normal perceivers in this context*. 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.

But that is not yet a fully satisfying explanation of color inaccuracy, because 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 (Figure 3), normal perceivers will perceive the left square as darker, and thus seem to misperceive at least one square. Likewise for the other illusions discussed in Section 4 of "Perceptual Variation and Relativism."

Another problem with this option 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.[4] As a result, the left square causes a phenomenal-dark-gray perception in at least one normal perceiver, and the right square causes a phenomenal-dark-gray perception in at least one normal perceiver. According to the option we're considering, phenomenal-dark-gray perceptions thus represent a color shared by the left and right square. For the same reason, phenomenal-medium-gray perceptions would represent a color shared by the left and right square. Thus, according to the option we're considering, your perceptions of the squares would leave open the possibility that they're the same color. More precisely, you would perceive colors such that it's possible for both squares to be exactly the same color. For example, given (2e), you might perceive

[4]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.

determinable colors with some of the same determinants (i.e., overlapping determinables). This is a problem, because you perceive the left square as *darker*, and when we perceive one object as darker than another, we're perceiving them as having incompatible colors, one darker than the other. Thus, if your perceptions of the squares leaves open the possibility that they're the same color, you aren't perceiving one as darker than the other. Returning to our analogy from "Perceptual Variation and Relativism," 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 option 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 perceiver, 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 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 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 1f: While the shaded region causes a phenomenal-dark-gray perception in you, you perceive the shaded region as instantiating *causes phenomenal-dark-gray perceptions in you in* ___.

> Proposal 2f: While the shaded region causes a phenomenal-dark-gray perception in you, 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 option is: *all normal contexts*. But this is a non-starter, because the same object can cause perceptions with different phenomenal characters in different normal contexts, for example under direct sunlight and under a shadow. Kitaoka's illusion is just one example. Thus, nothing causes phenomenal-dark-gray perceptions in *all* normal contexts.

Another option 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 you in a normal context, you'll accurately perceive that object as causing phenomenal-dark-gray perceptions in you in at least one normal context. However, the illusions we considered in Section 4 of "Perceptual Variation and Relativism" seem to involve misperception. Moreover, at least some of these illusions 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 option also has difficulty explaining our perceptions of color relations. As the illusions in Section 4 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 option we're considering, whenever you have a phenomenal-dark-gray perception, you're perceiving a color shared by both of these squares. For the same reason, whenever you have a phenomenal-medium-gray perception,

you're perceiving a color shared by both of these squares. But, for the reasons discussed in the last section, that makes it hard for the atomist to explain why you perceive the left square as darker.[5]

A third option 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 manuscript a). 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 option is to specify a phenomenal character for *each* context. For example, if we develop this proposal along the lines of (2f), 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 option 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

[5]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.

that specification.

This fourth option 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 (1f), so that the specifications are included in the color itself.

# 6 Other Relations

Finally, according to both of the proposals we considered, you perceive the shaded region because of a *causal* relation. What other 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 a). 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 the general problem with appealing to other relations mentioned in "Perceptual Variation and Relativism": 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.

# 7 Conclusion

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. I suggest perceptual structuralism (see my forthcoming a).

# References

Campbell, John (1993). A simple view of colour. In Haldane, John and Wright, Crispin, editors, *Reality, Representation, and Projection*, pages 257–268. Oxford University Press.

Chalmers, David J. (2006). Perception and the fall from Eden. In Gendler, Tamar Szabó and Hawthorne, John, editors, *Perceptual Experience*, pages 49–125. Oxford University Press.

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

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.

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

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

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

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

McLaughlin, Brian (2003). Colour, consciousness, and colour consciousness. In Smith, Quentin and Jokic, Aleksandar, editors, *Consciousness: New Philosophical Perspectives*, pages 97–154. Oxford University Press.

Morrison, John (forthcoming a). Perceptual variation and structuralism. *Noûs*. doi:10.1111/nous.12245.

Morrison, John (forthcoming b). Perceptual variation and relativism. In Vogt, Katja and Vlasits, Justin, editors, *Epistemology After Sextus Empiricus*. Oxford University Press.

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

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

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

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