

Anti-Atomism about Color Representation

JOHN MORRISON
 Barnard College, Columbia University

Abstract

According to anti-atomism, we represent color properties (e.g., *red*) in virtue of representing color relations (e.g., *redder than*). I motivate anti-atomism with a puzzle involving a series of pairwise indistinguishable chips. I then develop two versions of anti-atomism.

1. Introduction

Hold this piece of paper at arm's length.



Do these squares look exactly the same color (assuming gray is a color)? Intuitively: yes. Are you able to know this just by reflecting on your experiences — that is, introspectively? Intuitively: yes.

There is a puzzle that puts pressure on these and other core intuitions about how things look. According to the response I prefer, the fact that these squares look exactly the same color is more fundamental than the fact that each looks a particular shade of gray. In other words, our representing the squares as exactly the same color is more fundamental than our atomistically representing each as a particular shade of gray. The puzzle thereby motivates a new, anti-atomistic view of color representation. In the conclusion, I'll briefly consider whether anti-atomism extends to our representations of other kinds of properties, such as sounds and lengths.

To make the puzzle more concrete, I'm going to tell a story. My story invokes the notion of pairwise indiscriminability (hereafter just: indiscriminability). Nothing important will hinge on how this notion is understood. I just mean something like “can't tell them apart.”

Poor Jonah! He wanted to select exactly the right shade of paint for his new kitchen cabinets. Unfortunately, at the paint store he discovered that he was not up

to the task. Here's what happened: A few minutes after entering the store he was drawn to a charcoal-gray paint chip. He put it in his pocket and kept looking. A few minutes later he was drawn to another. And a few minutes later, yet another. He then pulled out the first two chips, placed them against a solid white background, and judged:

(1a) The first and second chips have indiscriminable colors in this context.

On this basis he inferred:

(1b) The first and second chips look exactly the same color in this context.

After returning those chips to his pocket he then pulled out the second and third chips, placed them against the same white background, and judged:

(2a) The second and third chips have indiscriminable colors in this context.

On this basis he inferred:

(2b) The second and third chips look exactly the same color in this context.

Concluding that all three chips must be coated with the same paint, he prepared to make his purchase and leave. But, to his great surprise, he then compared the first and third chips against the same white background and judged:

(3a) The first and third chips have discriminable colors in this context.

On this basis he inferred:

(3b) The first and third chips do not look exactly the same color in this context.

Just as two objects viewed simultaneously can look exactly the same color, so can objects viewed at different times, and so can one object viewed at different times. Jonah viewed the first chip twice, once when it was next to the second chip and once when it was next to the third chip. These two contexts were extremely similar: same lighting, same background, same viewing distance, and so on. Knowing this, Jonah inferred that the first chip looked exactly the same color both times he viewed it—its coloring looked constant. He likewise inferred that the color of the second and third chips looked constant. It will be helpful to express these inferences in a slightly different way. Let's say that a group of contexts is "stable" for an object if the object looks exactly the same color whenever viewed in those contexts. We can express Jonah's inference:

(4b) These three contexts were stable for these three chips.

He then recalled that as a schoolboy he was taught:

TRANSITIVITY

For all objects in stable contexts: if x and y look exactly the same color and y and z look exactly the same color, then x and z look exactly the same color.

But this principle is mutually inconsistent with (1b)–(4b). TRANSITIVITY, (1b), (2b), and (4b) together imply that the first and third chips look exactly the same color. But (3b) says that the first and third chips do *not* look exactly the same color. Confused by the inconsistency, Jonah went home empty-handed.

Jonah's predicament is familiar. It is also easy to replicate in a laboratory. Our job is to figure out which of the inconsistent claims is false and, if possible, explain how Jonah can know it is false. For now, let's not worry about whether Jonah's reasoning is representative of how people normally reason.

I think we should reject TRANSITIVITY.¹ An immediate problem with this response is that we have to give up a plausible way of thinking about color representation. Just as it is plausible that two objects weigh exactly the same in virtue of how much each object weighs, and just as it is plausible that two objects are exactly the same color in virtue of how each object is colored, it is plausible that two objects *look* exactly the same color in virtue of how each object *looks* to be colored. According to this way of thinking, what's fundamental is how each object looks, and these atomistic facts in turn fix whether two objects look exactly the same color. This way of thinking has an important consequence. As a first pass, we might express it:

Two objects look exactly the same color if and only if one object looks a color and the other object also looks that color.

This formulation is inadequate. Two objects might both look gray even though they do not look exactly the same color because one looks charcoal and the other looks ash. Here's a better formulation:²

Two objects look exactly the same color if and only if, for all colors, one object looks that color if and only if the other object looks that color.

By quantifying over all colors, including more exact/fine-grained/determinate colors like charcoal and ash, this formulation doesn't imply that two chips look exactly the same color even if one looks charcoal and the other looks ash.

TRANSITIVITY logically follows from this biconditional.^{3,4} An example involving numbers might help show why. Suppose you assign each of your shoes a number. If, as a result, a flipflop is the same number as a pump, then you must have assigned each of them the same number. Say: 2. Moreover, if the pump is the same number as a sneaker, then you must have assigned the same number to the sneaker—namely, 2. Therefore, the flipflop and the sneaker must be the same number. As this example makes clear, if our representations are atomistic, then our representations of sameness must be transitive.

If we want to reject TRANSITIVITY, we need to develop an alternative to this plausible, atomistic way of thinking about color representation. According to my preferred alternative, our experiences represent properties like *is charcoal-gray* in

virtue of representing relations like *is exactly the same color as* and *is not exactly the same color as*. For example, our experience represents a chip as charcoal-gray in virtue of representing it as exactly the same color as some objects, whiter than other objects, blacker than still other objects, and so on. Moreover, our experience fails to represent any determinate color when it represents mutually inconsistent relations, as when it represents: x is exactly the same color as y , y is exactly the same color as z , and x is not exactly the same color as z . Because this experience fails to represent any determinate color, there is no atomistic difference in how it represents x and z . Consequently, x looks some color if and only if z looks that color. Nonetheless, x and y look different. This undermines the right-to-left direction of the biconditional, because the right-hand condition is satisfied (x looks some color if and only if z looks that color) while the left-hand condition isn't satisfied (x and y look different). We are therefore free to reject TRANSITIVITY. I'm going to call this the **anti-atomistic response**.

In the next section I will motivate the anti-atomistic response by listing three other responses and their drawbacks. I will spend the rest of the paper introducing and developing anti-atomism.

But let's first clarify our puzzle. To start, it is not a traditional sorites paradox. Traditional sorites paradoxes depend on intuitions like: a paint chip is still red after it becomes incrementally yellower. They also depend on intuitions like: a paint chip still looks red after it is made to look incrementally yellower. More generally, traditional sorites paradoxes depend on the intuition that for higher-level predicates like 'is red' and 'looks red', if a change is small enough, then a thing satisfies the predicate at both times or satisfies the predicate at neither time. Our puzzle does not depend on a parallel intuition. Let's focus on TRANSITIVITY, the premise that seems most sorites-like. First, TRANSITIVITY is about cases when there seems to be no difference rather than about cases when there seems to be a small difference. As a result, one cannot respond to our puzzle by straightforwardly adapting responses to the sorites paradox that exploit small differences, such as degree-theoretic responses.⁵ Second, two paint chips can look exactly the same color even if we're uncomfortable categorizing them as looking red or as not looking red. As a result, one cannot respond to our puzzle by straightforwardly adapting supervenient or psychological responses to the sorites paradox, because these responses exploit our indifference or unwillingness to use higher-level predicates like 'is red' and 'looks red' to categorize things in one way or another.⁶

This isn't to deny that there are similarities between our puzzle and a sorites paradox. It's just to point out that the differences are great enough that we should approach our puzzle with fresh eyes.

It will also be helpful to clarify the relevant sense of 'looks'. Though they disagree about the details, many philosophers claim there are three or more different senses of 'looks'.⁷ We don't need to enter into that debate. All that's important is that 'looks' picks out a relation that satisfies the following four conditions: First, it relates us to paint chips and other external objects as opposed to inner, mental objects like experiences or sense-data. Second, it is representational in some sense—it corresponds to how things *seem* to us. For this reason, as another way of saying

that the chip looks charcoal-gray to Jonah, I will sometimes say that Jonah visually represents the chip as charcoal-gray. Third, it corresponds to a conscious mental state rather than unconscious information processing in that, for every change in how something looks, there is a corresponding change in our conscious mental state. Fourth, it is independent of our beliefs in that, for example, a wall can look gray even if we believe it is white.

We can accommodate different senses of ‘looks’ as long as they satisfy these conditions. Suppose you’re looking at a white wall that’s partially covered by a shadow. Arguably, there’s a sense in which the covered part of the wall looks darker, and there’s a sense in which the entire wall looks to be the same shade of white. Rather than choose between these two senses of ‘looks’, we can leave the discussion sufficiently general that it applies to both of them.

I also want to say something about how I’m using ‘look exactly the same color’. I’m using this phrase with a familiar and implicit restriction. We might think of the colors as forming a pyramid. Near the top of the pyramid are colors like red, yellow, and blue. Below them are colors like magenta, rose, canary, lemon, aqua, and navy. At the bottom are the most exact/fine-grained/determinate colors (hereafter just: determinate colors). There are some theories of color according to which the colors at the bottom are so determinate that no human being could ever detect a difference between sufficiently similar shades. Nonetheless, when we say that two things look exactly the same color we’re usually talking about the lowest level at which color differences are represented, and that’s how I’m using the expression. That is, I’m using ‘look exactly the same color’ so that it is restricted to the color differences that we can represent.

What is the lowest level at which color differences are represented? Just as there might be no moment at which the letters on an eye chart suddenly become invisible as it moves farther and farther away, so there might be no level in the color pyramid at which color differences suddenly become too fine-grained to be represented. My arguments will be compatible with this kind of fuzziness, though I will often talk as if it doesn’t exist.

Finally, a point about ‘phenomenal sorites’: A number of different puzzles are lumped together under this label. Russell and Goodman’s puzzle about sense-data is sometimes described as a phenomenal sorites.⁸ So are puzzles about the nature of demonstratives (‘that shade’) and high-level observational predicates (‘looks red’).⁹ Williamson’s anti-luminosity argument, which depends on a controversial safety condition on knowledge, is sometimes included in this group as well.¹⁰ Calling our puzzle a phenomenal sorites would falsely suggest that it turns on similar issues. It would also falsely suggest that our puzzle includes a traditional sorites premise, and, depending on one’s view of phenomenology, it might even falsely suggest that our puzzle is about phenomenology, because some people identify phenomenology with so-called qualitative feels and our puzzle is about how external objects look rather than about how any of our mental states “feel” to us. For all of these reasons, let’s not call our puzzle a phenomenal sorites.

There are a handful of papers that discuss puzzles that approximate our puzzle. But they often get entangled in distracting and inessential debates, use clunky and

unhelpful technical machinery, or fail to acknowledge the full range of responses. A subsidiary goal of this paper is to suggest a new framework.¹¹

2. Three Other Responses

In this section I will discuss three of the most straightforward responses to the puzzle: the roughness response, the contextualist response, and the inaccessibility response. As we'll see, each of these responses seems to undermine a core intuition about how things look (where anything that captures our natural, default perspective counts as an intuition). I'll later explain why the anti-atomistic response seems to preserve these core intuitions. While that's not a decisive reason to prefer the anti-atomistic response, it does make it more attractive.

The **roughness response** is that (1b) and (2b) are false and Jonah is in a position to know they are false because objects never look *exactly* the same color. Instead, they can only look *roughly* the same color. According to this response, Jonah shouldn't have inferred:

- (1b) The first and second chips look exactly the same color in this context.
- (2b) The second and third chips look exactly the same color in this context.

Instead, he should have inferred:

- (1c) The first and second chips look *roughly* the same color in this context.
- (2c) The second and third chips look *roughly* the same color in this context.

These beliefs don't generate a puzzle, because 'look roughly the same color' is not a transitive relation.

There are many ways to develop the roughness response.¹² But regardless of how it is developed, the roughness response abandons one of our core intuitions about how things look: that objects (and parts of objects) sometimes look exactly the same color. If you look at a uniformly colored wall under the right conditions, then, intuitively, different parts of the wall will look exactly the same color. Likewise, if you stare at a uniformly colored wall and nothing changes, then, intuitively, it will look exactly the same color from moment to moment. We know what it is like for two objects to look roughly the same color, and, intuitively, these aren't cases like that.

It might help to keep in mind how we're using 'looks exactly the same color'. If we were using it to talk about the bottom level of the color hierarchy, where the colors are so fine-grained that no human could ever detect a difference between them, it would be implausible to insist our visual experiences take a stand on whether things look exactly the same color. But that's not how we're using the expression. We're using it to talk about colors at the lowest level at which color differences are represented, which is why the roughness response is counterintuitive.

Proponents of the roughness response might argue that once we appreciate that color experiences are a kind of measurement and that all measurements are imprecise, we should give up this intuition. They might insist our experiences always represent a range of possible colors. However, this line of reasoning depends on a confusion. While it is true that measurements are imprecise in the sense that small errors are likely, it doesn't follow that measurements are imprecise in the sense that they represent a range of possibilities. For example, when an erratic speedometer says that you're moving at 60mph, it might be representing you as moving at exactly that speed even though, given what you know, you should have very little confidence that you're moving at 60mph rather than 55mph.¹³

Other proponents of the roughness response might argue that this isn't an intuition we should take too seriously. At the end of this section I'll explain why I think they're wrong. Regardless, it should be clear that, all other things being equal, we should prefer a response that preserves the intuition that things sometimes look exactly the same color.

The **contextualist response** is that (4b) is false. As everyone knows, the same object can look gray or brown in different contexts. The same chip can look gray when viewed in a room with normal lighting and look brown when viewed in a room with intense yellow lighting. Jonah, however, viewed the chips under the same lighting, against the same background, from the same distance, and so on. Knowing this, he inferred that each chip looked exactly the same color each time it was viewed. That is, he inferred:

(4b) These three contexts were stable for these three chips.

According to the contextualist response, that was his mistake.¹⁴

This response is more extreme than it might first appear. We can generate the puzzle even if Jonah is in a laboratory that allows us to make sure that the contexts aren't just similar but nearly identical. In addition to making sure that the background is exactly the same and the lighting is exactly the same, we can also make sure that Jonah's pupil size is exactly the same and there are no after-images. In fact, we can make sure that the only difference is that Jonah views the same chip alongside another chip. Accordingly, a shift in context can help explain why the first chip looked different only if an extremely small change in the coloring of a nearby chip is enough for that object to look different. Moreover, that extremely small change must *always* be enough for the first chip to look different, because otherwise the puzzle is simply relocated to one of the contexts where a small change in the coloring of a nearby chip isn't enough for that object to look different. And this goes far beyond the well-known fact that the same object can look gray or brown depending on the lighting.

Contextualism threatens to commit us to the view that any change in the visual scene, no matter how small, changes how every object looks. After all, there doesn't seem to be anything special about Jonah's situation. If a slight change in the coloring of a nearby chip is enough, then presumably so are the normal fluctuations of fluorescent lamps and the normal involuntary movements of the human eye.

Moreover, if a small change is enough, then presumably larger changes are also sufficient. The contextualist apparently must think that a brightly colored object moving rapidly through the visual scene changes how every object looks. The contextualist will therefore have a hard time maintaining that in ordinary contexts, where small changes and relatively large changes are ubiquitous, objects nonetheless look the same from moment to moment.

For this reason, the contextualist is pressured to give up another of our core intuitions about how things look: that objects (and parts of objects) often look exactly the same color from moment to moment despite small and large changes in their surroundings. That is, she must deny that there is color constancy. This leaves the contextualist in a hard spot. Ideally, she'd like to give a principled reason for insisting that small changes in Jonah's context are always enough for things to look different while denying that small and large changes in ordinary contexts are always enough for things to look different. But the prospects for finding such a reason seem dim; it seems improbable that a slight change in the color of a neighboring object would matter but other small changes, like the fluctuations in fluorescent lamps, and larger changes, like a brightly colored object moving through the scene, wouldn't matter.

Proponents of the contextualist response might try to undermine this intuition about constancy by pointing out that a chip can switch from looking gray_{11} to looking gray_{12} without our noticing the change.¹⁵ But there's usually a special explanation for why we fail to notice such changes, like inattention due to perceptual overload or selective focus, and none of these factors seems to explain why we would fail to notice that all objects change how they look from moment to moment, especially when, like Jonah, we're carefully examining objects against a solid white background. All other things being equal, it seems preferable to find a response that preserves the intuition that things often look exactly the same color from moment to moment in ordinary contexts.

Before turning to the next response, I want to separate the contextualist response from an increasingly popular view about the metaphysics of color. Some philosophers claim that the colors themselves are relative to contexts.¹⁶ According to these philosophers, nothing is gray_{11} . Instead, things are $\text{gray}_{11}\text{-in-}C$, $\text{gray}_{11}\text{-in-}C^*$, or $\text{gray}_{12}\text{-in-}C^*$. These philosophers might claim that the first chip looked $\text{gray}_{11}\text{-in-}C$ when compared to the second chip and then looked $\text{gray}_{12}\text{-in-}C^*$ when compared to the third chip, because the context changed. This might seem to commit them to the contextualist response. But it doesn't. Following Cohen [7], they might insist that things can look to have dispositional properties that aren't currently manifested, and therefore the first chip can still look $\text{gray}_{11}\text{-in-}C$ even after the context switched to C^* . The first chip would then look both $\text{gray}_{11}\text{-in-}C$ and $\text{gray}_{12}\text{-in-}C^*$, and, one might argue, that's all it takes for it to look exactly the same color in both contexts.¹⁷ Proponents of this view about the metaphysics of color can therefore reject the contextualist response and select any of the other responses we'll consider. For example, they can deny that things ever look exactly the same color, perhaps because things always look to have a range of possible colors, and thereby endorse the roughness response.¹⁸

The **inaccessibility response** is that the conjunction of (1b) and (2b) is false but that Jonah can't know which chips, if any, look exactly the same color. Recall:

- (1b) The first and second chips look exactly the same color.
- (2b) The second and third chips look exactly the same color.
- (3b) The first and third chips do not look exactly the same color.

According to the inaccessibility response, Jonah can know by introspection that (3b) is true and then infer that (1b)&(2b) is false. He can know that (3b) is true by introspection because if he figures out that two things have discriminable colors, then he can infer they do not look exactly the same color. He can then infer that (1b)&(2b) is false because TRANSITIVITY implies that (1b)–(3b) aren't all true. However, he cannot know by introspection whether (1b) is false or whether (2b) is false. More specifically, he can't tell whether $\neg(1b)\&(2b)$, $(1b)\&\neg(2b)$ or $\neg(1b)\&\neg(2b)$.

According to this response, here's how Jonah should proceed when trying to figure out whether two chips look exactly the same color: First, he should figure out whether they have discriminable colors. If they do, he should believe that they do not look exactly the same color. If they don't, he shouldn't believe anything, because, as we saw with (1b) and (2b), his evidence is consistent with both the possibility that they look exactly the same color and the possibility that they do not look exactly the same color. For this reason the inaccessibility theorist claims that Jonah can never know that two chips look exactly the same color—such facts are inaccessible by introspection.^{19,20}

By design, the inaccessibility response abandons another of our core intuitions about how things look: that we can know whether two objects look exactly the same color by introspection. In support of this intuition, consider that this is something that we seem to be able to know without asking other people, consulting scientific instruments, or relying on *a priori* principles. In contrast, it is far less intuitive that we can know by introspection whether objects *really are* exactly the same color, because, for all we know, our experiences might not be veridical. Nonetheless, at least under ordinary circumstances, we do seem able to tell by introspection that things *look* exactly the same color.

The intuition that we can know whether two objects look exactly the same color by introspection goes hand-in-hand with the first two intuitions, because those other intuitions seem to originate from our introspective access to our own experiences. Consider: How else might we know that objects sometimes look exactly the same color? Therefore giving up this third intuition seems to undermine the first two intuitions.

In further support of this intuition, remember that looks-talk is supposed to characterize how things seem to us, and there's something peculiar about cases in which x seems ϕ to us but we are unable to access this fact. One wonders: In what sense does x still seem ϕ to us? Perhaps the inaccessibility theorist can convince us that there's nothing peculiar about such cases.²¹ But, all other things being equal, it

still seems preferable to avoid postulating inaccessible facts about how things seem to us.

This is also a problem for the contextualist response. The contextualist claims that we can't know by introspection that a chip has switched from looking gray₁₁ to looking gray₁₂. Therefore, yet again, there are introspectively inaccessible facts about how things look.

I just argued that the three most straightforward responses abandon core intuitions about how things look: that objects sometimes look exactly the same color, that objects often look exactly the same color from moment to moment, and that we can know by introspection whether two objects look exactly the same color. There are some philosophers who won't take these intuitions seriously. These are the same philosophers who don't take seriously any of our other intuitions about the mind, including our intuitions about beliefs, desires, attention, and consciousness. I don't have anything original to say to them. It just seems to me that we should try as hard as possible to reconcile our natural, default conception of the mind with our best scientific theories, and this paper is written for those with a similar outlook.

Of course, some intuitions should be taken more seriously than others. How seriously should we take these three core intuitions? Very seriously. As a helpful contrast, consider the intuition, elicited by Ned Block in [2], that no matter how arranged or instructed, the citizens of China never suffice for a conscious being above and beyond the individual citizens. Some treat this intuition as an important problem for functionalism about the mind. No matter how seriously one takes this intuition, I think that one should take the three intuitions I listed above much more seriously. The key difference is that in the China example we have no direct access to whether there is or is not a superadded conscious being, whereas we apparently have a kind of direct access to our own minds that allows us to figure out how things *seem* to us.

There might be ways of improving these three responses so that they don't abandon core intuitions. Moreover, the fact that these responses force us to give up a core intuition isn't a decisive mark against them—it is only one consideration among many. Regardless, I think we should search for an alternative.²²

3. An Introduction to Anti-Atomism

Perhaps the problem isn't with (1b)–(3b), Jonah's pairwise comparisons between the chips, or with (4b), Jonah's assumption that the contexts are stable. Perhaps the problem is with TRANSITIVITY, the general principle that's responsible for their mutual inconsistency. Why not give up TRANSITIVITY?

As we noted earlier, there is a challenge for anyone who gives up TRANSITIVITY. TRANSITIVITY is entailed by:

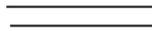
Two objects look exactly the same color if and only if, for all colors, one object looks that color if and only if the other object looks that color.

As we also noted earlier, this biconditional reflects a plausible way of thinking about color representation. The challenge is to develop an alternative.

Here's an alternative that I'm not going to pursue: that representing color relations is *independent* of representing monadic color properties (hereafter just: color properties). According to this view, the biconditional is false, because whether we represent two objects as exactly the same color is independent of how we represent each object. Jonah might represent the first chip as gray_{11} and the second chip as gray_{12} while representing them as exactly the same color. I'm not attracted to this view. As I said before, all other things being equal, I don't think we should postulate inaccessible representational differences, but, given that Jonah is unaware of any inconsistency in how he's representing the first and second chips, that's exactly what we'd be doing if we said that the first chip looked gray_{11} and that the second chip looked gray_{12} but that they nonetheless looked exactly the same color.

I'm instead going to focus on the possibility that we represent color properties *in virtue of* representing color relations—i.e., that representing color relations is *more fundamental than* representing color properties. In this section I'll introduce this anti-atomistic way of thinking, sketch two ways of developing it, and describe the corresponding senses of fundamentality. In the next section I'll show how it can be used to undermine TRANSITIVITY without abandoning any of our core intuitions. I'll end by considering three of its potential costs.

As a helpful model, let's consider two ways of thinking about our experiences of lengths, one atomistic and the other anti-atomistic. Consider the following lines:



The top line looks longer than the bottom line. It is natural to think that's because, for some lengths l_1 and l_2 , the top line looks l_1 , the bottom line looks l_2 , and l_1 is longer than l_2 . This is an atomistic way of thinking about our length experiences. To get a feel for an anti-atomistic alternative, consider the famous Müller-Lyer illusion:



The top line looks like it is longer than the bottom line. As before, it is natural to think that's because, for some lengths l_1 and l_2 , the top line looks l_1 , the bottom line looks l_2 , and l_1 is longer than l_2 .²³ However, given that in fact they are exactly the same length, this way of thinking leads to awkward questions like: Which line's length (if any) are you correctly representing? There might be a systematic way of saying exactly how long each line looks; perhaps it is a function of the angle formed by the arrow.²⁴ There might also be an unknown fact of the matter. But one could avoid questions like this altogether by claiming that the top line's looking longer than the bottom line is more fundamental than how each line looks. In that case, the awkward questions never arise, because, in virtue of representing one line as

longer than the other, you are misrepresenting the lengths of both lines. This is an anti-atomistic way of thinking about length experiences.

Returning to color representation, consider the following squares:



The left square looks darker than the right square. It is natural to think that's because, for some colors c_1 and c_2 , the left square looks c_1 , the right square looks c_2 , and c_1 is darker than c_2 . This is an atomistic way of thinking about our color experiences. To get a feel for an anti-atomistic alternative, consider Akiyoshi Kitaoka's [18] luminance-gradient-dependent lightness illusion:



The left square (hereafter: LS) looks darker than the right square (hereafter: RS). As before, it is natural to think that's because, for some colors c_1 and c_2 , LS looks c_1 , RS looks c_2 , and c_1 is darker than c_2 . However, given that in fact they are exactly the same color, that leads to awkward questions like: Which square's color (if any) are you correctly representing?

One might try to find a systematic way of saying which color is correctly perceived. But that project seems misguided, because it seems arbitrary to single out one of these viewing contexts as revealing the chip's true color. This is an instance of a general phenomenon. As Hardin [14, 67–76] and others have effectively argued, it seems arbitrary to say that an object's color is correctly perceived under certain conditions but not others: under natural sunlight rather than incandescent light, against a white background rather than against a gray background, from a distance of one meter rather than a distance of two meters, in a shadow rather than in direct light, and so on, even though an object can look different when viewed in these different contexts.²⁵ There might nonetheless be an unknown fact of the matter, and some are willing to bet that there is.²⁶ But I think it would be preferable to avoid the question.

The anti-atomistic alternative is that LS's looking darker than RS is more fundamental than facts about how each square looks. This avoids the question,

because in virtue of representing one as darker than the other, you are misrepresenting both.

Anti-atomism might initially seem strange. As a way of further softening this initial resistance, I think it's helpful to consider why we should expect there to be creatures with anti-atomistic systems of representation. While color perception has many objectives, three stand out, especially from an evolutionary perspective. First, color perception helps us pick out objects from their surroundings—it helps us see a puddle of water in a field of grass. Second, it helps us identify objects—it helps us recognize male parakeets by their plumage. Third, it helps us track objects as they move through their environment—it helps us follow a poisonous spider as it climbs up a tree. An anti-atomistic system of representation is capable of achieving all of these objectives, because it can represent the contrast between the color of the water and the color of the grass, the similarity between the coloring of this plumage and the colorings of plumages we've seen in the past, and the constancy of the spider's colors as it moves upward. We should therefore expect there to be creatures that have evolved anti-atomistic systems of representation, and we should be open the possibility that we're among them.²⁷

I'm going to focus on two ways of developing anti-atomism: **moderate anti-atomism** and **complete anti-atomism**.²⁸ Moderate and complete anti-atomists agree that we represent color properties in virtue of representing color relations. They disagree about which properties we represent. The moderate anti-atomist says we represent properties like *is charcoal-gray*, while the complete anti-atomist says we represent properties like *is exactly the same color as the second chip*. One difference between these properties is that *is exactly the same color as the second chip* can be instantiated only in worlds where the second chip exists. I'll say more about this disagreement later. Moderate and complete anti-atomists also disagree about the relevant sense of fundamentality. The moderate anti-atomist says it's roughly the sense in which 'Jonah's eyes are the color of the Mediterranean Sea' represents Jonah's eyes as turquoise in virtue of representing their relation to the Mediterranean Sea, and therefore more fundamentally represents their relation. The complete anti-atomist says it's roughly the sense in which the sentence 'Jonah's eyes are the color of the Mediterranean Sea' might be used to construct a sentence representing a property rather than a relation, like 'Jonah's eyes have the property of being the color of the Mediterranean Sea'. As a result of both disagreements, moderate and complete anti-atomists give different accounts of our representation of Kitaoka's lightness illusion and also generate different responses to our puzzle. I'll say more about their disagreements after I fill in some details.

3.1 Moderate Anti-Atomism

As a way of introducing moderate anti-atomism, let's again use spatial representation as a model. We might describe the location of the Empire State Building by describing its relations to other landmarks. We might say that it's on the same avenue as the New York Public Library and the same street as Penn Station. We'd then be referring to its location in virtue of describing its spatial relations; our

representing the Empire State Building's spatial relations would be more fundamental than our representing its location. According to the moderate anti-atomist, we similarly represent an object's color by representing its relations to other objects. We might represent a chip as charcoal-gray by representing it as exactly the same color as certain objects, darker than certain other objects, and lighter than still more objects. We'd then be representing its color in virtue of representing its color relations; the fact that we represent its color relations would be more fundamental than the fact that we represent its color.

There's an important complication. Just as a description of spatial relations can fail to pick out an exact/fine-grained/determinate location for the the Empire State Building, an experience of color relations can fail to pick out an exact/fine-grained/determinate color for a given object (keep in mind that I'm using 'determinate' to indicate a kind of property rather than a kind of representing or a kind of identity). What should the moderate anti-atomist say about such cases? With respect to the Empire State Building's location, a natural thing to say is that our description leaves open a number of different scenarios. In one scenario it's at the southwest corner of the relevant intersection, and in another scenario it's at the northeast corner. Despite these differences, there will be larger locations that it occupies in all the scenarios. Plausibly, our description picks out the smallest location that it occupies in all the scenarios.

Likewise, our experience of LS as darker than RS leaves open a scenario in which LS is gray₁₁ and a scenario in which LS is gray₁₄. Despite these differences, there will be a more determinable color that LS exemplifies in all the scenarios. Plausibly, we represent LS as the most determinate color that it exemplifies in all the relevant scenarios—we represent what the scenarios have in common. In this way the partial anti-atomist can use our experiences of color relations to explain how we represent more determinable color properties even if our experiences of color relations fail to pick out completely determinate color properties.²⁹

When we represent an object as a relatively determinate color (e.g., *charcoal-gray*) we also seem to represent more determinable colors (e.g., *dark gray*, *gray*). The moderate anti-atomist needs to explain how that's possible. I think she should appeal to more determinable color relations. In particular, I think she should say that when we represent a relatively determinate color relation between two objects (e.g., *exactly the same color*, *much darker than*) we thereby also represent more determinable color relations (e.g., *about the same color*, *darker than*). She should also say that we represent more determinate colors (e.g., *charcoal-gray*) in virtue of representing more determinate color relations and we represent more determinable colors (e.g., *dark gray*, *gray*) in virtue of representing more determinable color relations. Let's return to Kitaoka's lightness illusion. We represent determinate color relations between the squares like:

LS is much darker than RS, LS is exactly the same color as this, RS is much darker than . . .

This representation is consistent with some set of scenarios. The moderate anti-atomist should say that we represent LS as a relatively determinate color (e.g.,

charcoal-gray), because LS is that color in all the relevant scenarios. She should also say that we represent LS as a more determinable color (e.g., *dark gray*), because, now only taking into account our representations of more determinable color relations, we represent:

LS is darker than RS, LS is about the same color as this, RS is darker than . . .

This representation is consistent with even more scenarios, including scenarios where LS isn't charcoal-gray, and therefore we represent LS as a more determinable color (e.g., *dark gray*). In virtue of our representations of still more determinable color relations, we might represent LS as a still more determinable color (e.g., *gray*). In this way, moderate anti-atomism can be developed so that we represent both more determinate colors and more determinable colors. This will be a central part of the moderate anti-atomistic response to our puzzle.

As I said before, one of the reasons why anti-atomism is attractive is that we don't have to say whether we correctly perceive LS or RS. Instead, we misperceive both by misrepresenting LS as darker than RS. Further, we might misperceive both chips without misperceiving either of their individual colors. There might be a scenario in which LS has the correct shade of gray and another scenario in which RS has the correct shade. In that case, we wouldn't misperceive either square's individual color, because we would represent each as instantiating a determinable that includes its actual shade. But we'd still misperceive them by representing LS as darker than RS.³⁰

I'll fill in more details after I introduce another approach.

3.2 Complete Anti-Atomism

The complete anti-atomist claims that we represent only color properties like *is much darker than the second chip*. Such properties are standardly called impure, to distinguish them from properties like *is charcoal-gray*, which are standardly called pure. Here's how I think the view should be developed. Consider again how we represent Kitaoka's lightness illusion:

LS is much darker than RS, LS is exactly the same color as this, RS is much darker than . . .

This experience produces another experience, or experiential part, that represents LS as instantiating:

fine(x): x is much darker than RS, x is exactly the same color as this while RS is much darker than . . .

This second experience is constructed from the first experience in roughly the way the sentence 'Jonah's eyes have the property of being the color of the Mediterranean Sea' might be constructed from the sentence 'Jonah's eyes are the color of the Mediterranean Sea'. There's room for disagreement about how these kinds of experiences are constructed. One view is that it involves an automatic and unconscious computation or inference, perhaps triggered by attending to LS. Another view is that these experiences are so close that an experience that represents

the relevant color relations is non-computationally or non-inferentially sufficient for an experience that represents *is fine*, just as believing *that someone is kissing Jessica* is arguably sufficient for believing *that Jessica is being kissed* even though the first belief is about a relation (*is kissing*) and the second belief is about a property (*is being kissed*). This latter approach might not require the production of a new experience or experiential part; it depends on how one individuates experiences *qua* vehicles of representation. Let's set this entire disagreement aside.

As before, let's instead focus on our representations of more determinable colors, because it will be an important component of the complete anti-atomistic response to our puzzle. The complete anti-atomist should say that, in virtue of representing more determinable relations, we also represent LS as instantiating:

coarse(x): x is darker than RS, x is about the same color as this while RS is darker than . . .

This is a more determinable color because anything that instantiates *is fine* also instantiates *is coarse* but something can instantiate *is coarse* without instantiating *is fine*. Therefore, if complete anti-atomism is developed in the way I'm suggesting, then we represent both more determinate colors and more determinable colors. As I said, this will be a central part of the complete anti-atomistic response to our puzzle.

The complete anti-atomist can straightforwardly explain why we misperceive both LS and RS. She'll say that we misperceive LS because LS and RS are actually the same color and therefore LS doesn't actually instantiate *is fine*. For parallel reasons, she'll say that we misperceive RS. Thus, our misperception of both chips is symmetrical.

3.3 *Agreements and Disagreements*

While the moderate and complete anti-atomists agree that representing color relations is more fundamental than representing color properties, they disagree about which properties we represent. While objects instantiate both kinds of properties in the actual world, there are modal differences. Just as nothing instantiates the property *is the mother of Jonah* in worlds where Jonah doesn't exist, the complete anti-atomist says that nothing instantiates *is fine* in a world where RS doesn't exist. The moderate anti-atomist, however, says that we represent properties like *is charcoal-gray* that are instantiated in worlds where RS doesn't exist. While this is an important disagreement, I do not think it puts either view at a disadvantage. First, our experiences don't seem to tell us anything about the properties of objects in other possible worlds, and therefore our experiences don't seem to tell us that objects instantiate properties that are also instantiated in other possible worlds. Second, moderate and complete anti-atomism are specifically views about *visual representation*. They are not views about mental or linguistic representation. Complete anti-atomism is therefore consistent with the view that predicates like 'is red' and their correlates in thought refer to pure color properties like *is charcoal-gray*. We might have acquired the ability to linguistically and mentally represent these kinds of properties in any number of ways. People might have started using

‘is charcoal-gray’ to describe the monadic surface properties of objects that look roughly the same color as established paradigms. There might even be a chromatic analog to the standard meterstick. Alternately, people might have started demonstratively using ‘is charcoal-gray’ to pick out the property that causes objects to look grayer than most other objects but less gray than slabs of silver. As these proposals hopefully make clear, there are many ways in which we could refer to and think about pure color properties even if we visually represent only impure properties, like *is fine* and *is coarse*.

That being said, both moderate and complete anti-atomism have distinct advantages and disadvantages. Moderate anti-atomism better approximates our pre-theoretic view of color perception, because we still represent pure properties like *is charcoal-gray*. Complete anti-atomism is more elegant, and it gives us a more straightforward analysis of Kitaoka’s lightness illusion. Weighing their advantages and disadvantages is an important project, but in this paper I’m going to remain neutral. My goal is just to convince you that anti-atomism of some form gives us a promising solution to the puzzle.

The remainder of this section will fill in details that moderate and complete anti-atomists can agree about. According to both moderate and complete anti-atomists, we represent color properties in virtue of representing color relations. Which relations? Which objects? I’ll address these questions in turn.

Which relations? Relations like *is exactly the same color as*, *is redder than* and *is darker than*. There’s an interesting question about whether we also represent relations involving mixed colors like *is purpler than* or relations specifying exact degrees of difference like *is redder-to-degree- x than*. There’s also an interesting question about whether we can represent absolute maximums and minimums like *nothing is redder than*. Let’s set these questions aside for another occasion.

Which objects? The relevant objects presumably include nearby objects in the visual scene. For example, in the case of Kitaoka’s lightness illusion, we presumably represent the graded background. In addition, I think that we represent relations to remembered objects. If the suggestion that we represent relations to remembered objects seems strange, consider that it isn’t any stranger than the suggestion that, in experiences of recognition and in experiences of motion, one represents a relation between something present and something remembered. It might also help to note that we can represent an object as redder than a remembered object without representing anything else about that remembered object, such as its shape or when we saw it. One possibility is that, after several exposures to objects with the same color, we acquire the ability to represent them as a group without representing them individually, perhaps in the same way that superimposing a number of images of human faces might create an image that represents all the faces even if we can no longer tell them apart because they’ve been blurred together.

The anti-atomist owes us an account of what it takes for the object to be remembered in the relevant sense. On the one hand, it isn’t enough for there to be some information about the object even if one has trouble accessing it. On the other hand, one needn’t be consciously visualizing the object. The anti-atomist owes us an

intermediate sense of remembering, perhaps one that's close to what psychologists call "working memory." I'm optimistic about the anti-atomist's ability to fill in this detail.

One advantage of including remembered objects is that it allows us to distinguish intuitively different experiences. Suppose a person momentarily has a homogenous color experience produced by a constant green light and another person momentarily has a homogenous color experience produced by a constant red light. If they could represent relations only among objects in their visual scene, then it is unclear how we could differentiate their experiences. Each would just represent everything in the visual field as exactly the same color as everything else in the visual field. But if they also represent relations involving remembered objects, then we can easily differentiate their experiences.

Including remembered objects also has an interesting and plausible consequence. Plausibly, when we represent relations to remembered objects, we represent only more coarse-grained relations, like *is about the same color as*.³¹ Anti-atomism would then imply that when we consider an object by itself, we represent only more coarse-grained color relations and therefore represent only more coarse-grained colors. For example, when you consider a red paint chip by itself, it will look to have a more coarse-grained color than when you consider it next to chips with similar but different colors. This might explain why simultaneous comparisons are an important part of paint selection. When we simultaneously compare a paint chip to similarly colored chips, we see it as having a more exact color.

There are still a lot of details to fill in. Nonetheless, let's now apply anti-atomism to our puzzle.

4. The Anti-Atomistic Response

There are many ways to develop anti-atomism. If we wanted, we could develop it in a way that preserves TRANSITIVITY. But we can also develop it in ways that allow us to reject TRANSITIVITY. To reject TRANSITIVITY, we must reject:

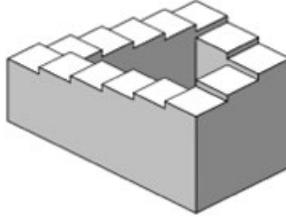
Two objects look exactly the same color if and only if, for all colors, one object looks that color if and only if the other object looks that color.

In this section I'll develop anti-atomism in ways that allow us to reject the right-to-left direction of this biconditional, freeing us to reject TRANSITIVITY. In the next section I'll consider objections.

Jonah's chips seem to bear inconsistent relations: the first chip looks exactly the same color as the second, the second chip looks exactly the same color as the third, but the first chip does not look exactly the same color as the third. Therefore, the chips look to bear relations that it would be impossible for three things to actually have. Because the anti-atomist says that Jonah represents colors in virtue of representing color relations, she might deny that the chips look any determinate color, which, as we'll see, undermines the right-to-left direction of the conditional, because the right-hand condition is satisfied (the first chip looks some color if and

only if the third chip looks that color) while the left-hand condition isn't satisfied (the first and third chips do not look exactly the same color).

As a helpful model, let's consider the parallel view of spatial experiences, which is the view that how we represent color relations is most fundamental to spatial experience. A proponent of this view needs to say something about L. Penrose and R. Penrose's [26] impossible staircase:



The stairs at the four corners seem to bear inconsistent relations: the far left stair looks like it is above the rear stair, the rear stair looks like it is above the far right stair, the far right stair looks like it is above the nearest stair, and the nearest stair looks like it is above the far left stair. Therefore, the stairs look to bear relations that it would be impossible for four things to actually have. If we represent locations in virtue of representing relations like *is above*, then, due to this inconsistency, we might deny that any of these stairs looks to be at a determinate location. Importantly, this wouldn't mean that we don't represent them at *any* location. For example, we consistently represent all of them as above the bottom of the figure. Likewise, some of the stairs look like they are at locations to the right or left of other stairs. Thus, they all look to be at locations. Nonetheless, they don't look to be at determinate locations, because we don't represent specific heights.

The anti-atomist might say the same thing happens in color perception. Because Jonah represents the chips in a mutually inconsistent way, the anti-atomist might deny that Jonah represents them as any determinate color. Instead, he represents only more determinable colors, like colors that are darker than the white background. In that case, even though, for all colors, the first chip looks that color if and only if the third chip looks that color, they still do not look exactly the same color, undermining the right-to-left direction of the biconditional.

Moderate and complete anti-atomists will give slightly different explanations. Moderate anti-atomists will say that, because Jonah represents determinate relations that are inconsistent, he represents relations that aren't compatible with any scenarios, and therefore he doesn't represent any determinate colors. Analogously, the description 'the mother of Jonah who is not a mother' does not describe anyone. But, the moderate anti-atomist will insist, Jonah also represents more determinable colors, because he also represents:

The first chip is roughly the same color as the second chip, the second chip is roughly the same color as the third chip, the third chip is roughly the same color as the first chip, all three chips are darker than the background. . .

This is a consistent representation and is therefore compatible with several scenarios. Thus, despite his failure to represent any determinate colors, Jonah still represents determinable colors.

The complete anti-atomist will say that, because Jonah represents determinate relations that are inconsistent, he doesn't represent a determinate color. I think that this reflects a natural, plausible principle about when our experiences represent impure properties. The underlying thought is that experiences of impure properties are *constructed* out of experiences of relations. Because Jonah's experience represents inconsistent relations, experiences constructed out of it are unable to represent anything.³² But Jonah still represents a more determinable color. Because he represents determinable color relations that are consistent (see above), he represents:

coarse2(x): x is about the same color as the second chip, x is about the same color as the third chip, x is darker than the background while the other chips are also darker than the background. . .

Thus, as before, despite his failure to represent determinate colors, Jonah still represents more determinable colors.

At this point, there's a natural question. Suppose we extend Jonah's series so that it begins with a pure white chip and ends with a pitch-black chip. Does anti-atomism imply that Jonah doesn't represent *any* color? No. It is psychologically impossible for someone to keep in mind so many relations, especially given that it isn't enough to merely store some information about each chip. Because we are able to keep in mind only a more limited number of relations, any inconsistencies will be local, like the four corner stairs in Penroses' impossible staircase, in which case we will represent determinable colors, because there is something that looks darker, redder, greener, etc., than all the other the colors we're considering.

Anti-atomism, whether moderate or complete, gives us a plausible way of thinking about color representation that undermines the biconditional that entails TRANSITIVITY. It therefore gives us a way of denying TRANSITIVITY. But does it preserve all three core intuitions? Yes!

The first intuition is that things sometimes look exactly the same color. According to the anti-atomist, these representations, along with our representations of other color relations, are at the foundation of our color experiences, and therefore the anti-atomist has no trouble preserving this intuition.

The second intuition is that things (and parts of things) sometimes look exactly the same color over time despite changes in context. The contextualist has trouble preserving this intuition, because she is pressured to claim that a thing doesn't look exactly the same color if we change the context. For example, if we merely change the coloring of a nearby chip, the second chip goes from looking charcoal-gray₁₁ to looking charcoal-gray₁₂ and therefore does not look exactly the same color as before. The anti-atomist isn't pressured to accept that conclusion. She can insist that the second chips looks exactly the same as before because we represent the second chip as bearing a certain relation to itself at an earlier time—that it is exactly the same color—despite changes in the context.

There is a point of agreement between the contextualist and anti-atomist. They both agree that whether an object looks exactly the same color over time depends on how we represent its relations to certain objects. The crux of the disagreement is about which objects. The contextualist is pressured to say that it is every other object in the visual scene, which is why any perceived change in the context results in a change in how the first chip looks. The anti-atomist, in contrast, says that the only relation that matters is its relation to itself at an earlier time. For example, the complete anti-atomist might claim we represent the first chip as instantiating:

constant(x): x is exactly the same color as itself three seconds ago, x is darker than the background. . .

Thus, the anti-atomist can insist that an object looks exactly the same color from moment to moment even if it alternates between looking lighter and darker than the object next to it. For this reason, the anti-atomist can accept (4b), the claim that Jonah's contexts were stable for all three chips.³³

The third intuition is that we can know by introspection whether two things look exactly the same color. The inaccessibility response abandoned this intuition. Yet again, the anti-atomist doesn't seem to have any trouble accommodating it. Consider Jonah's initial judgment:

(1a) The first and second chips have indiscriminable colors.

From this he inferred:

(1b) The first and second chips look exactly the same color.

The anti-atomist can say that this inference is cogent but unnatural because the natural order of inference is from (1b) to (1a). In particular, according to anti-atomism, we can judge (1b) directly from our experience, because that's at the foundation of how we represent the world. We would judge (1a) only subsequently, presumably inferring it from (1b). Further, because this kind of representation is foundational, it is plausible that we have immediate and reliable access to it, allowing us to know when we're representing things in this way. Thus we should be optimistic that the anti-atomist can explain our ability to know how things look.

I conclude that the anti-atomistic response preserves all three core intuitions.

There's a lot more to say about anti-atomism. Nonetheless, I think we've said enough to establish that it is an attractive way to preserve our intuitions about how things look and, as such, deserves our attention. Another attractive feature of anti-atomism is that it sidesteps the embarrassing question of whether we're correctly perceiving LS or RS in Kitaoka's lightness illusion. It also gives us a plausible explanation for why puzzles like this arise in the first place: when our visual system receives similar spectral information from two objects, it is biased towards representing them as exactly the same color. This bias would improve object recognition and object tracking, because minute color differences, real or

apparent, wouldn't be as distracting; it would be harder to track a beetle scampering across a leaf if the beetle and its background seemed to irregularly shift their colors. Jonah-type situations might be a harmless side effect, at least from a practical point of view.

5. Potential Costs of Anti-Atomism

In this section I will consider three potential costs of anti-atomism.

First, anti-atomism potentially implies that certain experiences don't represent colors. Consider a newborn, opening her eyes for the first time and looking directly at the uniform, white hospital ceiling. Anti-atomism potentially implies that her experience isn't representing the ceiling as white, because her experience is at most representing the relation *is exactly the same color as*. After she turns towards the nurse's blue scrubs, and then turns back again, the ceiling might look white, because she'll then be aware of the contrast. Some might consider this a cost of anti-atomism. But to the extent that we can speculate about what it's like to be a newborn, this strikes me as a plausible guess. You can't see a specific spatial location without appreciating its relations to other locations, paradigmatically the location of your body. Perhaps the newborn can't see a specific color without appreciating its relations to other colored objects. The newborn might just experience the ceiling as coated with a baffling and alien quality until she looks at the nurse's scrubs, at which point her experiences might become *oriented* in a color space, just as one's experiences are normally oriented in a physical space. The same considerations extend to other familiar characters, like the complete amnesiac, and Mary after she leaves her black and white laboratory.

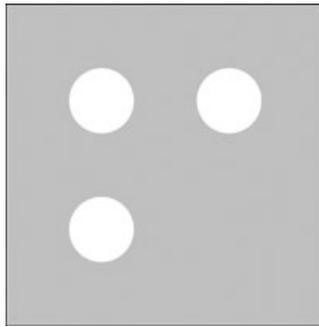
This isn't the anti-atomist's only response. Perhaps even a newborn appreciates that the ceiling could have looked darker and could have looked bluer, allowing her experience to represent relations to these non-actual alternatives, which might be enough for her experience to represent the ceiling as white. That would sidestep the implication.

Second, anti-atomism potentially forces us to give up the widespread view that there is a one-to-one correspondence between color phenomenology and color representation. In particular, it potentially forces us to give up the view that we represent charcoal-gray₁₁ whenever we have an experience with a certain color phenomenology, we represent charcoal-gray₁₂ whenever we have an experience with a certain other color phenomenology, and so on.³⁴ Consider our newborn again. As she's exposed to more objects that differ along more dimensions, her representations will become more specific. Assuming her color phenomenology remains the same whenever she looks at the ceiling, it follows that an experience with the same color phenomenology can represent a specific property, a general property, or no property at all, in which case there isn't a one-to-one correspondence between color phenomenology and color representation.

Building on what we said before, the anti-atomist could respond by insisting that the newborn's experience always represents the same shade of white, because from the very beginning she represents its relation to alternatives where it looks

darker, bluer, greener, etc. But that's not the anti-atomist's only option. The anti-atomist could also deny that the newborn's color phenomenology remains the same whenever she looks at the ceiling. Let's again use spatial representation as a guide. Suppose you're in a perfectly dark room when a spotlight illuminates a figurine. At first, the figurine might not seem to occupy any particular location—since you have no clues as to its size, it might not seem near or far. But as additional spotlights begin to illuminate other parts of the room, including the table underneath and the painting behind, the figurine will seem to occupy a more determinate location; it will seem to be at an increasingly exact location in the coordinate system centered between your eyes. Even if some aspects of your spatial phenomenology remain constant, representing additional spatial relations thereby changes your phenomenology with respect to the figurine, and this effect won't entirely disappear if the supplementary spotlights are turned off. Similarly, after the infant looks at the nurse's blue scrubs, at her mother's crimson lips, and then back at the ceiling, she'll situate the ceiling in a web of color relations, changing her color phenomenology. She'll have a new perspective on the ceiling; she'll become aware of it as whiter than the nurse's scrubs, and there will be a phenomenological component to this new awareness. The alternative, that her color phenomenology remains the same, falsely predicts that representing additional color relations has no phenomenological consequences. Of course, like before, it's impossible to verify this characterization of newborn phenomenology, but its plausibility is enough to relax any tension between anti-atomism and the view that there's a one-to-one correspondence between color phenomenology and color representation.

The third potential cost is that color illusions would result in widespread error. For example, suppose that three white dots are painted on a gray wall:



Someone then shines a light on the wall so that it looks like it has a fourth dot with exactly the same color as the original dots. According to the anti-atomist, you'd then be misrepresenting the three original dots, because you'd be misrepresenting their relations to the part of the wall where you see the illusory fourth dot.

This doesn't strike me as a genuine cost. Your experiences of the three original dots would be more accurate than your experience of the part of the wall where you see the illusory fourth dot, because you'd be accurately representing more of the relations involving the three original dots. In particular, you'd still be accurately

representing each of them as exactly the same color as the others. In contrast, your experience of the part of the wall where you see the illusory fourth dot would be almost completely inaccurate, because you'd be inaccurately representing almost all of its relations. Thus, while color illusions result in widespread error, the degree of error is uneven, and this strikes me as an intuitive result. As a helpful analogy, consider a map of tourist attractions in New York City. If it inaccurately located the Empire State Building next to the Metropolitan Museum of Art, then the map would inaccurately represent the locations of all the other tourist attractions, at least to some extent, because it would misrepresent their proximity to the Empire State Building. The map's misrepresentation of the Empire State Building would be most severe, because it would misrepresent all of its relations to the other landmarks.

There are, no doubt, other potential costs of endorsing anti-atomism. Whether those potential costs outweigh its benefits is a question we'll have to leave for another occasion.

6. Conclusion

Our puzzle was about color experiences. But there are symmetrical puzzles about other kinds of experiences, including our experiences of sounds, pains, lengths, angles, locations, and many other properties. Just as two objects can look exactly the same color, two notes can sound exactly as flat, two pains can feel exactly as intense, two lines can look exactly as long, two angles can look exactly as obtuse, and at two times an object can seem to occupy exactly the same location. Anti-atomism might give us a plausible response to all these puzzles.³⁵

Our experiences of some properties, like textures and odors, are too impoverished to represent exact sameness and therefore don't give rise to equivalent puzzles. It is impossible to imagine two blankets that feel *exactly* as soft or two perfumes that smell *exactly* as pungent. Nonetheless, there's nothing preventing us from extending anti-atomism to our representations of these other properties. Our representations of their relations would just be less precise than our representations of some color relations, a situation we already modeled in our discussion of more determinable colors.

Can we extend anti-atomism to all perceptual representations? Our experiences of shapes might seem recalcitrant. But there's a way. Arguably, a road sign looks square in virtue of how its sides and angles look. If so, we might represent its shape in virtue of representing the relations between its sides and angles, which is all that anti-atomism requires.

Given the widespread applicability of anti-atomism, it is natural to wonder whether it captures something essential about perceptual representation. While we might be able to talk and think about properties without representing any relations, perhaps there's something inextricably comparative about perceptual representation that makes it impossible for us to perceive things in this way. It's certainly a possibility worth exploring.³⁶

Notes

¹ Merely rejecting TRANSITIVITY actually isn't enough. The puzzle only requires the following principle, which is weaker but less straightforward: For all objects in stable contexts: if x and y look exactly the same color and y and z look exactly the same color, then x and z do *not* look *differently* colored. All the responses I'll consider give us reasons to reject this weaker claim as well as TRANSITIVITY.

² To keep things simple, I've left off a restriction. Consider two window panes. There is no color such that only one of them looks to have it. Therefore, as stated, the biconditional would imply that they look exactly the same color, which is the wrong result, because they look colorless. To avoid this consequence we need to restrict the biconditional to objects that look colored.

³ Formally, where Sxy translates ' x looks exactly the same color as y ' and Cx translates ' x looks color C ': $\forall xy[Sxy \leftrightarrow \forall C(Cx \leftrightarrow Cy)]$ entails $\forall xyz[(Sxy \& Syz) \rightarrow Sxz]$.

⁴ Though she formulates it differently, Fara [10, p. 914–915] says this biconditional is an undeniable truism and is therefore unwilling to give up TRANSITIVITY.

⁵ See Sanford [32].

⁶ See Fine [11] and Schiffer [33, p. 198–209], respectively.

⁷ See Breckenridge [3, p. 11–19] for an overview and discussion.

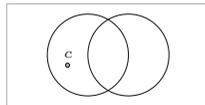
⁸ See Russell [31, p. 104–107], Goodman [12, p. 219–226], Armstrong [1, p. 218–221], and Jackson and Pinkerton [17]. Williamson [36] belongs in this group because he focuses on the phenomenal characters of experiences, like their painfulness.

⁹ Dummett [9], Wright [39], and Peacocke [24] focus on the predicate 'is red', which they consider an observational predicate like 'looks red'. McDowell [19, p. 170–171] and Pelling [25] focus on the demonstrative 'that shade' and the corresponding perceptual concept.

¹⁰ See Williamson [38, p. 93–113].

¹¹ See Chuard [5], Fara [10], Hellie [15], and Raffman [29].

¹² Hellie [15] is best understood as developing this response. As a helpful simplification, let's restrict ourselves to a certain two-dimensional slice through color space. According to Hellie, the first chip looks the color circumscribed by the left circle, and the second chip looks the color circumscribed by the right circle:



slice through color space

According to this diagram, Jonah's experience of the first chip is accurate if it is c , and Jonah's experience of the second chip is *inaccurate* if it is c . This asymmetry is inexplicable if the two objects look exactly the same color. Hellie should therefore say that the two chips look roughly the same color. (That's not, however, what he says—see p. 497)

Stepping back, we can appreciate why this diagram might correspond to the information carried by signals in Jonah's visual system. But that doesn't automatically imply that it is what Jonah's conscious color experiences represent. Jonah's visual system might be designed so that when it gets similar information about two surfaces, it represents them as exactly the same color. Proponents of roughness need to give us a reason to think otherwise.

¹³ This isn't the end of the discussion. Proponents of the roughness response might insist that a state represents that p only if it provides one with a sufficiently high degree of confidence that p . In that case, the speedometer does not represent you as moving at exactly 60mph. But this strikes me as an implausible principle, because a meteorologist and a soothsayer seem to represent the same thing when they say 'there will be a hail storm tonight' even though you shouldn't have the same degree of confidence in what they're saying. In any case, the confusion is in thinking that the mere fact that color experiences are a kind of measurement should lead us to accept the roughness response.

¹⁴ Raffman [29] is an especially clear statement of this kind of response.

¹⁵ Fara [10, p. 927–8] and Mills [21, p. 392] employ this tactic.

¹⁶ See Jackson and Pargetter [16], McLaughlin [20], and Cohen [6].

¹⁷ This would undermine the left-to-right direction of the biconditional unless it is restricted to a given time. Even if the first chip looked gray₁₁-in-*C* both times it was viewed, it presumably only looked gray₁₁-in-*C** the second time it was viewed, because Jonah hadn't been in *C** before.

¹⁸ It might also be helpful to distinguish this response from contextualism about vagueness, the view that vague expressions like 'look red' shift their content from one context to the next, analogous to the way indexicals shift their contents. See Raffman [28]. Unlike that contextualist, our contextualist does not claim that there are any shifts in the contents of linguistic expressions, including 'looks exactly the same color'.

¹⁹ We'd expect Peacocke [24] to endorse this response. He claims that sentences like '*a* is red' and '*b* is red' are true to varying degrees and that we can't know just by observing *a* and *b* which sentence, if either, is true to a greater degree.

²⁰ There's an obvious similarity between this response and epistemicism about vagueness: they both claim there are unknowable facts (see Williamson [37, p. 185–244]). But there are important differences. Epistemicists claim we're ignorant of what's represented by our *linguistic expressions*. Inaccessibility theorists claim we're ignorant of what's represented by our *visual experiences*. These two kinds of ignorance require different explanations. For example, while Williamson attributes our ignorance of linguistic meaning to our ignorance of the social factors that fix meanings, the inaccessibility theorist must appeal to other factors. In addition, our ignorance of what we visually represent would run deeper, because while the epistemicist denies that we can know what 'bald' means, she presumably still allows that we can know that 'bald' and 'bald' have the same meaning, whereas the inaccessibility theorist not only denies that we can know which property our experience represents but insists that we also can't know when two visual experiences represent the same property. Finally, epistemicism is motivated by the desire to preserve classical logic and to avoid higher-order vagueness, and neither applies to the inaccessibility response.

²¹ An anonymous reviewer suggests that we should expect such cases if introspection isn't a kind of "inner perception."

²² In my unpublished "Triangulating How Things Look" I discuss another alternative, the **triangulation response**. According to this response, (2a) and (3a) should have prompted Jonah to infer:

(1d) The first and second chips do not look exactly the same color.

And (1a) and (3a) should have prompted Jonah to infer:

(2d) The second and third chips do not look exactly the same color.

And, as before, (3a) alone should have prompted Jonah to infer:

(3d) The first and third chips do not look exactly the same color.

Because (1d) and (2d) are inconsistent with (1b) and (2b), the triangulation theorist rejects (1b) and (2b). In "Triangulating How Things Look" I argue that, due to visual noise, this view has the same drawback as the inaccessibility response.

²³ This can't be the whole story. The atomist must explain why representing one line as l_1 and the other line as l_2 is *sufficient* for representing one as longer than the other. This is a tricky question that anti-atomism helpfully sidesteps. Thanks to an anonymous reviewer and Nemira Gasiunas for pointing this out.

²⁴ See Pressey [27].

²⁵ These are cases of intrapersonal variation. Cohen [6], McLaughlin [20], and Jackson and Pargetter [16] use these cases to argue that colors are context-relative (I discussed this view in the last section). Anti-atomism gives us another approach to intrapersonal variation. Anti-atomism does not, however, give us another approach to other kinds of variation, namely interpersonal and interspecies variation.

²⁶ Byrne and Hilbert [4] and Tye [35] make this bet.

²⁷ There are other reasons why anti-atomism is attractive. In my [22] I developed a problem for all realist theories of color. While I didn't appreciate it at the time, my argument depended on atomism. Therefore, anti-atomism would make it easier to be a realist about color. In future work I plan to argue that anti-atomism also gives us the best account of peripheral color perception.

²⁸ According to Rosenthal [30], a mental quality represents charcoal-gray if we have the ability to use that mental quality to discriminate charcoal-gray objects from differently colored objects. Rosenthal's

view isn't anti-atomistic in my sense, because we might have such an ability even if we're not currently and consciously representing any relations to differently colored objects. But Rosenthal's view is anti-atomistic in another sense, since representing charcoal-gray is grounded in relations to other, differently colored objects.

²⁹ This approach has an obvious connection to supervaluationism, but there is an important difference. In the supervaluational framework, what varies between scenarios are the referents of our terms—the objects remain the same. In our framework, what varies are the properties of the relevant objects. One consequence is that our framework doesn't force us to acknowledge referential indeterminacy (though it doesn't rule it out, either) because we needn't think of our visual experiences as indeterminately representing each scenario. Another consequence is that our framework doesn't force us to acknowledge truth-value gaps (though it doesn't rule them out, either), because we're interested only in what's true in all scenarios.

There are other approaches. One alternative is more classically supervaluational: that we represent only one scenario but that there is no fact of the matter about which scenario we are representing. Another alternative is more epistemic: each scenario is epistemically possible, and together the scenarios exhaust the epistemic possibilities. Our representation would then be true if any of the scenarios is actual, otherwise false. I'm developing the approach I find most appealing.

³⁰ The story becomes more complicated when we include representations of relations to additional objects. Depending on the relations, it is possible that we'd end up correctly perceiving LS's color and misperceiving RS's color, because once those new relations are taken into account, LS has the correct shade of gray in all the compatible scenarios, while RS has the incorrect shade in those scenarios. That'd be an acceptable result, because we'd then have an explanation for the asymmetry.

³¹ There's empirical support for this claim. Collins [8], Hamwi and Landis [13], and Perez et al. [23] all demonstrate that we're worse at re-identifying colors after a delay. The best explanation is that we lose the ability to represent fine-grained relations to remembered objects. The alternative explanation is that we misremember colors as redder or greener—we misremember another, equally determinate color. But that's introspectively implausible. We aren't surprised by an object's color every time we look away and look back.

³² There is another option. Properties like *is a square circle* are inconsistent in that they combine incompatible properties. Jonah might be representing an inconsistent color like *is and is not exactly the same color as the second chip*.

³³ This potentially undermines the left-to-right direction of the biconditional that entails TRANSITIVITY. An object that was introduced into the visual scene three seconds ago might look to instantiate *is constant* only after three seconds have passed. Therefore, at least according to the complete anti-atomist, something can look exactly the same color as itself three seconds ago even if it didn't originally look to instantiate *is constant*. Unless the biconditional were restricted to a time, that's potentially a counterexample to its left-to-right direction. Other potential counterexamples include cases in which an object looks exactly the same color over time even if, at an earlier time, it looked darker than a neighboring object that was later removed and forgotten. There's another, more general reason why complete anti-atomists might be forced to give up the left-to-right direction of the biconditional. If we don't represent reflexive color relations, like *is exactly the same color as itself*, and we represent the left half of a wall as exactly the same color as the right half, then we'll represent the left half as instantiating *is exactly the same color as the right half* but we won't represent the right half as instantiating the same property. A moderate anti-atomist can also reject the biconditional's left-to-right direction but isn't similarly pressured towards that conclusion.

³⁴ Tye [34, p. 47–68] is one of the clearest statements of this view.

³⁵ Dummett [9, p. 316] and Fara [10, p. 926–927] discuss a parallel puzzle for sameness of location. It centers on an object whose motion is so slow that we can't notice any change in its position from one second to the next. Anti-atomism also gives us a response to this puzzle: deny that we see the object as having a different and determinate location at each moment.

³⁶ An exceptional group of people helped this paper develop from its first draft in January of 2010 to its final draft in December of 2012. With apologies to those I've forgotten, I'd like to alphabetically thank Ralph Bader, David Barnett ("the junior"), David Barnett ("the senior"), John Collins, Sinan Dogramaci, Nemira Gasiunas, Elizabeth Harman, Benj Hellie, Alex Madva, Colin Marshall, Stephen

Schiffer, James Stazicker, Sebastian Watzl, Bruno Whittle, Achille Varzi, Anubav Vasudevan, and an anonymous referee for this journal. I'd also like to thank the participants in my graduate seminar on perception (April 2012) and the Corridor Reading Group (June 2011), as well as audiences at CUNY (March 2011), Columbia (April 2011), and the Barnard-Columbia Perception Workshop (June 2011). Finally, Shamik Dasgupta deserves special mention for extensive comments on three drafts. Thanks, everyone!

References

- [1] David Armstrong. *A Materialist Theory of the Mind*. Routledge, New York, NY, 1968.
- [2] Ned Block. Troubles with functionalism. *Minnesota Studies in the Philosophy of Science*, 9:261–325, 1978.
- [3] Wylie Breckenridge. *The Meaning of 'Look'*. PhD thesis, New College, University of Oxford, Hilary Term 2007.
- [4] Alex Byrne and David R. Hilbert. Truest blue. *Analysis*, 67:87–92, 2007.
- [5] Philippe Chuard. Non-transitive looks & fallibilism. *Philosophical Studies*, 149:161–200, 2010.
- [6] Jonathan Cohen. Color properties and color ascriptions: A relationalist manifesto. *The Philosophical Review*, 113:451–506, 2004.
- [7] Jonathan Cohen. Color constancy as counterfactual. *Australasian Journal of Philosophy*, 86:61–92, 2008.
- [8] Mary Collins. Some observations on immediate colour memory. *British Journal of Psychology*, 22:344–352, 1932.
- [9] Michael Dummett. Wang's paradox. *Synthese*, 30:301–324, 1975.
- [10] Delia Fara. Phenomenal continua and the sorites. *Mind*, 110:905–935, 2001. Originally published under the name "Delia Graff".
- [11] Kit Fine. Vagueness, truth and logic. *Synthese*, 30:265–300, 1975.
- [12] Nelson Goodman. *The Structure of Appearance*. Harvard University Press, Cambridge, MA, 1951.
- [13] Violet Hamwi and Carney Landis. Memory for color. *Journal of Psychology*, 39:183–194, 1955.
- [14] C. L. Hardin. *Color for Philosophers: Unweaving the Rainbow*. Hackett, Indianapolis, IN, 1988.
- [15] Benj Hellie. Noise and perceptual indiscriminability. *Mind*, 114:481–508, 2005.
- [16] Frank Jackson and Robert Pargetter. An objectivist's guide to subjectivism about color. *Revue Internationale de Philosophie*, 41:127–41, 1987.
- [17] Frank Jackson and R. J. Pinkerton. On an argument against sensory items. *Mind*, 82:269–272, 1973.
- [18] Akiyoshi Kitaoka. Luminance-gradient-dependent lightness illusion, March 2006. <http://www.psy.ritsumeai.ac.jp/akitaoka/gilchrist2006mytalke.html#LGLI>. Reproduced with permission.
- [19] John McDowell. *Mind and World*. Harvard University Press, Cambridge, MA, 1994.
- [20] Brian McLaughlin. Colour, consciousness, and colour consciousness. In Quintin Smith and Aleksandar Jokic, editors, *New Essays on Consciousness*. Oxford University Press, 2003, 97–154.
- [21] Eugene Mills. Fallibility and the phenomenal sorites. *Noûs*, 36:384–407, 2002.
- [22] John Morrison. Colour in a physical world: A problem due to visual noise. *Mind*, 121:333–73, 2012.
- [23] Joaquín Pérez-Carpinell, Rosa Baldoví, M. Dolores de Fez, and José Castro. Color memory matching: Time effect and other factors. *Color Research and Application*, 23:234–247, 1998.
- [24] Christopher Peacocke. Are vague predicates incoherent? *Synthese*, 46:121–141, 1981.
- [25] Charles Pelling. Conceptualism and the (supposed) non-transitivity of colour indiscriminability. *Philosophical Studies*, 134:211–234, 2007.
- [26] Lionel Penrose and Roger Penrose. Impossible objects: A special type of visual illusion. *British Journal of Psychology*, 49:31–33, 1958.
- [27] Alexander Pressey. Effect of size of angle on the ambiguous Müller-Lyer illusion. *Acta Psychologica*, 38:401–404, 1974.
- [28] Diana Raffman. Vagueness without paradox. *The Philosophical Review*, 103:41–74, 1994.

- [29] Diana Raffman. Is perceptual indiscriminability nontransitive? *Philosophical Topics*, 28:153–175, 2000.
- [30] David Rosenthal. How to think about mental qualities. *Philosophical Issues*, 20:368–393, 2010.
- [31] Bertrand Russell. *An Inquiry into truth and meaning*. W. W. Norton & Company, New York, NY, 1940.
- [32] David Sanford. Borderline logic. *American Philosophical Quarterly*, 12:29–39, 1975.
- [33] Stephen Schiffer. *The Things We Mean*. Oxford University Press, New York, NY, 2003.
- [34] Michael Tye. *Consciousness, Color, and Content*. MIT Press, Cambridge, MA, 2000.
- [35] Michael Tye. The puzzle of true blue. *Analysis*, 66:173–178, 2006.
- [36] Timothy Williamson. *Identity and Discrimination*. Basil Blackwell, London, 1990.
- [37] Timothy Williamson. *Vagueness*. Routledge, Cambridge, MA, 1994.
- [38] Timothy Williamson. *Knowledge and its Limits*. Oxford University Press, New York, NY, 2000.
- [39] Crispin Wright. On the coherence of vague predicates. *Synthese*, 30:325–365, 1975.