

Color in a Physical World

John Morrison

October 18, 2009

1 Introduction

There is a tension between our perceptual experiences and our physical theories. Our perceptual experiences tell us about a world full of yellow dandelions, grey oysters, blue sapphires and green olives. Our best physical theories tell us about a world of spinning quarks, polarized leptons and energized bosons. Because we have little trouble imagining how large things could be constructed out of small things, we seem to have little trouble imagining how dandelions, oysters, sapphires and olives could be constructed out of quarks, leptons, and bosons, at least in rough outline. But we seem to have lots of trouble imagining how yellowness could relate to the spin, polarity, energy, etc., of a dandelion's fundamental particles. As a result we are led to wonder whether there is room for color in a physical world.

This paper is about **conciliatory thinking**, a widespread and influential way of thinking about colors and the way they fit into the physical world. There are two claims that are definitive of conciliatory thinking. The first claim is that the colors preserve the structure of our color experiences. What is the “structure” of our color experiences? One of the central tasks of this paper will be to get clear about this notion. But, for now, an example might help. Just as we experience objects as located in space and thereby experience them as closer or further from each other, we also experience objects as colored and thereby experience them as more or less similarly colored. For example, an experience of a yellow dandelion, an orange

tangerine and a blue sapphire also seems to be an experience of the dandelion and tangerine as more similarly colored than the dandelion and the sapphire. And this seems to be built into the structure of our color experiences. There are many other aspects of this structure and I will explore some of them later.

There are many ways to motivate the claim about structure. One way is to appeal to the intuition that the structure of our experiences reveals what is *essential* to the colors.¹ Another way is to appeal to the weaker intuition that the structure of our experiences is built into our color concepts, which implies that it is a *conceptual necessity* that the colors satisfy the structure of our experiences.² Yet another way is to argue for the modest view that the structure of our experiences places a constraint on which properties we represent in the actual world.³ Because a property should only count as a color if we actually represent it (or properties like it), it would follow that the colors at least *contingently* preserve the structure of our colors experiences. For the purposes of this paper it won't matter which

¹Many philosophers seem to share this intuition. Strawson [68] claims that the whole nature of the colors is “revealed” by experience. Harding [30] says that the nature of the colors is “laid bare” in experience. Campbell [13] claims that it is part of “common sense” that the essential facts about the colors are “transparent” to us. Johnston [41] claims that it is among our “core beliefs” that the colors are “perceptually available” in that we can know necessary truths about the colors by reflection on our experiences. Armstrong [2] thinks that our experiences can indicate what *follows from* the nature of the colors though he doesn't think that our experience can directly indicate anything about the nature of the colors. Shoemaker [66] thinks that we have this kind of access to what we might call ‘phenomenal colors’ and Chalmers [15] thinks we have this kind of access to what he calls ‘Edenic colors’. See also Hardin [28], who seems to endorse this kind of reasoning.

²A number of philosophers make claims in the vicinity of this claim. Wittgenstein [72] says that it is “unthinkable” that the colors could have different relations to one another. Broackes [7] talks about “phenomenal elements in our conception of the colors.” Yablo [73] says that to know the colors one must know them in a “subjectivity-involving way,” which might indicate a commitment to the second claim, depending on how we understand this expression. While they don't go into the details, Boghossian and Velleman [4] [5] say that it is “beyond question” that we can know necessary truths about the colors by introspection and Pautz [61] says that it is “obvious.”

³For reasons I will explore in a moment, it is sometimes difficult to determine whether someone accepts the claim about structure, and that's especially true for philosophers who say things that seem to correspond to this third motivation. Nonetheless, let's canvass some likely candidates. McLaughlin and Jakob [58] say that “color physicalism must be squared with the perceived similarity relations” and McLaughlin [57] builds these relations into the functional specification of the colors. So does Cohen [17], [16]. Tye and Bradley [6] also seem to belong to this group. Thompson [69] thinks that our theory of color must accommodate our pre-theoretic intuitions, and he seems to think that these intuitions reflect the structure of our experiences. Due to their reaction to Matthen [49], I suspect that Byrne and Hilbert [12] fall in this category.

motivation a conciliatory thinker endorses.

The second claim that is definitive of conciliatory thinking is that the colors *depend on* physical properties such as spin, polarization, energy and mass. There are many ways in which the colors might depend on these physical properties. For example, they might be identical to them, metaphysically necessitated by them, nomologically necessitated by them, have them as their dispositional grounds or just co-vary with them in the actual world. I will understand the claim about dependence so that it is neutral between these possibilities. So understood, this claim is endorsed by a wide swath of philosophers of color, including all color physicalists and many color dispositionalists and color primitivists.

Let's now turn to the claim about dependence. Like the claim about structure, there are many ways to motivate it. For brevity, I'm just going to focus on one of them: if the colors didn't depend on the properties of physics then it would be difficult, though not impossible, to maintain that anything is actually yellow, orange, red or green. To see why, suppose the redness of a stoplight causes a driver to press on her brake. If the redness of that stoplight did not depend on its physical properties then prior knowledge of the physical properties of the stoplight wouldn't have helped us predict that the car was about to slow down because, given that the colors wouldn't even co-vary with any physical properties, knowledge of the stoplight's physical properties would have been irrelevant information. That's a consequence that many people, including myself, are unwilling to accept. One way to avoid this consequence would be to deny the supposition that colors are causally efficacious, but that seems more unappealing than just denying that anything is really colored. Accordingly, people who think that things are actually colored ("realists") almost always maintain that the colors depend on physical properties, though they often disagree about the strength of that dependence, such as whether it is identity or mere nomological necessity.

To recap, the two definitive claims of conciliatory thinking are: the claim about structure (that the colors satisfy the structure of our color experiences) and the claim about dependence

(that the colors depend on physical properties). I think that it is helpful to see why these kinds of claims come in pairs. Suppose that in order for a property to count as yellowness it only needed to cause experiences of yellowness. In that case, it would be obvious, and perhaps even *a priori*, that yellowness is a physical property because it is obvious, and perhaps even *a priori*, that our experiences of yellowness have physical causes. In contrast, suppose that in order for a property to count as yellowness it needed to be an intrinsic and non-disjunctive property. In that case, it would be far less obvious whether yellowness is a physical property because it is far less obvious whether dandelions and lemons, for example, share any intrinsic and non-disjunctive physical properties. As these examples demonstrate, the claim about dependence needs to be paired with a claim about what conditions a property must satisfy in order to count as a color, and that's the role of the claim about structure. Thus we shouldn't think of conciliatory thinking as the mere disjunction of two independent claims. Instead, we should think of conciliatory thinking as an attempt to successfully integrate two commitments that press against each other.

In this paper I'm going to develop a new problem for conciliatory thinking. I will argue that the most natural and plausible strategy for integrating the claim about structure and the claim about dependence is unsuccessful, thereby calling into question the tenability of conciliatory thinking.

Of course, like all philosophical arguments, my arguments rely on claims that someone could reject. However, as we'll see, rejecting these claims would have some surprising, far-reaching and counter-intuitive consequences for one's theory of color and/or color experiences. A conciliatory thinker should therefore find my arguments challenging even if she isn't ultimately persuaded that her way of thinking is untenable.

Let's now address a complication. In large part because there is no standard vocabulary for talking about these matters, it can sometimes be difficult to determine whether a particular philosopher accepts the claim about structure. However, there is often plenty of indirect

evidence. Sometimes that evidence comes from the fact that the philosopher endorses one of the motivations listed above. Other times it comes from passages that seem equivalent:

Colors are the kinds of properties that fit together in characteristic ways to form structured color arrays, with a distinctive three-dimensional character. They are properties that as a group, form an internally related structure, built on the four unique, primary hues: green, red, blue and yellow, and related to the black/white pair. (Maud [52])

Like Maud, many people claim that structured colors arrays (also called ‘color solids’ or ‘color spaces’) not only capture the structure of our color experiences but also capture the structure of the colors themselves.⁴ Finally, evidence that a philosopher accepts the claim about structure sometimes comes from the way in which she responds to the **similarity objection**, which is an influential argument that targets one strand of conciliatory thinking. The people who develop this objection claim that we cannot *identify* the colors with physical properties because then there would be no genuine resemblance between red and orange, in which case, they conclude, the colors would not preserve the structure of our experiences. With few exceptions, philosophers of color respond in two ways. First, they endorse the conclusion and infer that colors aren’t identical to physical properties.⁵ Second, they argue that there is indeed a genuine similarity between the relevant physical properties, thereby undermining the argument’s crucial premise.⁶ Only a handful of philosophers respond by shrugging off the conclusion on the grounds that we shouldn’t expect the colors to preserve the structure of our experience.⁷ Because there doesn’t seem to be anything special about the similarities between colors, this suggests that most philosophers accept the more general claim that the colors satisfy the structure of our experiences.

I want to advertise two features of my arguments in advance. First, unlike existing arguments, my arguments target *all* strands of conciliatory thinking about color. In the

⁴See Tye and Bradley [6], McLaughlin [57] and Byrne [9].

⁵See Hardin [28] [29], Johnston [41], Maund [51], McGinn [55], Pautz [60] [61] and Thompson [69].

⁶See Byrne [9], Byrne and Hilbert [12], Bradley and Tye [6], Cohen [17] and McLaughlin [57].

⁷See Harman [31] and Matthen [50].

literature there are already arguments that target specific strands of conciliatory thinking. For example, as noted above, the similarity objection only targets the view that the colors are identical to physical properties, which is why it is often used to motivate other strands of conciliatory thinking, such as strands that merely insist that the colors supervene on physical properties.⁸ In contrast, my arguments target anyone who accepts the claim about dependence *regardless* of whether they think that the dependence is identity, metaphysical necessity, nomological necessity or even contingent co-variation. Moreover, the fact that people typically use the similarity objection, which depends on the claim about structure, as the primary motivation for these other strands helps give my arguments added force - I'm using one of their own premises against them.

Another example is a much-discussed argument targeting the strand of conciliatory thinking that categorizes colors as dispositions. Those who develop this objection allege that the colors aren't dispositions because our experiences reveal what is *essential* to the colors and our experiences don't reveal that the colors are dispositions.⁹ The success of this argument hinges on the claim that our experiences reveal what is essential to the colors, and that goes far beyond what the claim about structure requires. As I've formulated it, the claim about structure is extremely weak. For example, it doesn't entail that the colors must satisfy the structure of our color experiences in virtue of their essences, or even that the colors must satisfy the structure of our color experiences in all metaphysically possible worlds. It just entails that the colors at least satisfy the structure of our experiences in the actual world. Therefore, unless there is a stand-alone argument for this claim about what our experiences reveal, this other argument doesn't seem to extend to every strand of conciliatory

⁸Why aren't these strands threatened by the similarity objection? An analogy might help: if minds are not identical to brains then your mind might resemble an alien's mind when you're both thinking about dandelions even if your brain does not resemble the alien's brain.

⁹See Boghossian and Velleman [5], Campbell [13] [14], Harding [30], McGinn [55], and Strawson [68], among other places. See also Johnston [41], though he uses a slightly weaker premise. Boghossian and Velleman [5] and Pautz [60] use a similar premise to enhance the similarity objection.

thinking. Once again, what will set my arguments apart is their generality.

There’s another feature of my arguments that I want to advertise: one cannot respond to them by merely adapting standard responses to other problems. For example, philosophers who want to undermine the similarity objection often appeal to a claim that we can roughly state: an experience represents the disjunction of all the colors that reliably cause that experience. They use this claim to argue that redness and orangeness resemble one another even if they are identical to physical properties.¹⁰ Given the apparent usefulness of this claim in the context of the similarity objection, they might hope that it also gives them a way to undermine my arguments. But this claim is actually one of the premises that I’ll use to call into question the tenability of conciliatory thinking and thus it obviously cannot help a conciliatory thinker in this context. My arguments challenge conciliatory thinking in a novel way.

2 Preview

It may be helpful to have an idea of where we’re headed, so let’s use a toy model to roughly sketch the main argument

Suppose a collector of precious stones has a scale that displays one of four terms, ‘tiny’, ‘small’, ‘large’ or ‘huge’, depending on its measurement:

Display	Measurement (x)
tiny	$x \leq 10\text{mg}$
small	$10\text{mg} < x \leq 11\text{mg}$
large	$11\text{mg} < x \leq 12\text{mg}$
huge	$12\text{mg} < x$

¹⁰See Byrne [9], Byrne and Hilbert [12], Bradley and Tye [6] and Cohen [17]. Their arguments draw on Opponent Processing Theory.


Suppose also that the collector endorses two claims. The first claim is that each term represents a disjunction of weights. In particular, ‘tiny’ represents the disjunction of all the weights that reliably cause the scale to output ‘tiny’, ‘small’ represents the disjunction of all the weights that reliably cause the scale to output ‘small’, etc.

The second claim is that these terms represent properties that preserve the truth of certain conditionals, including:

- (a) If a stone isn’t tiny, small or large then it is huge.
- (b) If a stone is tiny and another stone is small then all stones with intermediate weights are tiny or small.

The collector might regard (a) and (b) as part of the structure of the scale’s system of representation.

While the first and second claim might at first seem independent, it turns out that they are inconsistent. Like all measuring instruments, the scale’s measurements will be “noisy.” For example, if a stone actually weighs eleven milligrams then the scale will nonetheless sometimes register more or less weight, and therefore will sometimes display ‘tiny’, ‘small’, ‘large’ and ‘huge’ (though not all at once). Noise often has the shape of a bell curve. For any stone that actually weighs eleven milligrams we might graph the probability that it will produce different measurements:



The area under the curve between ten and eleven is the probability that the scale will display ‘small’. Likewise for ‘tiny’, ‘large’ and ‘huge’. Because it is a *symmetric* probability distribution centered at eleven, and because it *asymptotically* approaches zero in both directions, all of these areas are less than one-half. Consequently, stones that weigh eleven milligrams have an equal likelihood of producing ‘small’ and ‘large’ as well as a non-zero likelihood of producing ‘tiny’ and ‘huge’. As a result, they do not *reliably* produce ‘tiny’, they do not *reliably* produce ‘small’, etc.

This example reveals something very general about measuring devices: due to the inevitable existence of noise, as long as the different types of stimuli fall along a continuum and there are only countably many outputs, some range of stimuli types will not *reliably* produce any particular output. In different devices the existence of noise will have different causes. For example, in metallic scales it might be caused by random variations in the temperature of the coil, and in the human eye it might be caused by random variations in blood pressure.

We’re now in a position to appreciate why the collector’s two claims are inconsistent. Due to the existence noise, the first claim entails that ‘tiny’, ‘small’, ‘large’ and ‘huge’ represent disjunctions of weights that do not include eleven milligrams. Therefore, a stone that is eleven milligrams is not tiny, small, large or huge, in which case (a) is false. Also, a stone that is nine milligrams will be tiny, a stone that is eleven and a half milligrams will be small, but a stone that is eleven milligrams will be neither tiny nor small, in which case (b) is false. Recall that the second claim is that ‘tiny’, ‘small’, ‘large’ and ‘huge’ represent properties that preserve the truth of (a) and (b). Therefore, the first claim and the second claim are inconsistent.

I’m going to argue that there is a similar problem for conciliatory thinking. In particular, I’m going to argue that the claim about structure implies that certain transitions in color are “full” in that there are no intermediate shades that aren’t represented. I will then argue

that, due to the existence of noise, the claim about dependence implies that there are always intermediate shades that aren't represented, in which case the two claims are inconsistent. That's obviously a problem for conciliatory thinkers. I will then canvass possible responses and argue that, at best, those responses would commit conciliatory thinkers to surprising, far-reaching and counter-intuitive claims.

Before I can state the full argument I need to identify the physical properties relevant to color and say more about the two definitive claims of conciliatory thinking.


3 Spectral Dispositions

There is no standard vocabulary for talking about the physical properties relevant to color and, moreover, the existing vocabularies are often skewed to particular views. As a remedy, I will introduce a number of new terms, including 'spectral dispositions', which is my term for the physical properties that the colors depend on.

A good place to start is Newton's [59] prism experiment. Newton cut a small hole in one of his laboratory's shutters, thereby allowing a narrow ray of sunlight to enter. He then used a prism to project a rainbow onto the opposite wall. Newton noted that he had different color experiences of different regions on the wall. In particular, he noted that he had different color experiences of regions with different angles of refraction relative to the original ray of sunlight. He inferred that there is a connection between a light ray's angle of refraction and an observer's color experiences. After Maxwell's [53] subsequent discovery that a ray's angle of refraction varies with its wavelength, many philosophers and scientists concluded that there is a connection between the wavelengths of light rays and an observer's color experiences. Even those who deny that actual objects are colored ("color eliminativists") can acknowledge this connection.

According to conciliatory thinking, the colors of objects depend on some of physical properties. The properties in virtue of which objects reflect, absorb and transmit light are among the properties that our best physical theories tell us about. Because they are also the properties in virtue of which objects cause color experiences, conciliatory thinkers theorize that there is a connection between the colors of objects and the properties of objects in virtue of which they reflect, absorb and transmit light. Let's explore this connection.

We need to start by introducing two technical terms. First, the **spectral power distribution** of a light ray is its intensity at each wavelength. We can graph the spectral power distribution of a light ray:




While objects do not themselves have spectral power distributions, they do reflect, emit and transmit rays of light with different spectral power distributions. In different contexts the spectral power distribution of those light rays will be different. For instance, in conditions of low illumination most objects will reflect, transmit or emit rays of light with low intensities at every wavelength. Leaving a more precise definition in a footnote, the **spectral disposition** of an object or volume is a function from each context to the spectral power distribution of the light rays reflected, transmitted or emitted by that surface or volume in that context.¹¹ Notably, even if two objects reflect, emit and transmit rays of light with

¹¹**Light** is electromagnetic radiation between 400nm and 700nm. **Spectral power distributions** are functions from all real numbers between 400 and 700 (wavelengths of light) to real numbers between 0 and

identical spectral power distributions in some contexts, they might nonetheless have different spectral dispositions due to their behavior in other contexts.

There is a widespread misconception about spectral dispositions that is due, in part, to a widespread misunderstanding of Newton’s prism experiment. Many people take Newton’s experiment to show that, for instance, whether we experience purple is exclusively a function of the intensity of light around 400nm, whether we experience green is exclusively a function of the intensity of light around 550nm, whether we experience red is exclusively a function of the intensity of light around 650nm, and so on. As a result, they think that there is a straightforward way of lining up spectral dispositions and our color experiences. However, that is not the case; there is no direct correlation between whether we experience purple, green or red and the intensity of light reflected at just those wavelengths.

For instance, suppose two objects reflect light rays with the following spectral power distributions:



∞ (intensity). The **total illuminant** is a specification of the angle of incidence, spectral power distribution, and polarity of the light that reaches the object. The **spectral disposition** of a surface or volume is the function from each total illuminant to the spectral power distribution and polarity of the light rays reflected, transmitted or emitted by that surface or volume to each location.

Note that by defining light as electromagnetic radiation with wavelengths between 400nm and 700nm I am following convention (see any introductory physics textbook). One shortcoming of this convention when discussing color is that strong electromagnetic radiation at 350nm or 800nm will produce, respectively, a weak purple experience and a weak red experience. Fortunately, nothing important will hinge on which bandwidths are relevant.

Traditionally, conciliatory thinkers have focused exclusively on reflected light, setting aside transmitted and emitted light and ignoring the complications due to polarity and angle of incidence, at least in their official formulations. But I can’t identify a good reason for this tradition.

Despite the different spectral power distributions of the reflected light rays we might nonetheless have indistinguishable color experiences. (This is due to the way in which the cones in our eyes respond to light - see Hurvich [36].) Accordingly, the relationship between spectral dispositions and color experiences is more complicated than Newton's experiment has led many to believe.

This has an important consequence: when associating color experiences with collections of spectral dispositions we can't rely on natural, human-independent divisions of spectral dispositions. Instead, just as in the toy model the collector relied on causal relations between the scale's outputs (e.g. 'tiny') and weights, we apparently need to rely on causal relations between our visual experiences and spectral dispositions. More on this later.

4 First Claim: Dependence

There are two claims that are definitive of conciliatory thinking, a claim about dependence and a claim about structure. The claim about dependence is that the colors we represent depend on collections of spectral dispositions. But which collections? And what kind of dependence? In order to precisely and concretely state the problem for conciliatory thinking it will be helpful to focus on particular answers to these questions.

Before we start, it is necessary to introduce a new term. We might think of the colors as forming a tree that is analogous to the zoological tree. Near the top of the tree are properties like red, yellow and green. Further down are properties like magenta, scarlet, amber, mustard, lime and jade. And at the bottom are the **foundational colors**, which are the most specific colors. By definition, objects with the same foundational color must instantiate *all* the same colors; if one is scarlet then the other is scarlet, if one is red then the other is red, and so on.

The best way to develop the claim about dependence is to break it down into four steps. The first step is to decide how fine-grained to make the foundational colors (which is analogous to deciding how fine-grained to make species classifications in the zoological tree). There is a range of possibilities. Towards one extreme is the possibility that the foundational colors are as fine-grained as the spectral dispositions. Towards the other extreme is the possibility that they are as coarse-grained as the discriminations of the human visual system. To keep things simple, I will suppose that the foundational colors are as fine-grained as the spectral dispositions. Because the human visual system cannot discriminate colors that are as fine-grained as the spectral dispositions, I will also suppose that the colors we represent are *non-foundational* colors. While most conciliatory thinkers will accept these suppositions, some will deny them. As a way of streamlining the discussion I will treat their responses as objections and consider them later.

The second step is to decide which spectral dispositions correspond to foundational colors (which is analogous to deciding which organisms correspond to species in the zoological tree - e.g. viruses? bacteria?). I will suppose that *every* spectral disposition corresponds to a foundational color. While some might reject this supposition, I will subsequently argue that this would have surprising and counter-intuitive consequences.

The third step is to decide how higher-up colors like red relate to lower-down colors like scarlet (which is analogous to deciding how genus relates to species in the zoological tree). To make my discussion more concrete I will suppose that the non-foundational colors are *disjunctions* of the foundational colors.¹² For example, where ‘red₁₆₇’ etc. are names for foundational colors, scarlet might be the disjunction $\text{red}_{167} \vee \text{red}_{168} \vee \dots \vee \text{red}_{199}$. Nothing will depend on this supposition.

The final step is to decide how the colors depend on the spectral dispositions. One

¹²Alternatively, the non-foundational colors might be determinables (see Yablo [73]) or sets (see Hilbert [34]).

possibility is that the colors are *identical* to spectral dispositions (or their shared bases).¹³ Another possibility is that the colors *supervene* in some way on the spectral dispositions so that they co-vary without being identical.¹⁴ A third possibility is that the colors are “response-dependent” in that they are constituted by relations between objects and observers that hold in virtue of the spectral disposition of the object as well as the make-up of the observer.^{15,16} I will focus on the weakest kind of dependence because that will demonstrate that it is possible to generate the problem for any strength of dependence. Let’s say that a foundational color **minimally depends** on a spectral disposition just in case in the actual world something instantiates that spectral disposition if and only if it instantiates that color. Similarly, let’s say that a non-foundational color **minimally depends** on a disjunction of spectral dispositions just in case in the actual world something instantiates one of those spectral dispositions if and only if it instantiates that color.

With this background we can state the claim about dependence as we’ve developed it:

DEPENDENCE

The colors that we represent minimally depend on disjunctions of spectral dispositions.

In the next section I will introduce the second claim. I will then argue against the most straightforward strategies for demonstrating that these claims are compatible. I will subsequently argue that other ways of developing conciliatory thinking confront the same problem.

¹³For example, see Armstrong [1], [2], Averill [3], Byrne and Hilbert [10] [12], Dretske [22], Hilbert [34], Lycan [48], Smart [67] and Tye [70] [71]. Note that Smart thinks that it is contingent identity. Jackson and Pargetter [40], Jackson [38] and McLaughlin [57] think that the colors are identical to the shared bases of certain spectral dispositions.

¹⁴For example, see Broackes [7], Campbell [13], McGinn [56] and Yablo [73]. Yablo thinks that it is local and metaphysical supervenience and Campbell seems to think that it is global and nomological supervenience.

¹⁵For example, see Cohen [18], Evans [24], Johnston [41], Kripke [43], McGinn [55] and Peacocke [63].

¹⁶At least when endorsed by a conciliatory thinker, these claims do not perfectly line up with so-called “physicalism,” “realist primitivism” and “dispositionalism.” For example, Broad [8] and Cornman [19] should probably be classified as “realist primitivists” but they would deny that the colors depend on any physical property. Also, McDowell [54] should be classified as a “dispositionalist” but he never commits himself to the view that the colors depend on spectral dispositions. This demonstrates that conciliatory thinking is somewhat orthogonal to traditional classifications of color theories. See Pautz [62] and Byrne and Hilbert [11] for the traditional classifications.

5 Second Commitment: Structure

The second claim is easy to state but difficult to understand:

STRUCTURE

The colors that we represent preserve the structure of our color experiences.

In this section I will develop an understanding of STRUCTURE that brings out what I think is essential to conciliatory thinking. In subsequent sections I will use it to develop a problem for conciliatory thinking. In response, a conciliatory thinker might suggest that we should understand it in some other way. However, for reasons I will introduce later, I doubt that this will allow her to avoid the problem.

In order to understand what I mean by ‘represent’, ‘preserve’ and ‘structure’, it is necessary first to understand what I mean when I say that an experience “tells us” something. I am using ‘tells us’ to pick out a phenomenon that we can grasp pre-theoretically but that admits of many different theoretical glosses. It roughly lines up with what people are talking about when they talk about “how things look,” “how things visually appear” and what is “experientially presented.” For present purposes, I will just highlight two of its characteristics. First, what our perceptual experiences tell us can be inconsistent with our background beliefs. For instance, even if we believe that the lines in the Müller-Lyer illusion are the same length, an experience of that illusion will still tell us that the lines are different lengths. Second, what our experience tells us must be accessible to conscious introspection. For instance, the mere fact that a sub-personal module in our visual system contains information about a thing is not enough for our visual experiences to tell us about that thing.

I am using ‘represent’ to pick out a related phenomenon. If an experience tells us *that something is red* then that experience has an intentional relationship to redness.¹⁷ I am

¹⁷The nature of this relationship depends on one’s view about the content of what our experiences tell us. For instance, if it is a Russellian proposition then redness is a constituent of what our experience tells us. Alternatively, if it is a Fregean proposition then redness is picked out by a constituent of what our experience tells us.

picking out that relationship by saying that the experience *represents* redness.¹⁸ Note that, while I am using ‘tells us’ and ‘represents’ to pick out different relations, it will not be crucial that you keep track of that difference.

Let’s now use examples to clarify what it means to say something belongs to the “structure” of our experiences. Consider an experience of both a cherry and a lemon. Our experience might tell us that the cherry is a certain shade of red and that the lemon is a certain shade of yellow, in which case our experience will thereby tell us that there is a difference between the color of the cherry and the color of the lemon. We might think of the relationship between the first thing (that the cherry is a certain shade of red and that the lemon is a certain shade of yellow) and what our experience thereby tells us (that there is a difference between the color of the cherry and the color of the lemon) as due to the *structure* of our experience. We might characterize the structure of our experience: if one thing is that shade of red and another thing is that shade of yellow then they are different colors. Therefore, those shades preserve the structure of our experience only if they preserve the truth of this

¹⁸There are many ways to talk about visual experiences. Those who are more familiar with other ways of talking might benefit from an explicit comparison. First, some people use ‘represent’ in a more restrictive sense. For instance, there is a debate about whether experiences represent properties, acquaint us with properties, or are merely “raw feels” that dispose us to form beliefs about properties. Participants in this debate must be using ‘represent’ in a more restrictive way than I am because, as long as everyone agrees that our visual experiences support beliefs in a consciously accessible way, everyone should agree that experiences represent properties in my sense of ‘represent’.

Second, while I am talking about the *properties* and *objects* that our experiences represent, some people talk about the *propositions* that our experiences represent (which are the “representational contents” of those experiences). They are using ‘represent’ in a way that is closer to how I am using ‘tells us’. But we shouldn’t take it for granted that their ‘represent’ and my ‘tells us’ are synonymous. One reason why it might be useful to distinguish the two relationships is that our experiences’ “representational contents” might be impoverished compared to what our experiences tell us. For example, someone who thinks that the representational contents of visual experiences do not involve motions should still agree that our experiences tell us about motions because our experience support beliefs about motion. Likewise, someone who denies that it is ever part of the “representational content” of visual experiences that the similarity between red and orange is greater than the similarity between red and yellow might still agree that our experiences tell us about those similarities because our experiences support beliefs about them in the right way.

Third, there are many senses of ‘looks red’ and one of them might roughly line up with my use of ‘represents red’. However, an important difference is that ‘looks’ always picks out a relationship involving *a person* and I am using ‘represent’ to pick out a relationship involving *an experience*. Plus, there are certainly other senses of ‘looks’ that do not line up my use of ‘represent’ (see Jackson [37] on some of the different senses of ‘looks’).

conditional.

Because in this example the first thing seems to *entail* the second thing, one may be tempted to think about the structure of our color experiences as a kind of *inferential* structure. But that would be a mistake. We do not need to *infer* anything in order for our experience to tell us that there is a difference between the color of the cherry and the color of the lemon or that there is a color shared by the apple and tomato. It is built into the way we represent those colors.

For our next example, consider an experience of a dandelion, tangerine and sapphire. If our experience tells us that the dandelion is a certain shade of yellow, the tangerine is a certain shade of orange and the sapphire is a certain shade of blue then it *thereby* tells us that the color of the dandelion resembles the color of the tangerine more than the color of a sapphire. Therefore, we might characterize the structure of our experiences: if one thing is that shade of yellow and another thing is that other shade of orange then they have colors that resemble each other more than that shade of blue.

It would be ideal if we had a complete theory about the structure of our experiences. But we don't, and it's beyond the scope of this paper to develop one. How, then, are we supposed to determine what is and what is not included in this structure? There is no algorithm. We must often rely on introspection to deliver a verdict about particular cases. Consequently, as much as possible, we should focus our attention on clear cases. And that's exactly what I'll try to do in the next section.

As noted in the introduction, there are a handful of philosophers of color who will reject STRUCTURE. Some of them will reject it because they merely think that the structure of our experiences provides a "folk theory" of the colors, and as such they will be comfortable with the existence of mismatch between the colors and the structure of our experiences.¹⁹

¹⁹Matthen [50] thinks about things in this way. Apparently, Harman [31] does as well because he thinks that apparent color similarities only reveal something about our physiology. Lewis [46] and Jackson [38] are tough cases. On the one hand, they frame the discussion around our folk theory of color and admit that

Others will reject it because they think that it is merely the phenomenal character of our color experiences, and not the colors themselves, that satisfies the structure of our color experiences.²⁰ Notably, these are the same philosophers who aren't troubled by the similarity objection. And, as I've set things up, these philosophers are not conciliatory thinkers because conciliatory thinkers do not allow for violations of STRUCTURE.

6 What Conciliatory Thinkers Must Demonstrate

The two definitive claims of conciliatory thinking are:

DEPENDENCE

The colors that we represent minimally depend on disjunctions of spectral dispositions.

STRUCTURE

The colors that we represent preserve the structure of our color experiences.

These claims interact. An analogy might help illustrate their interaction. Suppose one friend tells you that he's vacationing in either Paris or Berlin and another friend tells you that she's vacationing in either Paris or Tokyo. While each friend told you something different, what they told you leaves it open whether they are vacationing in the same city. Consequently, they didn't *thereby* tell you that they will vacation in the same city. Now suppose that they instead listed the cities such that they *thereby* told you that they weren't vacationing in the same city. In that case, they must have listed completely different cities.

Let's now return to the cherry and the lemon. Our experience of the cherry and the lemon is such that it thereby tells us that there is a color difference between the cherry and the lemon. As a result, it is part of the structure of our experience that: if something is that shade of red and something else is that shade of yellow then they are different

the folk theory of color might need to be revised. But, on the other hand, they insist that some elements of our folk theory cannot be compromised, and among those elements are the kinds of conditionals that characterize the structure of our experiences, including the conditionals about color similarities.

²⁰See Kalderon [42] and Hellie [33] for examples.

colors. Building on this, these shades preserve the structure of our color experiences only if nothing instantiates both shades. Therefore, given our assumption that non-foundational colors are disjunctions of foundational colors, STRUCTURE implies that these shades are non-overlapping disjunctions of foundational colors. Given DEPENDENCE, it follows that the shades minimally depend on non-overlapping disjunctions of spectral dispositions.

As this example demonstrates, DEPENDENCE and STRUCTURE together place demands on the relationship between the colors and the spectral dispositions. It is not obvious that all these demands can be met. Consequently, a conciliatory thinker must demonstrate that we represent properties that *both* minimally depend on disjunctions of spectral dispositions *and* preserve the structure of our color experiences. And, as best I can tell, that's only possible if she takes a stand on what spectral dispositions the colors we represent depend on.

Let's clarify what this *doesn't* require. It doesn't require a conciliatory thinker to develop a theory about *how*, or *in virtue of what*, our experiences represent properties that minimally depend on disjunctions of spectral dispositions. For instance, a conciliatory thinker does not need to choose between the theories of Dretske [21], Jackson [39], Peacocke [63] and Shoemaker [66].

In the next section I will introduce and criticize what I take to be the most promising strategy for satisfying this obligation, which I will call the 'reliable production strategy'. After responding to anticipated objections in section eight, in section nine I will consider an alternative strategy that might seem promising.

7 Reliable Production Strategy

Because we're supposing that the foundational colors are as fine-grained as spectral dispositions, and spectral dispositions are extremely fine-grained, we're also supposing that our experiences represent disjunctions of foundational colors. Which disjunctions? A plausible

principle is: an experience represents a disjunction of all the foundational colors that *reliably produce* that experience. The **reliable production strategy** centers on this principle.

An attractive feature of the reliable production strategy is that it coheres with popular theories about *how* our experiences represent the colors. According to these theories, our experiences represent colors *in virtue of* this very kind of causal co-variation.²¹ These theories dovetail nicely with conciliatory thinking about color because we might think of them as exemplifying conciliatory thinking about representation insofar as they try to show how representational properties depend on physical properties. Nonetheless, even those who are skeptical of these theories of representation might still endorse the reliable production strategy.

Another attractive feature of the reliable production strategy is that it is at the heart of many attempts to undermine the similarity objection. In particular, if the reliable production strategy is conjoined with an influential theory about how our visual system produces color experiences then, as Byrne [9], Byrne and Hilbert [12], Bradley and Tye [6] and Cohen [17] have all argued, there apparently must be some similarity between the spectral dispositions that produce similar colors experiences, which potentially establishes that there is a similarity between the colors that depend on those spectral dispositions. The details of their arguments are complicated and, fortunately, unnecessary for our purposes. All that's relevant is that many philosophers use the reliable production strategy in their response to the similarity objection and therefore they cannot give up the reliable production strategy in response to my arguments without developing another response to the similarity objection.

The problem with the reliable production strategy, which I previewed in section two, is that it sometimes fails to preserve the structure of our experiences. Here's a rough sketch of why. When conjoined with another claim, the principle at the center of the strategy implies that there are colors that our experiences *never* represent. In particular, it implies that there

²¹See Dretske [21] and Tye [70], among others, and see Loewer [47] for a helpful overview of these theories.

is a color that falls between orange and yellow such that our experiences never represent it on its own or as part of a disjunction of colors. But I'll argue that it is part of the structure of our color experiences that some colors together "fill" a path through color space in that there are no intermediate colors missing. I will conclude that the reliable production strategy does not preserve the structure of our experiences.

Here's a more precise statement of the principle at the center of the strategy:

RELIABLE PRODUCTION

An experience represents a disjunction of some foundational colors if and only if in a certain type of person in a certain type of context (i) exposure to each of the relevant foundational colors reliably produces that experience, and (ii) no additional foundational colors reliably produce that experience.

This formulation leaves room for debate about the relevant type of person and the relevant type of context. For instance, one might focus on ordinary contexts and on ordinary people. But one might also focus on idealizations of ordinary contexts and on idealizations of ordinary people. Fortunately, we can abstract away from these differences. This formulation also leaves room for debate about where to draw the line between reliable and unreliable production. But as long as the line is at or above fifty-percent, it won't matter. (In the next section I will consider what happens if we drop the threshold below fifty-percent.)


As I said, RELIABLE PRODUCTION has an unacceptable consequence when combined with another claim. For reasons that I am about to introduce, I do not think a conciliatory thinker should deny the other claim. Consequently, I will take the unacceptable consequence to provide a *reductio* of RELIABLE PRODUCTION, which forces the conciliatory thinker to find another way to integrate her commitments.

The other claim is a consequence of the fact that our visual system is noisy. In our toy example involving the scale, it was easy to imagine how noise would manifest itself - the scale would sometimes display different letters even though it is weighing the same stone. It is therefore natural to wonder, How does noise in the visual system manifest itself? Close your eyes and closely attend to your experience. At least if you're like me, you'll notice small,

constant changes in your experience’s phenomenal character; your experience will seem to “shift” between experiences with slightly different phenomenal characters. This “shifting” does not go away just because you open your eyes. It just becomes easier to ignore, which is why it often takes practice to notice it. Among its many biological causes are rapid variations in blood pressure, fluctuations in the temperature of the retina and interference along neural pathways. The existence of visual noise is well-established and much-discussed in the scientific literature. To a quote a vision scientist nearly at random, Dennis Pelli [64, p.12] observes that, “noise arises in virtually all neural elements of the visual system from photoreceptor to cortex.”

The existence of visual noise has an important consequence. Suppose that an object with a certain foundational color (orange_{256}) reliably produces a certain experience ($E1$). If we gradually change the foundational color of that object then at some point it will start producing that experience less often. At that point it will correspondingly start to produce another experience ($E2$) more often, until we reach a foundational color (orange_{343}) that reliably produces that other experience ($E2$).

I think that it helps to represent the situation with a graph. Where $E1$ is represented by the dashed line and $E2$ is represented by the dotted line, we might depict the situation:



Because noise is asymptotic, there is some probability that orange_{300} will produce other experiences ($E3$, $E4$, etc.), which is why both lines drop below fifty percent. As a result,

there will always be a foundational color like orange₃₀₀ that does not produce any *particular* experience more than fifty percent of the time.²² Hence:

NOISE

For any observer and any context: If two foundational colors reliably produce different experiences, then there is a foundational color between them that does not reliably produce any particular experience (even if it will always produce some experience or another).

Some conciliatory thinkers might hope to undermine NOISE by getting rid of all the biological causes of noise in the relevant observer's visual system. One strategy is to choose a perfectly ideal observer whose visual system is noise-free. Another strategy is to specify the context *so precisely* that all the biological factors that contribute to visual noise, including the observer's blood pressure, are all specified.²³ However, these strategies fail because quantum indeterminacies *by themselves* produce enough noise to guarantee the truth of NOISE, at least if we assume that the relevant observer's visual experience cannot take infinitely long measurements.²⁴

Of course, a conciliatory thinker could just deny that the relevant context is governed by the laws of quantum mechanics, or even that the observer can't take infinitely long measurements, but I think it is clear that this medicine is worse than the disease. It might help to note that, to the best of my knowledge, no existing theories of visual representation depend on such exotic contexts or observers. It might also help to note that if the conciliatory thinker starts tinkering with the laws of nature in the relevant context then there's no

²²Things get slightly more complicated if there are uncountably many visual experiences. But, in that case, the probability that a foundational color causes any particular visual experience is zero. Therefore, anyone attracted to reliable production would need to divide visual experiences into countably many groups. Setting aside whether there is any privileged way to group them, we could just think of 'E1' and 'E2' as names for groups of experiences, in which case the argument proceeds as before. Also, I doubt that there are uncountably many color experiences because I doubt that the brain could generate so many different experiences.

²³McLaughlin [57] adopts a strategy like this to solve a different problem - the problem of variation.

²⁴This is a result that I've proven in a technical report [XXX]. The proof utilizes two facts. First, due to quantum indeterminacies, measurements of intensity will be at least as noisy as the Poisson distribution (see Fox [26, p.76-80]). Second, due to the quantum indeterminacies responsible for Doppler shifts, measurements of wavelength will be at least as noisy as the Cauchy-Lorentz distribution (see Dickie [20]).

guarantee that objects in that context will have the same kinds of spectral dispositions as actual objects, in which case our color experiences won't represent colors that are actually instantiated, and that is one of the primary motivations for conciliatory thinking.

At this point, one might wonder, Why are we taking it for granted that orange_{e300}, and the properties like it, are colors? Recall that we're assuming that the foundational colors are as fine-grained as the spectral dispositions, and that all spectral dispositions correspond to foundational colors, in which case orange_{e300} is a color because it corresponds to a spectral disposition. In the next section I will consider what happens if we dispense with these assumptions.

Let's now identify a consequence of RELIABLE PRODUCTION and NOISE. NOISE implies that there are foundational colors, like orange_{e300}, that do not reliably cause any particular experience. RELIABLE PRODUCTION then implies that there is no experience that represents a disjunction that includes any of these foundational colors. Where 'carrot' picks out the color represented by *E1* and 'amber' picks out the color represented by *E2*, we might depict the situation:

$$\overbrace{\dots \vee \text{orange}_{256} \vee \dots \vee \text{orange}_{285}}^{\text{carrot}} \dots \text{orange}_{300} \dots \underbrace{\text{orange}_{315} \vee \dots \vee \text{orange}_{343} \vee \dots}_{\text{amber}}$$

More generally, RELIABLE PRODUCTION and NOISE logically entail:

MISSING

For any two color shades that we represent, there is a foundational color between them such that our experiences never represent a disjunction that includes it.

I will now argue that a conciliatory thinker cannot accommodate MISSING. Because NOISE seems beyond reproach, I will conclude that conciliatory thinkers should abandon RELIABLE PRODUCTION. In this way, I hope to provide a *reductio* of RELIABLE PRODUCTION, thereby forcing the conciliatory thinker to find another way to integrate her commitments.

MISSING will embarrass many conciliatory thinkers because it has the incredible consequence that there is a color between orange and yellow that is "invisible" in the sense

that we could never even represent a disjunction that includes that color. For example, an orange₃₀₀ wall would always be misperceived as carrot or amber. MISSING might even be incompatible with commitments shared by many conciliatory thinkers. For instance, it might be incompatible with the view that we can know everything about the nature of the colors by mere reflection on our visual experiences (e.g. see Cambell [13], Strawson [68] and to some extent Johnston [41]). It might also be incompatible with the view that ϕ is a color only if things sometimes look ϕ to us (e.g. see Peacocke [63], McGinn [55]). Nonetheless, I will identify a problem with MISSING that does not depend on whether one endorses these views; I will argue that it implies that the colors we represent *do not preserve the structure of our experiences*. I will conclude that conciliatory thinkers must reject MISSING, in which case they must also reject RELIABLE PRODUCTION.

As noted previously, we can think of the colors as filling a **color solid**, which is a three-dimensional solid where the axes correspond to hue, brightness and saturation. I will talk about lines through that solid as **paths through color space**.²⁵ Holding this paper at arm's length, consider your experience of the following strip:



Strip that transitions from orange to yellow

²⁵I'm ignoring a complication. If foundational colors are as fine-grained as spectral dispositions then they will fill an infinite dimensional space. As a result, there will be many different projections from that space to three dimensional spaces, leaving it unclear which if they are projections to color solids. Part of the problem is that there is controversy about how to give 'hue', 'brightness' and 'saturation' a precise physical meaning (see Land [44], Hulbert [35]), and therefore, because we're assuming that every spectral disposition corresponds to a foundational color, it is unclear which changes in foundational color correspond to changes in hue, brightness and saturation. Fortunately, my argument only requires that the relevant projection have the following property: letting x , y , and z be locations in the infinite dimensional space, and letting x' , y' , and z' be their projections in the color solid: if $x < y < z$ then $x' < y' < z'$. Note that if there were no such projection then the colors would not preserve the similarity orderings that are a part of the structure of our color experiences, and therefore conciliatory thinking would be a non-starter.

Your experience represents a series of colors, including (let's say) amber and carrot. I claim that your experience tells you that these colors together "fill" a path through color space so that there are no missing colors. This is not to say that your experience individually represents each intermediate foundational color. Given the number and grain of the foundational colors, that would be implausible. More plausibly, there just seem to be no gaps between amber, carrot and the other non-foundational, disjunctive colors that you represent.²⁶ This claim is supported by introspection. It might help to try to picture an intermediate color that is absent from the above strip. If you're like me, it's impossible. Of course, this doesn't *prove* your experience tells you that amber, carrot, etc. together fill a path through color space. But it is a consideration.

There is a useful parallel with another kind of experience. Consider your experience if you were to look at a brick under normal conditions. Your experience would tell you that the brick "fills" some region of space so that if the brick occupies two points then it also occupies all the intermediate points (at least in the sense in which a block of swiss cheese fills its own holes). This is not to say that your experience individually represents each point as occupied. Given the number and grain of those points, that would be implausible. More plausibly, there just seem to be no gaps between the sub-regions you represent. This claim is supported by introspection. It might help to picture an ordinary brick. Now try to imagine that it fills more of its volume. If you're like me, it's impossible.

I also claim that, in virtue of what your experience of the strip tells you about amber, carrot, etc., your experience *thereby* tells you that these colors together fill a path through color space. Consider that that it doesn't seem possible for your experience to represent the

²⁶Note that 'fill' is not synonymous with 'dense'. An analogy might help. One might represent the intervals $[0,1]$ and $[1,2]$ as together filling $[0,2]$ without representing any of these intervals as dense because one might not represent any relations between the points in the intervals. Additionally, while 'fill' picks out a *relation* between intervals, 'dense' picks out a monadic property of an interval. For similar reasons, 'fill' is not synonymous with 'continuous'. For example, unlike 'fill', 'continuity' picks out a property of a function from one interval to another.

same colors in the same order but fail to tell you that these colors together fill a path through color space. Consider also that you do not need to first check each color to make sure it is included; your experience immediately tells you that the colors together fill a region of color space - it seems to be built into your representation of amber, carrot, etc.²⁷

There is another helpful parallel with the brick experience. In virtue of what your experience would tell you about the sub-regions filled by the brick's parts, your experience would *thereby* tell you that the entire brick fills a spatial region. Consider that it doesn't seem possible for your experience to tell you that the brick's parts remain in the same position but the brick as a whole fails to fill some region. Consider also that you do not need to first check each point in each region; your experience would immediately tell you that the brick fills some region of space - it seems to be built into your representation of the brick's parts.

Note I'm not describing a *necessary* condition for having an experience of colors as filling a path through color space. For all I've said, it might be possible to have that kind of experience without representing any particular colors like amber or carrot. Instead, I'm describing a *sufficient* condition: if you represent amber, carrot, etc. in that order then you will have an experience of colors as filling a path through color space.

Putting this all together, it is part of the structure of your experience that: if amber, carrot, etc. are ordered in a certain way then these colors together fill a path through color space. Amber and carrot satisfy that conditional only if they are *adjacent* or *overlapping* regions of color space. Therefore, STRUCTURE implies that amber and carrot really are adjacent or overlapping regions of color space.

²⁷What exactly does your experience tell you about the strip's coloring? There is room for two views. On the one hand, your experience might represent parts of the strip in a way that does not distinguish the colors within them. From the perspective of your experience, each part might be completely homogenous with respect to color. For example, your experience might represent one part as homogeneously amber and another part as homogeneously carrot. But, on the other hand, it is also consistent with the possibility that your experience represents each part as continuously transitioning in color, so that, for example, one part is represented as continuously transitioning from amber to carrot. See Fara [25, 18-25] for a discussion of the tension between supposing that our experience represents continuous transitions and supposing that our experience represents homogenous regions.

The problem is that MISSING is incompatible with this implication. MISSING implies that amber and carrot are *not* overlapping or adjacent regions; MISSING implies that there is a color like orange₃₀₀ that is neither amber nor carrot. Therefore, MISSING implies that the colors you represent do not preserve the structure of your experience. As a result, conciliatory thinkers must reject MISSING - it is incompatible with STRUCTURE.

Recall that MISSING is a consequence of RELIABLE PRODUCTION and NOISE. Because NOISE follows from some extremely plausible assumptions and some well-known scientific facts, a conciliatory thinker should give up RELIABLE PRODUCTION. The conciliatory thinker is therefore forced to find another strategy for integrating her commitments.

Let's review. Conciliatory thinkers claim that the colors depend on physical properties. In particular, they think that the colors depend on what I'm calling 'spectral dispositions'. For reasons introduced in section five, they must show that the colors we represent both minimally depend on spectral dispositions and preserve the structure of our experiences. A promising strategy for satisfying this obligation is to appeal to a principle like: an experience represents a disjunction that includes all the foundational colors that reliably produce that experience. However, this strategy fails as the result of two considerations. On the one hand, due to the structure of our experiences, for each foundational color between orange and yellow we must at least be able to represent a *disjunction* that includes it. But, on the other hand, any principle of representation that centers around *reliable production* will preclude us from representing disjunctions that include foundational colors like orange₃₀₀. Conciliatory thinkers must therefore find another strategy.

One of the features of my argument that is worth emphasizing is that it applies to *all* contexts governed by the laws of quantum mechanics and *all* observers. Consequently, it is a problem even if one thinks that color representation is highly context-sensitive.

It is also worth re-emphasizing several features of my argument. One such feature is that my argument focuses on the weakest kind of dependence: extensional correspondence. It

therefore establishes a problem for conciliatory thinking regardless of how the colors depend on the spectral dispositions; among other possibilities, the colors might be identical to them, metaphysically necessitated by them, nomologically necessitated by them, or have them as their dispositional grounds.

Another feature of my argument worth re-emphasizing is that, unlike other arguments in the literature, it does not assume that the structure of our color experiences reveals anything *essential* about the colors. For instance, the structure might just be characterized by material conditionals (which permit there to be counterfactual worlds where the antecedent is true but the consequent is false.) In that case there is no reason to think that the structure reveals anything essential about the colors. Equivalently, descriptions like “Plato was the author of *The Republic*” might fix the referent of ‘Plato’ without revealing anything essential about Plato.

A third feature of my argument worth re-emphasizing is that a conciliatory thinker cannot respond merely by adapting standard responses to other problems. In particular, she cannot merely adapt the standard response to the similarity objection because that response hinges on RELIABLE PRODUCTION, which is one of the premises that I used to generate the problem.

In the next section I will consider eight responses and in the penultimate section I will consider an alternative strategy.

8 Responses and Consequences

Let’s begin by clarifying the dialectical situation. My argument is a *problem* for the reliable production strategy, not a *refutation*. One can respond by giving up conciliatory thinking, giving up the reliable production strategy or rejecting one of the assumptions that generate the problem. Our goal is to determine which option is best. Therefore the mere fact that it is *possible* to reject one of the assumptions that generate the problem has little signifi-

cance - after all, in philosophy it's possible to deny any assumption. What is significant is whether one can deny these assumptions without committing oneself to something equally problematic. I will argue that one cannot.

The first response targets my analysis of the experience of the color strip. I claim that your experience tells you: amber, carrot, etc. together fill a path through color space. It is tempting to respond that your experience only tells you something weaker: that amber, carrot, etc. together fill a path through the space *of all the colors that you can represent*. I think that this response is unattractive for two reasons. First, except perhaps in atypical circumstances, it is implausible to suppose your color experiences tell you about other color experiences, and, building on this, it is implausible to suppose that your color experiences tell you what can be represented by other color experiences, yet that's exactly what would be happening if your experience told you about the space of all the colors that you can represent. Second, it seems unmotivated. One might think it is motivated by the thought that if something is *in principle* not individually representable then it should be excluded from the domain of all our representations. But the missing fundamental color is *in principle* representable; just suppose that the relevant observer's experiences were calibrated so that orange₂₅₆ is the missing color instead of orange₃₀₀. In that case we would represent a non-foundational color that has orange₃₀₀ among its disjuncts. Further, our visual experiences presumably can't represent individual spatial points or individual foundational colors, yet we presumably can still represent collections of spatial points and collections of foundational colors. Thus, the fact that we cannot individually represent some foundational colors does not imply that they should be excluded from the domain of all our representations, which undermines the most natural motivation for this objection. A conciliatory thinker should therefore try to develop another response.

The second response is related. I claim that your experience tells you something about the relationships between amber, carrot and other colors. It is tempting to respond that

your experience merely tells you something about the relationships between the *phenomenal characters* of amber experiences, carrot experiences and other color experiences. In that case, perhaps your experience merely tells you that the phenomenal characters of these experiences together fill a path through *phenomenal space*. I think that this response is introspectively implausible; it seems clear to me that your experience tells you something about the relationships between amber, carrot, etc. Analogously, if you experience two objects as close together then your experience seems to be telling you something about those objects, and doesn't seem to be merely telling you something about your own phenomenology. Additionally, and perhaps most importantly, it is dialectically unsatisfying for a conciliatory thinker to insist that your experience only tells you something about your own experience's phenomenal characters. Just take anything that appears to be part of the structure of your color experiences. If in response to any incompatibility a conciliatory thinker can just respond that it is really a part of the structure of your experience's phenomenal characters then that would *trivialize* STRUCTURE. Yet STRUCTURE is supposed to be a substantive and interesting claim.

The third response targets the reasoning that led us to MISSING. For concreteness, let's again focus on orange₃₀₀. Orange₃₀₀ objects do not reliably produce amber experiences. Orange₃₀₀ objects also do not reliably produce carrot experiences. But orange₃₀₀ objects reliably produce *either* an amber experience *or* a carrot experience. That is, orange₃₀₀ objects reliably produce amber-or-carrot experiences. One might then think that RELIABLE PRODUCTION implies that amber-or-carrot experiences represent a disjunction that includes orange₃₀₀, which, when generalized, would establish that MISSING is false. The problem with this response is that it seems that the only way to have an amber-or-carrot experience is to have an amber experience or to have a carrot experience. If that's right then if neither amber experiences nor carrot experiences represent a disjunction that includes orange₃₀₀ then amber-or-carrot experiences do not represent a disjunction that includes orange₃₀₀.

Of course, a conciliatory thinker might just insist that one can have an amber-or-carrot experience without having either an amber experience or a carrot experience. The trouble is that it is hard to make sense of this possibility, at least if we're thinking about color experiences in the usual way. Perhaps there is another, better way of think about color experiences. But I'm skeptical; I suspect that any alternative will be unacceptably counter-intuitive. Regardless, if conciliatory thinkers were able to develop an alternative then that would a surprising and far-reaching outcome of my arguments.

The fourth response targets one of my background assumptions: that every spectral disposition corresponds to a foundational color. This assumption helps ground my claim that orange₃₀₀ is a color, which was a crucial premise in my argument for MISSING; consider that if orange₃₀₀ is not a color then my arguments do not establish that there is an intermediate color that is unrepresentable.

The best way to bring out the problem with this response is to consider a wall that is uniformly orange₃₀₀. Is that wall colored? Intuitively, it is. Consider an object that gradually transitions from orange₂₅₆ to orange₃₄₃. Intuitively, it is always colored; at no moment is it colorless. Therefore, intuitively, it is still colored even when it has a spectral disposition that will not reliably produce any particular color experience. In support of this intuition, keep in mind that even at those moments it still reliably produces some color experience or other, and therefore always looks colored. For example, a wall with that spectral disposition would always look carrot or amber, and it is extremely counter-intuitive to say that the wall is nonetheless colorless. Perhaps some will be willing to bite this bullet. But I suspect that most won't.

A conciliatory thinker might try to accommodate the intuition that the wall is colored. In particular, she might claim that the wall is colored without being any specific shade. The trouble is that it is hard to make metaphysical sense of this claim, at least at the level of the foundational colors. It seems just as paradoxical as the claim that there are polygons that

are not three-sided, four-sided, five-sided, and so on.

A conciliatory thinker might try to make metaphysical sense of her claim by invoking the concept of indeterminacy. Now letting ‘carrot’ and ‘amber’ pick out the most foundational colors that the wall might instantiate, a conciliatory thinker might claim that while it is determinate that the wall is orange, it is *indeterminate* whether it is amber or carrot. However, I don’t think that it is helpful to invoke indeterminacy of this kind. It is important to keep in mind that what’s at issue here is whether there is some indeterminacy in what properties are *instantiated*. What’s not at issue is whether there is some indeterminacy in what we *represent*, which is why standard theories of indeterminacy, such as those advanced by supervenientists, won’t help.

At bottom, the problem is that it is hard to make sense of how property instantiations could be indeterminate. For example, one possibility is that, just as some think that *parthood* comes in degrees so that some rocks are part of Mount Kilimanjaro to degrees less than one, perhaps *instantiation* also comes in degrees so that there are objects that instantiate both amber and carrot to degrees less than one. However, it is unclear whether this way of talking describes a metaphysically coherent possibility. For example, if we consider shapes, it is unclear what it would be for an object to instantiate squareness to a degree less than one, and if we consider mass, it is unclear what it would be for an object to instantiate ten milligrams to a degree less than one. Perhaps a conciliatory thinker can convince us that instantiation is all-or-nothing for shapes and masses while it comes in degrees for the colors. But that’s a heavy burden; it is hard to see what could ground such an asymmetry.

Another possibility is that, just as some think that clouds are “vague objects” in that it is indeterminate whether a given cloud is identical to a certain collection of water droplets, perhaps colors are “vague properties” in that, for example, it is indeterminate whether the wall’s color is identical to scarlet. One problem with this route is that a conciliatory

thinker would then need to overcome the standard objections to indeterminacy of this kind.²⁸ Another problem is that it is unclear how to think about vague properties at the foundational level. For example, we might naturally think of vague colors as disjunctions of foundational colors such that a foundational color can be a disjunct to a degree less than one. But that only gives us a way to think about *non-foundational* colors that are vague because only non-foundational colors are disjunctions. It does not give us a way to think about *foundational* colors that are vague, and that's what's needed.²⁹

Stepping back, I suspect that most conciliatory thinkers won't want their theory of color to depend on deep and controversial views about the nature of properties of their instantiations. Nonetheless, it would be a surprising, interesting and far-reaching outcome of my arguments if they led conciliatory thinkers to walk down either of these dark allies.

The fifth response targets another of my background assumptions. I assume that the foundational colors are as fine-grained as the spectral dispositions. This assumption also helps ground my claim that orange₃₀₀ is a color; consider that otherwise an orange₃₀₀ object might be the same color as an orange₂₈₅ object, and in our example orange₂₈₅ objects reliably cause a specific experience, in which case there is no intermediate color that is unrepresentable because RELIABLE PRODUCTION implies that we can represent the color of orange₂₉₀ objects.

One might reject this assumption. However, that would commit one to the view that colors are species-relative, which is a view many philosophers reject. To appreciate why rejecting my assumption has that consequence, consider any two objects with different spectral dispositions. Regardless of which objects we are considering, there is some possible creature that could reliably detect the difference between them,³⁰ and if colors are not species-relative

²⁸See Evans [23] and Lewis [45] on vague objects. Their arguments straightforwardly generalize to properties.

²⁹See Smith and Rosen [65] for a more sophisticated proposal that is unavailable to the conciliatory thinker for a similar reason.

³⁰This is a near-trivial consequence of signal detection theory. Just consider that if they reflect light

then we should think of this creature as discriminating the colors of the objects. More generally, if the colors are not species-relative then the foundational colors must be as fine-grained as spectral dispositions.

Some conciliatory thinkers will be happy with the implication that colors are species-relative, so let's briefly consider their view.³¹ If we develop it in the most natural way, what makes an object amber-for-humans is that it reliably produces amber experiences in humans. Likewise for carrot-for-humans, peach-for-humans, rust-for-humans and so on. In that case, objects that do not reliably cause any particular experience are not amber-for-humans, carrot-for-humans, peach-for-humans, rust-for-humans and so on. That leaves the conciliatory thinker in the same place as before; either she must deny that these objects are colored (which is counter-intuitive) or she must claim that they are colored without being any particular shade (which might not be metaphysically intelligible). Alternately, she might give some other account of what makes an object amber-for-human, though it isn't clear what such an account would look like.

The sixth response is that I've misunderstood STRUCTURE because, on a correct understanding, the fullness of some transitions in color is not part of the structure of our color experiences. However, when we think about the structure of our color experiences we are naturally led to think that it is captured, at least in part, by traditional color solids. Moreover, it seems to be constitutive of traditional color solid that not only do distances along a path correspond to degrees of similarity, but that all the paths through them are full, so that there are no missing shades. Equivalently, when we think about the structure of our spatial experiences, we are naturally led to think of it as something captured, at least in part, by

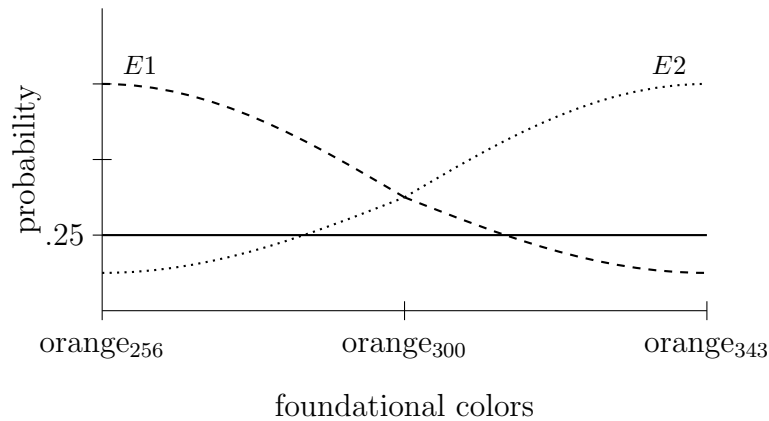
differently then, where d' is the distance between the peak of the two probability distributions at some wavelength, $d' \neq 0$. And if $d' \neq 0$ for two signals then there is always some detector that can reliably discriminate them. For more about signal detection theory, see Gescheider [27] and Heeger [32].

³¹See Cohen [18], Jackson and Pargetter [40], Kalderon [42] and McLaughlin [57]. For simplicity, I'm overlooking the fact that most color relativists think that the colors are relative to particular observers and particular contexts.

empty volumes, and it seems to be constitutive of these volumes that paths through them are full, so that there are no missing points. It might help to imagine a color solid in which each color is set apart from the others, like stars in the night sky. Intuitively, that color solid leaves something out; for lack of a better term, we might say that it fails to capture the *unity* of the colors. As I hope these considerations make clear, the similarity orderings and the fullness of some transitions go hand-in-hand, and it is typically taken for granted that similarity orderings belong to the structure of our color experiences. Thus, even if there is some other way of understanding STRUCTURE, I doubt that it will help the conciliatory circumvent the problem.

The seventh response is really just a modification of the reliable production strategy. The failure of the reliable production strategy is ultimately due to the existence of foundational colors that do not *reliably* cause any particular experience. A natural way to avoid this problem is to lower the threshold below fifty-percent.³² The problem with this modification is that the colors will no longer preserve the similarity orderings that are part of the structure of our color experiences. Let's again use a more abstract perspective to clarify the relevant issues. Consider the same series of foundational colors as before, orange₂₅₆ ... orange₃₄₃, and the same two experiences as before, *E1* and *E2*, but let's now lower the threshold:

³²I'm assuming for concreteness that for each color experience there is some foundational color that reliably produces it. However, nothing hinges on this assumption. Just note that the problems with the reliable production and this modification form a dilemma for any choice of thresholds: the threshold will either be so high that some foundational colors will fail to cross it for any experience, or the threshold will be so low that some foundational colors will cross it for two experiences.



A consequence of lowering the threshold is that these experiences represent overlapping disjunctions of foundational colors. Let's again use 'carrot' so that it picks out the color represented by $E1$ and 'amber' so that it picks out the color represented by $E2$. In that case we can depict the relationship between these colors:

$$\underbrace{\dots \vee \text{orange}_{256} \vee \dots \vee \text{orange}_{300} \vee \dots}_{\text{carrot}}$$

$$\underbrace{\dots \vee \text{orange}_{300} \vee \dots \vee \text{orange}_{343} \vee \dots}_{\text{amber}}$$

Here's the problem. Suppose you are having an amber experience of two paint chips and a carrot experience of a third paint chip. Your experience will thereby tell you that the color of the first paint chip is more similar to the color of the second paint chip than to the third paint chip. As a result, it is part of the structure of your experience: if two things are amber then they have colors that resemble each other more than carrot. The problem is that this conditional is false if we lower the threshold. Just note that if something is orange_{300} then it is both carrot and amber, in which case there is something that is amber that resembles carrot things just as much as other amber things.

As noted in the introduction, some people think that colors cannot be *identical* with disjunctions of spectral dispositions because the relevant disjunctions will be too heterogenous

to preserve the similarity orderings among the colors. These people think that this is a decisive problem for color physicalism and they use it to motivate other variants of conciliatory thinking. However, we've just established the surprising result that if a conciliatory thinker lowers the threshold below fifty-percent then she cannot preserve the similarity orderings among the colors even if the colors just minimally depend on spectral dispositions.

The eighth response is to invoke representational indeterminacy. Because there are so many theories of representational indeterminacy, and because it is hard to see how invoking representational indeterminacy could help the conciliatory thinker, this is a difficult response to evaluate. As a preliminary, I think it helps to note that many of the standard theories of representational indeterminacy were developed to undermine sorites paradoxes and, despite the existence of superficial similarities, like talk about where to place a cut-off in a series, the problem for conciliatory thinking is not an instance of a sorites paradox. To appreciate why, consider that there's no analog to DEPENDENCE or STRUCTURE in sorites paradoxes. Consider also that the problem for conciliatory thinking does not depend on a so-called 'sorites premise', which would have a form like: if a man with n hairs on his head is bald then a man with $n + 1$ hairs on his head is bald.

In any case, there are reasons to doubt that an appeal to representational indeterminacy will help the conciliatory thinker. Returning to our example, suppose that it is indeterminate which of the following properties a given experience represents:

$$\begin{array}{c} \text{carrot} \\ \underbrace{\dots \vee \text{orange}_{256} \vee \dots \vee \text{orange}_{285}} \\ \dots \vee \text{orange}_{256} \vee \dots \vee \text{orange}_{285} \vee \text{orange}_{286} \\ \underbrace{\hspace{10em}} \\ \text{carrot}' \end{array}$$

The problem is that, at least for most color experiences, it doesn't just seem that *there is some color that we're representing*. It also seems that we know which color it is, at least under one of its modes of presentation. Yet if it is indeterminate whether our experience is representing carrot or carrot' then we presumably can't know which color we're representing.

That's highly counter-intuitive.

Another problem is that if one accepts RELIABLE PRODUCTION then, even if there is some indeterminacy in which properties our experiences represent, it doesn't follow that there are no "missing" colors. Returning again to our example, it is determinate that orange₃₀₀ doesn't reliably produce any of our experiences. Therefore, given RELIABLE PRODUCTION, it is determinate that none of our experiences represent a disjunction that includes orange₃₀₀. As a result, even if there is some indeterminacy in which properties our experiences represent, perhaps as a result of indeterminacy in the extension of 'certain observer' and 'certain context', that doesn't automatically solve the problem unless one gives up RELIABLE PRODUCTION. And, as we've seen, it is unclear what could replace RELIABLE PRODUCTION.

There's another reason to be worried about the efficacy of this response. In our discussion of the reliable production strategy we established that if we draw the cut-off for representation above fifty-percent then the colors we represent will not satisfy the structure of our color experiences because amber, carrot, etc. will not fill a path through color space. In our discussion of the seventh response we then established that if we lower the cut-off for representation then the colors we represent will not satisfy the structure of our color experiences because amber, carrot, etc. will not have the correct similarity ordering. Therefore, given what we've established, if it was indeterminate where to draw the line then you'd expect it to be indeterminate whether amber, carrot, etc. fill a path through color space *and* indeterminate whether amber, carrot have the correct similarity ordering. And that certainly wouldn't help the conciliatory thinker.

Perhaps a conciliatory thinker can develop a clever, new theory of representational indeterminacy that sidesteps these problems. But it should at least embarrass the conciliatory thinker that her ability to integrate her commitments depends on it.

I just considered eight responses. There are doubtless other ways to defend or modify the reliable production strategy, and there might even be a way to develop one of these responses

into something less problematic. Regardless, it should be clear that the reliable production strategy and its most straightforward modification have a serious problem. Let's therefore investigate another strategy.

9 Alternative Strategy

The reliable production strategy relied on the probabilities linking color experiences and foundational colors to pick out the relevant color disjunctions. A natural response to the problem with this strategy is to instead include the probabilities in what is represented. The most plausible way to do this is to switch our focus from the probability that exposure to a foundational color will produce a certain experience (a likelihood approach) to the probability that an experience was produced by a certain foundational color (a posterior probability approach). The motivation for this strategy is the thought that the *causal information* a color experience provides about the environment is provided by these probabilities and, moreover, what a color experience tells us corresponds to the causal information it provides.

Let's develop this strategy with an example. Consider an arbitrary color experience as well as foundational colors orange_{256} , orange_{257} , etc. Suppose the probability an instance of this experience was produced by orange_{256} is .02, the probability it was produced by orange_{257} is .05, and there are similar probabilities for each of the other foundational colors. In that case, the information provided by that experience is: the probability that the relevant object has foundational color orange_{256} is .02, the probability that the relevant object has foundational color orange_{257} is .05, and so on. As a result, according to this strategy, the relevant experience represents a property like *having a probability of .02 of being color orange_{256} , having a probability of .05 of being color orange_{257} , etc.*

There are two ways to think about this property. According to the first way, we should think of *having a probability of .02 of being orange_{256} , etc.* as itself a color. Perhaps it is a

special kind of weighted disjunction. However, that would have absurd consequences. On the one hand, if we think of the probabilities as corresponding to *subjective* probabilities then it has the incredible consequence that if I change my credences about the foundational color of an object then the object thereby changes its color. On the other hand, if we think of the probabilities as corresponding to *objective* probabilities then, except in abnormal situations, nothing instantiates these properties. For instance, if an object is orange₂₅₆ then, if anything, it has the rival property *having a probability of 1.00 of being orange₂₅₆*. But one of the primary motivations for conciliatory thinking is the conviction that we represent colors that are actually instantiated.

The second way of thinking about this property is that it is not itself a color. In that case, the experience does not represent a non-foundational color that includes orange₂₅₆ and orange₂₅₇ as disjuncts. Instead, it merely tells us something about the probability that the object has various colors. However, that's implausible because at least some color experiences do not tell us anything *probabilistic*. Just look around. If your experiences are like mine then they'll tell you *which* surfaces have *which* colors, and there's nothing probabilistic about that. For instance, you might be tempted to say of an object's color, "That's my favorite color." But you'll rarely be tempted to say things like, "Look at one of the possible colors of this object" or "That object has some probability of having my favorite color." As this indicates, if we think about the property in this way that it will not preserve the structure of your experiences. For example, suppose you are having a uniform experience of a wall. Your experience might thereby tell you that the entire wall is the same color. But, if your experience merely represents the probability that each part of the wall is orange₂₅₆, orange₂₅₇, etc., and does not represent any less foundational color, then it would not thereby tell you that the entire wall is the same color. Instead, it would tell you the probability that the wall is the same color is less than one because there will always be some probability that the left-half is orange₂₅₆ and the right-half is orange₂₅₇. As a result, the colors you represent

would not preserve the structure of your color experiences.

I conclude that this strategy has significant problems.

10 Conclusion

The definitive claims of conciliatory thinking are the claim about structure (that the colors preserve the structure of our experiences) and the claim about dependence (that the colors depend on physical properties). The most straightforward strategy for integrating these claims is the reliable production strategy, but I argued that this strategy is incompatible with the claims it was supposed to integrate. How should the conciliatory thinker respond? If they want to remain conciliatory thinkers they must either replace the reliable production strategy or reject one of the assumptions that help generate the incompatibility. However, there is no obvious replacement for the reliable production strategy and rejecting any of the assumptions would at best have surprising, far-reaching and counter-intuitive consequences. That leaves the conciliatory thinker in a difficult spot.

References

- [1] Armstrong. *A materialist theory of the mind*. Routledge and K. Paul, 1968.
- [2] Armstrong. Smart and secondary qualities. In Pettit, Sylvan, and Norman, editors, *Metaphysics and Morality: Essays in Honour of J.J.C. Smart*. Blackwell Publishers, 1987. Reprinted in [?].
- [3] Averill. Color and the anthropocentric problem. *Journal of Philosophy*, 82:281–304, 1985. Reprinted in [?].
- [4] Boghossian and Velleman. Colour as a secondary quality. *Mind*, 98:67–106, 1989. Reprinted in [?].
- [5] Boghossian and Velleman. Physicalist theories of color. *Philosophical Review*, 100:81–103, 1991. Reprinted in [?].

-
- [6] Bradley and Tye. Of colors, kestrels, caterpillars, and leaves. *Journal of Philosophy*, 98:469–487, 2001.
- [7] Broackes. The autonomy of color. In Charles and Lennon, editors, *Reduction, Explanation, and Realism*. Oxford University Press, 1992. Reprinted in [?].
- [8] Broad. *Perception, Physics and Reality; An Enquiry into the Information that Physical Science can Supply about the Real*. Cambridge University Press, 1914.
- [9] Byrne. Color and similarity. *Philosophy and Phenomenological Research*, 66:641–665, 2003.
- [10] Byrne and Hilbert. Colors and reflectances. In Byrne and Hilbert, editors, *Readings on Color: The Philosophy of Color*, volume 1. Oxford University Press, 1997.
- [11] Byrne and Hilbert. Introduction. In Byrne and Hilbert, editors, *Readings on Color: The Philosophy of Color*, volume 1. Oxford University Press, 1997.
- [12] Byrne and Hilbert. Color realism and color science. *Behavioral and Brain Sciences*, 26:3–21, 2003.
- [13] Campbell. A simple view of color. In Haldane and Wright, editors, *Reality, Representation, and Projection*. Oxford University Press, 1993. Reprinted in [?].
- [14] Campbell, editor. *Reference and Consciousness*. Oxford University Press, 2002.
- [15] Chalmers. Perception and the fall from eden. In Gendler and Hawthorne, editors, *Perceptual Experience*. Oxford University Press, 2006.
- [16] Cohen. Color: A functionalist proposal. *Philosophical Studies*, 113:1–42, 2003.
- [17] Cohen. On the structural properties of the colors. *Australasian Journal of Philosophy*, 81:78–95, 2003.
- [18] Cohen. Color properties and color ascriptions: A relationalist manifesto. *The Philosophical Review*, 113:451–506, 2004.
- [19] Cornman. *Perception, Common Sense, and Science*. Yale University Press, 1975.
- [20] Dicke. The effect of collisions upon the doppler width of spectral lines. *Physical Review*, 89:472–473, 1953.
- [21] Dretske. *Knowledge and the Flow of Information*. MIT Press, 1981.
- [22] Dretske. *Naturalizing the Mind*. MIT Press, 1994.
- [23] Evans. Can there be vague objects? *Analysis*, 38:208, 1978.

-
- [24] Evans. Thing without the mind - a commentary upon chapter two of strawson's *individuals*. In Straaten, editor, *Philosophical Subjects: Essays Presented to P. F. Strawson*. Oxford University Peess, 1984.
- [25] Fara. Phenomenal continua and the sorites. *Mind*, 110:905–935, 2001. Published under the name “Delia Graff”.
- [26] Fox. *Quantum Optics: an introduction*. Oxford University Press, 2006.
- [27] Gescheider. *Psychophysics: The Fundamentals*. Lawrence Erlbaum Associates, 1997.
- [28] Hardin. *Color for Philosophers: Unweaving the Rainbow (expanded edition)*. Hackett, 1988.
- [29] Hardin. A spectral reflectance doth not a color make. *Journal of Philosophy*, 100:191–202, 2003.
- [30] Harding. Color and the mind-body problem. *Review of Metaphysics*, 45:289–307, 1991.
- [31] Harman. Explaining objective color in terms of subjective responses. In Villanueva, editor, *Philosophical Issues 7*, pages 1–17. Ridgeview Publishing, 1996. Reprinted in [?].
- [32] Heeger. Signal detection theory. Available at <http://www.cns.nyu.edu/~david/handouts/sdt/sdt.html>.
- [33] Hellie. Noise and perceptual indiscriminability. *Mind*, 114:481–508, 2005.
- [34] Hilbert. *Color and Color Perception: A Study in Anthropocentric Realism*. Standford: CSLI, 1987.
- [35] Hulbert. Formal connections between lightness algorithms. *Journal of the Optical Society of America A*, 3:1684–1693, 1986.
- [36] Hurvich. *Color vision: an introduction*. Sinauer Associates, 1981.
- [37] Jackson. *Perception*. Cambridge University Press, 1977.
- [38] Jackson. The primary quality view of color. *Philosophical Perspectives*, 10:199–219, 1996.
- [39] Jackson. Color and content. *Brain and Behavioral Sciences*, 26:34, 2003.
- [40] Jackson and Pargetter. An objectivist's guide to subjectivism about color. *Revue Internationale de Philosophie*, 41:127–41, 1987. Reprinted in [?].
- [41] Johnston. How to speak of the colors. *Philosophical Studies*, 68:221–63, 1992. Reprinted in [?].

-
- [42] Kalderon. Color pluralism. *Philosophical Review*, 116:563–601, 2007.
- [43] Kripke. *Naming and Necessity*. Harvard University Press, 1980.
- [44] Land. The retinex theory of color vision. *Scientific American*, 237:108–128, 1977.
- [45] Lewis. Vague identity: Evans misunderstood. *Analysis*, 48, 1988.
- [46] Lewis. Naming the colours. *Australasian Journal of Philosophy*, 75:325–342, 1997.
- [47] Loewer. A guide to naturalizing semantics. In Hale and Wright, editors, *A Companion to the Philosophy of Language*. Blackwell Publishing, 1997.
- [48] Lycan. *Consciousness and Experience*. MIT Press, 1996.
- [49] Matthen. Color vision: Content vs. experience. *Behavioral and Brain Sciences*, 15:46–47, 1992.
- [50] Matthen. The disunity of color. *Philosophical Review*, 108:47–84, 1999.
- [51] Maund. *Colours: Their Nature and Representation*. Cambridge University Press, 1995.
- [52] Maund. Color. In Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Stanford University, 2002. Available at <http://plato.stanford.edu/entries/color/>.
- [53] Maxwell. *A Treatise on Electricity and Magnetism*. Clarendon Press, 1881.
- [54] McDowell. Values and secondary qualities. In Honderich, editor, *Morality and Objectivity*. Routledge & Kegan Paul, 1985.
- [55] McGinn. *The Subjective View*. Oxford University Press, 1983.
- [56] McGinn. Another look at color. *Journal of Philosophy*, 93:537–553, 1996.
- [57] McLaughlin. The place of color in nature. In Mausfeld and Heyer, editors, *Colour: Connecting the Mind to the Physical World*. Oxford University Press, 2003.
- [58] McLaughlin and Jakab. Why not color physicalism without color absolutism? *Behavioral and Brain Sciences*, 26:34–35, 2003.
- [59] Newton. A letter of mr. isaac newton ... containing his new theory about light and colors. *Philosophical Transactions of the Royal Society*, 7:3075–3087, 1672. The experiment was conducted in 1666.
- [60] Pautz. Have byrne and hilbert answered hardin’s challenge? *Behavioral and Brain Sciences*, 26:44–45, 2003.
- [61] Pautz. Can the colour physicalist explain colour structure in terms of colour experience? *Australasian Journal of Philosophy*, 84:535–564, 2006.

-
- [62] Pautz. Philosophical issues concerning colour. In Cleeremans Bayne and Wilken, editors, *The Oxford Companion to Consciousness*. Oxford University Press, forthcoming. Draft available at <https://webpace.utexas.edu/arp424/www/papers.html>.
- [63] Peacocke. Colour concepts and colour experience. *Synthese*, 58:365–82, 1984. Reprinted in [?].
- [64] Pelli. The quantum efficiency of vision. In Blakemore, editor, *Vision: Coding and Efficiency*. Cambridge University Press, 1990.
- [65] Rosen and Smith. Worldly indeterminacy: A rough guide. *Australasian Journal of Philosophy*, 82:185–198, 2004.
- [66] Shoemaker. Phenomenal character. *Nous*, 1994:21–38, 1994. Reprinted in [?].
- [67] Smart. On some criticisms of a physicalist theory of colors. In Cheng, editor, *Philosophical Aspects of the Mind-Body Problem*. University Press of Hawaii, 1975. Reprinted in [?].
- [68] Strawson. Red and ‘red’. *Synthese*, 78:193–232, 1989.
- [69] Thompson. *Colour Vision: A Study in Cognitive Science and the Philosophy of Perception*. Routledge Press, 1995.
- [70] Tye. *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind*. MIT Press, 1995.
- [71] Tye. *Consciousness, Color, and Content*. MIT Press, 2002.
- [72] Wittgenstein. *Tractatus Logico-philosophicus*. Routledge, 1981. First published in 1921.
- [73] Yablo. Singling out properties. *Philosophical Perspectives*, 9:477–502, 1995.