Are Colors Real?

Emiliano Boccardi
(Centre in Metaphysics of the University of Geneva)
emiliano.boccardi@gmail.com

Introduction

Do colors exist in the world as mind-independent properties or, as many have argued, are they *virtual properties*: properties the world *might* have instantiated, but in fact doesn't? I am here going to assume that this problem, also known as the *problem of color realism*, concerns the existence and nature of color properties as they are *represented* by visual experience. It is natural to think that, within this framework, the starting point for a discussion of color realism would be some theory of the representational content of perceptions. This is not supposed to be the result of a conceptual analysis of color concepts, or at least not only. It is something that we supposedly know by pure introspection on the content of our color experiences. But is it true? If our color experiences have a determinate content, and if they are (at least sometimes) veridical, this is a fact that falls beyond the grasp of our a-priori reason. It is possible, at least in an epistemic sense of the word, that color perceptions systematically fail to have determinate contents. It is also (epistemically) possible that, although they have determinate content, color perceptions systematically fail to be veridical, under a metaphysically thick notion of truth, or correctness. This, incidentally, opens the logical space for so called eliminativism about colors: the thesis that nothing in the world is really colored.

On the other hand, and for the same reasons, if at least some color perceptions do have veridical contents, this bounds the meaning of color perceptions at all possible worlds. In other words, if we discover that color properties are type-X properties of our world, this fixes the content of color perceptions once and for all, in spite of the fact that there might be (counterfactually) possible worlds where color perceptions systematically fail to have determinate contents or to be veridical. I think the best way to describe this situation is through a two-dimensional theory of color perceptions. On one side is what I shall call the *character* of color perceptions. What are colors? If we are to be guided at all in answering this question, I argue, it must be the character of color perceptions that guides us. The character of perceptual experiences, as I see it, is a map from contexts to contents. It is the aspect of meaning that guides our enquiry into the reality of color properties.

You're looking at a red tomato on the table, and your perception seems to have the (propositional) content that there is a red tomato on the table. The event of your looking, and that particular tomato, (partly) constitute the token-context of your perception. The character of the perception fixes a content for that particular token color experience. The character trades with "representations" and their semantic properties, while the content, if it exists at all, only trades with material objects and their properties.

In a (Kripkean) sense, whatever the content of your experience happens to be, it must be a necessary intentional property of your particular experience. This property is necessary, but it is so a-posteriori: it is open to discovery what that particular content in fact is, if at all.This distinction between the character and the content of color perceptions, moreover, is what explains the significance of our enquiry. A philosophical theory of colors is cognitively revealing, it is informative, only if there is a difference between the character and the content of our color experiences. We have a-priori access to the character of color experiences, but not to their contents. Another way to express this thought, is to say that the character of color experiences, which is cognitively accessible for us, gives us (implicitly) a descriptive knowledge of the content of our experiences. This "description", if we are lucky (i.e. if at least some of our color perceptions are ever veridical, hence if they ever have determinate contents), must be sufficiently strict so as to fix a map from contexts to contents.

I think we are moderately lucky. I shall argue that the character of color perception, and the particular nature of our world, justifies some degree of optimism for the realist. There is at least one kind of properties that could constitute the content of veridical color perceptions. Such properties, however, I shall argue, are irreducibly extrinsic. So how does the character of color perceptions constrain the individuation of their contents? We said that the character consists of a map from context to contents. How does this map work? How would we describe such map in a meta-language that contains both perceptual terms and terms for describing physical facts? Before answering this question, let me say something about the content of perceptual experiences in general. How are perceptual contents fixed, in general?

This question has occupied an entire sub-industry of philosophical enquiry for quite a long time now. So, if we are wondering how the character of color perceptions manage to fix their contents, it sounds like a good place to start would be a good theory of content. Moreover, as most scholars involved in the discussion share physicalistic intuitions (especially with regards to what fixes the content of *perceptual* experiences) a good place to start would be a *naturalistic* theory of content. In our case, a good naturalistic theory of content would provide us with a description of the mechanisms that underlie the character of color perceptions, viz. those mechanisms that constitute the mapping from perceptual contexts to perceptual contents. Many authors, however, have been discouraged from adopting this strategy because they have a poor opinion of the achievements of naturalistic theories of content so far. D. Hilbert, for example, concedes that "one way of settling the problem of color realism would be via some naturalistic theory of content". However, he goes on to argue, "none of these theories is well-enough developed to allow this sort of argument to be formulated in the required details".[1]

Now, while this is certainly true, I think that without some intuitions about what could help to fix perceptual content, we would be incapable of setting the whole enquiry about colors off the ground. What are we looking for, when we ask for the "real" content of color experiences? Moreover, suppose that we did succeed at individuating the "real" content of color experiences; how could we know that we have so succeeded, if we have no previous knowledge of what that content should be like to start with? As a matter of fact, I think that most debates assume, more or less implicitly, some restriction or other on the notion of perceptual content. These restrictions, moreover, function as tests for the adequacy of the various proposals. This goes relatively unnoticed because there is sufficient amount of agreement about what helps to fix contents in general.

If there is no consensus about the details of a naturalistic theory of content, in fact, there is wide agreement about a number of *necessary* conditions for something to be the content of a perceptual representation. I

[1] Byrne & Hilbert, 2003, 8.

think that these conditions are often implicitly at work in framing the debates and the various arguments for and against color realism. Among these presuppositions, for example, is the claim that the content of a *veridical* perceptual experience should be (at least in part) the *cause* of that experience, and that such causal relation contributes to individuating the content itself. In other words, many authors, in line with their physicalistic intuitions, assume a causal theory of content.

There is also wide agreement about the fact that the contents of perceptual experiences (e.g. color experiences), should mirror, at least to some extent, their phenomenal structure. Many arguments in the literature heavily depend on similar assumptions.[2] I propose to start by making these (purely semantic) presuppositions explicit (section 1.1). The peculiar character of color experiences is widely (and reasonably) believed to impose further, color-specific conditions on what fixes their contents. Among these, the intuition that physical objects should be the proper bearers of color properties, if anything is. In sections 1.2–1.7 I critically discuss a number of these further restrictions. With these restrictions in place, I go on to discuss the limits and scope of objectivist theories of color. I conclude that the content of color experiences must contain a relational property of objects. In particular, I argue that such relational property is the instantiation of the projection of the space of spectral reflectances (the distal stimuli) onto the 3-dimensional space of retinal (proximal) stimuli.

This has the somewhat unwanted consequence that part of the content of color experiences (on top of physical objects) are retinas. Retinas are part of the observers, if anything is, so, like most relationalist proposals, mine faces the challenge of mind-independence: under most understandings, if an object has the property of being (literally) colored, then whatever fact makes true the proposition that the object is colored must be a mind-independent fact, whatever this means.

The mind-independence restriction requires more then the mere objective nature of color properties. After all, the properties of our retinas are objective just as much as those of physical objects are. Retinas *are* physical objects! So, one may try to cheat, "although color properties are relational, and although the retinas of the observers are among the relata of color properties, such properties are nonetheless objective properties". This would be cheating

[2] As I shall discuss in some details, for example, Hilbert's account of metamers, or his contention that objects are represented as having proportions of hue magnitudes, implicitly draw on both of the above mentioned conditions.

because the rationale behind the mind-independence requirement is that it is to make room for faulty disagreements. What is red-relative-to-my-retina, may not be red-relative-to-yours, and this appears to block a-priori the possibility of error, or disagreement.

A second standard objection against relationalist theories of colors, is that part of what we perceive, when we perceive a colored object, is that color is a property that the object has monadically, not relationally. Monadic properties are typically conceived as necessarily intrinsic to their bearers, so that relationalist accounts appear to fly in the face of the very character of color experiences.

I believe my brand of relationalism has the resources to tackle both challenges. In section 1.5 I argue that the monadic character of color properties is relative to the mode of presentation of these properties in perception, and that it doesn't impose any restriction as to the extrinsic or intrinsic nature of the properties that are to be identified with colors. The character of color experiences, in other words, present color properties as monadic (it is the tomato that is red, not a system that includes something else, a part for the tomato), but this, we shall see, only entails that color properties must be *describable* as monadic, not that they must be intrinsic to their bearers.

My response to the observer-independence challenge is two-fold. According to my proposal, the character of color experiences only places second-order constraints on their contents.[3] This has the consequence that when I veridically perceive a red tomato on the table, what fixes the content of my representation is not some relation that only THAT tomato bears to MY retina. What fixes the content of the experience is the fact that THAT tomato and MY retina instantiate a second-order relational property: a property that could be instantiated by other (sufficiently similar) tomatos and other (sufficiently similar) retinas.[4] THAT particular tomato and MY particular retina right then, at most, make THAT experience veridical.

After presenting a formal toy model of color perception (section 2), I go on to consider a number of possible variants of my proposal (sections 3.1-3.3), testing them against standard anti-relationalist arguments. In particular, I test them against the threat of faultless disagreement, that hangs as a sword of

[3] In this respect, my proposal has a lot in common with functionalist accounts.
[4] My brand of relationalism, however, does have the consequence that if my retina was substantially different (as is the case with some non-human species), then the color properties of THAT tomato might turn out to be different.

Damocles over the heads of all non-physicalist accounts of colors. I conclude that a viable candidate is what I call a teleological relationalist theory of colors (section 3.3). According to this variant, what robustly fixes the content of color experiences, and makes room for genuine disagreement, is a teleological ingredient. Put crudely, according to this variant of my account, the character of any given color experience contains reference to *what would have had to have been the case, had the perceptual system actually harboring that experience instantiated it when functioning properly.*[5]

The account, then, rests on some naturalistic notion of proper function. I briefly mention a few alternative options as to how one may hope to naturalize functions (section 3.3), but ultimately I am interested in the viability of my account as a philosophical theory of colors.[6] I argue that color properties are objective mind-independent properties of physical objects. If I'm right, then, we can say that the world is populated by objectively (albeit relationally) colored objects. Some, however, will insist that objects are not *really* colored, if colors are not basic, intrinsic properties of them.

Here enters the second part of my response (section 4.1). I argue that all color experiences, independently of the particular makeup of the respective perceptual apparatuses, share the same character. Now, relational properties often possess "narrow correlates". The narrow correlates of a relational property are the properties of an object in virtue of which that object participate to the instantiation of the property (sec. 1.3-1.5). In the case of the relational property of weight, for example, the narrow correlate is mass. Mass is the basic intrinsic property in virtue of which material objects possess a weight, under suitable circumstances. According to my account, all color properties, independently on the observer, share the same narrow correlates, viz. the reflectance profiles of their bearers. This feature of my account, I shall argue, allows for as much room for disagreement as any other objectivist theory.

Because of the peculiar nature of the restrictions imposed on content by my account, moreover, all color properties, regardless of their observers, can be compared (metrically) with their common narrow correlates (reflectance profiles). Although reflectance profiles do not constitute, alone, the content of color experiences (they are not *the* colors), they have an essential role to play in any explanation of why we developed the capacity to perceive colors, or,

[5] I borrowed this way of expressing the teleological ingredient from Ruth Millikan.
[6] This will depend on the viability of some naturalistic theory of proper functions.

which is the same, in any explanation of why color perceptions can be so useful. The possibility to compare the contents of various color experiences as to how *accurately* they approximate the reflectance profiles of the bearers of color properties, therefore, provides us with a notion of relative "accuracy" of our perceptions. Once we know that a given color perception is veridical, according to my account, in fact, we can further ask how "accurate" it is. The perception is veridical iff its apparent bearer instantiates the relevant kind of relational properties. Such properties may approximate more or less accurately the reflectance profile of the object (i.e. the narrow correlate of the color property). My account, as we shall see (sections 3.2-3.3) allows for a quantitative notion of "accuracy". Depending on how much accurate the property in question is (in this technical sense), the correspondent perception will be said to be more or less accurate.

Because we can measure the distance of our color perceptions from "ideally accurate" color perceptions, moreover, we can judge how much our physical world is far from instantiating ideally accurate color contents. My verdict is: not much! We live in a quasi-colorful world (section 4.2). In the limit, as the degree of accuracy of various color experiences increases, I argue, my relationalist account conflates with standard realist accounts (such as Hilbert's), according to which colors are to be identified with reflectance properties of objects.. This, however, does not have the consequence that colors are, *really,* reflectance profiles. What colors *really* are depends solely on the character of color experiences in our world, and on how our physical world happens to be. Although what colors really are is a matter open for empirical discovery, I repeat, it is a *necessary* a-posteriori matter of fact. Whatever color properties turn out to be in this world (if anything does at all), those properties will be "the colors" at all other nomologically possible worlds.

Even if some creatures had retinas capable of discriminating and individuating single reflectance profiles of objects, this would not entail that what these creatures would perceive would be the "true" colors. Of course, the colors that these creatures would see would extensionally coincide with reflectance profiles. And of course the perceptions of these creatures would be much more accurate then ours (in the technical sense mentioned above). But this would not entail that the colors these creatures would perceive are the *true* colors. I consider my account to be a physicalist account of the nature of colors. My considerations, I hope, will allow us to avoid the consequence that if one rejects standard physicalist theories of colors, then one is committed to think that nothing in the world is really colored: an admittedly embarrassing consequence.

Let us begin to make explicit the constraints that the character of color perceptions places upon their contents.

1. Constraints on the content of color experiences

1.1. Semantic desiderata

The following are widely accepted conditions that a property must satisfy for it to be (part of) the content of a perceptual experience.

1) Co-variation condition. Veridical color experiences form a domain whose (phenomenal) structure is (at least) omeomorphic to that of their contents.

This assumption derives from widely shared epistemological tenets. Both those who believe that sense-data mediate our experience of the external world, and those who believe that we have direct experience of the external world, will claim that we have (direct or indirect) experience also of the *structure* of the causes of our perceptions. Moreover, most philosophers find it plausible that such structure is (at least partly) captured by the phenomenal structure of our experiences.

So, for example, our auditory experiences of certain sounds can be arranged according to their pitches (e.g. Do, Re, Mi, Fa, Sol, La, Si). Call the structure determined by the relations of perceptual pitch similarities among these experiences: P_{phen}. According to the co-variation assumption, perceptual auditory experiences represent the world as instantiating (at least) the structure P_{phen}. It follows that P_{phen} consists of veridical auditory experiences only if a portion (D) of the world (W) is such that there exists relations defined on D such that their structure is homeomorphic to P_{phen}:

For some $D \subseteq W$ there are $R_1, R_2,...R_n$ on D such that $(D; R_1, R_2,...R_n) \cong P_{phen}$

Although this condition is the trademark of internalist theories of representational content (e.g. conceptual role theories, or Cummins' theory of content), most philosophers in the "causal camp" also sympathize with it. This is how Dretske expresses this requirement, for example.

> The fundamental idea is that a system, S, represents a property, F, if and only if S has the function of indicating (providing information about) the F of a certain domain of objects. The way S performs its function (when it performs it) is by occupying different states s1, s2,...,sn corresponding to the different determinate values f1, f2,...,fn,of F.[7]

Millikan is even more explicit on this point.

> [R]epresented conditions are conditions that vary, depending on the form of the representation, in accordance with specifiable correspondence rules that give the semantics for the relevant system of representation.[8]

2) Causality condition. The content of a veridical perceptual experience must be part of the cause of that experience, at least under some epistemically salient conditions.

As Hilbert points out, "any plausible version of physicalism will identify colors with physical properties implicated in the causal process that underlies the perception of colors". I would add that this is a desideratum of any non-eliminativist theory of colors that wishes to comply with physicalistic intuitions, and not only of the brand of physicalism advocated by Hilbert. The caveat on "epistemically salient conditions" is to avoid a vacuous notion of content, or, if you wish, it is to make room for epistemic error. More about this later (section 3.1-3.3).

3) Asymmetric dependence condition. If it is (nomologically) possible for a given non-veridical perceptual experience to be veridical, then its causes (qua causes of that experience) asymmetrically depend on the causes that the experience would have had, had it been veridical.

Fodor notoriously proposed a causal theory of content whose essential ingredient is the asymmetric dependence of the causes of non-veridical perceptual experiences on the causes of veridical ones. While it is still controversial whether this places sufficient (or substantial) constraints on the individuation of content, it seams to me safe to claim that any causal theory of content should be such as to have this condition come out true.[9]

4) Robustness condition. The content of a given perceptual experience must be robustly the same, regardless of whether the experience is veridical or not.

---

[7] Dretske, 1995, 2.
[8] Millikan, 1990, 224.
[9] As we shall see, however, nothing in my arguments hinges on the assumption that this condition holds.

We shall discuss this condition at length. For the moment, it suffices to say that the condition, among other things, is to make room for disagreement. If I say that a certain object is red, and you think I'm wrong, then we better mean the same thing by "red", otherwise our disagreement would be only apparent. More on this later.

1.2. Color experiences

Let us apply these general constraints to the case of the content of color perception. Let C_{phen} be the phenomenal structure of color experiences as of their hues. It consists, suppose, of the structure of similarities among them, plus the opponent structure. Let $C_{sim\text{-}phen}$ and $C_{op\text{-}phen}$ name respectively the similarity substructure and the opponent substructure. The above mentioned conditions on the individuation of perceptual content, then, allow us to say that color experiences are veridical if:

1. For some domain $D \subseteq W$, there are relations $(S_1, S_2, ... S_n)$ on D, such that: $(D; S_1, S_2, ... S_n) \cong C_{sim-phen}$ and $(D; S_1, S_2, ... S_n) \cong C_{op-phen}$
2. Under epistemically salient conditions, the instantiation of $(D; S_1, S_2, ... S_n)$ causes the instantiation of $C_{sim-phen}$ and $C_{op-phen}$
3. If an instance of C_{phen} is non veridical, its cause must depend asymmetrically on the relation that obtains between C_{phen} and its causes when C_{phen} is veridical.
4. The contents of C_{phen} would have been the same, even if the experiences in C_{phen} had not been veridical.

If we assume these conditions, then they provide us with constraints on what colors may be taken to be (if they exist at all): if the content of color perception is ever veridical, colors must (at least) be properties satisfying conditions 1-4. These constraints derive from the assumption that color perception is a representational phenomenon (i.e. that it involves tokening representational properties), plus the thesis that color properties are part of the content of color perceptions. It is easy to realize, however, that these conditions place very weak constraints, by themselves. In fact, without filling in the details of their interpretation, the constraints are compatible with virtually every representational account of colors of which I'm aware of. My thesis, I anticipate, is that under the only sensible interpretation, these constraints are sufficient to rule out all but a relationalist accounts of colors.

The further constraints that need to be added, to individuate what kind of properties colors are (if they exist at all) come from conceptual considerations about the particular nature of color, as well as from our extensive knowledge of optics, colorimetry and the neurophysiology of color perception.

1.3. What are the bearers of color properties?

Notice, first, that the four conditions given above, short of further indications as to how one should interpret them, leave open the question of what portions of the world are to provide for the class of possible instantiations of the domain $D \subseteq W$. Should the portion D of the world include the brain of the perceiver? Should it also include the whole environment? Or should it only include the (supposedly) colored objects? One possible restriction can be justified by the following argument. If the account is to construe of colors as observer-independent properties, the domain D should be taken as excluding at least our brains (and our retinas). This, as we shall see, does not, by itself, commit us to say that the properties instantiated by D must not be ultimately related to the brain. It just means that the bearers (if at all) of the color properties represented by color experiences are to be found outside of the brain of the perceivers (if any). The intuitive argument for this thesis seams to be the following.

A minimal requirement for a property to be mind-independent (however one wants to construe this notion), is for it to be a property that is not necessarily co-instantiated with any mental property. Necessary co-instantiation, in fact, is a sign of "dependence", under all sensible understandings of the word "dependence". Since, presumably, the brain instantiates mental properties, the requirement that color and mental properties are never necessarily co-instantiated entails that the bearers of color properties must be found entirely outside of the brain. I will come back to mental independence later (section 4). We shall call this restriction the:

5) Externality condition. The bearers of color properties, if any, must be physically disjoint from the brain of their potential perceivers.

Another line of argument that places a-priori constraints on the suitable instantiating domain comes from our intuitive conceptual knowledge of colors. One may reason as follows. According to our pre-theoretical understanding of color concepts, colors, if they exist at all, must be properties *of the objects* that we perceive (or of their surfaces). Forget about what *kind* of properties colors are for the moment: whatever they are, they certainly must be properties *of the objects*! Should it turn out that, under closer scientific scrutiny, there are no properties *of the objects* that comply with conditions 1-5, then, too bad for real colors! In that case one should say that we perceive the world *as if* objects instantiated color properties, when in fact they don't.

The intuition that a mere introspective scrutiny of color perceptions will reveal something about the metaphysics that they presuppose, is very strong, and indeed very widely held. We could try to explain this intuition by saying that perceptual experiences have, among their properties, a formal, predicative structure. If I perceive a red object, I come to believe that I am in front a red object: "the proposition that there is a red bulgy object on the table is part of the content of the subject's experience", says Hilbert for example.[10] If perceptual experiences have (also) a propositional content, one cannot, supposedly, have a visual experience, without thereby coming to know its propositional content. Propositional contents, in turn, have a predicative structure,[11] whence the metaphysical presuppositions. Let us call this:

The Propositional Content Assumption. The content of perceptual experiences consists partly of structured propositions. By perceiving a visual scene, subjects also perceive the predicative structure of these propositions.

Setting aside the question of where these presuppositions come from, let us now turn to the consequences they would presumably have for a theory of color. Not only does perception present objects as colored, but perception also presents what these colors are like.

> When [a person] perceives a blue bead, not only does he perceive the bead to be blue, but he perceives what blue is like. The qualitative nature of the colors is manifest to us in our perception of them. Objects are perceived to instantiate color properties, and these color properties are perceived to instantiate higher-order properties that constitute their qualitative character. So, not only does

[10] Byrne & Hilbert, 2003, 5.

[11] By this I mean that grasping a proposition entails, eo ipso, grasping its surface logical structure, viz. grasping what is predicated of what.

> color perception present the existence and distribution of the colors, but it also presents their nature.[12]

Both eliminativists and realists about colors may sympathize with this line of argument.

The eliminativist master argument is that if colors cannot be thought of as properties that inhere in the objects and that cause our color experiences (in the counterfactually strong sense described above), then, we must conclude that nothing is really colored.

Most realists would also find this argument convincing. Hilbert, for example, argues that the representationalist theory of color perception entails that "the view that no physical objects are colored is equivalent to the view that the contents distinctive of color experiences (for example, that there is a red bulgy object on the table), are uniformly false".[13] On similar grounds, many typically discard as inadequate the idea the colors may be properties of light. Let us call this:

6) The proper subject condition. The proper subject of color ascriptions are physical objects

I think condition 6 is essentially correct, but that it hides a potential unwarrented presupposition: that if an object possesses a certain real property, it must possess it in and of itself. Before turning back to this important point, let me continue with our analysis of how we should fill in the details left open by the four semantic conditions on color content.

1.4. Are color properties relational?

Conditions 1-6 leave open what sort of properties are to constitute the domain of instantiation of the structure $(D; S_1, S_2, ...S_n)$. What kind of properties are colors? Are they extrinsic or monadic? Dispositional or non dispositional? And if they are dispositional, do they involve a relation to the (cognitive system of the) perceivers or not? Is there any a-priori argument that could help us to individuate the kind of property that colors are, if anything is a color property? As we have already seen, a part from the restrictions imposed by conditions 1-4, one can try to place constraints derived from our intuitive notion of color, or from the alleged metaphysical presuppositions of color

[12] Kalderon, 2007, 563.
[13] Byrne & Hilbert, 2003, 5.

experiences (conditions 5-6). In the previous section, for example, we argued that the proper subjects of color ascriptions must be physical objects. Does this place constraints on the kind of properties colors might be (if they exist at all)? Prima facie, I think, we would answer in the affirmative. This would be the argument.

If colors are properties of physical objects, and if they (or rather their instantiations) must be causally efficacious (condition 2), then colors must be *physical* properties of physical objects. "[I]t is of course *the object* that looks colored […]", says Hilbert for example, "and so the relevant physical property must be a property *of* objects".[14] Now, at a first glance, it seems that if this reasoning is sound, then we should restrict the domain of instantiation of color properties to *monadic*, intrinsic physical properties of material objects. But this is certainly wrong. No one thinks that this is what has been shown (not even Hilbert). But I think it is interesting to see what is wrong with this conclusion.

Consider the following example. The physical world appears populated by more or less heavy objects. When someone has a tactile experience, the tactile scene appears to the subject to be one way or another. Just like the proposition that there is a red bulgy object on the table is part of the content of a visual perceptual experience, the proposition that there is a heavy object in your hand, is part of the content of your tactile experience. Now, everything that we said about colors (conditions 1-6), also apply to this case. If your experience is to count as veridical (at least possibly veridical), then it must be taken as representing the world as populated by objects that possess the property of being heavy. A line of reasoning virtually identical to the one that we have seen above, lead us to conclude (correctly, I think) that the bearers of this property, if any, must be material objects. The co-variation condition on the content of representations, moreover, leads us to conclude that, if our perceptual experiences are ever to be veridical, the property in question must be a magnitude of some kind. The causality constraints, finally, entail that such property must be a physical magnitude instantiated by material objects.

It is rather straightforward, given our background knowledge of physics, to conclude that the property represented by this experience is weight. The property of having a certain weight, in fact, complies with the three desiderata on the content of representations (1-4), and with our pre-theoretical intuitions as to what kind of property it is, as well as to what entities could bear it (5-6). In this case, it is clear that these considerations, by themselves, do not entail

[14] Byrne & Hilbert, 2003, 9, my emphasis.

anything about the particular nature of weight. We know, on independent grounds, that weight is a relational property: it is a property that material objects have relative to the earth.[15] But nothing to this effect follows from a priori arguments.

It is worth to pause a moment to think about the representation of relational properties. First, does the fact that weight is not an intrinsic property mean that weight is not *really* a property possessed *by* material objects? Should we say that, strictly speaking, the property is *really* possessed, say, only by a system that comprises the object and the earth? If so, we should conclude that the proper subjects of weight ascriptions should be entire astronomical systems. But this is certainly wrong! The system that comprises the object that you're holding in your hand and the planet beneath your feet, is not the proper bearer of the property, as this is represented by the predicate *is heavy*! The object is the bearer of the property (relative to its mode of representation) regardless of whether the property is intrinsic or relational.

The property of being a hundred meters away from a plumber is clearly a relational property. But if you are a hundred meters away from a plumber, it is *you* who are a hundred meters away from a plumber! In this case, because of the predicative structure picked up by this particular representation of that property (viz. its conceptual content), *you* are the proper subject. Notice, however, that the same (relational) property can be presented in such a way that its proper bearer is, instead, the plumber. If the predicative structure intrinsic to a representation of that property had the plumber as its proper subject, then it would be *the plumber* that has the property of being a hundred meters away from you. Similarly, the same property can be seen as an intrinsic property of a pair constituted by you and the plumber. In this case, the proper subject of ascription of the property would be the pair constituted by you and the plumber.

Distinguishing intrinsic and extrinsic, or monadic and relational properties, is notoriously a tricky task. I do not wish to delve into the details of these distinctions here, but some clarification is in order. Let me introduce some useful concepts and distinctions. First, intuitively, whether a property is relational or not, seems to be a matter of objective fact, that can be subject to rational and empirical scrutiny. Given what we know about physics, for example, it seems that weight is unquestionably and objectively a relational property of material objects. This "fact" appears not to be relative to a

[15] Even Aristotle thought that, though for different reasons.

particular way of picking up (or of representing) the property. Yet one may reason as follows.

We think that weights are relational properties of material objects because their instantiations are always conditional on the presence, feature and distribution of other (astronomical) material objects. Why do we think that? Because we believe that the weight of an object is due to the gravitational force exerted on it by the presence of other massive objects. Change the distribution or the masses of these other objects, and weight changes accordingly; whence the idea that weight cannot be an intrinsic property. But why do we think that weight is *due to* gravitational forces, rather then thinking that it *consists of* gravitational forces? After all the weight of an object is nothing but one manifestation of the gravitational forces exerted upon it. So why not say that weight is the same property as (rather then being caused by) gravitational attraction, under certain circumstances?

But if weight is nothing but gravitational forces, then whether it is a monadic property or not depends on what portion of the world we take as its relevant bearer. Gravitational attraction, in fact, is a monadic property of the system that comprises the object and the planet. It is relational only if its bearer is taken to be the object alone. I think that it is safe to conclude from this example that whether a property is relational or not, is a matter that is relative to factors that do not depend on its intrinsic nature. One and the same property has different "modes of presentation", as it were, depending on how it is picked up by its representations. Presented as a property of the object, weight is relational, while presented as a property of a larger system, it is monadic.

We argued in the preceding section that the proper subject of color ascriptions must be material objects (condition 6). Now we can see that this condition, by itself, does not constrain the metaphysical nature of color properties. It constrains the nature of color properties only relative to our mode of representing them. Following these considerations, from now on, in stead of saying that a property *is* relational, we shall say it is *relationally fixed* (by a given representation). We should then better express condition 6 as follows:

6*: Proper subject condition. The proper bearers of color properties (as these are fixed by our color perceptions), are physical objects

1.5. Relational properties and their narrow correlates

Some relational properties can be thought of as relating a narrow correlate (relatum) with a wider correlate. The narrow correlate of a relational property R of an entity, is the intrinsic property (or properties) of that entity in virtue of which the entity contributes to the instantiation of R. In the case of weight the narrow correlate is mass. An object possesses the weight that it does in virtue of having a certain mass. Mass is (a) one of the relata that constitute the property of weight (the other relata being all the relevant celestial bodies in the surroundings); mass also happens to be (b) an intrinsic property of the proper subject of weight ascriptions (physical bodies).[16] These two features of mass make of it the "narrow correlate" of weight. I am going to argue that color properties must be relational, and that the reflectance profiles of their bearers are their narrow correlates.

A more precise definition of narrow correlate requires that we distinguish basic from non-basic properties. Intuitively, a property is non-basic if it is instantiated (when it is instantiated), in virtue of the instantiation of some other property. It is basic otherwise. To pin down this notion, I shall introduce the following:

Substitutivity Test. For any property P and a pair of objects x and y, it is true that, when x is in a nomologically possible context that fixes that x is P, had y been in that context instead of x, y would also have had P. Basic properties, intuitively, are those such that two objects sharing them pass the substitutivity test:

A *basic property* is a property that belongs to the minimal subset B of the properties of the world that satisfies the following requirement: every two objects that share all their B properties pass the substitutivity test. The notion of basic property can be used to define narrow correlates in a more precise fashion. The *narrow correlate* of a relationally fixed property R of an object Q, is the single minimal set of basic properties of Q by virtue of which it has the ability to contribute to the tokening of R. It is interesting for our discussion of color properties, I think, to make a few further remarks about narrow correlates. As I have already anticipated, I am going to argue that colors are relational properties, whose narrow correlates are reflectance profiles. What

[16] For the sake of the example I am pretending that our physical world is non relativistic.

we represent in our color perceptions, I shall argue, are these relational properties, and not their narrow correlates (the reflectances).

Some might have the (fallacious) intuition that the real content of a veridical perceptual experience of a property is always a narrow correlate. If colors are objective, mind independent properties of objects, one may think, they better be intrinsic properties of objects! Narrow correlates are intrinsic properties of the bearers of color properties, and they must be (at least partly) causally responsible for our perceptions. So why not think that it is the narrow correlates (the reflectances) that we represent? I think that this intuition derives from the predicative structure of our perceptions. We ascribe color properties to physical objects, by attaching to their names/descriptions monadic predicates such as “is red”. It could be argued that this predicative structure (subject/monadic-predicate) is part of the implicit content of color perceptions. In other words, we instinctively think of monadic properties as inherent to their bearers, whence the intuition. It is interesting, for the purpose of exposing my thesis, to see how far this intuition can get. It can be spelled out as follows.

Suppose that a phenomenal structure W_{phen} represents the physical structure $(D; R_1, R_2, \ldots R_n)$. Suppose further that a candidate for the instantiation of $(D; R_1, R_2, \ldots R_n)$ is a certain class of relationally fixed, physical properties. As we said these relational properties must have narrow correlates. If this is so, should we not conclude (a priori) that the properties that are *really* represented by W_{phen} are these narrow correlates? Consider again our example. From the fact that weight has a narrow correlate, does it follow that what you are representing when you experience a heavy object in your hand, is, *really*, its mass? This is a tricky question. Notice, in fact, that the magnitude mass appears to comply with all the relevant desiderata, just as well as weight does. As I shall argue, however, the magnitude mass fails to comply with the robustness condition, hence the existence of narrow correlates will not affect, by itself, a given metaphysical account of perceptual representation. First, let us try to push the case for narrow correlates as far as it can get.

Co-variation condition. The instantiations of the magnitude mass can be arranged so as to have a structure that mirrors perfectly well those of the magnitude weight. If weights instantiate $(D; R_1, R_2, \ldots R_n)$, then so do masses. Hence mass complies with condition 1 on the individuation of content.

Causality condition. If instantiating certain weight properties causes a perceiver to instantiate W_{phen}, then, a fortiori, so does instantiating their respective narrow correlates. After all, it is the instantiation of a certain masses that cause the instantiation of a certain weights, which in turn cause the instantiation of W_{phen}. Hence mass complies also with condition 2.

Asymmetric dependence condition. Remind that the asymmetric dependence condition states that: if it is (nomologically) possible for a given non-veridical perceptual experience to be veridical, then its causes (qua causes of that experience) asymmetrically depend on the causes that the experience would have had, had it been veridical.

Thus, suppose that you are hallucinating holding various heavy objects in your hand. This means that you instantiate the phenomenal structure W_{phen}, although there is nothing heavy in your hand. The cause of this instantiation is, say, that some scientist stimulates your neurons in the appropriate way. Strictly speaking, the proximal cause of the instantiation is a certain pattern of stimulation.The rationale behind the asymmetric dependence condition is that we would like the following hypothetical conditional to come out true. Had not the presence of heavy objects caused the instantiation of W_{phen} in the past, then the same pattern of stimulation that now causes the instantiation of W_{phen}, would not be causing it. This is the essence of the "dependence" condition in question. Of course, the above conditional may turn out to be vacuously true in the case that there exist no heavy objects in reality. So, if the condition is to cut some ice, it must be understood under the assumption that weight perceptions are, some times at least, literally veridical. Let us turn back to our question: do narrow correlates always also comply with the asymmetric dependence condition? It appears that they do. Suppose we take the magnitude mass (and not weight) to be part of the content of the proposition that you are holding a heavy object. The asymmetric dependence condition, then, would take the following form: had not the instantiation of mass caused the instantiation of W_{phen} in your past, then the same pattern of stimulation that now causes the instantiation of W_{phen}, would not be causing it. It is easy to realize that if weight complies with this condition, then so will the magnitude mass.

Robustness condition. I shall argue that narrow correlates sometimes fail to comply with the robustness condition. This is the case, for example, I argue, of colors. The robustness condition states that the content of a perceptual experience must be the same, regardless of whether the experience is veridical or not. So, if the content of a veridical experience to the effect that you're holding a heavy object, is that there is an object with a given mass in your hand, then this should be the content of your experience, also in cases in which the experience is non veridical.

Now, imagine holding the same object in outer space. If, under these circumstances, you were nevertheless to experience the presence of a heavy object in your hand, your experience would not be veridical.[17] But the robustness condition imposes that the content in the two circumstances be the same. Hence, also now that you are hallucinating weight in outer space, the content of your experience is that there is an object with a given mass in your hand. But it is true that there is an object with a given mass in your hand! So your experience must be veridical, contrary to the hypothesis. This is enough, I believe, to conclude that the narrow correlates of relational properties that comply with the relevant desiderata for being the content of a given perceptual experience, are not necessarily the "true" contents of that experience.

A second remark on narrow correlates is in order. Which relational properties possess narrow correlates, and which don't? I don't have a full answer to this question, but it seems reasonable to assume that if a (relational) property is to have autonomous casual powers, as is the case with color properties, then it must have a narrow correlate. If this proves to be correct, then the causal condition on the individuation of the content of color experiences entails that, if colors are relational properties, they must have a narrow correlate. Before applying all that was said to the problem of color realism, let me make some further remarks about the conditions for identifying perceptual content.

1.6. The causality condition

There is an ambiguity in the expression of the causality condition, as expressed above. If the instantiation of the structure $(D; S_1, S_2, \ldots S_n)$ is the content of veridical color experiences, we said, it must *cause* the instantiation of structure C_{phen}. Now, there are infinitely many ways in which the world may instantiate both structures. So, to say that the instantiations of token-

[17] Imagine, for example, that while (really) holding the object in outer space, you find yourself in the same mad scientist scenario as before.

structures of type $(D; S_1, S_2, \ldots S_n)$ cause the instantiation of token-structures of type C_{phen} can be taken to mean either of the following:

1. It can be (minimally) taken to mean that each token of the structure $(D; S_1, S_2, \ldots S_n)$ causes a token of the structure C_{phen}, but no counterfactual causal conditional holds between the former and the latter. Or, maximally,
2. It can be taken to mean that the causal relations among the tokens of the two structures hold *because* there exist a law that causally connects the instantiations of $(D; S_1, S_2, \ldots S_n)$ with the instantiations of C_{phen}.

The essential difference between these two interpretations is that according to the second the causal relation between the two structures supports counterfactual conditionals, while according to the first it consists of mere material conditionals. Which of the two interpretations is most sensible? Remind that the purpose of the causal requirement is to participate in the individuation of content. Now, content (if color experiences have contents at all) must be robustly the same at different times and under different circumstances (robustness condition). The causal requirement, then, must be interpreted as supporting counterfactual claims. Not only must be the case that tokens of type $(D; S_1, S_2, \ldots S_n)$ accidentally happen to cause tokens of type C_{phen} under given circumstances. In a case where tokens of type $(D; S_1, S_2, \ldots S_n)$ are not instantiated, it must still be true that any token of type $(D; S_1, S_2, \ldots S_n)$ *would have* caused a token of type C_{phen}, had the former been instantiated.

In short, if we take the first interpretation of the causality condition to be the correct one, then the condition would merely suffice to say that certain properties cause our *allucinating* color properties, while we want the casual condition to help us grounding veridical color *representations*. This requires, as we said, that the members of the instantiation basis for a given color property share some relevant second-order properties, over and above the accidental fact of causing the same perceptual experience. In other words, the instantiation bases of color properties must carve nature at its joints. This leads us to opt for the second interpretation. I shall argue that, under this interpretation, standard versions of color physicalism are not compatible with the causal condition.

1.7. A posteriori constraints on color properties

Finally, a number of constraints on what colors may reasonably taken to be come from our impressive body of knowledge about color processing in the visual system, from psychological data about color perception, from the optical properties of physical objects, and from how these may be recovered by our perceptual systems. Psychological data, for example, show a certain degree of constancy in the perception of colors. Objects appear to retain their color properties under very different environmental conditions. In particular, they appear to retain their color properties in spite of significant changes in illumination conditions. This suggests that colors, whatever they may be, should be properties that do not depend (to a too great extent) on illumination conditions:

6. Color properties must be retained under significant changes in the spectral power distribution and wave-length composition of the illuminant.

Finally, as noted by various authors, colors must be properties that can (at least under certain ideal conditions) be recovered by our perceptual apparatuses. We know that all the information about color properties is processed in the human brain from the patterns of stimulation of three types of photoreceptors in the retina: the L-, M- and S- cones. Light of various wavelengths stimulate these types of cells to varying degrees. Red light, for example, stimulates the L-cones much more than the M-cones, and it hardly has any effect on S-cones. This suggests that colors satisfy also the following desideratum:

8) Recoverabilty condition. Color properties, whatever they are, must be such as to be (at least approximately) discernable and identifiable by processing information that consists solely of patterns of stimulation of the three types of cones in the retina.

2. The geometry of color perceptions

In this section I introduce some formal properties of color perception. My aim is to provide a toy (formal) model of color perception. It should not be taken as a realistic model: its purpose is simply that of clarifying my relationalist proposal.

2.1. Geometry of color stimuli

Physical color stimuli can be represented by functions *C*(*w*) from the range of visible wavelengths (represented by the interval of real numbers $I = [W_{\min}, W_{\max}]$) to the real numbers. In the intended interpretation, these functions assign to each wavelength $w \in I$ its intensity *C*(*w*). Each of these functions is a (linear) combination of pure "spectral color stimuli", i.e. stimuli whose intensity is non zero only for one wavelength value $\overline{w} \in \mathrm{I}$. Physical color stimuli, thus represented, are elements of a Hilbert space of square-integrable functions: *H(I)*.

As we said, stimuli of various wavelengths stimulate the three types of photoreceptive cells in the retina to varying degrees. Such "degrees" can be represented by three functions: s(w), m(w) and l(w). The "extent" to which a given physical color C(w) stimulates each of these receptors, can thus be calculated, respectively, as:

$$\int_{W_{\min}}^{W_{\max}} C(w)s(w)dw, \quad \int_{W_{\min}}^{W_{\max}} C(w)m(w)dw, \text{ and } \int_{W_{\min}}^{W_{\max}} C(w)l(w)dw.$$

Perceived colors can then be represented as points in a three-dimensional space: R^3_{color}. The relations between these points and the functions in *H (I)* will be crucial for our proposal. Let me spell them out in some more details. For reasons to be discussed later, the "human" case of a 3-dimensional perceptual space will be generalized to an N-dimensional space. Given the Hilbert space of physical color stimuli *H (I)*, we select an N-dimensional subspace: $H_N(I)$. We introduce, for $H_N(I)$, a basis: $\mathrm{b_n(w), n = 0, \dots N}$. Each element $C \in H(I)$, can be approximated by its orthogonal projection onto $H_N(I)$. If we indicate the projection operator with O:

$$C(w) \approx O \cdot C(w) = C_N(w) = \sum_{\mathrm{n=0...N}} \beta_n b_n(w)$$

The coefficients are calculated as follows: $\beta_n = \langle C, b_n \rangle$, where $\langle .,. \rangle$ is the scalar product of *H(I)*.[18] Let S be the subset of *H(I)* that represents the color stimuli. Among these, the monochromatic stimuli, $m_{w_0}(w)$, are defined as those stimuli that are concentrated at some wavelength $w_0 \in I$. Call "black" the function $B \in S : B(w) = 0$, for all $w \in I$. And call "white" the function $W \in S : W(w) = 1$ for all $w \in I$.The line connecting the monochromatic stimulus $m_{w_0}(w)$ with the white point *W(w)* crosses the boundary of **S** at $m_{w_0}(w)$, hence the half-line $\{c \cdot m_{w_0}(w_0), c \geq 0\}$ lies at the boundary of S.

It follows that, given any two stimuli $C_i(w) \in S : i \in \{1,2\}$, their linear combinations $c \cdot C_1(w) + (1-c) \cdot C_2(w)$ are also stimuli, for all $0 \leq c \leq 1$. Thus S is convex. More precisely, it is the convex closure of the set of monochromatic stimuli.[19]

The image $S_N = \{O \cdot C : C \in S\}$ of S is a subset of $H_N(I)$ (also known as the "spectral locus"). The projection operator (O) can be chosen so that the spectral locus lies at the boundary of S_N. This matches the fact that monochromatic spectra lies at the boundary of *S*. The line in $H_N(I)$ connecting the projections of the limit points $m_{w_{\min}}(w)$ and $m_{w_{\max}}(w)$, viz. the line connecting $O \cdot m_{w_{\min}}(w)$ and $O \cdot m_{w_{\max}}(w)$, is called the "purple line". Both S and its image S_N consist of half-lines departing from the black point. They both form a (mathematical) cone, whose vertices are the spectral colors and whose apex is the black point. Each half-line in the S-cone represents a given color stimulus. Receding from the apex (the black point), the stimulus retains its chromaticity, while increasing its intensity.

As we said, human color space is three-dimensional because our eyes contain three types of receptors, each with its own type of spectral response. The projector operator, in the human case, maps the set of stimuli $S \subseteq H(I)$ onto a subset of a 3-dimensional space $S_3 \subseteq H_3(I)$. The choice of a basis for this space is rather arbitrary. At the beginning of this section, we suggested that a basis could match the fundamental response functions of the receptors in the eye. This, however, is not imposed upon us. Any three linearly

[18] This ensures that the basis is ortonormal
[19] From Grassmann's laws.

independent combinations of these bases will constitute a suitable basis for the same color space.

A concrete manifestation of this arbitrariness is the fact that color-matching data from normal individuals underdetermine the eye's primary response functions. Indeed, "[a]ll the colors of the spectrum […] can be mimicked by combinations of different intensities of […] blue, green, and red".[20] The amounts of the three primaries required to match a given color are called its "tristimulus values". Because any three linearly independent combinations of these color-matching functions is also a triplet of color-matching functions, the choice of "primaries" is arbitrary, so long as their vectors in color space are not coplanar. The projector vector O, can only be determined by empirical investigations performed on human (or other) observers. The characteristics of vector O depend on what sets of spectral stimuli are visually identical to a given subject. Any two stimuli belonging to such a set, are called metamers. Mathematically, metamers are stimuli mapped onto each other by functions whose projection under O is null.

2.2. The structure of phenomenal colors

In the previous paragraph we described some geometric properties of color stimuli. These stimuli are processed and modified by our perceptual apparatus shortly after being input to the cognitive system. Whatever the processes involved in this information processing, the result of them is the phenomenal structure of color properties as we experience them. There are several ways in which we can investigate empirically this structure. Probably the best known is the "Munsell color system". It is a 3-dimensional color space based on the phenomenal dimensions of hue, value (lightness), and chroma (color purity). It was introduced by Albert H. Munsell at the beginning of the twentieth Century, and was improved in the following decades through extensive (psychological) experimental studies. For the purposes of this paper, the details of Munsell color space are not relevant. It suffices to note the following.

Munsell color space can be represented cylindrically in a 3-dimensional space as an irregular color solid. For the purposes of my argument, as I said, it is irrelevant whether the details of this particular solid accurately represent

[20] Malin and Murdin,1984, 35, 60-61.

phenomenal color space. It will be here taken to represent the structure of phenomenal color experiences, whatever they are. What I mean, by this, is that minor changes in the detailed structure of Munsell cylinder won't affect the strength of my argument. In the notation introduced at the beginning of this paper (§ 1.2), Munsell color cylinder will be taken to be the structure C_{phen}. As noted in paragraph 1.2, the contents of color experiences, if they are ever veridical, must be such that:

For some portion of the world (subdomain $D \subseteq W$), there are relations ($C_1, C_2, ...C_n$) on D, such that: $(D; C_1, C_2, ...C_n) \cong C_{phen}$.

I shall argue that, contrary to most physicalist proposals, such structure must be homeomorphic to the N-dimensional subspace of the Hilbert space of color stimuli introduced in the previous section. I shall further argue that such structure can only be instantiated if colors are taken to be relational properties. According to my proposal, I anticipate, the reflectance profiles of the surfaces of material bodies are the narrow correlates of these relational properties.

3. Colors as instantiations of orthogonal projectors

Given the restrictions placed on color properties by the character of color perceptions (conditions 1-8), it follows that the properties represented by color experiences cannot be the spectral reflectances of the surfaces of objects. Spectral reflectances, in fact, instantiate at best the structure $S \subseteq H(I)$ described in section 2.1. The entities in this subset do not naturally instantiate the structure $C_{phen,}$ as required by the co-variation condition (condition 1). Intuitively, this means that spectral reflectances do not stand to each other in the right similarity relationships. If we identify colors with spectral reflectances, for example, then two reflectances belonging to a metameric pair should count as two different colors, while they appear to be exactly the same to all normal observers.

While, as we shall see, some authors are prepared to bite the bullet on this point, I think there are reasons to think that this is a drawback of standard physicalist accounts. More strikingly still, phenomenological colors that correspond to monochromatic stimuli ($m_{w_0}(w)$) vary continuously as w_0 takes up increasing or decreasing values within the visible spectrum, but tend toward the same color (puple/magenta) at both opposite extremes:

respectively in correspondence with w_0 = .40μm and w_0 = .70μm. There is no property of the vectors $m_{w_0}(w)$ in H(I) that correspond to this fact.

Now, while the structure S of distal stimuli is not homeomorphic to the space of phenomenal colors C_{phen}, the physical causal properties that instantiate the projection operator can be argued to be. In the case of humans, for example, there is a homeomorphic mapping from the Munsell color cylinder (the human C_{phen}) to the cone represented by S_N. My proposal is to identify colors with the relational properties that instantiate the projector operators. This ensures that the content of color experiences satisfies the co-variation condition.[21] As we shall see, the proposal can be argued to be immune to standard objections to relationalism.

Before exposing my proposal, it is interesting to consider Hilbert's response to the objection raised above. Hilbert proposes to identify colors with specific reflectances of physical surfaces. He is well aware of the above mentioned potential objection: "[d]eterminate colors", he writes, "cannot be identified with specific reflectances because there will typically be (indefinitely) many reflectances that result in the appearance of a given determinate color, and no motivation for choosing between them." (Byrne & Hilbert, 2003, 13) Here is how Hilbert proposes to amend his theory to meet this objection:

> The solution to this problem is clear: we can identify the determinable colors with reflectance types (or sets of reflectances) rather than with the specific reflectances themselves. For example, the property purple, on this modified account, is a type of reflectance rather than a specific reflectance. As a bonus, this proposal also solves the problem of metamers (and so it is not really an additional problem): both determinable and determinate colors are reflectance-types. Metameric surfaces are, according to the revised theory, the same in

[21] It may be objected that such mapping is not complete, or that it is not a "perfect" homeomorphism. Topographically, the two structures match pretty well. They are both 3-dimensional, they agree on conflating metameric pairs, and, finally, most phenomenical similarity relations are preserved. But not all! Human phenomenal color space is metrically distorted in ways that are not always matched by the cone S3. There are some qualitative properties expressed by the Munsell color system that have no match in the triple-reflectance color space. Similarity relations along the dimensions of brightness and saturation, for example, have a different character from changes of hue from unique green to unique yellow to red to blue. Such distortion of the color cylinder have no correspondence in SN. These differences, however, are minor, and do not play any significant role in standard color perception. Those who sympathize with my proposal, will have to bite the bullet. They will have to accept that there are (few) properties of color experiences that have no correspondence in reality, although most of them do.

> determinate color in spite of their physical differences (Byrne & Hilbert, 1997a; Hilbert, 1987).

The resulting account is known as "Anthropocentric Realism". Real colors, according to this view, are spectral reflectances. Then there are anthropocentric colors, identified with *groups* of spectral reflectances. Falk talk of colors, according to this view, refer to anthropocentric colors, while what is really represented in color experiences, are real colors. Now, if the considerations exposed in section 1.6 are sound, then neither "real" nor "anthropocentric" colors could be the content of color experiences. According to the causality condition, if the instantiation of a structure $(D; S_1, S_2, ...S_n)$ is the content of veridical color experiences, it must *cause* the instantiation of structure C_{phen}. As noted in section 1.6, this condition must be taken to entail that the causal relations among the tokens of the two structures must hold *because* there exist a law that causally connects the instantiations of $(D; S_1, S_2, ...S_n)$ with the instantiations of C_{phen}.

Hilbert concedes that "the reflectance-types that we identify with the colors will be quite uninteresting from the point of view of physics or any other branch of science unconcerned with the reactions of human perceivers". However, he continues, "[t]his fact does not […] imply that these categories are unreal or somehow subjective (Hilbert, 1987, 11)". I agree that the fact that these properties are "uninteresting" does not entail that they are unreal. But, given our causality condition, this is not enough. If they are to constitute the content of veridical color experiences, these properties must be projectible, i.e. they must be (jointly) capable of supporting inductive reasoning, quite apart from inductions related to the response of perceivers. So, if by saying that they are "uninteresting" Hilbert means that the only inductions that these properties support are related to color perceivers, then the causality condition rules them out as candidates for the content of color experiences. More on this point later. If spectral reflectances (or classes thereby) cannot be identified with colors, however, they certainly have a lot to do with them. For example, it is unquestionable that we would not perceive any colors, if it wasn't for them. It could be argued even that we could not even hallucinate colors, if it wasn't for them (asymmetric dependence condition). So what's the role of spectral reflectances in color perception? As I have already anticipated, I argue that spectral reflectances are the narrow correlates of color properties.

3.1. Teleological relationism

3.1.1. The character of veridical color experiences

Let me briefly summarize what we said about color perceptions. The distal stimuli that cause color perceptions form a structure that can be represented by $S \subseteq H(I)$, as described in section 2.1 above. The stimuli undergo two transformations.

The first formally consists of an (orthogonal) projection that "squeezes" the space of spectral stimuli into a N-dimensional subspace of *H(I)*: $H_N(I)$. The resulting structure is the structure of proximal stimuli: a convex subset S_N of $H_N(I)$. The dimensionality of $H_N(I)$ depends on the number of types of photoreceptive cells in the perceptual apparatus of the perceiver.[22]

Such projection is (formally) realized by the projection operator O, so that the space of proximal stimuli is the image of S under O: $S_N = \{O \cdot C : C \in S\}$. The causal chain that links the reflectance properties of objects to color perceptions, must therefore instantiate the projector O. This causal chain, in humans, is realized by the reflectance properties of objects and by the three types of photoreceptive cells in the retina, resulting in a 3-dimensional space of proximal stimuli.

The second transformation is realized by the brain alone, and it leads to the instantiation of a structure that we called "phenomenal color space": C_{phen}. Here is a sketch of my proposal. I propose that color properties, i.e. the properties represented by veridical color experiences, should be identified with the physical properties that instantiate the projection operator O, whatever they are. As we said in the introduction, the character of color experiences is a map from perceptual contexts to perceptual contents. Now we can say how the character map works, i.e. how it assigns contents to various contexts.

[22] I deliberately leave open this dimensionality, to allow for color experiences in creatures whose visual apparatus is different from that of normal humans.

3.1.2. The context of perceptual token-experiences

Perceptual token-contexts are constituted by (1) an individual object (or surface), (2) an environment and (3) an individual perceptual apparatus. For the purpose of the individuation of content (as specified below), token-contexts belong to non-conventional types. Such types are individuated by the formal properties of their instantiations, and do not depend on the properties of the phenomenal color space. More precisely, two token-contexts belong to the same type (equivalence class) iff they instantiate all and only the same projection operators.[23] The character of color perceptions (the map) works according to the following instructions:

Rule 1. The reflectance profile of the perceived bearer of the color (the object) is always one of the relata of the color property represented by a veridical color perception. More precisely, it is the narrow correlate of the color property.

Before the distal stimuli are processed by the perceptual apparatus, the space of reflectances is projected onto a finite-dimensional space: the space of proximal stimuli.[24] Such projection can be represented by an orthogonal projection operator, O.

Rule 2. The content of a veridical color experience, given the context (as specified in step 1), is the relational property of the colored object in virtue of which the object and the perceptual apparatus co-participate in the instantiation of the projection operator O.

Notice that not all the physical details of a given token-context are relevant for applying rule 2. What a given color property is, is insensitive to changes to the properties of a context that leave unaltered the formal properties required to instantiate the projection operator. As we said, two token-contexts belong to the same type iff they instantiate all and only the same projection operators. Rule 2, then, is only sensible to the *types* to which a given context belongs. My proposal entails that individual colors are identified with relational properties. In the case of humans, for example, each color will be identified with a physical relational property whose *relata* are at least (1) the physical object (the bearer of the color property) and (2) the retina. Should we conclude that it is part of the essence of colors to be relational properties? If

[23] The idea of generalizing the relevant normative contexts to solve the problem of faultless disagreement has been defended in Cohen, 2004.

[24] Subsequent processing consists in further transforming these stimuli so as to construct the phenomenal color space

so, why isn't this transparently part of their characters? Why, that is, did we have to look at how our world is, to figure out that colors are relational properties, when supposedly this is a consequence of the *character* of their representation, and character is the cognitively accessible semantic dimension of representation? Could it have turned out that colors are intrinsic properties of their bearers, or that they are their reflectance profiles?

Yes and no. In other counterfactually possible worlds, the answer to the last question is yes: colors could have turned out to be reflectance profiles. But under a Krepkean notion of possibility the answer is no. The character of color perceptions is whatever allows perceivers to go from perceptual contexts to perceptual contents. As it happens, this map, as we described it, is "world specific", i.e. it allows to *successfully* individuate content (if at all) only at worlds sufficiently similar to ours. Once the character of color experiences is individuated (in our world), however, it remains robustly the same at all other possible worlds, like water remains robustly identical with H_2O at all possible worlds. A consequence of this is that the same character that individuates what colors are in our world, might not be successful at individuating contents at all in worlds nomologically very different from ours.

3.1.3. The character of non-veridical color experiences

Notice that rules 1 and 2 only provides us with means for fixing the content of color representations when (and if) they are veridical. To complete the identification of the character of color perceptions, then, we need to add another rule that fixes the content of non-veridical experiences in a robust way. Intuitively, such content will be individuated by those properties that *would* instantiate the relevant projection operator, in *that* context, if the experience *were* veridical.

This, however, is highly problematic, for it threatens my account to beg the relevant question. Suppose in fact that your retina starts to dysfunction (or to function differently), so that the same tomato that appeared red to you this morning, now appears to be blue. Remind that, for our purposes, the functioning of the retina is completely captured by the "degrees" to which a given physical color C(w) stimulates each of the three receptors. So the assumption that your retina functions differently this evening effectively means that (at least) one of the three response functions (s(w), m(w) and l(w)) has changed. This would lead inevitably to three different coordinates:

$$\int_{W_{\min}}^{W_{\max}} C(w)s(w)dw, \int_{W_{\min}}^{W_{\max}} C(w)m(w)dw, \text{ and } \int_{W_{\min}}^{W_{\max}} C(w)l(w)dw.$$

In sum, the 3-dimensional projection of the Hilbert space of distal stimuli will be different. It follows that also the projection operator would be different from the one you and the tomato instantiated this morning! If I don't add anything to the account, this would have the consequence that this evening the content of your experience is different from the content that your experience of the same tomato had this morning, in spite of the fact that the tomato hasn't changed at all. Worst still, the two experiences will be (necessarily) equally veridical! This is a typical drawback of relationalist accounts. I believe that my framework has the resources to tackle this problem, but I will have to make relevant concessions to ecological theories of color. The secret, I believe, is in the relation between the manifest bearers of color properties and their narrow correlates. Let us resume our discussion of narrow correlates (sec. 1.5).

Consider again the example of weights. Weights, we said, are relational properties of familiar material objects and other astronomical objects. Our representations present the familiar objects as the proper bearers of weight properties. I (usually) weigh 75 kilos. It is *I* who weigh 75 kilos: not a system that includes the earth! We noticed, however, that this is not an irreducible feature of weight. Weights, in fact, are irreducibly extrinsic properties. What happens is that our representations pick up these properties in a monadic mode, as it were. This is why we ascribe weights to people and objects, and not to astronomical systems. Now, because the property is to be causally efficacious, we expect weights to have a narrow correlate. As it happens, this is mass. I have argued that, in spite of this, our weight representations do not have the magnitude mass as their sole content. However, the mass of an object, being a narrow correlate, plays a (Krepkean-) necessary role in the individuation of the content of veridical weight perceptions. Indeed, it could be argued that mass is what interests us, in making weight judgments, although it is not the content of weight perceptions.

What to make of false (or incorrect) weight perceptions? Suppose you wanted to buy a 1 kilo beefsteak. And suppose that the shopper tricks you in the following way. When he weighs the beefsteak in front of you, he activates an elevator that accelerates upwards the whole shop. As a result, you will get less meet then you expect. Yet both the scale and your perceptions would agree that you're in front of 1 kilo beefsteak! What to make of this? Is your perception non veridical, in the elevator? If we specify the character of weight perceptions only making references to familiar and astronomical objects, like we have done, then there is no way to say that the shopper is wrong. According to the only available notion of weight, he's absolutely right: the beefsteak weighs 1 kilo!

Yet something must be missing from our specification of the character of weights, such that, if we took it into account, we could explain why the shopper is cheating, and why the weight perception is non-veridical. I think that part of the character of weight experiences, in fact, is that their contents correlate with mass (intuitively: quantity of matter). Mind it, I said that the contents *correlate* with mass, not that they *are* masses. When we experience a given weight, I expect to be experiencing a given quantity of matter. This is why you would be surprised if you were still hungry after eating your elevator beefsteak. I think that a reference to masses should therefore be inbuilt in the character of weight experiences. A weight experience is veridical (among other things), if it gives us an optimal idea of the mass of its proper subject.

Something analogous, I believe, happens to the character of colors. We said that reflectance profiles are the narrow correlates of color properties. This explains why color perceptions exist at all. It is by latching to reflectance profiles, that color properties convey information about the physical characteristics of objects. If it is true that people's hair tends to turn grey with age, and if ripening bananas and pears tend to turn yellow, and if it is true that red striped spiders are venomous, this is because color properties latch onto reflectance profiles. This explains why color properties are projectible, to some degree, and why we expect them to be found out there in the external world, independent from our perceptions.

The adequacy of such latching, I submit, must then be inbuilt into the character of color experiences. In particular, I propose that it should be relevant in fixing the content of false (or incorrect) color experiences. But how can we do that, without concluding that nothing is really objectively colored? Notice in fact that this notion of "adequacy", is relative to token-contexts: my perception of the tomato is "adequate" only relative to the present conditions of my perceptual apparatus. So, if we in-build the notion of proper functioning into the character of color perceptions, we seem to be confronted with the following dilemma.

3.1.4. The problem of error

In saying that my apparatus “dysfunctions”, or that it “doesn’t perform at its best”, we appear to be saying either of these things:

1. Either we are saying that it “dysfunctions” in the sense that it fails to capture the exact reflectance of the tomato, in which case ALL possible perceptual apparatuses dysfunction.
2. Or we say that it dysfunctions in the sense that it is not “performing at its best”, whatever this means.

In the first case, if proper functioning is in-build into the notion of veridical color perceptions as I have suggested, we will conclude that there are no veridical color perceptions after all, i.e. that nothing is really colored in the relevant sense (eliminativism). However, if we opt for the second interpretation, we fall into the relativist horn of the dilemma. It seems that a retina can only be “performing at its best” (or fail to do so) relative to itself. In fact, if we said that a retina is not performing at its best relative to a “healthy” retina, we must be referring to the first interpretation of “dysfunction”. A “healthy retina” can only be (1) a statistically typical retina, in which case the epistemically normative character of the notion is lost; or (2) a retina that optimally approximates reflectance properties of objects. But “optimally” relative to what other possible retinas? This has the absurd consequence that, under the second interpretation, no retina could possibly dysfunction. A retina, in fact, can do nothing but follow the laws of physics. How could it possibly go wrong about that? How could you blame a retina for following the laws of physics?

This is the good old problem of error. Where are we to find room for epistemic error in a world that submissively obeys to the laws of physics? In a nutshell, this is the problem. If there is a sense in which a given color perception is non veridical, there must be a sense in which, in that context, that perception *could* have been veridical. Hence there must be a sense in which, in that context, the perception *could* have been different from what it is. Nothing empirical can be false, if it *could not* have been true!

Now, these modal notions must be understood in a nomological sense. This is because, if we strip the physical details from the context of a given perceptual experience, it is not clear anymore that it is THAT perceptual experience that COULD have been true. Let me be more precise about the problem of error. Remember that each element $C \in H(I)$ from the distal stimuli (the reflectance profiles), can be approximated by its orthogonal projection onto $H_N(I)$. This, we have seen (section 2.1), can be expressed by:

$$C(w) \approx O \cdot C(w) = C_N(w) = \sum_{\mathrm{n=0...N}} \beta_n b_n(w)$$

How "good" is this approximation? Could a different choice of response functions make this approximation better? Is there any other function in $H_N(I)$ that approximates $C \in H(I)$ better then $O \cdot C(w) = C_N(w)$ does? The following are standard mathematical notions that will help us to answer these questions.Given any two functions $C^1(w), C^2(w) \in H(I)$, define their inner product as:

$$\left\langle C^1(w) \mid C^2(w) \right\rangle = \int_{W_{\min}}^{W_{\max}} C^1(w) C^2(w) dw$$

This allows us to define a positive definite norm for each vector $C(w) \in H(I)$ in the Hilbert space of stimuli:

$$\|C\| =_{def} \sqrt{\left[\langle C \mid C \rangle\right]}$$

With this norm we can define a "distance" between any two functions of the space. Such distance turns our space into a metric space.

$$d(C^1, C^2) =_{def} \left\|C^1 - C^2\right\| = \sqrt{\int_{W_{\min}}^{W_{\max}} \left[C^1(w) - C^2(w)\right]^2 dw}$$

Now we can give a precise definition of what it means to say that a given projection $O \cdot C(w) = C_N(w)$ "approximates" the stimulus *C(w)*. We shall say that the projected vector $O \cdot C(w) = C_N(w)$ approximates *C(w)* to a degree of accuracy that is measured by the distance:

$$d(C(w), C_N(w)) =_{def} \|C(w) - C_N(w)\| = \sqrt{\int_{W_{\min}}^{W_{\max}} \left[C(w) - C_N(w)\right]^2 dw}$$

Now, given a Hilbert space $H(I)$ and a point (vector/function) in it, $C \in H(I)$, and given a non-empty closed convex subset, such as $\mathrm{S_N} \subseteq H(I)$, there exists a unique point $C_N^{Best} \in S_N$ which minimizes the distance between *C* and the points in tri-stimulus space S_N:[25]

$$C_N^{Best} \in S_N \,,\ \left\| C - C_N^{Best} \right\| = d(C, S_N) = \min\left\{ d(C, C_N^i) : C_N^i \in S_N \right\}$$

The existence of vector $C_N^{Best} \in S_N$ suggests that we may use it to ground the normative character of color perceptions. We could, for example, in-build a reference to it among the features that individuate the content of color experiences (relative to a given context), along the following lines:

The content of veridical color perceptions (proposal 1)

A given perceptual token-context fixes a tri-stimulus space $\mathrm{S_N} \subseteq H(I)$ and a projection operator O. The content of color experiences are the properties that instantiate O. The character of these experiences determines the conditions under which their content is veridical:

> *The properties that instantiate a given projection operator O_i are the content of a veridical color experience only if O_i is such that, for any possible stimulus $C \in H(I)$, the image of C under O_i, $C_{N_i}(w) = O_i \cdot C(w)$, is the best approximation of C, relative to O_i:*
> $C_{N_i}(w) = C_{N_i}^{Best}$

Now the problem expressed above is quite clear. Call the projection operators that the tomato and your retina instantiated this morning and this evening, respectively, $O_{morning}$ and $O_{evening}$. The same reflectance profile of the tomato has two images in the two different tristimulus spaces: $C_{N_{morning}}(w) = O_{morning} \cdot C(w)$ and $C_{N_{evening}}(w) = O_{evening} \cdot C(w)$. The condition we placed above consequently splits into the following two conditions:

$$C_{N_{morning}}(w) = C_{N_{morning}}^{Best} \text{ and } C_{N_{evening}}(w) = C_{N_{evening}}^{Best}$$

Suppose that these conditions apply to our case. They express the fact that your retina performed “at its best” both this morning and this evening. The

[25] Rudin, 1987, theorem 4.10

retina performed "at its best" relatively to what it could (nomologically) have done, given its current properties at the time of assessment. If these are the conditions for a given color representation to be true, then we will have to say that the tomato *was* red this morning and *blue* this evening. If you and I instantiated respectively $O_{morning}$ and $O_{evening}$ in front of the same tomato, moreover, according to my proposal the tomato would then be red for me and blue for you. And that's that: no possible disagreement! This is the second horn of the dilemma.

On the other hand, if we required that the projection operator be such as to capture *exactly* the reflectance properties of the tomato, that is if O is required to be an identity operator, then neither $O_{morning}$ nor $O_{evening}$ could be the content of a true color experience: hence, strictly speaking, the tomato would be neither red nor blue. This is the first horn of the dilemma.

3.2. True colors relative to the dimensionality of color space?

Another option comes to mind. Perceptual contexts, as defined above, fix the relevant color spaces (hence also the projection operators) in two ways. First, they determine the *dimensionality* of the projection. For us trichromats, for example, this dimensionality is 3. Other perceptual contexts (in non-human animals or in anormal humans), will fix color spaces and projection operators differently.

Secondly, perceptual contexts fix the detailed shape of proximal perceptual spaces. This is determined, in the case of humans, by the response functions s(w), m(w) and l(w). So far we have proposed to make optimal performance of visual experiences relative to a given triplet of response functions. This created the problem of error as explained above. Could we not have fixed the normative notion of optimal performance relative to a given dimensionality, rather then to a specific triplet of response functions? To say that the tomato is red because this is the best I could do, given the current conditions, makes color properties relative to idiosyncratic visual conditions and to the current physical properties of the retina. But what if we define red relative to the best I could do *as a trichromat*, rather then relative to the best I could do as Emiliano (my name) this morning?

Technically, this is what the proposal would look like. Given a certain dimensionality (N), there are uncountably many N-tuplets of response functions, corresponding to as many projection operators. Let us confine ourselves to the case of trichromats, for simplicity. Given a distal stimulus $C \in H(I)$, to each triplet $s_i(w)$, $m_i(w)$ and $l_i(w)$ in the space of possible projections, there corresponds an optimal approximation vector $C_i^{Best} \in S_3$. We can partition the class of triplets into those subclasses that share the same optimal approximation vector, so that now there is a one-one correspondence between 3-D projection operators belonging to the same equivalence class, and optimal approximation vectors. The proposal sketched above is then to select the normatively relevant (equivalence class of) projection operators so as to optimize the distance from the distal stimulus C.

Let T_i be the triplet $\langle \mathrm{s_i(w), m_i(w)\,and\,l_i(w)} \rangle$, and let T be the class of all these triplets. $C_i^{Best} \in S_3$ is the best approximation vector relative to T_i.

Now, there exists a vector $\mathrm{C^{Best}} \in H(I)$ such that

$$\left\| C - C^{Best} \right\| = \min \left\{ d(C, C_i^{Best}) : C_i^{Best} \in T \right\}.$$

Notice that the identity of $\mathrm{C^{Best}} \in H(I)$, hence of the correlated projection operator O^{Best} depends solely on the *dimensionality* of the projection, and on no other idiosyncratic feature of the perceptual context. Let us call O^{Best} the "3-best projection operator" (generalizing dimensionality: the N-best operator). The amended proposal, then, would be the following:

The content of veridical color perceptions (proposal 2)

> *The properties that instantiate a given projection operator O_i are the content of a veridical color experience only if O_i is such that, for any possible stimulus $C \in H(I)$, the image of C under O_i, $C_{N_i}(w) = O_i \cdot C(w)$, is the N_i-best approximation of C.*

The projection operator that wins this game will be called the N-best operator. Analogously, we define the "N-best tri-stimulus-space" and the "N-best N-tuple of response functions". The predictable complication with this version of my account is that it is very likely that, within the same dimensionality, there will be N-tuples that "win the game" relative to certain distal stimuli, while others win it relative to other stimuli. If this is the case, as I think it is, then my restriction won't suffice to establish a relation of total order among N-tuples/spaces/operators. One needs a *total* order because the normative character of color perceptions now hinges upon the possibility to compare N-tuples/spaces/operators with N-best N-tuples/ N-best spaces/ N-best operators.[26]

What to do? One could of course try to further restrict the condition, for example by averaging for accuracy among the various possible projectors. The epistemically relevant order relation in the space of projectors would be then established by comparison with "average best performance", rather then with "best performance" *simpliciter*. I won't pursue this solution, however, because I think there is a general serious problem with my strategy that has to be tackled first. It is possible (if not likely), that under this restriction most ordinary color perceptions would come out false. Who guarantees that that the typical human retina is averagely the N-Best receptive apparatus? Indeed, one may have reasonable doubts about whether the typical human retina is an N-best retina (averagely or not averagely) with respect to any stimulus.

So eliminativists will probably sympathize with this relationalist proposal. Nothing, or nearly nothing, would be colored, in our world, if this theory of color were correct. Here goes my last proposal, as far as this paper is concerned. Essentially, the idea is to replace the restriction exposed in this section with a *teleological* restriction, thus weakening it substantially.

3.3. Teleological relationalism

My favorite version of the relationalist account advocated in this paper, is a teleological version. It is "teleological" because the epistemologically normative ingredient is a naturalized notion of purpose, or function, rather then the technical notion of best performance introduced above.

In a nutshell, this is the proposal:

The content of color perceptions (teleological proposal)

> *The content of color experiences are the projections that would have had to have been instantiated, had their respective perceptual systems instantiated that experience when* functioning properly*.*

What makes this proposal "teleological" is the fact that the epistemologically normative ingredient is a naturalized notion of purpose, or

[26] The instantiation of a projection operator, under this proposal, is the content of a veridical N-dimensional color experience only if it is an N-best operator.

function. The notion of “functioning”, for example, could be borrowed from biology:

> The concept of a biological function is defined in terms of natural selection (Wright [[92]], Neander [[58]]) along the following lines: it is the function of biological system S in members of species Sp to F iff S was selected by natural selection because it Fs. S was selected by natural selection because it Fs just in case S would not have been present (to the extent it is) among members of Sp had it not increased fitness (i.e. the capacity to produce progeny) in the ancestors of members of Sp. [27]

We are not forced, however, to adopt this particular reductionist strategy. In fact, maybe we shouldn't. As etiological accounts of function cannot be cashed out in terms of the present state of the instantiating system, some might worry (with reason, I think) that these are causally epiphenomenal, i.e., causally inert. Remind that we are after *real* causal constraints on representations, so if this difficulty cannot be amended, this fact could threaten our proposal. Some authors suggest that biological function could be cashed out in non-etiological and non-teleological terms. Here it suffices to say that, while teleological functions are often considered as selected effects, they can also be considered as selected *dispositions*: certain traits are selected because they produce certain effects in response to certain causes.[28] Moreover, there is hope that one could define *proper functions* in non-etiological, and non-biological terms.[29] What's good about teleological solutions, is that they can be adapted to various theories of content to block the problem of error.

> An appeal to teleological functions can be combined with various ideas to form hybrid theories. [...] it's worth mentioning that such an appeal can also be combined with isomorphism theories (e.g. Cummings 1996). If we combine the idea that representations are isomorphic with their representeds with idea that psychosemantic norms depends on the norms of proper functioning, we can generate several proposals: for example, the proposal that the relevant mappings are those that the systems were designed to exploit... [30]

Here I do not wish to argue in favor or against of any particular teleological theory of content. For our purposes, what counts is that if any of these theories proves to be sound, it would allow us to induce an externalist

[27] Loewe, 127.
[28] This, of course, does not make teleological functions a set of current dispositions, but a set of selected dispositions.
[29] See for example Bickhard, 1991.
[30] Neander, 2004.

restriction to the contents of color perceptions, thus bypassing the relationalist dilemma. Notice that this restriction is weaker then that imposed by requiring that the contents of veridical perceptions be N-Best operators. A given projection can be the one *that would have had to have been instantiated, had the respective perceptual systems instantiated that experience when functioning properly* (I bet you can't say it without breathing), even if it is not an N-Best operator. Natural selection is very clever at designing solutions, but it is not perfect!

This proposal, however, is similar to the previous one in a relevant respect. Under any understanding of proper functioning, a perceptual apparatus functions properly only if it exploits all (and only) the photoreceptors that natural selection has designed for it. So, implicitely, this proposal also makes color perceptions true relative to the dimensionality of phenomenal color space. In fact, it does more: it makes them relative to specific *kinds* of perceptual systems.[31]

4. It's a quasi-colorful world

According to teleological relationalism, there will be as many classes of color properties as there are *kinds* of perceptual systems. So much the worse, I say, for the intuition that we and the bees, for example, represent (exactly) the same properties of a flower when we're looking at one. One advantage of my account, I think, is this. As I shall argue, although two different creatures might be representing different properties of a flower, when looking at it, it is still possible to say that these properties belong to a common natural kind. This is exactly as it should be, if the different properties in question are to deserve the name of colors.

Colors, according to teleological relationism, are relational properties of physical objects and perceptual apparatuses. These properties are represented in such a way that their proper bearers, relative to these representations, are the physical objects. Empirical discoveries allow us to say that color properties have a narrow correlate (reflectance profiles). Very different color representations (for example harbored by very different creatures), may represent different color properties of the same physical object, under the same environmental conditions.

[31] Not necessarily, as I have said, these kinds must be biological kinds.

Very likely, the subjective experience of these representations will also be very different. This is, I think, what should be expected. Any realist theory of colors also has the consequence that there are color properties that we humans cannot represent. Therefore, I don't take this to be a peculiar drawback of my account. Nor I think that this is a serious objection for anyone.

Up to a certain extent, whether a certain property is a *color* property, is a terminological issue. In my account, what all color properties have in common, is the *character* of their representations. My representation of a flower and a bee's representation of the same flower share the same character. This is to say that both I and the bee fix the content of color representations in the same way. Character, remember, is the map from context to content. My context is different from that of a bee's, whence the fact that we represent different properties. The content of my representation and that of the bee also (necessarily) share a common relatum: the colored object. Moreover, the two perceptions represent color properties that also have their narrow correlates in common: the reflectance profiles. All color properties (regardless of their class of provenience), are arranged in a metrical space that allows us to say which color property is the more accurate approximation to its correspondent narrow correlate.

Finally, most cases of disagreement that one may want to accommodate, e.g. the case of the red vs/ blue tomato of our example, can be easily accommodated by teleological relationalism. If the tomato looks blue to me, then my retina is not functioning properly, hence the content of my experience is (robustly) that the tomato is blue, while in fact it is red. This is as much room for error as teleological relationalism can afford. I think it is enough room. This is as far as the substantial, non-terminological dispute can go, I believe. Whether we want to call properties and perceptions that have that much in common "colors", I submit, is now a terminological issue.

4.1. So, is the world colored?

Suppose I'm right about what properties we represent when we have color experiences. What should we make of the claim that the world is objectively colored? Retinas necessarily contribute to instantiating color properties. Just as one can pick up color properties so that physical objects are their proper bearers (our representations do), one can also pick them up so that retinas are the proper bearers. I am doing it now while writing, and you are doing it while reading these words. Doesn't this make color properties mind-dependent? Doesn't it violate the externality condition?

I think not. It is important to distinguish mere environmental differentiators from mental representations. Any metal bar, for example, implicitly categorizes environments that have the same temperature, because its length co-varies with temperature in a lawful way. However, we would not say that any metal bar *represents* the temperature of the environment. I will say that metal bars are environmental differentiators. Whether a metal bar also represents temperature, depends on whether an organism (or cognitive system) *uses it* to represent temperature. It is undeniable that certain metals are particularly apt to be so used. This is why we can build thermometers exploiting this property. However, thermometers only represent temperatures relative to our using them as representations. A thermometer, in and of itself, is a mere environmental differentiator.

While being a representation is certainly a "mental property", being an environmental differentiator is not. Now, according to my proposal, retinas (and similar perceptual apparatuses) are necessarily among the relata of color properties. Seen "from the side of the retina", so to speak, color properties are properties of the retinas. The narrow correlate of color properties, when these are viewed "from the side of the retina", are all those physical properties in virtue of which retinas act as environmental differentiators. As I noted above, however, these properties are not mental properties. Plausibly, a property is mind-dependent only if it is necessarily co-instantiated with some mental properties (whatever these are). The instantiations of the projector operators in no ways entail the (co-)instantiation of mental properties. In fact, we have seen, they only entail the (co-)instantiation of *environmental differentiators.* It follows that color properties, under my account, are not mind-dependent. Having said that, we can conclude that if I'm right, the world is indeed a colorful place, for color perceptions are often veridical!

Yet someone might still be perplexed at this solution. "All right", my detractor could concede, "the properties you call colors are not mind-dependent, but they are certainly different from the brain-independent properties that we were expecting!" I have already noted that the idea that objects should be the proper subjects of color ascriptions is due to the particular mode of presentation of color properties in our phenomenal world. We have seen how the ultimately extrinsic or intrinsic nature of color properties is an empirical question, and not one that could be accessible to phenomenological introspection.This should be enough to dispel the impression that my account entails some form of eliminativism. However, I think that one can say more to diffuse this worry.

5. Conclusions

Color properties, according to the view put forward here, are objective properties that we use to gather information about distal stimuli. The properties that we represent in color perceptions (i.e. colors) are very similar to their narrow correlates (i.e. reflectances). Such similarity, we have seen, can be measured. The similarity explains why our color perceptions can be used in (approximately) sound inductive reasoning about properties of the objects that are not themselves relative to retinas. Being a ripe banana, or a venomous spider, for example, are certainly not properties that depend in any way on our retinas, let alone on some mental properties. What explains our capacity to infer retina-independent properties of bananas and spiders, I submit, is the (measurable) degree of similarity between the contents of our perceptions and their narrow correlates.

Summing up, colors are not basic properties of the world (see the definition of basic property in section 1.5), but they are extremely close to some basic properties of the world. Those readers who insist that the world can only be said to be *really* colored if colors are basic properties of objects, will have to content themselves with the claim that the world is quasi-colored: colors are quasi-basic properties. I have argued on a number of grounds that the properties that we represent in our color experiences should best be thought of as relational properties of physical objects and perceptual apparatuses. In particular, I have argued that color properties are those that instantiate the operators that projects the infinite-dimensional space of spectral reflectances onto the finite color spaces that organisms perceive.

Colors, under this account, are objective, mind-independent properties of the world. Teleological relationalism, that is, allows us to claim that the world is populated by objectively colored objects[32], and that most of our color perceptions are veridical. The account has been shown to be immune from

[32] As I said, those particularly picky about real colors being basic properties (see section 1.5), will have to content themselves with saying that the world is populated by quasi-colored objects.

standard objections to relationalism. In particular, it has been argued to resist standard faultless disagreement counterarguments.

References

Bickhard, M. H., 1991, How to build a machine with emergent representational content, *CogSci News* 4(1), 1–8.

Campbell, J., 1994, A Simple View of Color, in Haldane, John, and Wright, Crispin (eds.) (1994), *Reality, Representation and Projection*, Oxford: Clarendon Press: 257-69.

Cohen, J., 2004, Color Properties and Color Ascriptions: a Relationalist Manifesto, Philosophical Review, 113(4):451-506

Dretske, F., 1995, *Naturalizing the mind*, Mit Press, Cambridge, Ma.

Hardin, C. L., 1988/1993, *Color for Philosophers*, Indianapolis, Ind.: Hackett.
--- 2003, A Reflectance Doth Not a Color Make, *The Journal of Philosophy*, 100: 191-202.

Hilbert, D. R., 1987, *Color and Color Perception*, Stanford, Calif.: C.S.L.I.

Byrne, A. and Hilbert D., 2003, Color Realism and Color Science, *Behavioural and Brain Sciences*, 26: 3-21.

Millikan, R., 1990, *Biosemantics, Mind and Cognition* (W. G. Lycan, ed.), Blackwell: 221–230.

Rudin, W., 1987, *Real and Complex Analysis*, McGraw-Hill