

Perceptual Variation, Color Language, and Reference Fixing. An Objectivist Account*

Mario Gómez-Torrente
Universidad Nacional Autónoma de México

Abstract. I offer a new objectivist theory of the contents of color language and color experience, intended especially as an account of what normal intersubjective variation in color perception and classification shows about those contents. First I explain an abstract account of the contents of (utterances of) color and other gradable adjectives; on the account, these contents are certain objective properties constituted in part by contextually intended standards of application, which are in turn values in the dimensions of variation associated with the adjectives. Then I propose an explanation of normal variations in linguistic color classification; these are postulated to be effects of differences in intended standards for the color adjectives appearing in the classifying predicates. Next, I consider a potential objection to this explanation, based on the suggestion that contextual content should be more accessible than the explanation predicts. In reply, I point out that contextual content is occasionally opaque to unsophisticated reflection. Finally, I present a companion account of the contents of color experiences on which these represent an object as lying on certain salient intervals in the chromatic dimensions, and I show how the account accommodates intersubjective variation in color perception.

Introduction

Subjectivist accounts of color language and color experience postulate that the color properties and phenomena referred to by color words and represented by chromatic experiences are metaphysically constituted by perceptual relations between objects and perceiving subjects. Objectivist accounts postulate that such color properties and phenomena are not metaphysically constituted by relations of this kind.¹ Recent

* I am grateful to many people for their reactions to this material in question sessions, conversations or e-mail exchanges. Among those who had an influence on the text I especially remember Alexis Burgess, Eleonora Cresto, Manuel García-Carpintero, David Hilbert, Mark Eli Kalderon, Carlos López Beltrán, Adam Pautz, Jesse Prinz, Olivier Roy and Pedro Stepanenko. Axel Barceló, Jonathan Cohen and an anonymous referee for *Noûs* deserve special thanks for their extensive sets of comments. My thanks also to Alberto Barrañón for his help with the figures. Parts of this research were supported by the Mexican CONACyT (CCB 2011 166502) and by the Spanish MICINN and MINECO (research projects FFI2008-04263 and FFI2011-25626). A generous sabbatical leave grant from the DGAPA-UNAM allowed me to complete work on this paper.

¹ Jonathan Cohen (2004, 2009) has recently used the label “relationalist” for the views I call “subjectivist”. An error-theoretic view of a certain property, on the other hand, denies that this property is instantiated in the actual world. An objectivist view of the color properties can in principle be coupled with an error theory of them, but I will often reserve “objectivist” for non-error-theoretic views. Classical dispositionalist theories such as those of McGinn (1983), Peacocke (1984), Wright (1988) and Johnston (1992) and recent theories such as Cohen’s come out subjectivist on the present understanding. Physicalist theories such as those of Smart (1963), Armstrong (1968), Hilbert (1987), Matthen (1988), Byrne and Hilbert (1997, 2003) and Tye (2000) come out objectivist. Error-theoretic views include those of Hardin (1988), Boghossian and Velleman (1989), Maund (1995) and Averill (2005).

consideration of facts of normal intersubjective variation in color perception has led to the formulation of difficult problems for existing (non-error-theoretic) objectivist theories of the contents of color language and color experience. The present paper offers a new (non-error-theoretic) objectivist theory that aims to compare favorably to existing theories, especially as an account of what intersubjective variation shows about those contents.²

The theory is inspired by some natural considerations on contextual reference and the accessibility of content derived from well-known insights from the philosophy of language, considerations which initially suggest an account of the contents of color expressions and expression utterances, but which indirectly recommend an allied account of a relevant part of the contents of chromatic experiences. These two accounts will be presented below.³ While the problem of the content of color experience is in a number of ways more fundamental than the corresponding problem about color language, it is natural to seek cues for a resolution of the former problem by considering potential solutions to the latter problem. For it is natural to suppose, at least cautiously, and along with most of the literature on the philosophy of color, that the color properties signified by color words are represented by color experience.⁴ While the route to theories of the content of color

² Besides the intersubjective variation arguments that are the focus of the considerations in this paper, the recent subjectivist literature develops *intrasubjective* and *interspecies* variation arguments (see e.g. Cohen 2009, ch. 2). In my view, these different kinds of arguments clearly exploit very different phenomena and intuitions, and objectivist responses are bound to be based on different considerations in each case. In fact, I think that existing objectivist theories already have essentially appropriate responses to intrasubjective and interspecies variation arguments, and I would be happy to adapt some versions of those responses to the present account. For some of those responses see e.g. Tye (2007), ch. 7, and Byrne and Hilbert (2003), 13-16.

³ Questions concerning the contents of color experience will be dealt with in a comparatively more superficial way than questions concerning the contents of color language. The paper focuses on some of the highly complex issues concerning these latter questions, which require their own extended independent treatment. Some issues concerning the present account of the contents of color experiences not examined here will be treated elsewhere.

⁴ See e.g. Byrne and Hilbert (2003), 4: “it is natural to suppose that there are intimate connections between a certain salient property that tomatoes appear to have, the word ‘red’ and the concept RED; in particular, the word ‘red’ refers to this property”. This is not to say that the natural supposition cannot be questioned; and it has been questioned, in particular by philosophers who think that many words for sensible properties and phenomena stand for physical properties of certain kinds but experience represents objects as having properties that are somehow “purer” or “simpler” than any physical properties of those kinds. (See e.g. Gendler and Hawthorne 2006, 7ff. for an introduction.) But even philosophers who think that color experiences represent certain broadly chromatic properties other than those named in language typically think also that color experience does represent the named properties as well, and thus do not reject the

experience via theories of the content of color language may have been somewhat neglected, it clearly seems good methodology to take the route and see where it leads us.⁵ The considerations in this paper suggest that this route has much to offer.

Section 1 describes one main argument against (non-error-theoretic) objectivism based on facts of normal intersubjective variation in color perception and linguistic classification. The argument concludes that these variations are compatible only with error-theoretic or subjectivist views of the contents of color language and color experience. The section also explains some of the unsatisfactory aspects of several objectivist attempts to deal with this argument. Section 2 explains a basic ingredient of the present theory, an abstract account of the contents of (utterances of) color and other gradable adjectives. On the account, these contents are certain objective properties constituted in part by contextually intended standards of application; and these standards are in turn identified with values in the dimensions of variation associated with the adjectives. Section 3 describes the main part of the theory, an explanation of normal variations in linguistic color classification; the theory explains these as effects of differences in intended standards for the color adjectives appearing in the classifying predicates. The section also argues that the resulting objectivist rebuttal of the argument of section 1 compares favorably to the existing objectivist reactions also discussed there. Section 4 considers and rejects a potential objection to the theory. The theory traces some variations in color classification to differences in intended standards, differences that could only be guessed on the basis of moderately sophisticated knowledge of facts about color, unavailable to a typical current speaker. The objection is then that this possibility is at odds with the essence of context-sensitive resources, which might be suggested to lie in part in their interpretability without help of sophisticated knowledge. The objection is defused noting that, in several cases, including the case of color adjectives, contextual content can occasionally be opaque to unsophisticated reflection, and in fact as opaque as the content of natural kind terms is susceptible of being under well-known and plausible accounts of reference fixing. The final section 5 presents the objectivist account of the contents of color experience naturally suggested by the preceding account of the contents of color language. On this account, color experience represents an object as lying on certain intervals in the chromatic dimensions, as occupying a particular locus in those dimensions that stands in certain relations to other significant loci that mark the bounds of those intervals. The facts

natural supposition (see e.g. Shoemaker 1994, 2006, and Chalmers 2006). It is worth stressing also that the mentioned way of questioning the natural supposition does not spare sensible properties traditionally considered as primary and objective, such as shape, and is thus not directly relevant to a resolution of the debate between (non-error-theoretic) objectivism and (error-theoretic views and) subjectivism.

⁵ The neglect is of course not universal. The “language to experience” route has been traveled at least by Jackson (1996, 1998), Lewis (1997), and arguably Kripke (see Gómez-Torrente 2011).

concerning normal intersubjective variation in color perception again receive a natural explanation on this account.

1. An intersubjective perceptual variation argument against color objectivism

In a certain psychological experiment, a large group of normal subjects⁶ are presented with several chips that when illuminated with some fixed normal illuminant reflect lights of distinct dominant wavelengths in the 490-520nm zone of the spectrum, and are asked to point to the one that is (closest to) unique green, i.e. a shade of green that is neither somewhat yellow nor somewhat blue, and classify the others as either somewhat yellow or somewhat blue.⁷ The answers vary considerably. There are for example chips that a considerable group of normal subjects classify as green but neither somewhat yellow nor somewhat blue, and another considerable group of normal subjects classify as green but somewhat blue; there is even a large group of subjects who classify those same chips as blue but neither somewhat green nor somewhat red (i.e. as unique blue).⁸ All this suggests that there are phenomenal differences among the experiences that different normal subjects have of the same chip in the same normal viewing conditions, which in turn induce the differences in their linguistic classifications. Those phenomenal differences are in all probability caused by a number of physiological factors resulting in minute interpersonal differences in the absorption curves of cones (see e.g. De Valois and De Valois 2000, 135ff.).

⁶ The tests of normality for chromatic perception (such as the Farnsworth test) are based essentially on the idea that the vast majority of subjects (the normal subjects) ought to order in the same way a series of objects along the dimensions of their apparent color (especially hue). (Not all subjects who are normal by these tests are able to make side by side color distinctions with the same apparent acuity. However, the differences in linguistic classification about to be mentioned are robustly independent of the apparent acuity of the subjects.)

⁷ “Yellowish” is also used instead of “somewhat yellow” (and similarly for the other color adjectives). I take these to be synonymous in the relevant contexts.

⁸ The classic reference for experiments of this kind is Hurvich *et al.* (1968). There are especially accessible descriptions in Hardin (1988), 79ff. and Block (1999); see also Kuehni (2014) for a somewhat technical presentation of a large range of experiments of this kind. In a related kind of experiment (Malkoc *et al.* 2005), subjects are asked to classify chips as either “red”, “orange”, “yellow”, “green-yellow”, “green”, “blue-green”, “blue” or “purple”; some classify some chips as red and not orange and others classify those same chips as orange but not red; some classify some chips as orange but not yellow and others classify those same chips as yellow but not orange, and so on. (For discussion see Tye 2006, Cohen *et al.* 2006 and Byrne and Hilbert 2007.) Although I will focus on the classifications involving the unique hues, the theory in this paper is designed as an account of the classifications in this latter experiment as well.

One main argument against non-error-theoretic objectivism based on these facts goes as follows:

It seems natural to think that, if χ is a chip that a subject UG perceives and classifies as green but neither somewhat yellow nor somewhat blue and another subject BG perceives and classifies as green but somewhat blue, there is no reason why one of these experiences or attributions should be veridical or correct to the exclusion of the other. After all, the subjects are both normal according to reasonable tests, and the differences between them appear to be traceable merely to their cones having slightly different absorption curves, neither of which would seem to be “better” in any sense than the other. It does seem natural to suppose that either both UG and BG have veridical experiences and make true attributions or both have falsidical experiences and make false attributions. If the latter, then an error theory must be correct. And if the former, then the properties of being green but neither somewhat yellow nor somewhat blue and of being green but somewhat blue must be subjective. For “green but neither somewhat yellow nor somewhat blue” and “green but somewhat blue” are incompatible unless they signify subjective properties. If we suppose that they do, then both BG’s and UG’s experiences and attributions can be taken to be correct: BG perceives and implicitly classifies χ as something like *looking green but somewhat blue to BG* and UG perceives and implicitly classifies χ as something like *looking green but neither somewhat yellow nor somewhat blue to UG*. There is no contradiction between UG and BG because the properties they talk about come constituted by perceptual relations involving UG and BG; if we don’t assume such constitutive relationality, e.g. if we suppose that those properties are objective, the contradiction seems inevitable.⁹

To be sure, this argument can seem convincing (or at least convincing at first sight) only to moderately sophisticated language users with knowledge of the spectrum of wavelengths, of cone absorption curves, of what counts as a normal perceiver according to reasonable tests, etc. At least some language users lacking this information or unable to digest it think, when presented with the raw data of the subjects’ classifications in the mentioned psychological experiments, that at most one of UG or BG can be correct. We will come back to this in section 4.

As I see things, this argument doesn’t merely suggest the dilemma error theory/subjectivism, but favors subjectivism over an error theory, because the experiment appealed to doesn’t simply generate an impression in sufficiently sophisticated readers that UG and BG are equally right or wrong: it also generates an impression that they are equally right. But subjectivism is not without its counterintuitive consequences. One of them lies in what I take to be an extended further impression that the color properties are physical properties “out there”, not involving subjects or perceptual relations at all. A companion

⁹ For places where forms of this argument are developed and embraced, see Cohen (2004), (2009, ch. 2), Hardin (1988, 79ff.), (2003). Cohen chooses the subjectivist horn, Hardin the error-theoretic one.

difficulty involves a related idea that can be somewhat more precisely stated, the idea that color predicates have what we might call “objective modal stability”, a property closely related to rigidity. Say that a (possibly context-sensitive) monadic predicate P is objectively modally stable iff, for any given object o and context c, there are no two different worlds w and w' which are the same as to the objective properties of o but are such that o is in the extension of P with respect to (c, w) but fails to be in the extension of P with respect to (c, w'). If UG says *If my eyes were just a bit different, χ would stop being green and neither somewhat yellow nor somewhat blue*, what he says will probably sound false to him (it did sound false to all my informants). This is plausibly explained under the supposition that “green and neither somewhat yellow nor somewhat blue” is objectively modally stable: the antecedent of the counterfactual takes UG to a world w' that is at most subjectively different from the actual world w, and UG thinks that in w' (and with respect to the same context of his utterance) “green and neither somewhat yellow nor somewhat blue” doesn't change extension. But UG's utterance is true on a standard unsophisticated subjectivist view, on which what he says is something like *If my eyes were just a bit different, χ would stop being green-but-neither-somewhat-yellow-nor-somewhat-blue-looking (to me)*. On a standard unsophisticated subjectivist view, “green but neither somewhat yellow nor somewhat blue” is not objectively modally stable.¹⁰

Problematic as error theories and subjectivist theories may be, existing (non-error theoretic) objectivist theories and their reactions to intersubjective variation arguments seem to me to incur unreasonable costs. Let me now briefly review some of these theories and the costs they incur, so that I can later indicate how the account suggested in this paper avoids those costs. The theories in question obviously cannot be dismissed merely on the basis of the following brief discussion. But my main aim in this paper will not be to give knockdown arguments against rival views, but just to note that my view compares favorably to them in a few key respects.

¹⁰ On some sophisticated subjectivist views, color predicates are objectively modally stable. “Rigidified” properties like *being disposed to look green but neither somewhat yellow nor somewhat blue to UG as actually constituted in the actual experimental conditions C* are subjective, but a predicate that signifies this property is presumably modally stable. (See e.g. Shoemaker 1986, or Cohen 2009, ch. 7 on “role-functionalism”, for essentially similar proposals.) But an explanation on which “green but neither somewhat yellow nor somewhat blue” stands for a non-rigidified objective property is arguably simpler (*ceteris paribus*), much in the same way in which an explanation of the rigidity of proper names in terms of their being semantically unstructured is simpler (*ceteris paribus*) than one on which their content is given by sophisticated rigidified descriptions. Further, the “rigidifying” explanation needs to postulate that speakers with stability intuitions who do not have at their disposal rigidifying concepts are not really competent users of color concepts; I think this is implausible. See also note 16 and the surrounding text below for further criticism of related views.

A prominent reaction can be found in Alex Byrne and David Hilbert (2003, 2004) and Michael Tye (2006), among other authors. They reject the idea that either both UG and BG must have veridical perceptions or both must have falsidical perceptions, and thus the assumption that either both UG and BG must make true attributions or both must make false attributions. On this view, it may well be that at most one of either BG or UG has a veridical perception of χ 's objective chromatic property, and thus that at most one of either BG or UG makes a true attribution about χ . It may well be that we just don't know at present, and perhaps will never know, who among BG or UG or the other normal subjects with phenomenally different experiences is right about χ . It seems fair to say that this sort of reply is very unconvincing in view of what we know about the physical and physiological phenomena underlying color perception. We have no clue as to what fact of the matter, if this is to be describable in terms of the relevant known physics and physiology, could make BG or UG or some other subject be in the right to the exclusion of the others. Moreover, in part because of this, but also insofar as purely linguistic judgments are concerned, neither BG nor UG would seem to have made any mistake in their use of color words. It is hard to avoid the conclusion that objectivist accounts that respected the impression that all the subjects are equally right would enjoy a definite advantage over this kind of reaction.

One such objectivist account has been proposed by Mark Eli Kalderon (2007).¹¹ He concedes that both UG and BG have veridical perceptions and make true attributions, but rejects the assumption that "green but neither somewhat yellow nor somewhat blue" and "green but somewhat blue" as used by UG and BG must signify incompatible properties unless they signify subjective properties. Kalderon claims that being green but neither somewhat yellow nor somewhat blue and being green but somewhat blue must in some way be compatible properties, and that while UG is able to see how χ possesses the first of these properties but not how it possesses the second, BG is able to see how χ possesses the second but not how it possesses the first. I agree that the weakest spot in the argument is the transition from the compatibility of the attributed properties to their subjectivity, and the theory below can be seen as one possible way of fleshing out Kalderon's abstract idea, as it will propose that UG and BG do indeed refer to two different chromatic properties of χ . A main problem with Kalderon's abstract idea, however, is how it can adequately deal with the fact that "not somewhat blue" in UG's attribution *χ is green but neither somewhat yellow nor somewhat blue* and "somewhat blue" in BG's attribution *χ is green but somewhat blue* are contradictory predicates.¹² Kalderon's proposal doesn't appeal to any sort of relativization that could make the properties signified by those predicates compatible. But if no relativization of any sort is invoked, it's hard to see what fact of the matter could make

¹¹ See also Mizrahi (2007) and Allen (2009).

¹² See the related objection in Hardin (2006, 343).

it the case that “not somewhat blue” in the first attribution and “somewhat blue” in the second are not contradictory.¹³

An objectivist reaction that invokes one kind of relativization is one favorably considered by Brian McLaughlin (2003).¹⁴ The idea is that the reference of the predicate “somewhat blue” when used by UG is fixed by the a priori equivalent descriptive predicate “has the physical property that in normal conditions makes something look somewhat blue to UG” (mutatis mutandis for BG)¹⁵, but this doesn’t imply that that predicate stands for the subjective property *having the physical property that in normal conditions makes something look somewhat blue to UG*; in fact, it stands for the physical property that in normal conditions makes something look somewhat blue to UG, which we can suppose to be objective (mutatis mutandis for BG). This is consistent with UG’s attribution *χ is green but neither somewhat yellow nor somewhat blue* and BG’s attribution *χ is green but somewhat blue* being true attributions of compatible objective properties. And it suggests an explanation of why “not somewhat blue” in UG’s attribution *χ is green but neither somewhat yellow nor somewhat blue* and “somewhat blue” in BG’s attribution *χ is green but somewhat blue* are not contradictory predicates, as “somewhat blue” is postulated to be a priori equivalent with different descriptive predicates in the two cases. But one serious problem with this view is that “somewhat blue” and “has the physical property that in normal conditions makes something look somewhat blue to UG” are not a priori equivalent. It’s imaginable to UG that some things that are somewhat blue are in some way precluded from looking somewhat blue to him in certain conditions he could not have excluded a priori as

¹³ Cohen (2009, 81), on behalf of Kalderon, considers as a way of avoiding the contradiction the possibility of denying that “not somewhat blue” in the first attribution and “somewhat blue” in the second are both intended to apply to χ . But I take this way out to lack independent motivation, and should hope that Kalderon would not be tempted by it.

¹⁴ It is also closely related to an objectivist reaction to other subjectivist considerations, favorably considered by Frank Jackson (1996). Compare the discussion of these views, under the label “realizer functionalism”, in Cohen (2009), ch. 7.

¹⁵ Actually, in McLaughlin’s *preferred* view, the idea is that the predicate “somewhat blue” when used by UG *means* something like “has the physical property that in normal conditions makes something look somewhat blue to UG”. One problem with this is that it implies that UG’s predicate “somewhat blue” is not objectively modally stable. McLaughlin is not too worried about this, as he doesn’t take the intuition of objective modal stability to be robust. Another problem is that this view implies that a color predicate like “somewhat blue” means different things in the idiolects of different speakers, which I take to be fairly implausible. But there is another, in my view more basic, problem with all views of this sort, as noted in the text.

abnormal.¹⁶ It's just implausible that the conventions governing a color predicate should stipulate an a priori equivalence between it and a predicate ascribing the association in normal conditions of a certain perceptual quality with a certain property in a particular person, even if this person is the speaker. An account that used some kind of relativization without yielding such equivalences would thus be preferable.

2. Reference fixing for color and other gradable adjectives

Let me begin the exposition of the account of this paper with one of its main ingredients, an account of reference fixing for color adjectives. This relies on some by now fairly standard ideas concerning the ways in which gradable adjectives in general are plausibly context-sensitive. A gradable adjective like “hot” or “big” seems susceptible of different standards of application along one or several assumed dimensions of variation.¹⁷ On a standard appropriate when we are thinking of washing some clothes, the heated water that comes out of the tap counts as hot, but on a standard appropriate when we are thinking of scalding a chicken, that same water doesn't count as hot and even counts as cold. On a set of standards appropriate when I'm thinking of an apartment where I can live with my big family, the micro-apartment I once occupied in a visit to Seoul counts as not big and even as (very) small, but on a standard appropriate when I'm thinking of a short touristic visit to that crowded city, that same apartment counts as big. In the case of “hot” the relevant dimension is the dimension of amounts of heat that we have come to identify with the physical dimension of temperature, and in the case of “big” the relevant dimensions (at least in many uses) are those of height, width and depth.

The relevant standards are usually taken to be fixed by contextual aspects of the use of the gradable adjective, fundamentally by the utterer's generally vague and often merely implicit intentions to set such standards. What apparently happens is that in a particular context of use of “hot”, a thermal property gets contextually determined as the content of “hot” with the help of some implicit rule that operates on the contextually intended standard and delivers a corresponding property, which will be constituted in part by that standard.¹⁸ In the case of “hot” such a rule might be similar to this: *An utterance of “hot” refers*

¹⁶ This objection is Kripke's; see Gómez-Torrente (2011). Accounts such as the subjectivist “rigidifying” accounts mentioned above in note 10 are liable to a related criticism.

¹⁷ See e.g. Kennedy (2007) for a recent example of use of standards in a semantic theory of gradable adjectives, and for references to the technical semantic literature. In the following presentation I attempt to make my use of standards fairly non-technical by comparison.

¹⁸ These are orthodox assumptions I share with most of the literature alluded to in the preceding note. More specifically, I will be assuming that the content of a gradable adjective at a context of utterance (of a sentence S containing the adjective) is a property which is partially constituted by standards determined by the context of utterance. This property then partly constitutes a proposition (expressed by S at the context)

to the property of having heat in at least the minimum rough amount intended by the utterer, in the case of “big” the rule (at least in many uses) might be similar to this: *An utterance of “big” refers to the property of having a height, a width and a depth of at least the rough values intended by the utterer.* Expanding the Kaplanian terminology for pure indexicals and demonstratives, we may call these rules the *characters* for the adjectives. The character for “hot”, for example, when coupled with the standard 30°C for “hot” intended in a particular context, delivers as the content of “hot” in that context the property of having heat in at least the minimum rough amount 30°C. One further presumable fact is that there are meaning restrictions on how the properties signified by uses in the same context of “hot” and “cold” (or “big” and “small”) must be related.¹⁹ Fundamentally, those properties must be incompatible, which imposes, say, the restriction that the contextual standard for “hot” must be a temperature value higher than the temperature value that is the standard for “cold”.

The possibility of different contextual standards is compatible with the objectivity of the properties characterized and constituted by such standards. To begin with, the determinable properties of hotness and bigness are presumably not constituted by perceptual relations between objects and subjects. Furthermore, the more determinate properties signified by contextual uses of “hot” and “big” are presumably objective, even if standards of application for “hot” and “big” are somehow selected by subjects (utterers). Moreover, and crucially, these determinate properties are presumably objective even though these standards are often selected by an utterer via intentions formed with the *help* of *perceptual qualities* as felt, remembered or imagined by him.²⁰ For example, the utterer may

which is evaluable at indices partly constituted by times and worlds. Though this format is the one I favor for the influence of context on content in the case of standards for gradable adjectives, some of the views in this paper might conceivably be applied under other formats for such influence explored in the recent literature. (See e.g. Lasersohn (2005) and MacFarlane (2009) for representative examples.) For example, standards might appear along times and worlds as elements of the indices with respect to which the proposition is evaluated; and they might be postulated to be provided by some context different from the context of utterance. Here I cannot fully argue that my orthodox framework is more adequate than other alternatives, but see note 24 below.

¹⁹ A presumable character for the adjective “cold” is *An utterance of “cold” refers to the property of having cold in at least the minimum rough amount intended by the utterer* (where the “cold” that appears in the character is the noun, not the adjective); here the assumed dimension of variation is the dimension of amounts of cold, that we have *also* come to identify with the physical dimension of temperature. A presumable character for “small” is *An utterance of “small” refers to the property of having a height, a width and a depth of at most the rough values intended by the utterer.*

²⁰ Often but certainly not always. Nowadays most of us feel comfortable hearing astronomers speak of stars as hot when their temperature is above 10000°C, and cold when it is below 4000°C (which in usual

implicitly intend the contextual reference of “hot” to be fixed by “the property of having heat in at least *this* amount”, where “this amount” picks out a temperature value via a deferred demonstration where the directly demonstrated thing is a remembered thermal perceptual quality. But this is no obstacle to the objectivity of the selected property: the determinate hotness property intended by the utterer in such a case presumably consists in the possession of a temperature equal to or greater than a certain rough temperature value (the temperature value *corresponding to* the relevant remembered quality), and is thus a property uncontroversially not constituted by perceptual relations. In the case of “big”, a determinate bigness property presumably consists (in many contexts) in the possession of a height, width and depth equaling or surpassing certain numerical values of height, width and depth, and is thus a property uncontroversially not constituted by perceptual relations. And this is so even if speakers can refer to such properties via the mechanisms for fixing size standards, and even if these mechanisms allow them to fix the relevant standards via appeals to the ways those standards are sensed, remembered or imagined by them.

The fact that the contents of utterances of adjectives like “hot” and “big” are objective properties, even though the content-fixing mechanism for those utterances involves subjective aspects of the utterers, deserves emphasis. An analogy with a simpler case may help clarify this crucial distinction. Suppose you and I are looking at a group of people, and that a certain man in the group appears happy. If I then say *He is happy*, the content of my utterance of “he” (at least under the most widely accepted views on the contents of indexicals of this kind) is simply the man I intend to call attention to, the object itself, and does not involve aspects of my subjective life. However, the way in which the content of that utterance is fixed does involve aspects of my subjectivity. The conventions that determine what the content of such an utterance can be say, roughly, that this content is to be the male the utterer intends as he or she makes the utterance. And my intention to refer to the happy man is not the only aspect of my subjective life that plays a role in the fixing of a content for my utterance of “he”; for example, and especially, I form this intention with the help of a previous or concurrent perception of the man in question. But this intention and this perception (at least under the most widely accepted views on the contents of indexicals of this kind) do not make it into the content of my utterance of “he”. Analogously, the conventions that determine the content of an utterance of the adjectives “hot” or “big” require this content to be fixed with the intervention of intentions, and specifically of intentions typically formed with the help of previous or concurrent perceptions. These intentions, in particular, fix the objective standards that partially constitute the properties that provide the contents of utterances of “hot” and “big”; and the intentions are formed with the help of ways of identifying these standards in

contexts would count as hot). These standards are intended via theoretical descriptions of those temperatures, not via felt, remembered, or even imagined perceptual qualities.

perception. However, the intentions themselves and the perceptions that contribute to their formation are not constitutive of the contents of those utterances. These contents are objective thermal and size properties.

If the determinate properties referred to by utterances of “hot” and “big” are objective, then it is to be expected that both “hot” and “big” are objectively modally stable. And this is in fact the case. If I am in a context in which the water is not hot according to my washing standard of 30°C, and I say *If my nerve endings were more sensitive, the water would be hot*, what I say sounds false. And in any typical context where I say of a small apartment *If my vision were microscopic, this apartment would be big*, what I say will equally sound false.²¹ This strongly confirms the idea that the contents of these utterances of “hot” and “big” are objective properties, not constituted by the subjective aspects that play a role in reference fixing.

Like “hot” (and “cold”) and “big” (and “small”), also “green”, “blue”, “red” and “yellow”²² are susceptible of different standards of application along the dimensions of variation presupposed by our use of them, especially hue.²³ On some standards along the dimension of hue, an aquamarine jewel counts as green, but on others it counts as blue. Imagine that you are shopping for a jewel you want to give as a present to a friend who loves all kinds of greenish jewelry. If you ask the jeweler *Show me some green jewels*, you intend your utterance of “green” to apply both to emeralds and to aquamarine beryl gems of substantively different hues. However, if you are shopping for a jewel for a pickier friend, who likes only the pure greens of the gem stones that you often see in the jewels he wears, then for you given your interests, the aquamarines count as not green and even blue, while the emeralds count as green and not blue. Plausibly, you intend matching standards in the physical dimension of hue, via some seen or remembered perceptual qualities that you take

²¹ But think of a context in which the water is not hot on my standards for *taking a bath*. Doesn't it sound true to say here *If my nerve endings were more sensitive, the water would be hot?* Maybe, but this is probably due to the fact that here “hot” is used to mean (non-semantically) something like “causing impressions as of heat”, presumably because in this context I'm especially interested in effects on my nerve endings. There are similar examples with “big”. Suppose I'm calmly contemplating a small apartment of reasonable height, width and depth, without any very practical purpose in mind such as calculating if my old furniture or my big family will fit inside, but rather with some vague aesthetic attitude. Then it sounds true to say *If my vision were microscopic, this apartment would be big*.

²² Much of what I will say will hold also, with appropriate modifications, for many words for non-unique colors, like “orange”, “olive”, etc. But I will only deal explicitly with terms for the unique colors.

²³ It may not be clear that the dimension of variation is the same for “red” and “green”, or for “hot” and “cold”. We wouldn't normally say that water frozen at 0°C is *hotter* than water frozen at -10°C, or that tomatoes are *redder* than geranium leaves. But there is a clear continuity between the corresponding dimensions of variation, and they jointly form all-embracing dimensions.

to correspond to those standards. These intentions are again aspects of the context of use of the color adjective, and the corresponding intended standards are presumably constituents of the contextually determined color properties.²⁴

There are of course important differences between the use of contextually set standards for “hot” and their use for “green”. A use of “hot” appears to presuppose only one contextual standard, a rough value in the dimension of temperature starting from which things count as hot. The use of “green” requires the operation, however implicit, of a set consisting of more than one standard, in fact a pair of rough levels in the dimension of hue, a rough level in the dimension of saturation, and a pair of rough levels in the dimension of brightness. Two rough standards in the dimension of hue are required establishing dividing lines between blue and green and between green and yellow; a rough level in the dimension of saturation is required establishing a dividing line between green and white (or grey or black); and a pair of rough levels in the dimension of brightness are required establishing dividing lines between white and (bright) green and between (dark) green and black. Presumably, something counts as green in a context when its hue is between the two contextual hue standards or equals one of them, its saturation meets at least the contextual saturation standard, and its brightness is between the two contextual brightness standards or equals one of them.²⁵

²⁴ That standards of application are constitutive of the properties signified by utterances of color and other gradable adjectives is also made compelling by a number of tests for the existence of covert syntactic arguments. Standards for color adjectives arguably pass, e.g., the binding, licensing and sluicing tests. (See e.g. the useful presentation in Schaffer 2011.) It sounds natural to say *At this jeweler’s, you can find something green for all standards of greenness*, which suggests a bindable tuple of standards argument projected by “green”. It also sounds natural to say *Aquamarines are green on these standards*, which indicates that the existence of a tuple of standards argument licenses the explicitation of some such standards; further, “on these standards” is arguably not an optional prepositional adjunct, as naturalness is not preserved when a clear adjunct is interposed: *Aquamarines are green when not industrially manipulated, on these standards* sounds odd, as opposed to *Aquamarines are green on these standards, when not industrially manipulated*. Finally, sluicing constructions targeting standards arguments are natural, as in *If aquamarines are green, on what standards is this so?*

By contrast, as Jonathan Cohen has remarked to me, *perceiver* arguments for color adjectives fail these tests. To use an example of Schaffer’s, to say *The stinky cheese is blue, but to whom?* sounds odd, which suggests failure of the sluicing test. Cohen says this has led him to abandon his subjectivist account as an account of ordinary color language, and to think of it as a revisionary proposal on which the subjectivist contents of color experience are properly captured. This seems to me to detract from the value of his proposal, for, as noted in the Introduction, there is a widely accepted natural assumption that the color properties signified by color words are represented by color experience.

²⁵ This may be the place to add that I see little justification in recent attempts to argue that color adjectives are sensitive to a contextually filled argument specifying the *part* of an object that the adjective is intended to

Thus, analogously to “hot” and “big”, in contexts of use of “green” the particular property that gets contextually determined is presumably picked by some such implicit character as *An utterance of “green” refers to the property of having a hue of the color green between (or equaling one of) the rough hue values intended by the utterer, a saturation of the color green of at least the rough saturation value intended by the utterer, and a brightness of the color green between (or equaling one of) the rough brightness values intended by the utterer.* There are also some presumable complex meaning restrictions on how the properties signified by uses in the same context of “green”, “blue”, “yellow”, etc. must be related. Fundamentally, the properties signified by “blue” and “yellow” must be incompatible, as must be the properties signified by “red” and “green”; the property signified by “green” can be compatible with the property signified by “blue” and the property signified by “yellow”, but it may also be compatible with just one of these, or with neither (in contexts where “green” is intended to stand for the same property as “green but neither somewhat yellow nor somewhat blue”); and so on.²⁶

What are the dimensions presupposed by our use of “green” and the other color adjectives? And what is the color green mentioned in its character? (And what are the other colors mentioned in the characters of other color adjectives?) The present account will work under any reasonable answer to these questions, provided we grant some suppositions that will sound plausible to any objectivist: (1) The supposition that the relevant content- or reference-fixing mechanisms manage to determine that the dimensions in question are certain objective physical dimensions—just as those mechanisms presumably determine that the dimension presupposed by our use of “hot”, “cold”, “warm”, etc. is temperature, and that the dimensions presupposed by our use of “big”, “small”, etc. are the dimensions of physical height, width and depth. (2) The supposition that the relevant content- or reference-fixing mechanisms manage to determine that the property signified by a use of the adjective “green” in a context is the property of *having a*

apply to (see e.g. Szabó 2001 and Rothschild and Segal 2009), or that they are semantically ambiguous in some related way (see e.g. Kennedy and McNally 2010 and Hansen 2011). I cannot go into this here, but let me stress that these proposals have nothing to do with the more established idea that standards for gradable adjectives contextually fill syntactically real arguments. (Perhaps I should also note that I see equally little justification in other recent views, such as that in Travis 1994, on which the data that motivate the mentioned proposals raise a challenge to orthodox truth-conditional semantics.)

²⁶ This imposes the restrictions that the standards for “blue” and the standards for “yellow” must delimit disjoint hue intervals, and likewise for “red” and “green”; that the standards for “green” in a context can in principle delimit a hue interval that overlaps the interval delimited by the hue standards for “blue”, or the interval delimited by the hue standards for “yellow”, or both, or neither; and so on. These restrictions arguably reflect certain structural characteristics of chromatic representation in perception, which plausibly induce isomorphic structural characteristics in the perceptually represented properties and indirectly the linguistic restrictions in question. See section 5 below for further discussion of structural characteristics.

physical hue value between (or equaling one of) two rough standards set by contextual utterer's intentions on the postulated dimension of hue, a physical saturation value greater than or equal to the rough standard set by contextual utterer's intentions on the postulated dimension of saturation, and a physical brightness value between (or equaling one of) the two rough standards set by contextual utterer's intentions on the postulated dimension of brightness—where the relation of *having* at stake is objective. Finally, (3) the supposition that the relevant content- or reference-fixing mechanisms manage to determine that the thing signified by (a use of) the noun “green”, or the noun phrase “the color green” (as it appears in the conjectured character for the adjective “green”) is a certain objective phenomenon. The main ideas of the present account will be applicable under any objectivist conception of the dimensions of hue, saturation and brightness, of the properties of having values in (intervals of) these dimensions, and of the phenomenon that the color green is, provided it is a conception that satisfies these suppositions. In this sense, the basic ideas of the present account will be just as abstract as the basic ideas of the objectivist accounts reviewed in section 1.

This is not to say that I don't have a preferred view as to the nature of the dimensions of hue, saturation and brightness, as to the properties of having values in these dimensions, and as to what the color phenomena are. For the record, let me say that the dimensions in question must be closely related to the colorimetric dimensions of dominant (and complementary) wavelength, purity and luminance factor, though literally understood these ultimately qualify as subjective dimensions, for they are defined by reference to perceivers. One objectivist possibility, though a rather artificial one, is to identify, e.g., the hue dimension with the dimension of values singled out by the notion of dominant (and complementary) wavelength, taken as a purely extensional function from types of light to real numbers (or tagged real numbers, in the case of complementary wavelength), and similarly for saturation and purity, and for brightness and luminance factor.²⁷ There are less artificial possibilities, however, as it's not so difficult to define in purely objective physical terms functions with values in approximate correlation with the standard colorimetric dimensions. (But it would take us too long to do this here.) On the other hand, the colors, or equivalently, the significations of the *nouns* “green”, “blue”, etc. may all have turned out to be the same thing, namely light (electromagnetic radiation), there being no natural distinction between blue and green, or between green and yellow, etc. (in the same way that heat is presumably just energy and cold is also just energy, there being no natural distinction between heat and cold). And for a thing to “have” a particular hue value (and a particular saturation value, and a particular brightness value) of (say) the color green, may just be for the thing to have the disposition to irradiate light with precisely the

²⁷ Perhaps this purely tabular understanding is the way some philosophers of color and color scientists actually understand the colorimetric dimensions.

corresponding values in the appropriate physical dimensions, in the condition of illumination which is typical or paradigmatic for that thing, which may be fairly vaguely specified.²⁸

As already noted, and again analogously to the case of “hot”, the two rough hue standards, the rough saturation standard and the two rough brightness standards, generally vaguely and often merely implicitly intended by the utterer in a context of use of “green”, will typically be set by the utterer with the help of some seen, remembered or imagined chromatic perceptual qualities. The intention to fix the relevant standards as those corresponding to the perceptual qualities the utterer has in mind is again a contextual aspect that together with the character for “green” fixes a fairly determinate chromatic property as the content of the adjective “green” in the context. But crucially, these determinate properties are presumably objective even if the standards set by an utterer are often selected via intentions formed with the help of perceptual qualities. For example, the utterer may implicitly intend the contextual reference of “green” to be fixed by “the property of having a hue of the color green between (or equaling one of) *these*, a saturation of at least *this*, and a brightness between (or equaling one of) *these*”, where the demonstratives pick out hue, saturation and brightness values via a deferred demonstration in which the directly demonstrated things are certain remembered perceptual qualities. But this is presumably no obstacle to the objectivity of the selected property. For the determinate greenness property fixed as the content of “green” in the context presumably consists in the possession of a physical hue value between (or equaling one of) the two hue standards the utterer intends, a physical saturation value greater than or equal to the saturation standard the utterer intends, and a physical brightness value between (or equaling one of) the two brightness standards the utterer intends (whatever the hue, saturation and brightness dimensions turn out to be). This is thus a property that (under basic objectivist assumptions) is not constituted by perceptual relations. As in the case of “hot” and “big”, although some aspects of the utterer’s subjective life—including especially certain intentions and perceptions contributing to their formation—are plausibly involved in the process of content fixing for utterances of “green”, those aspects do not make it into the content of these utterances.

²⁸ One motivation for this set of hypotheses comes from the suspicion that it compares favorably to reflectance proposals as an account also of the contents of color experiences. It’s implausible to suppose that color experience represents the full spectral reflectance information of an object (let alone the full information concerning its possession of one of the reflectances in a complex set of reflectances). By contrast, I see great plausibility in the hypothesis that a color experience represents an object as having values in certain intervals in the chromatic dimensions, as argued in section 5.

3. Perceptual variation and the contents of thermal and color predicates

Now that we have a fairly developed (if abstract) picture of reference fixing for (utterances of) color adjectives, we can begin to see how it throws light on intersubjective variation arguments for the subjectivity of the contents of color language. In order to see this, it's heuristically useful to construct and consider first a perceptual variation argument analogous to the argument of section 1 but about hotness and coldness.

We could conduct the following experiment on a large group of subjects, all with unimpaired thermal perception on any reasonable criteria. We would present them with several objects of different temperatures around 33°C (roughly the average temperature of the human skin); and we would ask them to touch the objects and point to the one that is thermally neutral, i.e. the one that is (closest to being) neither somewhat cold nor somewhat hot, and classify the others as either somewhat hot or somewhat cold. The answers will predictably vary a great deal,²⁹ to all appearances as a function of a number of intersubjective physiological variations, including minute differences in the activation thresholds of the coldness and hotness thermoreceptors of the subjects.³⁰ It seems natural to suppose that, if ψ is an object that a normal subject N perceives and classifies as neither somewhat hot nor somewhat cold, and another normal subject H perceives and classifies as somewhat hot and not somewhat cold, then either both N and H have veridical perceptions of ψ and make true attributions about it or both have erroneous perceptions and make false attributions. If the latter, then an error theory must hold. If the former, H must perceive and implicitly classify ψ as something like *feeling somewhat hot and not somewhat cold to H* and N as *feeling neither somewhat hot nor somewhat cold to N*.

Presumably this thermal perceptual variation argument will not shake the by now apparently majoritarian view that thermal properties are objective. Why? The reason will certainly not be that most philosophers will be inclined to embrace an (objectivist) error theory or a theory that at most one of N and H is right. Nor will it be that philosophers will suppose that “hot” is a priori equivalent with “feels hot to me” in N’s and H’s idiolects. I submit that the reason why this argument will not impress objectivists about thermal properties will be captured by the account of the semantics of thermal adjectives of section 2. A thermal adjective is susceptible of being used by an utterer together with default (implicit) intentions fixing on (vague) contextual standards for the adjective which

²⁹ “Predictably”, because a variety of experiments have shown that there are robust, wide intersubjective variations in thermal sensation thresholds even among subjects who have unimpaired thermal perception on any reasonable criteria. (It may be noted that the differences are so wide in normal subjects that they cannot be attributed to differences in capacities for thermal discrimination, which are very fine-grained in subjects with unimpaired thermal perception.) See e.g. Fruhstorfer *et al.* (1976) and Harju (2002).

³⁰ Coldness thermoreceptors seem to be the free endings of thin myelinated (A δ) fibers, and hotness thermoreceptors the free endings of unmyelinated (C) fibers (see e.g. Brodal 2010, 169).

may diverge from the standards fixed on by default (vaguely and implicitly) by another person, including listeners of that utterer. It is this fact that makes certain superficially incompatible attributions involving complex thermal predicates composed by thermal adjectives compatible, and also compatible with objectivism. Let's see this in a bit more detail, drawing on the account of section 2.

Suppose that the property signified by an utterance of “(somewhat) hot” in a context is implicitly picked on the basis of the character for “hot” in section 2.³¹ This plausibly happens through the utterer's implicit association of his utterance with a description licensed by that character, e.g. “the property of having heat in at least *this* amount”, taken together with a (possibly implicit) intention fixing on a certain rough degree of heat, a degree that the utterer will typically deferentially demonstrate via a certain felt, remembered or imagined perceptual quality. And similarly for “(somewhat) cold”. In particular, N's attribution of the complex predicate “neither somewhat hot nor somewhat cold” is presumably made as he implicitly fixes on a set of default standards that he demonstrates via certain remembered perceptual qualities of “minimal hotness” and “minimal coldness”. But what standards correspond to these qualities in N's perception is something determined by peculiarities of N's thermoceptive apparatus, in particular by the specific tuning of his thermoreceptors. The default standards for “hot” and “cold” that correspond to the qualities of “minimal hotness” and “minimal coldness” as felt by N are simply the activation thresholds t_1 and t_2 of N's hotness and coldness thermoreceptors: t_1 is the minimum temperature a thing must have to feel hot to N (and below which a thing does not feel hot), and t_2 is the maximum temperature a thing must have to feel cold to N (and above which a thing does not feel cold). The rough temperature level t of ψ that N feels as neither somewhat hot nor somewhat cold will then be precisely between t_1 and t_2 . (See

³¹ I will assume that “somewhat” is a “vacuous” modifier, in the sense that it doesn't affect the extension of a gradable adjective it modifies, though its use will often pragmatically convey something beyond the satisfaction conditions of the adjective. If we say that something is somewhat hot, we typically mean that it's not very hot. Similarly, with *Some philosophers from the university protested against budget cuts* we say the same as with *Philosophers from the university protested against budget cuts*; but we would avoid saying either of those things if we could have said *All the philosophers from the university protested against budget cuts*. If I say *Some philosophers from the university protested against budget cuts: all did*, I'm not contradicting myself. Similarly, if I say *This is somewhat hot: it is very hot*, I'm not contradicting myself. Alternatively, “somewhat” might be treated as non-vacuous, as “very” often is: an object counts as “very hot” in a context when its degree of heat is not only at least the degree contextually intended for “hot”, but also above some other higher standard contextually associated with “very” (see e.g. Kennedy and McNally 2005). On a treatment of this sort, the semantics of “somewhat” would require the introduction of further standards, in this case of closeness to the thermally neutral. In any case, the account in the text could be straightforwardly adapted to this more complex semantics.

figure 1.) So the property picked out by N's use of "(somewhat) hot" is the property of having heat in at least the amount t_1 , i.e. the property of having a temperature of at least t_1 ; and the property picked out by N's use of "(somewhat) cold" is the property of having cold in at least the amount t_2 , i.e. the property of having a temperature of at most t_2 . And by simple compositionality and intuitive meaning restrictions, the property picked out by N's use of "neither somewhat hot nor somewhat cold" is the property of having a temperature between t_1 and t_2 . Given N's default standards, then, ψ counts as neither (somewhat) hot nor (somewhat) cold, and N's attribution is true.³²

In H's case, by contrast, the set of standards fixed on by default is induced by H's thermoceptive apparatus, which again surely determines his default intentions. Here the standard that is the minimum temperature a thing must have to count as (somewhat) hot is a rough temperature level t_3 below t_2 , t and t_1 ; and the standard that is the maximum temperature a thing must have to count as (somewhat) cold is a rough temperature level t_4 slightly below t_3 . (See again figure 1.) So the property picked by H's use of "somewhat hot" is the property of having a temperature greater than or equal to t_3 , and the property picked by H's use of "somewhat cold" is the property of having a temperature lesser than or equal to t_4 . And by simple compositionality and intuitive meaning restrictions, the property picked out by H's use of "somewhat hot and not somewhat cold" is the property of having a temperature greater than or equal to t_3 . On H's default standards, then, ψ counts as (somewhat) hot and not (somewhat) cold, and H's attribution is true.

³² Note that the predicate "neither somewhat hot nor somewhat cold" (or equivalently, the adjectival predicate "thermally neutral") is not a gradable predicate, but still its conditions of application and its extension obviously vary as a function of the standards of application for the adjectives that compose it. For another example of non-gradability accompanied by dependence on standards induced by the dependence on standards of the composing adjectives, compare the predicate "neither big nor small".

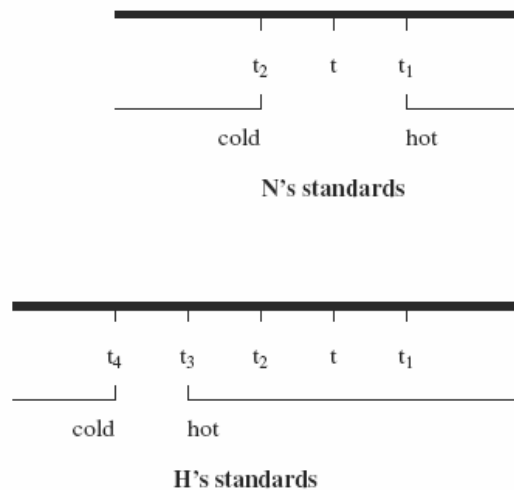


Figure 1

As stressed in section 2 for the case of many simple adjectives, the temperature properties that this account assigns to thermal adjectives and to many predicates composed of them are obviously objective, even though the properties can be different for (utterances of) different subjects intending their own default standards, as illustrated in the case of N and H. A further key to seeing why the picture is not some kind of “thermal subjectivism” is appreciating the distinction between adjectives or predicates (or utterances thereof) signifying subjective properties *and* adjectives or predicates that merely get to signify certain (objective) properties as a function of aspects of the occasion of use which may be prompted in turn by peculiar properties of subjects (e.g. through default contextual intentions formed with the help of demonstrated perceptual qualities). The latter adjectives or predicates are simply ones with a character such that for some object *o* and some contexts *c*, *c'*, *o* is in the extension of the property that that character delivers for *c* but fails to be in the extension of the property that it delivers for *c'*. Again an indication of the objectivist nature of the account is that this is compatible with “hot”, “cold”, “neither hot nor cold”, etc. being objectively modally stable, as opposed to the prediction of an (unsophisticated) subjectivist account on which (an utterance by H of) “hot”, say, signifies something like the property of *feeling hot to H*. Given any context *c*, there are no two different worlds *w* and *w'* which are the same as to the objective properties of *o* but are such that *o* is in the extension of the property of having a temperature greater than or equal to the contextually set standard with respect to (*c*, *w*) but fails to be in its extension with respect to (*c*, *w'*).

It may be good to stress again at this point that N and H do fix on different standards for “hot” and “cold”, even if in a certain purely phenomenological sense they might have been loosely said to intend “the same standards”: N and H both intend to call “hot” those things which feel hot, and presumably the “hotness feeling” is

phenomenologically the same for both. But this is indeed only a loose criterion for identity of standards. The present account, like usual theories of standards for gradable adjectives, takes these standards to be values in certain dimensions which need not be the immediate phenomenological dimensions used in content-fixing: N intends to call “hot” the things with the property of having heat in at least *this* amount, where “*this* amount” refers to a degree t_1 that N is disposed to think of via the remembered phenomenological quality (of “minimal hotness”) that he takes to correspond to that degree; by contrast, H intends to call “hot” the things with the property of having heat in at least *this* amount, where “*this* amount” refers to a degree t_3 *different* from t_1 even if H is disposed to think of it via the same phenomenologically felt quality (of “minimal hotness”) that N associates with t_1 . Given these intentions, N and H fix on different objective physical standards for “hot”, even if they form their corresponding intentions via the same remembered quality.

It may also be good to stress already at this point that the physical standards that speakers fix on for “hot” and “cold” need not be publicly available, or be part of a “context” understood as an already publicly available common ground. The notion of contextual determinant at play is a Kaplanian one on which such things need not be immediately publicly available, in the same way in which the intended referent of an utterance of a demonstrative need not be immediately publicly available. Of course, if there is to be successful communication between N and H (and between each of them and the experimenter or other language users) they will each have to guess in some way which thermal properties and consequently which physical standards the other intended. But these guesses of standards, of referential intentions and, in general, of contextual determinants and of contextual aspects fixing the contents of context-sensitive words, need not always be successful, and may be prevented by false beliefs of several kinds. In fact, I will show in section 4 below that the case of intersubjective perceptual variation in linguistic thermal and color classifications, as well as other cases involving context-sensitive expressions, involve special obstacles to the correct guessing of contextual determinants and thus of contextual content.

Now, an account analogous to this account of superficially incompatible thermal attributions can also be reasonably given in the case of superficially incompatible color attributions. First, recall that the property signified by an utterance of, say, “(somewhat) green” in a context is presumably implicitly picked on the basis of the character for “green” above.³³ This plausibly happens through the utterer’s implicit association of his utterance

³³ Again I am assuming that “somewhat” is a “vacuous” modifier (see note 31). And again on an alternative view, “somewhat” would be treated as a non-vacuous modifier like “very”. Something would count as “very green” in a context when its degree in the dimension of hue levels is not only between (or equaling one of) the standards contextually intended for “green”, but also between (or equaling one of) some other

with a description licensed by that character, e.g. “the property of having a hue of the color green between (or equaling one of) *these*”,³⁴ taken together with a (possibly implicit) intention fixing on a couple of rough physical hue values of the color green, values that the utterer will typically demonstrate via certain sensed, remembered or imagined chromatic perceptual qualities. And similarly for “(somewhat) blue”, “(somewhat) yellow”, etc.

Second, turn to the UG and BG case. Just as utterers’ default standards in the thermal case were induced by thermoreceptor activation thresholds, UG’s default demonstrative intentions for setting standards, and thus the standards themselves, are presumably induced by his visual apparatus, and in particular by the idiosyncratic absorption curve of his cones. In his case the hue value h of χ is precisely between the “lower” hue standard h_1 for “yellow” from which things start looking (somewhat) yellow to UG and the “upper” hue standard h_2 for “blue” from which things start looking (somewhat) blue to UG;³⁵ and all these standards lie between the two contextual hue standards, h_3 and h_4 , for “green”. (See figure 2.) So the property picked out by UG’s use of “(somewhat) yellow” is a property consisting in part in the possession of a hue value of at least h_1 , the property picked out by his use of “(somewhat) blue” is a property consisting in part in the possession of a hue value of at most h_2 , and the property picked out by his use of “green” is the property of having a hue value between (or equaling one of) the hue values h_3 and h_4 . And by simple compositionality and intuitive meaning restrictions, the property picked out by UG’s use of “green but neither somewhat yellow nor somewhat blue” is the property of having a hue value between h_1 and h_2 . So, in context, UG’s attribution is true.³⁶

By contrast, BG’s attribution is made as he implicitly intends default standards induced by *his* visual apparatus, by the idiosyncratic absorption curve of *his* cones. This is

standards closer to unique green and contextually associated with “very” as applied to “green”. But again the account in the text could be easily adapted to this more complex semantics.

³⁴ For the sake of brevity, I will henceforth often omit mention of saturation and brightness.

³⁵ Since the dimension of hues presumably constitutes a circumference, talk of “lower” and “upper” standards is strictly speaking misleading. But it is simpler to visualize, and readily intelligible when only the non-purple hues are involved. More exactly expressed, and assuming that in the hue circumference the blue, green, yellow and red hues appear in this order in clockwise direction, the idea is that h_1 is the initial standard for “yellow” in this direction, and that h_2 is the initial standard for “blue” in counterclockwise direction. *Mutatis mutandis* for the standards mentioned later. See section 5 for a more accurate representation of the dimension of hues.

³⁶ Much as before, it may be noted that the predicate “green but neither somewhat yellow nor somewhat blue” (or equivalently, the adjectival predicate “unique green”) is not a gradable predicate, but still its conditions of application and its extension obviously vary as a function of the standards of application for the adjectives that compose it.

such that the hue value h of χ is below the “upper” hue standard h_5 for “blue” from which things start looking (somewhat) blue to BG, below the “lower” hue standard h_6 for “yellow” from which things start looking (somewhat) yellow to BG, and finally also between the two contextual hue standards, h_7 and h_8 , for “green”. (See again figure 2.) So the property picked out by BG’s use of “(somewhat) yellow” is a property consisting in part in the possession of a hue value of at least h_6 , the property picked out by his use of “(somewhat) blue” is a property consisting in part in the possession of a hue value of at most h_5 , and the property picked out by his use of “green” is the property of having a hue value between (or equaling one of) the hue values h_7 and h_8 . And by simple compositionality and intuitive meaning restrictions, the property picked out by BG’s use of “green but somewhat blue” is the property of having a hue value between (or equaling one of) h_5 and h_8 . So, in context, BG’s attribution is true. Just as in the thermal case, no implication that the relevant properties must be subjective is warranted, provided we grant our assumption that the dimension of hues is objective and that the relation of having a certain hue is also objective.³⁷

³⁷ It may be good to note that, while this account handles the normal intersubjective variations in color ascriptions among speakers of the same language, the account is not intended as an explanation of certain cross-linguistic variations in color classification. For example, it is widely accepted that languages differ as to how many terms for what Berlin and Kay (1969) called “basic color categories” they contain; see also Kay and McDaniell (1978). Among speakers of a type IV language in Kay and McDaniell’s (1978) taxonomy, with only terms for the basic categories “white”, “red”, “yellow”, “green or blue” and “black”, we will presumably not find differing linguistic classifications of chip χ (at least using terms for basic color categories), unlike among their English-speaking counterparts. But this linguistic difference between these two language populations surely cannot be explained by appeal to differences in intended standards of application for shared color terms. Note for example that among speakers of English the differences in intended default standards are plausibly ultimately induced by physiological differences in levels of cone absorption. But the mentioned linguistic difference between the two language populations does not indicate in any way that the physiological differences in cone absorption detected in English speakers must not have the same incidence in a population of speakers of a type IV language. In fact, the linguistic difference between the two populations is presumably just a function of the quantities of terms for basic color categories within the populations’ languages. Such quantities are in turn presumably a function of a multiplicity of historical, social, cultural and environmental factors, but probably not of physiological factors. (In fact, Berlin, Kay and McDaniell argue from the existence of certain regularities in the distribution of terms for basic color categories across languages to the conclusion that there are underlying common physiological constraints on the amount of variation in color vocabulary one can find in different languages. This conclusion is widely, though not universally accepted.)




Figure 2

As in the thermal case, it may be good to stress that UG and BG do fix on different standards for their color adjectives and the complex color predicates composed of them, even if in a loose descriptive sense they might be said to intend “the same standards”: UG and BG both intend to call “green but neither somewhat yellow nor somewhat blue” those things which look green and do not look yellow or blue at all, and presumably the looks of the green that is neither somewhat yellow nor somewhat blue are the same for both. However, UG intends to call “green but neither somewhat yellow nor somewhat blue” those things with the property of having a hue value between *this* and *that*, where *this* hue value is a value h_1 that UG is disposed to think of via the remembered phenomenological quality (of “greenness with minimal yellowness”) that he takes to correspond to that value, and where *that* hue value is a value h_2 that UG is disposed to think of via the remembered phenomenological quality (of “greenness with minimal blueness”) that he takes to correspond to that value. By contrast, BG intends to call “green but neither somewhat yellow nor somewhat blue” those things with the property of having a hue value between *this* and *that*, where *this* hue value is a value h_6 *different* from h_1 that BG is disposed to think of via the same remembered phenomenological quality (of “greenness with minimal yellowness”) that UG associates with h_1 , and where *that* hue value is a value h_5 *different* from h_2 that BG is disposed to think of via the same remembered phenomenological quality (of “greenness with minimal blueness”) that UG associates with h_2 . Given these intentions, UG and BG fix on different physical standards determining different extensions for “green

but neither somewhat yellow nor somewhat blue”, even if they do so via the same remembered qualities.³⁸

Before turning to a potential objection to the present account, let me offer some further support for the account by observing how it respects the desiderata violated by the objectivist theories mentioned in section 1.

First, the account does respect the impression that all the attributions involved in the argument of section 1 are true, unlike the accounts of Byrne and Hilbert and Tye. And it does so respecting objectivism too. The fact that what is responsible for the truth of UG’s and BG’s attributions is the implicit operation of two sets of contextually intended standards, ultimately induced by peculiarities of the visual apparatuses of the two subjects, doesn’t imply the subjectivity of the properties they refer to, just as the fact that “hot” and “big” apply or fail to apply under the operation of (sets of) standards selected by (and in some cases induced by peculiarities of the perceptual apparatuses of) subjects doesn’t imply that the properties signified by (utterances of) “hot” or “big” are subjective. Note also that on the present account neither UG nor BG (nor the experimenter) make any mistake, either purely linguistic or factual, when offering their classifications; they just set their standards via their default standards-setting intentions, as they are surely allowed to do by the conventional semantics for context-sensitive adjectives.

Second, the account respects the desideratum that Kalderon’s abstract account did not deal with. “Not somewhat blue” in UG’s attribution *χ is green but neither somewhat yellow nor somewhat blue* and “somewhat blue” in BG’s attribution *χ is green but somewhat blue* do signify compatible properties in virtue of a relativization having to do with contextual effects. Since the standards involved are different in the two cases, “somewhat blue” in UG’s attribution and “somewhat blue” in BG’s attribution stand for different chromatic properties. Hence there is no problem in holding that “not somewhat blue” in the first attribution and “somewhat blue” in the second are not contradictory.³⁹

Third, unlike the account viewed favorably by McLaughlin, on which UG’s predicate “(somewhat) blue” is a priori equivalent with the reference-fixing descriptive predicate “has the physical property that in normal conditions makes something look (somewhat) blue to UG”, the present account is compatible with the absence of a priori connections between being of a certain color and looking that color to a particular person. A speaker may implicitly apply the character of “(somewhat) blue” and the description “the property of

³⁸ Note that, given that the standards for color adjectives are physical values, every difference in truth value between attributions of the same adjective to the same object must be due to a difference in the intended standards. This is of course compatible with speakers intending different physical standards via identical phenomenological characterizations.

³⁹ And in particular, no problem in doing this without denying that both are predicates that apply to χ , as in Cohen’s half-hearted supplement to Kalderon’s theory.

having a hue of (somewhat) blue between (or equaling one of) these” intending some hue standards with the help of demonstrated perceptual qualities. But this doesn’t require it to be a priori that these qualities are associated with the corresponding property in the utterer’s perception, even in normal conditions (though they may be so associated as a pure matter of fact). Reference fixing via a description licensed by the relevant character will be successful when some feasible hue standards do actually correspond to the perceptual qualities demonstrated by the utterer in context; but the existence of the correspondence need not be a priori for the mechanism to work.⁴⁰

4. Misidentifying contextual content: an objection defused

An objection to the account of section 3 might be as follows. The account traces some apparent disagreements between normal users of color language to differences in certain contextual determinants of their utterances. However, the existence of these differences, and as a consequence the differences in the contextual contents of many of those utterances, could only be ascertained on the basis of moderately sophisticated knowledge of scientific facts about color, unavailable to a current typical speaker. Moderately sophisticated speakers exposed to perceptual variation arguments do certainly get a *prima facie* impression that UG’s and BG’s utterances are compatible—and, if the account in section 3 is correct, those speakers see in some implicit way UG and BG as relying on different but feasible chromatic standards. However, at least some current typical speakers (including all my informants) think of UG’s and BG’s utterances as incompatible, and thus at best think of UG or BG as relying on unfeasible standards. And they will not be able to think of those utterances as compatible and based on feasible standards unless they digest the moderately sophisticated information concerning the physics and physiology of color presupposed in perceptual variation arguments. The objection is then that this is at odds with a thesis one might take to be of the essence of context-sensitive linguistic resources: the content of a context-sensitive expression must be identifiable without help from sophisticated scientific knowledge—while misidentifications of contextual content might and surely do occur, these ought to be impossible if the hearer is appraised of all relevant unsophisticated facts.

If this thesis is right, it puts some pressure on the suggestion that the differences in color classifications appearing in intersubjective arguments are to be explained as differences in contextual content. But we will see that the thesis is not right. I will show

⁴⁰ Compare the case of the fixing of the reference of an utterance of “he” by means of (an implicit appeal to) the description “the male I intend to refer to”. It is surely not a priori that my intention does actually single out a man (I might be deluded, e.g. because the experience that contributes to the formation of my intention is not perceptual but hallucinatory). But if my intention does single out a certain man, the man in question will be the referent of my utterance of “he”.

that phenomena entirely similar to the variations in linguistic color classification and the accompanying difficulties for an unsophisticated identification of contextual content emerge for paradigmatic indexicals and for context-sensitive predicates for other sensible properties. In fact, these phenomena, even if rare, are perfectly expectable under the plausible prevailing conception of the mechanisms of reference fixing for (utterances of) context-sensitive expressions. I will develop these points in this section.

One of the most enduring insights of twentieth-century philosophy of language is, I believe, that in many cases, with proper names and natural kind terms being paradigmatic instances, the content of a word in a speaker's idiolect is fixed in part via non-descriptive links between speaker and content. As a consequence, the content of a word in such cases gets fixed independently of whether it satisfies all the descriptive characterizations that the speaker and other language users are disposed to take confidently as reference-fixing for the word, or "authoritative" characterizations, as I will call them.⁴¹ This may get in the way of an unsophisticated identification of content, and even cause its misidentification, in cases where some incorrect but well established general beliefs are in place. If such beliefs could only be corrected in some sophisticated way, this may lead a language user to think wrongly that different uses of a word with associated descriptively similar reference-fixing material must get the same content. For the relevant non-descriptive reference-fixing links may be different for different uses, determining different contents for them, and yet this fact may be opaque in the absence of sophisticated knowledge.

Think of the following familiar example. The content of "water" in the idiolect of a terrestrial visitor to Putnam's Twin Earth is H₂O, and the content of "water" in the idiolect of his twin inhabitant of Twin Earth is XYZ, even if in both idiolects the reference-fixing material for "water" is descriptively similar. The terrestrial visitor, one can plausibly suppose with Kripke (1972), initially established in some implicit way the reference of his uses of "water" via something like a demonstrative description such as "the liquid, transparent substance over here", and similarly for his twin. But while in the case of the terrestrial visitor the content of uses of "water" is H₂O because "over here" there was H₂O, in the case of the twin the content of uses of "water" is XYZ for the analogous reason. The terrestrial visitor will nevertheless authoritatively associate with the twin's uses of "water" a use of the description "the liquid, transparent substance over here", made

⁴¹ I take the convenient terminology of "authoritative association" from Braun (2006). By a descriptive characterization D "authoritatively associated" by a language user U with an expression E Braun more strongly understands one that U very robustly associates with E, that U thinks can be confidently used to determine the reference of E, and that it is such that U will generally take utterances of a sentence "E is P" to communicate the proposition to the effect that D is P. But Braun's stronger characterization would be equally suitable for my purposes. See the next-to-last paragraph of subsection 4.1 below for examples of descriptions authoritatively and non-authoritatively associated with uses of "today".

with the intention to refer to the space occupied by a terrestrial lake, say. The terrestrial visitor is wrongly confident that such a use of the description applies to the samples of liquid in Twin Earth, as he wrongly thinks that it is for all essential purposes the same as a use fixing the reference of “water” in the twin’s idiolect. As a result of this, he misidentifies the content of “water” in the mouth of his twin.

Now a related but neglected phenomenon concerns the contextual content of utterances of some context-sensitive words. As is widely accepted, in many cases the contextual content of a word utterance gets fixed in part via non-descriptive links between utterer and eventual content, including some facts about utterer identity and spatiotemporal location in the case of pure indexicals and certain non-descriptively identifying intentions on the part of the utterer in the case of impure indexicals. What is perhaps not so widely recognized is that, as a consequence, the contextual content of utterances of some context-sensitive expressions gets fixed independently of whether it satisfies all the descriptive characterizations that language users are disposed to take confidently as reference-fixing for those utterances. This creates obstacles of principle to the identification of contextual content, and in fact incites its misidentification, in cases where language users have wrong but well established general beliefs that could only be corrected through the acquisition of sophisticated knowledge. Specifically, it leads to problems of identification, and to misidentification, in cases where the relevant wrong general beliefs lead language users to wrongly think that utterances with associated descriptively similar reference-fixing material must get the same content. For these utterances may differ as to the relevant non-descriptive reference-fixing links between utterer and content, thus coming to have different contextual contents.

Before we see how this phenomenon emerges for color language, it will be good to illustrate how it emerges in a relatively simple case, the case of “today”. (Similar points could be made about other indexicals, pure and impure.)

4.1. Misidentifying the content of “today”

Consider the following story, which gives rise to a situation arguably analogous to the one in the variation argument of section 1. Imagine that Jill is a member of an isolated English-speaking community living in an island somewhere in the British North Sea. As used to be the convention, for this community a day begins and ends at sunset. More remarkably, Jill and her fellow isolated islanders don’t have a modern understanding of the astronomical and geographical reality behind sunset and sunrise. For them the Earth is flat and small and the Sun goes round it. Nothing makes them think that the Sun sets at substantively different moments on different Earth locations. At a time of economic crisis, Jill’s husband Jack is offered a job as a shepherd in an island of the Sea of Japan, and he accepts it. Jack’s employers give cell phones to Jack and Jill, so that they can stay in touch. Jack spends his night flight to the Asian island sleeping, and wakes to his following Asian mornings

thinking that his jet-lag is due to the uncomfortable trip. Their first call takes place shortly after a sunset in the Sea of Japan, thus a few hours before sunset in the North Sea. After a conversation nothing in which makes them suspect that the Sun sets at different times on the two islands, Jack says to Jill *I'll call you again later today*. A skeptical Jill says *I won't get any more calls from you today*. Shortly before the next Sea of Japan sunset, well after the Sun has set on the North Sea, Jack calls Jill and in triumph notes that he was right and she was wrong. Jill is baffled. She thinks of what she said as true, and of what Jack said as false—and *mutatis mutandis* for Jack. On the other hand, on what I take to be the likely reaction of non-isolated readers of the story, both spoke the truth.

On what I take to be a reasonable view that accounts for these reactions, both the islanders and non-isolated readers implicitly associate with “today” a Kaplanian character such as *An utterance of “today” refers to the day of the utterance*. On the basis of this character they can assign to each utterance of “today” particular (uses of) reference-fixing descriptions licensed by it, such as a use of “the day of this utterance”. Besides this, the islanders and non-isolated readers also associate with “day” certain implicit, well established general postulates and beliefs, e.g. the postulate that a day is a time period between sunsets.⁴² Some of the beliefs that the islanders associate with “sunset”, and hence with “day”, are wrong, such as the belief that sunset takes place at the same time on all Earth locations. But even though the islanders have wrong beliefs about sunset, and hence also wrong beliefs about what a day is, any utterance of “today” by one of them gets a contextual content, namely the day (sunset-to-sunset time period) of the utterance, as picked e.g. by an implicit use of a description such as “the day of this utterance” (or by similar descriptions licensed by the character of “today”). As it turns out, the day of Jack’s utterance is a certain sunset-to-sunset time period, while the day of Jill’s utterance is a certain different sunset-to-sunset time period, so the contents of their utterances of “today” are different. If this is so, both utterances are true, and non-isolated readers of the story are correct.

Whence then the appearance of contradiction for Jack and Jill? Note that Jill (*mutatis mutandis* for Jack), besides thinking of her utterance of “today” as having its content conventionally fixed by an implicit use of “the day of this utterance” made while thinking of her utterance, also thinks (correctly) of Jack’s utterance of “today” as having its content conventionally fixed by a slightly different implicit use of “the day of this utterance”, one made while thinking of Jack’s utterance. These two uses of “the day of this utterance” are descriptively essentially similar from the point of view of Jill: they both point to utterances of what is auditorily clearly the same word made almost simultaneously, and in fact made in the same sunset-to-sunset time period from Jill’s point of view. But because of this similarity, Jill’s wrong belief that sunset takes place at the same time on all Earth locations leads her to take those uses to be interchangeable for all essential purposes. For example,

⁴² The non-isolated readers embrace this ancient convention while they are reading the Jill and Jack story.

Jill is ready to authoritatively associate with Jack's utterance of "today" a use of "the day of this utterance" made as she thinks of her own utterance. Jill is thus naturally led to expect her utterance of "today" and Jack's to have the same content. This makes it unfeasible for Jill to see how differences in utterer spatio-temporal location operate so as to fix different contents for her utterance and Jack's. She misidentifies the contextual content of Jack's utterance of "today"—though she will think she doesn't, and will probably come to think, after the second conversation, that Jack must be confused or deluded.

It should be noted that the Jill and Jack case contrasts with mundane cases in which someone is wrong about the content of an utterance of "today". In mundane cases someone has correct general beliefs about what a day is, but simply lacks the appropriate unsophisticated information about the particulars of when or where the utterance was made, information which is not authoritatively associated with the utterance. For example, a listener may be wrong about the reference of an utterance of "today" in an old tape merely due to a misleading label about when it was recorded. He may associate with the utterance the description "the 29th day of February, 1988" appearing on the label, but (typically) will not do so authoritatively. The islanders, however, know when and where Jack and Jill are talking, and derive their bafflement (in part) from an authoritatively associated but wrong general belief about what a day is, a belief that could only be corrected by imparting to them some relatively sophisticated knowledge.

On the other hand, non-isolated readers know all the relevant astronomical and geographical facts. They of course know (in some implicit way) that an utterance of "today" by Jack or Jill gets as contextual content the day picked by a relevant use of, say, the description "the day of this utterance". But their comparatively sophisticated knowledge of the relevant astronomical and geographical facts lets them know that those uses pick out different time periods because the days of utterance are different. This is what accounts for the non-isolated readers' impression that both utterances are true.⁴³

⁴³ While this account of the Jill and Jack story seems the most reasonable one to me, it might of course be questioned in not unreasonable ways, some of which parallel some of the different responses to the basic perceptual variation argument of section 1. For example, it might be postulated that the contents of utterances of "today" are not objective time periods, but subjective descriptive contents constituted in part by perceptual relations between utterers and time periods, which in turn determine time periods. It might be postulated that an utterance of "today" has a descriptive content such as the property of *being the day to which the utterer perceives the moment of utterance as belonging*. On this view, both Jack's utterance and Jill's are true. But on this view "today" can shift reference from one circumstance of evaluation to another, even if we keep the context fixed, and this leads to odd results, related to the oddness of modally unstable views of color adjectives: an utterance of *If I were now in the North Sea Island, I would not be speaking to you today* made by Jack in the first conversation would count as true. But it clearly sounds false. Compare this result with the implications of the view favored in the text. On that view, "today" is a typical indexical whose content at a

4.2. Misidentifying the content of thermal predicates

The story of Jack and Jill arguably provides a case analogous to our case of variation in linguistic color classification, and clearly indicates how general difficulties for an unsophisticated correct identification of utterance content can arise even in the case of a paradigmatic indexical. The cases of thermal and chromatic adjectives are not substantively different from the case of “today”—though, as I will presently note, they differ from “today” and between themselves as to the different levels of persistence of traditional misconceptions and of general availability of the scientific information relevant for a correct identification of contextual content.

At earlier stages of our use of thermal words people in general presumably thought of standards different from their default ones as intuitively unfeasible. The intuitive feasibility of thermal standards arguably depends on one’s beliefs about the heat mentioned in the character for “hot”—and similarly for “cold”, “warm”, etc. Some of these beliefs were presumably always true or satisfiable, such as the belief that heat and cold are phenomena exemplified in boiling water and frozen water, say. However, other relevant early beliefs were false. For example, one early belief was apparently to the effect that heat and cold were non-anthropocentric varieties in the dimension of amounts of heat and cold, varieties not singled out merely by the idiosyncrasies of the human perception of this dimension, but by differences in the physical natures of two elemental substances or phenomena.⁴⁴ Another early belief was then presumably to the effect that a normal person’s perception delivers a correct indication of the (net) amount of the non-anthropocentric phenomena of heat or cold present in a certain object, and more specifically that there is a single correct correspondence between thermal perceptual qualities and physical levels in the objective dimension of degrees of heat and cold.

If these are still the preconceptions of a given language user U, he all the same knows how the content for utterances of “hot” is fixed via applications of its character (*mutatis mutandis* for “cold”, here and in what follows). In the case of his own uses, U implicitly fixes the content of an utterance of “hot” by means of a use of a description licensed by the character of “hot”, typically with an accompanying non-descriptively identifying intention of his—e.g. a use of “the property of having heat in at least *this* amount” taken together with an intention to refer to an amount of heat demonstrated by him via a certain felt, remembered or imagined thermal perceptual quality; and the property picked out is the

context is an object (time period), unmediated by a descriptive content, and hence the reference of “today” is rigid, cannot shift from one circumstance of evaluation to another. So on this view Jack’s counterfactual utterance comes out false for natural reasons.

⁴⁴ The intuitive view of heat and cold as different substances or phenomena is also easily elicited in contemporary children and adolescents. See e.g. Erickson (1979) and Lewis and Linn (1994).

property of having a temperature greater than or equal to *that* contextually intended standard. In the case of another speaker's uses, U implicitly knows that the content of an utterance of "hot" is analogously typically fixed by means of a use of that description taken together with an intention on the part of *that* other speaker to refer to an amount of heat demonstrated by *her* via a certain felt, remembered or imagined thermal perceptual quality.

However, some of U's well established general beliefs about the heat mentioned in the character of "hot" are false. U has preconceptions to the effect that if normal, non-deluded speakers are in some situation expected or asked to intend the default or laxest standard for "hot", they will all fix on the same standard corresponding to a minimum (positive net) amount of heat, as he thinks that in normal speakers there is one correct correspondence between thermal perceptual qualities and physical levels of heat. In the N and H setup U has no reason to think of either N or H as abnormal or deluded, so he expects them to fix on the same standard for "hot" via similar descriptive material, e.g. uses of "the property of having heat in at least this amount", and thus to fix on the same determinate hotness property as the content for their utterances of "hot". Because of the similarity in the reference-fixing descriptive material, U's wrong beliefs lead him to authoritatively associate with N's and H's utterances of "hot" and with his own utterances (if he is not either N or H) a use of "the property of having heat in at least *this* amount" taken together with his own default intention to refer to an amount of heat presented via the felt, remembered or imagined thermal perceptual quality *he* takes to correspond to a minimal amount of heat. U is thus naturally led to expect all the relevant utterances of "hot" to have the same contents. Faced with N's and H's actual utterances, then, he is likely to doubt that they are both normal perceivers, or that neither of them is deluded, and thus that both of their utterances can be true.

While there may still exist speakers like U, at least moderately sophisticated speakers are by now aware that there is no such thing as the elemental heat and cold, and that there need not be a privileged correspondence between perceptual qualities and physical levels of heat and cold. Moderately sophisticated speakers know (in some implicit way) that the contextual contents of utterances of "hot" and "cold" in contexts where the utterers intend default standards are a function of the assumed correspondences between perceptual qualities and physical temperature levels, which need not be the same for all utterers. So moderately sophisticated speakers also presume that the implicit uses of the characters of "hot" and "cold" by N and H determine different properties. Hence they also presume a difference in the contents of N's and H's utterances of "hot", and they consequently presume that both N's and H's utterances are true.

4.3. Misidentifying the content of color predicates

The situation in the color case is essentially similar to that in the thermal case. The two cases share an essential core of common relevant features. Nevertheless, there are differences, arguably due to historical differences in the levels of availability of the relevant scientific information and of persistence of traditional misconceptions.

Analogously to the thermal case, the intuitive feasibility of chromatic standards depends on one's beliefs about (say) the color green mentioned in the presumable character for "green" postulated above—and similarly for "yellow", "blue", etc. Some of these beliefs are presumably true or satisfiable from an objectivist viewpoint, but other relevant beliefs are false. For example, our initial preconceptions are to the effect that the colors correspond to non-anthropocentric varieties in the chromatic dimensions, especially hue. The color blue, we initially think, doesn't differ from the color green merely in that some human idiosyncrasy establishes an anthropocentric division in the physical dimensions along which color varies. There is surely some preconceived anthropocentricity due to the relatively well-understood effects of the use of (intuitively feasible) contextual standards, but by and large the basis for the division is preconceived as non-anthropocentric: if something counts as blue, it's of a shade where some non-anthropocentric phenomenon we call the (pure) color blue predominates. Another belief is then presumably to the effect that a normal person's perception delivers a correct indication of the (net) amount of the non-anthropocentric phenomena called the color blue, the color green, etc. present in objects, and more specifically that there is one correct correspondence between chromatic perceptual qualities and physical levels in the objective dimension of hue. Preconceptions of this kind are unsatisfiable from an objectivist viewpoint, for (as the more informed among us now know or at least suspect) the divisions imposed by perception on any dimension of wavelengths plausibly corresponding to perceived hue will be thoroughly anthropocentric, and will exhibit differences from normal subject to normal subject as a function of (great or small) physiological differences. But it is natural to conjecture that typical language users still have these preconceptions about color. And yet, typical language users must all the same know how the content for utterances of "blue", "green", etc. is fixed via applications of their characters. In the case of his own uses, a language user implicitly fixes the content of an utterance of "blue" by means of a use of a description licensed by its character, typically with an accompanying non-descriptively identifying intention of his—e.g. a use of "the property of having a hue of the color blue between (or equaling one of) *these*" taken together with an intention to refer to a couple of hues demonstrated via certain sensed, remembered or imagined chromatic perceptual qualities; and the property picked out is the property of having a hue value between (or equaling one

of) *those* contextually intended standards.⁴⁵ In the case of another speaker's uses, the language user implicitly knows that the procedure is the same, *mutatis mutandis*.

I hypothesize that false preconceptions and conventional linguistic knowledge combine to make a typical language user think (in some implicit way) that if normal, non-deluded speakers are in some situation asked or expected to intend the default or laxest standards for "blue", they will all fix on the same standards corresponding to minimum (positive net) amounts of the color blue. In the UG and BG setup such a language user (possibly identical with either UG or BG) has no reason to think of either UG or BG as abnormal or deluded, so he expects them to fix on the same standards for "blue" via similar descriptive material (e.g. uses of "the property of having a hue of the color blue between (or equaling one of) *these*"), and via demonstrative reference to the same perceptual qualities supposed to correspond to the same hue levels. So he also expects UG and BG to fix on the same determinate blueness property as the content for their utterances of "blue", and thus not to differ in their classifications involving "blue". Faced with BG's and UG's actual utterances, then, he is likely to doubt that they are both normal perceivers, or that neither of them is deluded, and thus that both of their utterances can be genuinely true.

Whence then the impression generated by the perceptual variation argument of section 1, that UG and BG are both right (or at least that neither is more correct than the other)? On the present hypothesis, this impression is elicited in *readers* of intersubjective perceptual variation arguments, precisely because they are among the few of us who already know that our relevant preconceptions of chromatic non-anthropocentricity are false, and specifically that our perceptual partition of the spectrum of wavelengths is thoroughly anthropocentric and there is really no privileged correspondence between perceptual qualities and hue values. People able to digest the arguments know (in some implicit way) that the contextual contents of utterances of color predicates (in contexts where the speakers intend default standards) are a function of the assumed correspondences between perceptual qualities and physical hue values. So they also presume that uses of reference-fixing descriptions licensed by the character of, say, "blue", by people like UG and BG need not single out the same property. Hence they presume a difference in the contents of UG's and BG's utterances of "blue", and also that both utterances can be true.⁴⁶

⁴⁵ Again I omit the specification of the contextually intended saturation standard, and of the contextually intended brightness standards, for brevity.

⁴⁶ Note that although some readers of perceptual variation arguments ultimately reject these intuitions, and in a way stick to the intuitions of the unsophisticated speaker (e.g. Byrne and Hilbert and Tye), this is surely as a result of ulterior philosophical considerations. As far as I can tell, none of them denies that they feel the initial pull of the arguments.

We have thus seen that a plausible view of reference fixing for context-sensitive expressions, as being in part a function of non-descriptive links between utterer and eventual content, makes expectable and accounts for the difficulties in the correct unsophisticated identification of contextual content for utterances of those expressions. And we have also seen that the view receives direct support from the unified explanation it provides for the judgments of different kinds of speakers concerning the scenarios featured in intersubjective variation arguments. We have thus defused the objection that contextual content should be more accessible than the present account postulates, which I take to be the main potential objection to the account.

5. The contents of color experience

The present account of the contents of color language suggests a companion account of a relevant part of the contents of color experiences, an account that I think is also susceptible of independent motivation. On this account, a typical color experience represents at least some objective chromatic properties, including the properties that are the contents of color expressions in the typical color attributions that could be made by the experiencer on the basis of the experience.⁴⁷ The account postulates, in particular, that a color experience by a

⁴⁷ I say “at least some”, because I think it likely that experience represents both objective and subjective chromatic properties, thus agreeing with Shoemaker (1994, 2006). I think it likely, for example, that all the relevant facts are compatible with UG’s experience representing both objective properties of the sort postulated above and below and subjective properties such as the property of *looking this way to me (UG) now*. I also view with sympathy other theories that postulate a plurality of kinds of chromatic contents in color experience, such as the view explored in Chalmers (2006). (I am not so sympathetic with other “pluralist” views which, unlike Shoemaker’s and Chalmers’s, postulate that color experience represents, besides chromatic properties instantiated by extra-mental objects, also chromatic properties instantiated by mental objects such as experiences and sense-data (see e.g. Maund 1981, Peacocke 1984, Rosenthal 1999 and Brown 2006). I think that the postulation of “mentalist” chromatic properties is fraught with difficulties, and that it seems fair to say that it is widely rejected. However, the present account is not necessarily incompatible with the existence of such properties or with their being represented in color experience, as long as they are represented along with (objective) chromatic properties instantiated by extra-mental objects.)

By contrast, I think that our implicit beliefs and postulations relevant to content fixing are sufficiently rich to exclude the possibility that UG’s predicate “green”, say, signifies a subjective property. For example, the fact that our implicit convention is that “green” is objectively modally stable excludes the possibility that “green” signifies the property of *looking this way to me (UG) now*. (But the question of whether color predicates have subjectivist or even mentalistic acceptations would ideally require an extended treatment that admittedly cannot be provided here.) If all this is so, then no anti-subjectivist result about color perception can emerge from the considerations pertinent to the content of color expressions. But, if the

typical perceiver represents an object as lying on certain hue intervals; that if I am a typical perceiver, I see an object not merely as occupying a particular hue locus, but also as occupying a locus between, or above, or below certain other significant hue loci (*mutatis mutandis* for the other dimensions). The crucial aspect of the proposal is its postulate that membership in *intervals* in the chromatic dimensions is represented in a certain way by typical chromatic experiences. On the account, a typical chromatic experience does this in turn by representing in a certain way the bounds of those intervals, even if those bounds are not being “seen” in some direct way at the time of the experience, or at any rate it does this provided the experiencer has some previous representations of those bounds in her experiential memory or her empirically informed experiential imagination.

If I am a typical perceiver, I possess previous representations of minimally green blue, of slightly green blue, of slightly blue green, of minimally blue green, of pure green, and so on. But if this is so, then I possess previous representations of my default bounds for green—these are just my representations of minimally green blue and minimally green yellow. And I also possess previous representations of other bounds that I could especially have in mind when I intend linguistic standards for “green” other than the default ones provided by my perceptual partition of the physical dimension of hues, e.g. representations of the bounds constituted by slightly blue green and slightly yellow green. My possession of these previous representations, and sometimes my concurrently bearing them in mind as I have a particular chromatic experience of a certain object *o* that I see as green, does seem to influence the phenomenal character, and indirectly the content, of my experience of *o*. I don’t merely see *o* as having the shade it has, as it were in isolation from other shades; I see *o* as occupying a hue locus between minimally green blue and minimally green yellow; or (inclusive “or”) I see *o* as occupying a hue locus intermediate between, say, the loci of slightly blue green and of slightly yellow green. And this is plausibly rooted in qualitative aspects of my experience induced by the existence of my previous representations of those loci. In the end, these qualitative aspects are responsible for making a typical experience represent *o* as lying on the corresponding hue intervals. (*Mutatis mutandis* for saturation and brightness.)

This will sound plausible to those who think it plausible that background mental representations and abilities influence the phenomenology and content of experience, following the lead of Siewert (1998, ch. 7) and Siegel (2006, 2010, chs. 4ff.), among others. (Though, as far as I can see, these authors do not consider the case of the representation of membership in intervals in the color dimensions). A quote from Siewert describes the general phenomenon that on the present proposal is instantiated in the case of properties characterized by bounds in the chromatic dimensions:

account of this section is correct, then it is also the case that no anti-objectivist result about color language can emerge from considerations pertinent to the content of color experiences.

This kind of variation in the phenomenal character of visual experience, where what one sees is recognizable as of a certain type (...) shades into, and is perhaps never wholly separable from, a sort that relates to differences in what ‘lies beyond’ what one sees—that is, both what literally is *located* somewhere else, and what lies beyond the *actual* condition of what one sees, in *possibilities* in which it figures. (1998, 257)

On the present account of (some of) the chromatic properties represented in visual experience, in particular, when an object *o* is seen as having a certain color property of the kind described in sections 2 and 3, this seeing amounts to placing *o* in a certain hue locus (say, of *green but neither slightly blue nor slightly yellow*) that is partly individuated by reference to other hue loci that are “beyond” what one directly sees. But these loci, or rather their representations, nevertheless influence the phenomenal character of one’s experience of *o*.

In the case of UG and BG, the present account postulates that their experiences possess different contents involving the corresponding mentioned properties of χ . In particular, their visual experiences represent χ as lying on the relevant hue intervals. One reason for this is that UG and BG possess, however implicitly, certain pre-existing representations of the levels that fix the bounds of those intervals, and this influences the phenomenal character, and indirectly the content, of their experiences of χ . UG is certainly disposed to recognize χ as having a hue “below” h_1 , a hue previously represented as marking the starting point for the (somewhat) yellow (strictly speaking, the minimally yellow green), and also as having a hue “above” h_2 , a hue previously represented as marking the starting point for the (somewhat) blue (strictly speaking, the minimally blue green), and finally also as having a hue strictly between h_3 and h_4 , the starting points for the (somewhat) green (strictly speaking, the minimally green yellow and the minimally green blue). These dispositions go hand in hand with χ being presented to him in a certain way qualitatively different from the way it would be presented to him if he had no memory or imaginative knowledge of h_1 , h_2 , h_3 and h_4 . BG’s experience, by contrast, arguably presents χ to him as having a hue “below” h_5 , a hue previously represented as marking the starting point for the (somewhat) blue (strictly speaking, the minimally blue green), and also as having a hue “above” h_6 , a hue previously represented as marking the starting point for the (somewhat) yellow (strictly speaking, the minimally yellow green), and finally also as having a hue strictly between h_7 and h_8 , the starting points for the (somewhat) green (strictly speaking, the minimally green yellow and the minimally green blue). It is thus that it is legitimate to postulate that UG’s and BG’s visual experiences represent χ as lying on different hue intervals and possessing different color properties characterized by membership in those intervals, on which they place χ by means of their respective attributions. On this hypothesis, UG’s and BG’s experiences have compatible and in fact

veridical objective contents.⁴⁸ (However, UG need not know or even suspect that BG's experience has such a different content from his (and vice versa).)

This account of the relevant part of the content of chromatic experiences can be used to provide an objectivist explanation of certain presumable facts about the structural relations of the colors. Critics of objectivism have often sought to challenge its ability to account for apparently obvious structural facts such as that green is a unique color while orange is a binary color, and that red is more similar to orange than it is to yellow. The criticism is that the physical properties that have typically been identified with the colors, e.g. reflectance properties, do not appear to stand in any comparable structural relations.⁴⁹ But we will now see that the properties that the present account identifies with the objective chromatic properties represented by perceptual experience and by color language, i.e. properties of membership in intervals along the chromatic dimensions, do stand in the appropriate structural relations. Thus, under the present account structural facts are susceptible of receiving a reasonable objectivist explanation.⁵⁰

Focus on UG's phenomenal partition of the objective dimension of hues. This partition can be seen as a multiplicity of divisions of the hue circumference into intervals whose bounds are marked by perceptually significant loci. Among these intervals, eight stand out, as depicted in figure 3: the closed interval bounded by UG's default standards

⁴⁸ N's and H's experiences also presumably possess different contents that contain as ingredients the corresponding mentioned properties of ψ . In particular, their tactile experiences presumably represent ψ as lying on the relevant temperature intervals. For example, ψ is plausibly presented to N in a certain way qualitatively different from the way it would be presented to him if he had no memory or knowledge of (his representation of) t_1 and t_2 . Mutatis mutandis for H.

⁴⁹ For places where versions of this criticism are developed, see e.g. Hardin (1988), Boghossian and Velleman (1991), Johnston (1992), Maund (1995) and Pautz (2006) (the latter contains an especially clear presentation of challenges of this sort and of various dialectical moves from the two sides).

⁵⁰ This explanation is in the same spirit as the one postulated by Byrne and Hilbert (2003) and Byrne (2003) in terms of their theory of "hue magnitudes". This is a theory I feel uncomfortable with, in part because it's not entirely clear to me from Byrne and Hilbert's compressed presentation what objective physical magnitudes are to be identified with their hue magnitudes (see Pautz 2003 for some worries in this direction). By contrast, the theory in the text appeals only to (a perfectly clear objective identification of) color properties. Another, not implausible objectivist reaction to the criticisms based on the apparent existence of structural relations between the colors is to deny that these relations are more than apparent; Matthen (1999) and Gómez-Torrente (2011) give voice to reactions of this sort. It may be worth noting in this connection that thermal properties do also stand in intuitive structural relations: the property of being thermally neutral is "unique", in the sense of not being intuitively a mixture of being hot and being cold; being warm is more similar to being hot than to being cold. However, philosophers don't appear to think that this is an objection to the idea that objective temperature properties are represented in perception.

for “green”, the hues corresponding to minimally green blue (h_4) and minimally green yellow (h_3);⁵¹ the closed interval bounded by UG’s default standards for “yellow”, the hues corresponding to minimally yellow green (h_1) and minimally yellow red (h_9); the closed interval bounded by UG’s default standards for “red”, the hues corresponding to minimally red yellow (h_{10}) and minimally red blue (h_{11}); the closed interval bounded by UG’s default standards for “blue”, the hues corresponding to minimally blue red (h_{12}) and minimally blue green (h_2); the closed interval bounded by UG’s default standards for “olive”, the hues corresponding to minimally yellow green (h_1) and minimally green yellow (h_3); the closed interval bounded by UG’s default standards for “orange”, the hues corresponding to minimally red yellow (h_{10}) and minimally yellow red (h_9); the closed interval bounded by UG’s default standards for “purple”, the hues corresponding to minimally blue red (h_{12}) and minimally red blue (h_{11}); and the closed interval bounded by UG’s default standards for “turquoise”, the hues corresponding to minimally green blue (h_4) and minimally blue green (h_2).⁵²

The presumable fact that orange (as perceptually represented and referred to by UG) is a binary color can be explained as the fact that objects having a hue in the default “orange” interval $[h_{10}, h_9]$ ⁵³ always also have a hue both in the default “yellow” interval $[h_1, h_9]$ and in the default “red” interval $[h_{10}, h_{11}]$. Or, strictly speaking, as the fact that the property of membership in the “orange” interval is in a certain relation to the property of membership in the “yellow” interval and the property of membership in the “red” interval: the “orange” interval that partly constitutes the first property is the intersection of the “yellow” and the “red” intervals that partly constitute the latter two properties. The presumable fact that green (as perceptually represented and referred to by UG) is a unique color can be explained as the fact that objects having a hue in the default “green” interval

⁵¹ In describing intervals in what follows, the bound mentioned first is to be understood as initial and the bound mentioned second as final in the clockwise direction.

⁵² A referee has insisted to me that we standardly think of the hue dimension as consisting of things, the hues, which by their nature are not susceptible of being variously partitioned by different speakers or perceivers—unlike the interval membership properties in the objective hue dimension postulated here, which are so susceptible in virtue of their indirect relation to phenomenal color qualities. If respecting this idea about the nature of the hues is a desideratum on a theory of them, it must be acknowledged that the present theory does not meet it. However, I don’t think that this should be a desideratum on a theory of the hues. Note that we similarly do not initially think of thermal properties as susceptible of being variously partitioned by different speakers or perceivers, and yet our mature theory of them postulates the existence of an objective dimension of properties which are so susceptible, and identifies them with the thermal properties.

⁵³ In what follows, “[a, b]” and “(a, b)” stand respectively for the closed and open intervals with a as initial and b as final bounds in the clockwise direction.

$[h_4, h_3]$ sometimes do not have a hue in any of the default “yellow” or “blue” intervals $[h_1, h_9]$ and $[h_{12}, h_2]$. Or, strictly speaking, as the fact that the property of membership in the “green” interval is in a certain relation to the property of membership in the “yellow” interval and the property of membership in the “blue” interval: the “green” interval that partly constitutes the first property has a subinterval (i.e. (h_2, h_1)) disjoint from both the “yellow” and the “blue” intervals that partly constitute the latter two properties. The fact that red (as perceptually represented and referred to by UG) is more similar to orange (as perceptually represented and referred to by UG) than it is to yellow (as perceptually represented and referred to by UG) can be explained as the fact that objects having a hue in the default “orange” interval $[h_{10}, h_9]$ always have also a hue in the default “red” interval $[h_{10}, h_{11}]$, but objects having a hue in the default “yellow” interval $[h_1, h_9]$ sometimes do not have a hue also in the default “red” interval $[h_{10}, h_{11}]$. Or, strictly speaking, as the fact that the property of membership in the “red” interval is in a certain relation to the property of membership in the “orange” interval and the property of membership in the “yellow” interval: the “red” interval that partly constitutes the first property includes the “orange” interval that partly constitutes the first of the latter two properties, but does not include the “yellow” interval that partly constitutes the second. Mutatis mutandis for the analogous facts about the other unique and binary colors and for relations of similarity between other triads of colors.⁵⁴

⁵⁴ It is sometimes claimed that green is more similar to blue than it is to yellow, or that red is more similar to yellow than it is to blue. I don’t share these intuitions, but if they are genuine and require an explanation in terms of structural relations between objective color properties, such an explanation would need to appeal to a different apparatus. Alternatively, these might be illusory intuitions in a sense in which more intuitive claims of similarity may not be.

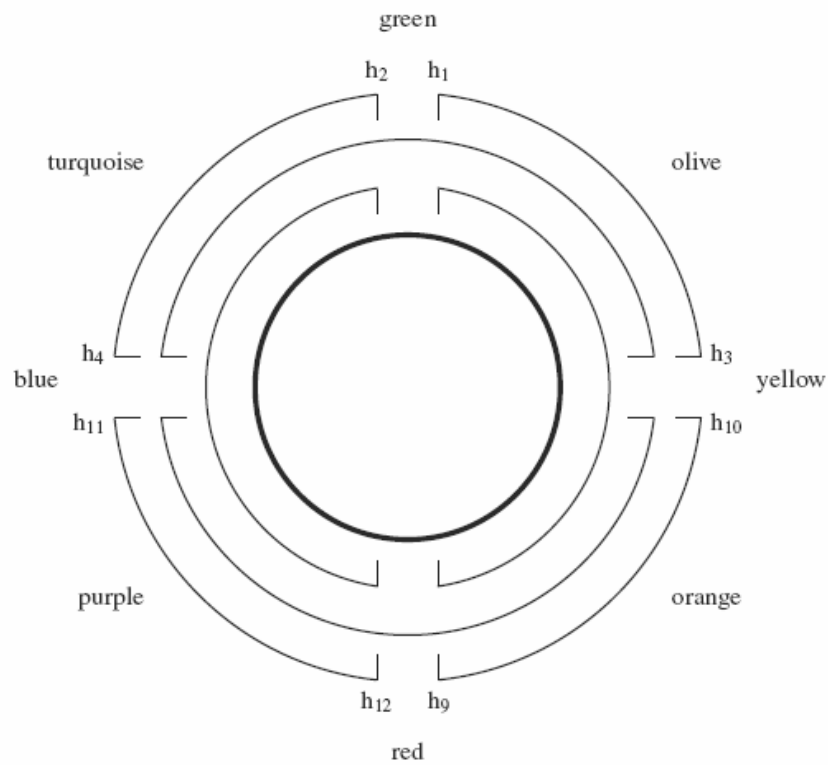


Figure 3

The considerations in this paper show, I think, that the present two allied accounts, of the contents of color language and color experience, have much to recommend them. One consideration that was only hinted at in the Introduction, however, seems to me to lend particular weight to the accounts. It is undeniably natural to suppose that the color properties signified by color words are represented by color experiences. If this natural supposition is correct, it strongly advises us to heed the desideratum that an appropriate account of the contents of color experience must mesh well with an appropriate account of the contents of color language. However, the problem of the contents of color experience has been predominantly considered in isolation from the corresponding problem about color language. The accounts in this paper, by contrast, are closely interrelated, and arguably satisfy the mentioned desideratum in a harmonious way. I think that the existence of these close connections between the accounts, just as much as the other independent motivations for them also considered here, constitutes in itself an important consideration in their favor.

References

- Allen, K. (2009), "Inter-Species Variation in Colour Perception", *Philosophical Studies* 142, 197-220.
- Armstrong, D.M. (1968), *A Materialist Theory of the Mind*, Routledge and Kegan Paul, London.
- Averill, E.W. (2005), "Toward a Projectivist Account of Color", *Journal of Philosophy* 102, 217-234.

- Berlin, B. and P. Kay (1969), *Basic Color Terms*, University of California Press, Berkeley.
- Block, N. (1999), "Sexism, Racism, Ageism and the Nature of Consciousness", *Philosophical Topics* 26, 39-70.
- Boghossian, P. and D. Velleman (1989), "Colour as a Secondary Quality", *Mind* 98, 81-103.
- (1991), "Physicalist Theories of Color", *Philosophical Review* 100, 67-106.
- Braun, D. (2006), "Names and Natural Kind Terms", in E. Lepore and B. C. Smith (eds.), *The Oxford Handbook of Philosophy of Language*, Oxford UP, New York, 490-515.
- Brodal, P. (2010), *The Central Nervous System. Structure and Function*, 4th edn., Oxford UP, Oxford.
- Brown, D.H. (2006), "On the Dual Referent Approach to Colour Theory", *Philosophical Quarterly* 56, 96-113.
- Byrne, A. (2003) "Color and Similarity", *Philosophy and Phenomenological Research* 66, 641-665.
- Byrne, A. and D. Hilbert (1997), "Colors and Reflectances", in Byrne and Hilbert (eds.), *Readings on Color, Vol. 1. The Philosophy of Color*, MIT Press, Cambridge (Mass.), 1997, 263-288.
- (2003), "Color Realism and Color Science", *Behavioral and Brain Sciences* 26, 3-21.
- (2004), "Hardin, Tye, and Color Physicalism", *Journal of Philosophy* 101, 37-43.
- (2007), "Truest Blue", *Analysis* 67, 87-92.
- Chalmers, D. J. (2006), "Perception and the Fall from Eden", in T. S. Gendler and J. Hawthorne (eds.), *Perceptual Experience*, Clarendon Press, Oxford, 49-125.
- Cohen, J. (2004), "Color Properties and Color Ascriptions: A Relationalist Manifesto", *Philosophical Review* 113, 451-506.
- (2009), *The Red and the Real*, Oxford UP, New York.
- Cohen, J., C. L. Hardin and B. P. McLaughlin (2006), "True Colours", *Analysis* 66, 335-340.
- De Valois, K. K. and R. L. De Valois (2000), "Color Vision", in K. K. De Valois (ed.), *Seeing*, Academic Press, London, 129-175.
- Erickson, G. L. (1979), "Children's Conceptions of Heat and Temperature", *Science Education* 63, 221-230.
- Fruhstorfer, H., U. Lindblom and W. G. Schmidt (1976), "Method for quantitative estimation of thermal thresholds in patients", *Journal of Neurology, Neurosurgery, and Psychiatry* 39, 1071-1075.
- Gendler, T. S. and J. Hawthorne (2006), "Introduction: Perceptual Experience", in T. S. Gendler and J. Hawthorne (eds.), *Perceptual Experience*, Clarendon Press, Oxford, 1-30.
- Gómez-Torrente, M. (2011), "Kripke on Color Words and the Primary-Secondary Quality Distinction", in A. Berger (ed.), *Saul Kripke*, Cambridge UP, Cambridge, 290-323.
- Hansen, N. (2011), "Color Adjectives and Radical Contextualism", *Linguistics and Philosophy* 34, 201-221.
- Hardin, C.L. (1988), *Color for Philosophers. Unweaving the Rainbow*, Hackett, Indianapolis. Expanded edition: Hackett, Indianapolis, 1993.
- (2003), "A Spectral Reflectance Doth Not a Color Make", *Journal of Philosophy* 100, 191-202.
- (2006), "Comments", *Dialectica* 60, 341-345.
- Harju, E.-L. (2002), "Cold and Warmth Perception Mapped for Age, Gender, and Body Area", *Somatosensory & Motor Research* 19, 61-75.
- Hilbert, D. (1987), *Color and Color Perception*, CSLI, Stanford.

- Hurvich, L. M., D. Jameson and J. D. Cohen (1968), "The Experimental Determination of Unique Green in the Spectrum", *Perceptual Psychophysics* 4, 65-68.
- Jackson, F. (1996), "The Primary Quality View of Color", *Philosophical Perspectives* 10, 199-219.
- (1998), *From Metaphysics to Ethics. A Defence of Conceptual Analysis*, Clarendon Press, Oxford.
- Johnston, M. (1992), "How to Speak of the Colors", *Philosophical Studies* 68, 221-263.
- Kalderon, M. E. (2007), "Color Pluralism", *Philosophical Review* 116, 563-601.
- Kay, P. and C. K. McDaniel (1978), "The Linguistic Significance of the Meanings of Basic Color Terms", *Language* 54, 610-646.
- Kennedy, C. (2007), "Vagueness and Grammar: the Semantics of Relative and Absolute Gradable Adjectives", *Linguistics and Philosophy* 30, 1-45.
- Kennedy, C. and L. McNally (2005), "The Syntax and Semantics of Multiple Degree Modification in English", in S. Müller (ed.), *Proceedings of the 12th International Conference on Head-Driven Phrase Structure Grammar*, CSLI, Stanford, 178-191.
- (2010), "Color, Context, and Compositionality", *Synthese* 174, 79-98.
- Kripke, S. (1972), "Naming and Necessity", in D. Davidson and G. Harman (eds.), *Semantics of Natural Language*, Reidel, Dordrecht, 1972, 253-355.
- Kuehni, R. G. (2014), "Unique Hues and their Stimuli—State of the Art", *Color Research & Application* 39, 279-287.
- Lasersohn, P. (2005), "Context Dependence, Disagreement, and Predicates of Personal Taste", *Linguistics and Philosophy* 28, 643-686.
- Lewis, D. K. (1997), "Naming the Colours", *Australasian Journal of Philosophy* 75, 325-342.
- Lewis, E. L. and M. C. Linn (1994), "Heat, Energy and Temperature Concepts of Adolescents, Adults and Experts: Implications for Curricular Improvements", *Journal of Research in Science Teaching* 31, 657-677.
- MacFarlane, J. (2009), "Nonindexical Contextualism", *Synthese* 166, 231-250.
- McGinn, C. (1983), *The Subjective View. Secondary Qualities and Indexical Thoughts*, Clarendon Press, Oxford.
- McLaughlin, B. P. (2003), "Color, Consciousness, and Color Consciousness", in Q. Smith and A. Jolic (eds.), *Consciousness: New Philosophical Perspectives*, Oxford UP, New York, 97-154.
- Malkoc, G., P. Kay and M. A. Webster (2005), "Variations in Normal Color Vision. IV. Binary Hues and Hue Scaling", *Journal of the Optical Society of America A* 22, 2154-2168.
- Matthen, M. (1988), "Biological Functions and Perceptual Content", *Journal of Philosophy* 85, 5-27.
- (1999), "The Disunity of Color", *Philosophical Review* 108, 47-84.
- Maud, J. B. (1981), "Colour—A Case for Conceptual Fission", *Australasian Journal of Philosophy* 59, 308-322.
- (1995), *Colours: Their Nature and Representation*, Cambridge UP, Cambridge.
- Mizrahi, V. (2007), "Color Objectivism and Color Pluralism", *Dialectica* 60, 283-306.
- Pautz, A. (2003), "Have Byrne and Hilbert Answered Hardin's Challenge?", *Behavioral and Brain Sciences* 26, 44-45.
- (2006), "Can the Physicalist Explain Colour Structure in Terms of Colour Experience?", *Australasian Journal of Philosophy* 84, 535-564.

- Peacocke, C. (1984), "Colour Concepts and Colour Experience", *Synthese* 68, 365-382.
- Rosenthal, D. M. (1999), "The Colours and Shapes of Visual Experiences", in D. Fissette (ed.), *Consciousness and Intentionality: Models and Modalities of Attribution*, Kluwer, Dordrecht, 95-118.
- Rothschild, D. and G. Segal (2009), "Indexical Predicates", *Mind and Language* 24, 467-493.
- Schaffer, J. (2011), "Perspective in Taste Predicates and Epistemic Modals", in A. Egan and B. Weatherston (eds.), *Epistemic Modality*, Oxford UP, New York, 179-226.
- Shoemaker, S. (1986), Review of McGinn (1983), *Journal of Philosophy* 83, 407-413.
- (1994), "Phenomenal Character", *Noûs* 28, 21-38.
- (2006), "On the Ways Things Appear", in T. S. Gendler and J. Hawthorne (eds.), *Perceptual Experience*, Clarendon Press, Oxford, 461-480.
- Siegel, S. (2006), "Which Properties Are Represented in Perception?", in T. S. Gendler and J. Hawthorne (eds.), *Perceptual Experience*, Clarendon Press, Oxford, 481-503.
- (2010), *The Contents of Visual Experience*, Oxford UP, New York.
- Siewert, C. (1998), *The Significance of Consciousness*, Princeton UP, Princeton.
- Smart, J.J.C. (1963), *Philosophy and Scientific Realism*, Routledge, London.
- Szabó, Z. G. (2001), "Adjectives in Context", in I. Kenesei and R. M. Harnish (eds.), *Perspectives on Semantics, Pragmatics, and Discourse*, John Benjamins, Amsterdam.
- Travis, C. (1994), "On Constraints of Generality", *Proceedings of the Aristotelian Society* n.s. 44, 165-188.
- Tye, M. (2000), *Consciousness, Color and Content*, MIT Press, Cambridge (Mass.).
- (2006), "The Puzzle of True Blue", *Analysis* 66, 173-178.
- Wright, C. (1988), "Moral Values, Projection and Secondary Qualities", *Proceedings of the Aristotelian Society*, supp. vol. 62, 1-26.