THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Structural realism for secondary qualities

Citation for published version:

Isaac, AMC 2014, 'Structural realism for secondary qualities', *Erkenntnis*, vol. 79, no. 3, pp. 481-510.
<https://doi.org/10.1007/s10670-013-9519-3>

Digital Object Identifier (DOI):

[10.1007/s10670-013-9519-3](https://doi.org/10.1007/s10670-013-9519-3)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Erkenntnis

Publisher Rights Statement:

The final publication is available at Springer via <http://dx.doi.org/10.1007/s10670-013-9519-3>

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Structural Realism for Secondary Qualities

Alistair M. C. Isaac

July 27, 2013

Abstract

This paper outlines and defends a novel position in the color realism debate, namely *structural realism*. This position is novel in that it dissociates the veridicality of color attributions from the claim that physical objects are themselves colored. Thus, it is realist about color in both the semantic and epistemic senses, but not the ontic sense. The generality of this position is demonstrated by applying it to other “secondary qualities,” including heat, musical pitch, and odor. The basic argument proceeds by analogy with the theory of measurement. I argue that perceptual experiences are analogous to numerical structures in that they are suitable for measurement, but only report measured values after they have been linked to states of a measurement device via *calibration*. Since the calibration of our sensory apparatus varies with context, it is inappropriate to identify specific experiences with specific properties in the world. Rather, it is structural relations between possible experiences which represent relations between possible external properties, and it is at the level of these structural relationships that veridicality is appropriately assessed.

1 Introduction

What is the relationship between the world *as we experience it* and the world *as it is*? This is the question of perceptual realism, the question of whether or not the properties attributed to the world in our experience are in fact the properties of the world. This question has been most hotly debated for the property of color, but the considerations raised in that literature apply equally to the rest of the so-called “secondary qualities.”¹ My aim here is to defend a novel position in this debate, namely *structural realism*.

The basic idea behind structural realism is that our experience of secondary qualities conveys only relational information to us about the world.² An experience of *this much warmth* does not convey an absolute value of *this much temperature*. Rather, it conveys to us the *difference in temperature* between the warmth inducing stimulus and some baseline. This insight motivates a novel interpretation of some well known phenomena. For example, if one hand is cooled while

¹I use the term “secondary quality” throughout to refer merely to a well-known category of epistemologically worrisome properties. I do not intend to thereby endorse any substantive theory of the primary / secondary quality distinction.

²This view should not be confused with the view that secondary qualities are “relational” in the sense that they are defined in terms of the relation between the organism and the environment.

the other is warmed, then both hands are thrust into a lukewarm bucket of water, the cool hand will sense the water as warm, while the warm hand will sense the water as cool. On the account developed here, these apparently “contradictory” sensations may both be veridical.

I develop this view by analogy with the theory of measurement. When a measurement is performed, a correspondence is established between some quantity in the world and a numerical value. The veridicality of the assignment of a particular number to a particular quantity in the world depends crucially on the *calibration* of the measuring device. The assignment of 86° to today’s temperature in Houston may be correct if the thermometer is calibrated for degrees Fahrenheit, but incorrect if it is calibrated for degrees Celsius. Our sensations of warmth or coolness are analogous here to the numbers of the real line. They hold the potential for representing external temperature, but they cannot actually perform that function until a correspondence is established through an act of calibration. In the case of the lukewarm water, each hand has been calibrated differently. Just as there is not *necessarily* a contradiction between a thermometric reading of 86° and one of 30° , there is no necessary conflict between the assignments of warm and cool delivered by the differently calibrated hands, nor is there one between an assignment of orange and one of brown to the same surface when viewed under different lighting conditions. Although one *could* stipulate that veridicality is fixed to a single baseline, e.g. a “standard lighting condition” or “typical body temperature”, to do so would misrepresent the information conveyed by experience, just as insisting that only the Celsius, but not the Fahrenheit, scale is “correct” would misrepresent the information conveyed by practical thermometry.³

This account is *structural* in the sense that the correct analysis of the relationship between the world *as we experience it* and the world *as it is* is one of *structural correspondence*. It is a form of *realism* in the sense that sensations may be evaluated for their veridicality. This veridicality rests not on the correct representation of properties in the world, however, but of *relations* between properties. In the case of colors, it is not particular surface properties of objects which particular color sensations represent, but rather relations between surface properties are represented by relations between sensations. Consequently, we may maintain both that surfaces are not in fact “colored” (in the sense that we experience color properties) but also that attributions of colors to surfaces are typically veridical. The veridicality of these attributions does not rest on objects in the world having the properties we experience them as having, just as my success in referring to trees with the word “tree” does not rest on trees in the world comprising four letters, or a single syllable. An important starting point for this position is the individuation of perceptual properties by their phenomenology, not their content. This is because the representational content of a perceptual property will in general change across different calibrations.

This point illustrates a striking advantage of the structural realist position: it characterizes the *general presuppositions* of the methods by which perceptual experience is scientifically studied. These methods presuppose independent individuations of sensations and of the physical correlates

³Although the analogy between perception and measurement has been much discussed, including the comparison between thermometry and color vision (e.g. Tye, 2006), these discussions differ radically from that presented here. In particular, even when calibration has been discussed (e.g. Matthen, 2005, 260f), its significance for revealing the context relativity of measurement values is not recognized (Section 4.3.1).

of these sensations. This is in sharp contrast with physicalist analyses of color, which rest upon *particular theses* of perceptual science, and are thus contingent on the outcome of scientific disputes (Section 4.2). This advantage becomes especially clear when we consider secondary qualities which are more poorly understood than color. Whereas a physicalist analysis of olfactory qualities must wait for a more mature theory of the physical correlates of smell, the structural realist analysis can already provide a general answer to the epistemic question of the relationship between smells as experienced and smells in the world. Furthermore, it can explain the practices of olfactory research, which we observe struggling to characterize (i) the structure of our experience of smell; (ii) the structure of the physical correlates of smell; and (iii) the process of calibration which explains the shifting correspondence between these two across different contexts (Section 5.2).

This notion of structural realism shares some features with that which has recently become popular in philosophy of science. When considering an answer to the realism question for scientific theories, philosophers have struggled to find a middle ground between the Scylla of the “No Miracles” argument and the Charybdis of the “Pessimistic Induction.” The issue here is how successive stages in the theoretical development of a field which are apparently contradictory in the properties they ascribe to objects (e.g. Newtonian and Einsteinian theories of gravity) may nevertheless both be “true.” The answer provided by the structural realist has been first, that we should assess theories for the veridicality of the structural relations they ascribe to the world, not the intrinsic properties they ascribe to particular objects (Worrall, 1989), and second, that in order to be effective, this structural realism must not be merely *epistemic* (“all we can know is structure”), but *ontic* (“all that there is is structure”) (Ladyman, 1998). Although the view that the relationship between theory and world is properly understood as one of structural correspondence has a long history (a full account would include Poincaré, Hertz, Carnap, Grover Maxwell, and others), the contemporary literature has supplemented this long tradition with the stronger ontic claim, and a flurry of results relating to fundamental issues in philosophy of physics.

The view defended here is an epistemic rather than an ontic structural realism. This is in part because the problem addressed is somewhat different for science than for perception. Structural realists in philosophy of science worry about how successive theoretical structures may correspond veridically to a presumably unchanging external world, which is itself accessible only through science. In philosophy of perception, it has become customary to allow oneself the physical description of the external world and ask merely how experience relates to that description. So, a potentially dangerous circularity faces realists in philosophy of science which simply does not arise for philosophy of perception. Furthermore, the problem for philosophy of science is one of preservation of veridicality across directed, diachronic theory change. As presented below, the problem for philosophy of perception is one of preservation of veridicality across undirected, yet frequent, contextual changes.

In order to motivate structural realism, we begin in Section 2 with a discussion of the basic features of measurement. Two fundamental concepts for analyzing the epistemic status of perceptual experience are developed here: i) the concept of *calibration*, and ii) the distinction between *artifactual* and *representational* structure. The latter distinction will help dissolve some long standing

problems for realism about perceptual qualities. For instance, color similarity as assessed in experience does not appear to correspond to any physical similarity between the physical correlates of color. While this observation is usually taken as a challenge to realists, here it merely constitutes evidence that the similarity structure amongst color experiences is an artifact, with no representational content. I develop this argument in more detail in Section 3, which applies these concepts to the example of color perception.

In Section 4, I discuss the relationship between the position developed here and the rest of the color realism literature, focusing primarily on those views which are most closely related to my own. In general, structural realism is distinguished in this debate by insisting first, that color experiences not be individuated by their content, and second, that the attribution of color properties to surfaces in experience does not represent surfaces as having those very same properties. Nevertheless, the insistence on the distinction between color properties as experienced and those properties in the world with which they correlate does not imply widespread error. This issue is clarified by introducing a distinction from the philosophy of science literature between ontic, epistemic, and semantic interpretations of the realism question. While much of the color realism debate conflates these questions, the structural realist teases them apart, denying ontic realism, while endorsing semantic and epistemic realisms. Section 5 concludes with a quick survey of structural realist interpretations of odor and pitch perception. Here the close relationship between this view and the presuppositions behind perceptual science is illustrated by demonstrating that the epistemic assumption of structural realism explains the practice of psychologists working in these fields.

2 Basic Features of Measurement

The theory of measurement analyzes the relationship between (i) a measured space and (ii) a measuring space, as established by (iii) a measurement procedure. In the case of (typical) thermometry, for example, the measuring space is the real line, the measured space is the space of possible mean kinetic energies in the object, and the measurement procedure consists in holding a thermometer up against the object. After working through this example in more detail, I argue that the representational relationship between measuring and measured spaces is best thought of as a structural one. I conclude the section by applying this analysis to the sensation of heat, arguing that sensations of heat should be interpreted as measuring temperature. Since the representational content of sensations of heat stands in a structure preserving relationship to temperatures in the world, this analysis implies that the relationship between heat as experienced and the physical correlates of heat in the world is one of structural realism.

2.1 Measuring Temperature and Calibration

The standard view in the theory of measurement is that a foundation for the assignment of numerical values to the outcomes of a measurement procedure is provided by demonstrating an isomorphism from a structure axiomatically defined by the qualitative assumptions underlying the procedure into a numerical structure such as the real line. Such a “representation theorem”

demonstrates that for any model which satisfies the qualitative axioms there exists an isomorphic numerical structure, and this thereby legitimates our use of a numerical structure to represent all such models (Krantz et al., 1971).⁴ Let’s look at how this strategy applies to a specific example.

In the case of temperature, the qualitative measurement procedure involves holding a thermometer—for concreteness, a glass tube filled with mercury—up against various surfaces and noting the position of the height of the column of mercury. This procedure assumes that the values being measured can be linearly ordered, just as the relative heights of the column may be linearly ordered. Note, however, that the column is always there, it just responds differently to different surfaces. Because the column is always present, the physical behavior of the thermometer does not by itself imply a natural zero point.⁵ Because the column varies continuously in height, it does not imply any preferred unit size. But if we want to use our thermometer to assign numbers to our measurements, we need to specify a way to map mercury heights into the real line. In order to do this, then, we must pick a zero point and a unit size. Although the term has other technical meanings, for the purpose of the present discussion, we’ll call this process *calibration*.

Calibration – The establishment of a baseline correspondence between states of the measurement device and points in the measuring space such that each state of the device determines a unique point.

By writing a scale on the side of our thermometer, we calibrate it, thereby establishing a map from the property being measured into the real line.

In the contemporary theory of temperature, we analyze that which is measured in this case as mean molecular motion. But the theory of temperature as mean molecular motion is independent of the qualitative assumption that whatever is being measured can be linearly ordered. Consider, for example, the caloric theory, which analyzed temperature in terms of a special intermolecular fluid, caloric. Levels of (free) caloric in a body can also be ordered linearly. The naïve practice of thermometry just described does not make assumptions about whether that which is measured is mean molecular motion or free caloric, it only assumes that some property of the body varies linearly and that different values of this property affect the height of the column of mercury differentially.

And this is why the correct interpretation of the relationship between the numerical values assigned by a thermometric measurement and the measured property of the body is one of structural correspondence. The properties of caloric, mean molecular motion, and numbers are largely disjoint. One is a concrete substance, the other a summing over behavior, the third a set of abstract objects. What all three systems share, however, is the structural feature of being organized

⁴Recently, the standard view has come under fire from philosophers who argue that a full theory of measurement must also take into account the intentions of the scientist developing the measurement procedure (van Fraassen, 2008) or more details of the empirical procedure itself than just the axiomatic characterization of its presuppositions (Frigerio et al., 2010). I set these subtleties aside here, but note in passing that the development of the analogous measurement procedure in perceptual systems took place on an evolutionary time scale, and the appropriate analog to *scientist’s intentions* on the picture offered here would thereby be something like *selective evolutionary pressures*.

⁵It is important not to confuse the presuppositions of our current theoretical analysis of temperature with the presuppositions of naïve thermometry as described in this section. While the former imply a natural zero point, the latter do not. This issue is discussed further in the following section, which contrasts naïve thermometry with other types of measurement, including thermometric measurement in degrees Kelvin.

linearly: there can be more or less caloric, more or less mean molecular motion, greater or lesser numbers.

To summarize: in the case of simple thermometry, the measuring space is the space of possible real numbers. Today, we interpret the measured space as the space of possible mean molecular motions. The only property of this space which is represented in the measuring space, however, is the relational property that mean molecular motions can be linearly ordered. The correspondence between the space of mean molecular motions and the space of numbers is established by a combination of a physical process which responds differentially to the measured space, namely the equilibrium height of a column of mercury when a thermometer is held against a body, and a calibration procedure, which establishes a convention for assigning relative heights unique numbers.

2.2 Artfactual and Representational Structure

So, a measurement procedure establishes a structural correspondence between two spaces. What can we learn about the measured space by examining the measuring space? The answer to this question depends upon the nature of the calibration procedure.

The most important thing to notice is that a measuring space must have antecedent structure. We already ordered real numbers linearly before Galileo suggested to Sagredo that he write numbers on the side of a tube filled with spirits and lower it down a well.⁶ But because the structure of the real numbers is present antecedent to its use as a measuring space, there is no guarantee that it will all correspond to structure in the measured space. In fact, one of the central practices of measurement theory is the categorization of scales in terms of those aspects of the antecedent structure of the real line to which they assign representational content. Correct categorization is crucial for determining whether a particular relationship definable in terms of numbers is meaningful when those numbers are interpreted as measurement values.

Consider for example ratios between real numbers. If x and y are real numbers, then x/y is a meaningful quantity and we can meaningfully assert, for example, that if $x/y = 1/2$, then y is twice the value of x . If it is 100° Fahrenheit in Houston, Texas, and 50° Fahrenheit in Anchorage, Alaska, is it meaningful to say that it is twice as hot today in Houston as it is in Anchorage? No, and the answer can readily be seen by translating the temperatures in Houston and Anchorage into Celsius, namely 37.8° and 10° respectively: $50/100 = 1/2 \neq 10/37.8$.

The problem here is that ratios inherit their meaning from the fixity of the zero point. Since our thermometric calibration procedure set a zero point arbitrarily, that structure in the real line which depends upon the significance of the zero point does not represent anything about the structural relationship between temperatures. We can see this in the formula for converting Fahrenheit into Celsius. If y is degrees in Fahrenheit and x degrees in Celsius, then the two values are related by the formula

$$y = x\left(\frac{9}{5}\right) + 32.$$

⁶Obviously, this is a caricature of the history, but the basic point is correct. Thermometry begins around 1600 and the differences in calibration procedures across researchers meant a) that only claims of relative temperature could be communicated, and b) that the establishment of fixed point standards for calibration became the primary goal of early thermometry (Chang, 2004, Ch. 1).

This is an affine transformation: it both changes zero point (by adding 32) and unit size (by multiplying by 9/5). Only structural features of the real line which are invariant across affine transformations are meaningful as representations of structure of the measured space of possible temperatures.

These considerations motivate a distinction between representational structure and artifactual structure:

Representational Structure – Those structural features of a measuring space which are invariant across all mappings from a model of the qualitative assumptions of the measurement procedure.

Artifactual Structure – Those structural features of a measuring space which are *not* invariant across all mappings from a model of the qualitative assumptions of the measurement procedure.

The theory of measurement organizes numerical scales in terms of their relative proportions of representational to artifactual structure. In *ratio scales*, i.e. those with meaningful zero points, such as the measurement of length or weight, ratios are meaningful—it makes sense to say of this board that it is twice as long as that board, or of this baby that it weighs twice as much as that baby. In *interval scales*, such as those resulting from simple thermometry or IQ testing, ratios between values are not meaningful, but ratios between intervals are, e.g. if x_1 , x_2 , x_3 , and x_4 are temperature measurements in degrees Fahrenheit, then the value

$$\frac{x_1 - x_2}{x_3 - x_4}$$

is meaningful because it is invariant across affine transformations (Krantz et al., 1971, Ch. 1; see also Luce et al., 1990, Ch. 22).

For example, suppose it is 90°F in Houston today at 9 a.m. and 45°F in Anchorage also at 9 a.m. The temperature at both locations is measured again at noon, and it is then 100°F in Houston and 50°F in Anchorage. The claim that between 9 and 12 this morning, the temperature increased twice as much in Houston as in Anchorage is meaningful because it is a claim about the ratio between intervals and, consequently, invariant across affine transformations such as conversion to Celsius.

$$\frac{100 - 90}{50 - 45} \text{ (Fahrenheit)} = \frac{2}{1} = \frac{37.8 - 32.2}{10 - 7.2} \text{ (Celsius)}$$

Ordinal scales are invariant under any monotonic increasing function, consequently they use even less of the structure of the real line as representational structure. Only the ordering of assigned values is meaningful for such a scale, not the distances between them. Since the only representational structure we need is an ordering in the case of an ordinal scale, we could easily use some measuring space other than the real line, so long as it has an antecedent ordering defined over it. Consider, for example, the Mohs hardness scale, which orders minerals by the hardest sample substance they can scratch. Distance relations in this scale are not meaningful, only the ordering it produces. Traditionally we use natural numbers to represent this ordering, but we could

just as easily use letters of the alphabet, days of the week, or middle names of presidents of the United States. Any structure with an antecedent ordering could represent the exact same structure as the natural numbers typically do in this case. An example of this practice for a non-scientific ordinal scale is the use of letters to represent notes in a musical scale.

In the limit, simple categorization via some specified procedure is also a form of measurement (the *nominal scale* of Stevens, 1946). Here the relation between measuring and measured spaces is still “structural” although very little structure is preserved, merely difference in category membership. Consider, for example, the classification of brainwaves into Gamma, Alpha, Beta, and Theta waves. In this example, there is antecedent structure in the measuring space, namely the standard ordering of letters in the Greek alphabet, but this ordering does not correspond to an ordering in the measured space—if we order brainwaves by frequency, from greater to lesser, we get $\text{Gamma} > \text{Beta} > \text{Alpha} > \text{Theta} > \text{Delta}$.

It is important to notice that the distinction between representational and artifactual structure depends on the assumptions of the measurement procedure *not* on absolute properties of the measured space. Once we interpret temperature as mean molecular motion, we can theoretically define a meaningful zero point (zero motion) and thereby establish a ratio scale for temperature, such as degrees Kelvin. This does not make ratios between temperatures *as measured by simple thermometry* meaningful, however. If we convert a measurement made in Fahrenheit into Kelvin (by first transforming it into Celsius, then adding 273, an instance of an affine transformation), this new value is properly understood as the outcome of a new measurement procedure, one with qualitatively different assumptions than simple thermometry. The assumptions of this new procedure are those of simple thermometry *plus* the theoretical assumptions which motivate the analysis of the zero point for mean molecular motion.

This also illustrates a final point: measurement procedures can be arbitrarily complex and theory laden. Often (typically!) measurements are “indirect,” “derived,” or made *by proxy*.⁷ Measurement devices can be arbitrarily complex (think of a Geiger counter, or the detectors used in a particle accelerator), yet still be understood from a measurement theoretic standpoint as making relatively simple qualitative assumptions about the measured domain. An illustration of this can be seen in the recent controversy over the measurement of neutrino speed (Cartlidge, 2012). Speed is assumed to be organized into a ratio scale just like length, but the device which delivers numerical values on this scale for neutrinos is large, complex, and depends upon many more theoretical assumptions and technical details than the simple practice of holding a ruler against an object.

2.3 The Sensation of Heat

I claim that the sensation of heat stands in the same relationship to temperature as the outcome of a simple thermometric measurement. Sensations of heat are linearly ordered against a neutral

⁷Arguably, the only direct form of scientific measurement is length measurement (George Smith, personal communication); we measure the height (length) of the column of mercury directly, but only in an indirect way measure some value of the body against which we hold the thermometer. For a detailed discussion of measurement by proxy using the example of Thomson’s measurement of the charge of the electron, see Smith, 2001.

baseline. Above this baseline, we describe them as more or less warm, below the baseline, more or less cold. Just as a single thermometer can be calibrated differently to deliver numbers on either the Fahrenheit or Celsius scale, so also the same physiological apparatus may be calibrated against different baselines when generating sensations of heat. If I hold ice in one of my hands for five minutes but not the other, then plunge both hands into a bucket of water, the hand which formerly held ice will sense the water as warmer than the hand which did not. The sensory (measurement) devices in the two hands have been calibrated differently.

It is important to note that we need not equate the measurement procedure corresponding to sensation with just the interaction between sensor and stimulus at the surface of the skin. Just as a simple thermometric measurement may be combined with theoretical assumptions to produce a value on the Kelvin scale, or the sensor activity in a particle accelerator may undergo complex processing in order to deliver a value for neutrino speed, the neural processing of the signal returned from the nerve endings at the surface of the skin may be arbitrarily complex. Whatever the neural correlates here may be, from the standpoint of phenomenal experience, we clearly sense objects as more or less warm. Our ability to compare these phenomenal sensations indicates that they stand in some structural relationship to each other, and our experimentally confirmed ability to order objects linearly with respect to the sensations they produce (think Goldilocks and the three bowls of porridge) confirms that our possible sensations of heat are linearly ordered.

The calibration of sensory experience is largely opaque. In fact, the history of early thermometry involves a sequence of discoveries about the heretofore unforeseen degree to which sensations of heat were subject to calibration. In 1615, for example, Sagredo wrote to Galileo in excitement at his discovery that “well-water is colder in winter than in summer ... although our senses tell differently.”⁸ Well water feels cooler to us in summer than in winter because the baseline ambient temperature calibrates our sensations, but the fact of this calibration is not transparent. Only with an external instrument, one subject to a different calibration process, could Sagredo discover the extent to which our sensations of warmth and cold depend upon a variable baseline.

If we accept this analogy, then it appears that the relationship between heat *as we experience it* and heat *as it is in the world* is one of structural correspondence. The physical causes of experiences of heat are linearly ordered, as are our sensations of heat. When a baseline is fixed, then our sensations of greater or lesser heat veridically represent the relative ordering of these physical causes. If we attempt to compare sensations of heat across different calibrations, however, as when we compare our experience of the coolness of well water in the summer to that of its warmth in the winter, we may arrive at false conclusions. The error here is analogous to the error of comparing thermometric measurements across different calibrations, for instance if we conclude it is warmer today in Anchorage than in Saskatoon since it is 50° in the former and 11° in the latter, neglecting the fact that the former measurement is in Fahrenheit while the latter is in Celsius.

Moreover, the observation that our sensations are subject to contextual calibration demonstrates why it would be incorrect to identify the content of a particular heat sensation with a particular (range of) mean molecular motions *simpliciter*. There is a double dissociation between

⁸Translated in Müller and Weiss (2005, 224).

heat sensations and temperatures: once we consider the structural correspondence between sensations and temperatures across different contexts, we realize that different heat sensations can veridically represent the same temperature, and the same heat sensation may veridically represent different temperatures. Consequently, despite the typical veridicality of heat sensations, our realism about them must remain at the structural level.

3 Structural Realism for Color

The realism debate about secondary qualities has been most extensive for the example of color. In this section, I defend the positive proposal for structural realism about color, with a special focus on the importance of the distinction between representational and artifactual structure for clarifying questions about how color experience represents the world. I will reserve discussion of the relationship between structural realism and other positions in the color realism debate for the following section.

3.1 Color Vision as Measurement

We perceive surfaces as colored, but how does our experience of surfaces *as colored* relate to the properties of surfaces *as they are*? I argue that color sensations *measure* properties in the world, typically surface properties, but also properties of translucent solids and light sources. The argument for this position is a demonstration that the relationship between color sensations and properties in the world is analogous to the relationship between a measuring space and a measured space. If this argument is correct, then we should endorse structural realism about color.

Our possible experiences of color are organized into a geometrical space commonly called the color solid. The representation of this space most familiar to philosophers is as a spindle with a vertical axis (lightness), a radial axis (saturation), and a circular axis (hue). The color solid characterizes the relative distances between possible experiences of color, as determined through assessments of color similarity. Although these distances vary across subjects, and even within subjects from day to day, they are fixed enough that they can be determined with a high degree of precision by psychophysical methods such as color matching experiments. The asymmetries of this space and the distances within it are remarkably robust across observers.⁹

From the standpoint of physics, the property of a surface which determines the color which will be attributed to it is its surface spectral reflectance profile (SSR), this is the percentile for each possible wavelength of light in the visible range (roughly 400–700 nm) with which that wavelength is reflected when incident on the surface. An illuminant is characterized by its spectral power distribution (SPD), a function which gives the strength of each wavelength in the emitted light. The color signal which arrives at the retina after light from an illuminant I has bounced off a surface S is then given by $SSR_S \times SPD_I$.

⁹For a survey of the many proposed color solids and a discussion of their respective virtues, see Kuehni and Schwarz, 2008. For a discussion of the history of experimental methods for investigating the color solid and an assessment of the evidence for similarities in color experience across observers, see Isaac, 2013.

The “measurement procedure” for color perception involves the transduction of the color signal at the retina by photoreceptor cells (rods and cones) plus later processing of this signal in the retina, the lateral geniculate nucleus, and further cortical regions in the visual processing chain. The crucial fact about this procedure for the present discussion is that the assignment of color values is calibrated by the SPD of the illuminant (as well as other contextual features of the scene). This calibration effect generates the paradoxical phenomena of “color constancy.” On the one hand, our assignment of baseline categories such as neutral white is relatively fixed across gross changes in illuminant. On the other hand, close attendance to the phenomenal features of our experience of a particular surface demonstrates variation across these changes. Our implicit knowledge of the context sensitive calibration of color vision allows us to conclude that surface properties nevertheless remain fixed. I can notice that the carpet looks yellowish orange in the sunlight, but dark red in shadow, yet still maintain that its surface is uniformly characterized by the same set of properties. Control of contextual effects in artificial images allows these features of phenomenal color experience to be systematically manipulated. These manipulations demonstrate that we may attribute the same color to surfaces with different reflectance properties, but also different colors to surfaces with the same reflectance properties. Thus, color sensation exhibits the same double dissociation from surface properties as heat sensation does from temperature.

Just as in the case of heat perception, the exact details of the processing involved in the physiology of color perception are not fully known (although they are much better understood for color than for heat!). We do know that the illuminant calibrates this procedure, however, because models for predicting color appearance must take into account not only the color signal incident at the retina, but also the illuminance level (and, in more elaborate models, other facts about context as well, for instance absolute luminance, SSR’s for surround and background, and even the spatial organization of the scene, Fairchild, 2005, 184). The basic fact here is summarized succinctly by Wandell: “The appearance of an object in a scene is generally predicted somewhat better by the tendency of the *surfaces to reflect light* rather than by the actual light arriving at the eye” (1989, 187).

These facts about color vision are easily explained by interpreting the color solid as a measuring space. It measures some property of surfaces, plausibly surface spectral reflectance profiles. The measurement procedure involves not only the transduction of the color signal at the retina, but also complex processing of this signal before the neural correlates of color experience (whatever they may be) are triggered. This measurement procedure is calibrated during this processing in a manner controlled by the illuminant and other features of the scene. The effect of calibration is a relative fixity in assessments of the relative differences between surfaces independent of the SPD of the illuminant.

The properties of surfaces measured by color sensations need not be SSRs, they may instead be chemical or ecological properties (see discussion in Sections 3.2 and 4.3 below). What is clear, however, is that these properties are not themselves colors *as we experience them*, any more than mean molecular motions are themselves numbers. Nevertheless, modulo a particular calibration, color attributions may be assessed for veridicality. This is because the attribution of a color to a

surface does not depend upon the surface itself being colored, but rather on the structural relations between that surface’s measured property and other possible values of that property. There are three important reasons why we must resist the urge to make the further claim that these measured properties are just colors. First, to do so would conflate conceptually distinct categories (c.f. Section 4.2). Second, to do so would ignore the fact that color experience may measure different types of property in different circumstances (e.g. surface properties versus properties of translucent solids). Third, assignment of color experience to surface property will in general differ with different calibrations (a color swatch experienced as forest green under this lighting may be experienced as hunter green under that lighting). Maintaining a distinction between colors as experienced and surface properties of objects reconciles this phenomenological fact with the intuitive judgment that our color attributions are not typically in error.

3.2 Artifactual Structure in the Color Solid

I believe the reason that structural realism is not a standard position in the color realism debate is the observation that the structural relations between colors do not seem to correspond to any physically interesting structural relations between surface spectral reflectance profiles. A straightforward application of the distinction between representational and artifactual structure demonstrates that this inference is fallacious. In fact, structural realism handles this apparent discrepancy between experience and the physical better than other forms of realism: it transforms an ontological embarrassment into an evidentiary virtue.

Two surfaces are *metamers* if they are physically different but perceived as perceptually identical under a fixed illumination.¹⁰ In general, the SSRs of surface metamers are not “similar” in any physical sense. Furthermore, SSRs which correlate with similar color experiences need not be “similar” or “close” in any obvious physically specifiable way.¹¹ More subtle structural relations between color experiences also fail to have any obvious physical correlates, for instance color opponency phenomena such as the apparent opposition between yellow and blue or green and red (Hardin, 1988, 45–58, 121–7).

The first point to note here is that, even if none of the qualitative relations between colors as experienced are mirrored in qualitative relations between SSRs, the interpretation of color experience as measurement of SSR, and consequently the structural realist view, is not thereby undermined. The assignment of different instances of a measured domain to different points within a measuring domain is an act of measurement even if all the remaining structure of the domain is artifactual (this is the case with categorization of brain waves into Gamma, Alpha, etc.). Furthermore, the assessment of such a procedure as an act of measurement is not undermined if the categories in the measured domain turn out to lack significance, i.e. if the categories of SSRs corresponding to par-

¹⁰Note that for any pair of surface metamers, there will be *some* illuminant under which they appear different. To see this, note that in order to be physically different, they must differ with respect to the reflectance of at least one wavelength of light. Now consider the surfaces as illuminated by monochromatic light at precisely this wavelength. The one which reflects more will appear lighter. In general, even relatively minor changes in illuminant are enough to distinguish formerly metameric surface pairs.

¹¹*Pace* the attempt by Churchland (2007) to provide such a specification; for a rebuttal see Kuehni and Hardin (2010).

ticular colors turn out to share no feature in common other than that they are categorized together by human experience. Consider, for example, the phrenologist’s measurement and categorization of bumps on the skull. We now judge these bumps to be of no theoretical significance, but so long as there is consistency in her procedure, the phrenologist still satisfies the logical requirements for performing a measurement.

Nevertheless, the conclusion that all structural relations between color categories are artifactual is way too strong. At the very least, the ordering of hues around the color solid corresponds to the ordering of homogeneous lights by wavelength. Furthermore, if SSR_1 and SSR_2 are “similar” in the sense that their curves are very close together, the corresponding color sensations will also be close. This follows immediately from the fact that a continuous change in the spectral power distribution incident at the retina results in a continuous change in color experience (a fact utilized to great effect in color matching experiments). So, even if we accept that the existence of metamers demonstrates that some structure in the color solid is artifactual, it does not follow that all structure is artifactual.

Furthermore, although we have adopted the working hypothesis that color experience measures SSR, other interpretations of the measured space are consistent with the structural realist view. We can conclude that *some* property of surfaces is being measured by color experience from the relative regularity with which we assign colors to surfaces. But this regularity by itself does not tell us *which* property of the surface is being measured. Taking the evolutionary perspective, we might ask: *given* the structural features of the color solid, which features in the world are plausible candidates for a measured space? From this perspective, it is not physical, but rather biological or ecological properties which are more plausible candidates.

Many of the apparently arbitrary features of color vision can be explained once one takes the ecological perspective. For instance, Kurt Nassau has emphasized that the range of wavelengths to which the human eye is sensitive is precisely that at which the interaction between radiation and molecules is substantive, but not destructive, making it ideal for the detection of chemical properties of surfaces (2001, 31). Furthermore, the three dimensionality of color space seems much less restrictive once we note that most naturally occurring SSRs are smooth curves, and can be recovered through the linear combination of relatively few basis curves (Maloney, 1986). For a sustained analysis of this kind of consideration, see Shepard (1992).

4 The Rest of the Realism Debate

I have postponed discussion of the color realism debate because the view advanced here does not fall within any of the broad categories of response within that debate as typically construed. This is because of the novel feature of structural realism, namely the dissociation of the veridicality of color attributions from the claim that objects themselves are colored. After discussion of the basic shape of this debate and my position within it, I contrast my position first with the physicalism of Byrne and Hilbert, and finally with some specific ecological views which share many features in common with structural realism.

4.1 Disentangling Three Types of Realism

In typical presentations of the question of color realism, it is framed as a question about the properties of physical objects: “Are physical objects colored?” Answers fall into three broad categories: eliminativism, relationalism, and realism. The eliminativist takes colors to be purely subjective features of experience, thereby “eliminating” them from the physical world. The realist takes colors to be objective properties in the world, and then must face the further question: which properties in the world? The relationalist walks a middle ground by affirming that colors indeed exist as properties of the world (thereby avoiding eliminativism), but insisting that colors are a special type of property, defined in terms of the relation between observer and object.¹²

If one endorses realism, then one might take colors to be primitive properties of surfaces (*primitivism*, e.g. Campbell, 1993), or one might identify them with (metameric sets of) SSRs (*physicalism*, e.g. Byrne and Hilbert, 2003a; Churchland, 2007). A common form of relationalism is *dispositionalism*, the view that colors are defined in terms of the dispositions of surfaces to cause color experiences in standard observers. Although this view is frequently identified with Locke, there is some controversy over how exactly to interpret his position (for a modern example of dispositionalism, see Johnston, 1992). *Eliminativists* tend to emphasize the claim that science has demonstrated that no physical feature of surfaces exhibits the right properties (e.g. similarity relations) to count as colors, and therefore there are no colors (Hardin, 1988). The *ecological* view takes colors to be properties of relevance to the organism on an evolutionary timescale; this position can be cashed out in either realist or relationalist terms (see Section 4.3).

However, we have been careful to distinguish in the above discussion two senses of “color”: first, colors as experienced, individuated by their phenomenal character; second, colors in the sense of properties in the world measured by experience. If we interpret “color” in the standard question of color realism in the second sense, the answer appears to be simply analytic: do surfaces have whatever surface properties our experience of color measures? Of course. But if we replace the question of whether or not surfaces have the property of color with the first sense of color, the question becomes conceptually problematic. Do physical surfaces have whatever phenomenal qualities characterize color as experienced? At best the question seems misguided, at worst an outright category mistake.¹³

The problem here is that the question of color realism is deeply entangled with a debate about the metaphysics of experience. If one has a view on which properties as they are can be experienced *directly*, or that a supervenience relation holds between properties of the world represented in experience and the phenomenal character of those experiences, then the view that phenomenal character reveals physical properties becomes plausible. These views contrast with those theories of experience on which properties as experienced are either themselves properties of

¹²For versions of this taxonomy see, e.g. Byrne and Hilbert, 1997, or Hatfield, 2003.

¹³In fact, the accusation that realists (in particular, physicalists) are guilty of a “Rylean category mistake” has been made by MacLeod (2003). Specifically, he argues that talk of estimation and recovery (see Section 4.2) encourages a category mistake since “the ‘estimated’ quantity may have no simple and well-defined physical referent” (433). He is commenting here on Mausfeld (2003), who criticizes not only physicalism, but also analogies with measurement in general. Mausfeld’s critique, however, applies only to measurement analogies which fail to make the representational / artifactual distinction, and so does not undermine structural realism.

experience (such as *adverbialism*) and those on which some third thing (whether it be cashed out as *sense data*, *qualia*, or something else) mediates between properties of the world and experience. On these latter views, while the phenomenal character of color experience *may* reveal the nature of properties in the world, such revelation requires separate argument due to the indirect nature of perceptual representation. Structural realism is consistent with any of the views in this latter category.

If we countenance a view on which our experience of the world represents its properties indirectly, then there is room to import a distinction from the debate on scientific realism which I think clarifies what is at stake in the color realism debate: this is the contrast between realism understood as an ontic, an epistemic, or a semantic position (Psillos, 1999, xix). The *ontic* realist about x makes a claim about the metaphysical status of x . If the traditional color realism question is interpreted in the ontic way, it asks of physical objects in the world whether they have that very property of the world as we experience it which we call color. This reading is perhaps most natural on views for which experience has direct access to properties of the external world, but it is by no means inconsistent with indirect views.

Nevertheless, if experience may only indirectly access properties of the world, we may also read the question as semantic or epistemic. The *semantic* reading takes the question “Are physical objects colored?” to be a question about the truth value of attributions of colors to objects: may I truthfully utter “that chair is red”? The *epistemic* reading asks whether color attributions to objects constitute a display of knowledge about the world. On this reading, we demonstrate that we know *something* about the world when we say “that chair is red”, even if what we know is not best characterized by the ontological claim that the chair has the property of redness. The structural realist is perfectly happy with asserting that physical objects are colored if this claim is understood on either the semantic or epistemic readings.

The break that the structural realist makes with the traditional debate is crystalized in two readings of the following claim:

CR – Sensations of objects as *colored*₁ represent objects in the world as *colored*₂.

If **CR** is interpreted such that *colored*₁ and *colored*₂ refer to the same property, then the epistemic and semantic questions reduce to the metaphysical question. The structural realist denies this interpretation, and endorses **CR** only if *colored*₁ and *colored*₂ are interpreted as referring to different properties. In their survey of the contemporary color literature, Byrne and Hilbert (1997) can identify only a single paper which rejects the one property reading of **CR**.¹⁴ The sole dissenting voice is Tolliver (1994) whose view is that

¹⁴See p. xiv: “We can all agree that, at least typically, a red-feeling experience is red-representing, and conversely.” I believe that, on a broader historical analysis, the endorsement of this claim no longer appears so universal, and there are in fact many antecedents to the structural realist position, see for instance Köhler: “I cannot identify the final products, the things and events of my experience, with the physical objects from which the influences come. If a wound is not the gun which emitted the projectile, then the things which I have before me, which I see and feel, cannot be identical with the corresponding physical objects” (1947, 22). Most influential on my own thinking here has been Hermann von Helmholtz. In the interests of brevity, however, I shall set aside further historical discussion for a future venue.

[S]ensuous color properties are part of an internal code for the type-individuation of visual representations, i.e. color experience is part of a system of internal bookkeeping. Any content our color experiences have is best thought of as information content rather than representational content. (412)

Although I do not agree with the particulars of Tolliver’s view (for example, I endorse a different theory of informational content), it is basically consonant with structural realism. In particular, he also emphasizes the fundamental point that, although “the property revealed [by visual sensations] is not a property shared by external physical things”, nevertheless this does not imply “systematic error” nor, crucially, is it “necessary for maintaining a distinction between veridical and illusory color perceptions” (412).

If we allow ourselves to endorse the semantic and epistemic readings of the claim that physical objects are colored, why not go all the way and embrace the ontic reading as well? I hope I have demonstrated above that the structural realist position is the one which fits most closely with the facts of color science. Other philosophers who have looked closely at those facts, however, have drawn quite different conclusions. Where exactly do we differ? For the sake of specificity, I will focus here on one of the most prominent and well-developed positions to purport to follow from color science, the physicalism of Byrne and Hilbert (2003a), who identify colors with sets of metameric SSRs.¹⁵

4.2 Which Realism is Implied by Color Science?

I have argued that the structural realist position characterizes the presuppositions behind perceptual science as a research program, and is thus independent of any particular theory within that practice. Just as the analysis of a measurement procedure breaks down into three components: i) the measuring space; ii) the measured space; and iii) the process of calibration which links the two; so also the science of any type of perceptual experience seeks to characterize three things: i) the space of possible experiences; ii) the space of possible external correlates to experience; and iii) the process by which the two are linked (the analysis of which proceeds in both physiological and functional terms). By insisting that (i) and (ii) collapse, the physicalist erases a distinction crucial for perceptual science.

In fact, Teller (2003) levels exactly this criticism against Byrne and Hilbert (2003a):

Now, as far as I can see, color realism is the view that of the vision scientist’s three entities—surface spectral reflectance, neural signals, and perceived color—one *is* color, and the other two are not. But if you ask a color scientist which of the three entities *is* color, she will answer that the question is ill-posed. We need all three concepts, and we need a conceptual framework and a terminology that makes it easy to separate the three, so that we can talk about the mappings among them. Color physicalists can

¹⁵Modulo some additional nuances and qualifications which are irrelevant for the present discussion. The following section draws heavily on Wright, 2010, Section 5. Although I find Wright’s criticism of physicalist accounts very compelling, his own view (as he acknowledges) is not yet fully worked out, so I refrain from discussing that here (if anything, however, it falls closest to the ecological accounts discussed in Section 4.3).

call surface spectral reflectance *physical color* if they want to, although *surface spectral reflectance* is a more precise term. But to call it *color* (unmodified) is just confusing and counterproductive, because for us the physical properties of stimuli stand as only one of three coequal entities. (2003, 48)

In their response to Teller, Byrne and Hilbert (2003b) accuse her of ignoring the importance of intentionality for making sense of color sensation:

We conjecture that the reason **Teller** sees only a tedious squabble about words is that she fails to recognize fully the *intentionality*, or representational nature, of visual experience. . . . Once we have accounted for the “regularities” between external stimuli and color experiences, it is hard to see why there would be a further question about whether color experiences represent the world as it really is. (52)

Now I take it that the move from the claim that color experiences *represent* SSRs to the claim that colors *just are* SSRs depends crucially on the single property reading of **CR**. As argued above, the veridicality of color experience does not turn on the identification of colors as experienced with external properties. Likewise, it is perfectly coherent to admit that color experiences correlate with, or are caused by, SSRs (as I do and I’m sure Teller does as well) without making the further move of identifying color properties as we experience them with SSRs.

Wright (2010) helpfully diagnoses a key misunderstanding in this exchange by identifying Byrne and Hilbert’s implicit reliance on the thesis that the goal of color vision is to “estimate” or “recover” surface reflectance properties (27). If this thesis is correct, then color science appears to motivate the conflation of the color properties assigned by experience with the surface properties they aim to recover, typically construed as SSRs (28). After referencing MacLeod’s accusation of category mistake (see footnote 13 above), Wright himself levels several pragmatic criticisms against any single-minded focus on “finding physical counterparts. . . for perceptual qualities”:

[T]here are two main ways in which this mindset threatens to harm inquiry: it can lead to mischaracterizations of perceptual phenomena by trying to force upon them a vocabulary derived from physical theory or it might encourage mistaken attributions of features that are only present in experience to the stimulus. The former is Mausfeld’s (2002) “physicalist trap” and the latter is Köhler’s (1947) “experience error.” Relatedly, this outlook might set off a quest to discover (or stitch together) “natural kinds” that correspond to perceptual categories, but which turn out to be so fractured, vacuous, or *ad hoc* that they are of no aid to inquiries into the nature of perception. (28)¹⁶

If Wright is correct that Byrne and Hilbert take their view to follow from the particular (albeit popular) thesis that the goal of color perception is to recover SSR,¹⁷ then I think a stronger criticism

¹⁶This last sentence is clearly intended as a dig at the “unknowable” (21) and physical, but “uninteresting from the point of view of physics” (11) color categories of Byrne and Hilbert (2003a).

¹⁷Alternatively, they may be antecedently committed to collapsing semantic and epistemic questions about realism into metaphysical ones via their endorsement of the single property reading of **CR**. If this is the case, then the arguments above apply, and the correct interpretation of vision science is simply not at issue.

can be made against them than the pragmatic ones he levels here. In particular, physicalism does not follow from the practice of even those vision scientists who endorse this thesis, *a fortiori* it does not follow from vision science in general. To see this, let's examine the case of Brian Wandell.

Wandell (e.g. 1989) is one of the champions of an approach to color vision which interprets the goal of color perception as one of recovering or estimating SSR. Nevertheless, I believe an examination of his reasoned views on color (e.g. Wandell, 1995) does not support identifying color sensation with SSR. The reason is already present in Teller's quote: "We need all three concepts ... so that we can talk about the mappings among them." Only by keeping color as experienced and SSR conceptually and empirically distinct can Wandell marshal and evaluate evidence in favor of the thesis that color sensations recover SSR. In Wandell (1995), he puts great weight on precise discussions of such mappings, emphasizing both the value of searching for them and the importance of acknowledging discrepancies when they appear.

For instance, he spends several paragraphs (95–97) discussing a mapping which shows there are equivalent amounts of information in the retinal signal and in color experience as measured through color matching experiments. He emphasizes that this demonstration was only possible as the outcome of a prolonged process of "trying to recast our experiments using different methods until the relationships become evident" (96). The upshot is that something is learned, but its value depends upon sensitivity to the difficulties involved in getting there. His continued emphasis in later passages on the discrepancies between color experience as measured through psychophysical experiments and the neural processing of the color signal in the brain demonstrates that he strongly resists the conclusion that such mappings imply reductionism.

Similar considerations shape Wandell's discussion of the relationship between SSR and color experience. Now, we should admit at the outset that he does in some passages attribute colors directly to surfaces, although this appears to be a presupposition rather than a result, e.g. "If color must describe a property of an object, the nervous system must *interpret* the mosaic of photopigment absorptions and estimate something about the surface-reflectance function" (295). However, if we were to look for the endorsement of a particular realist thesis in Wandell, there is perhaps even greater support for structural realism: "The defining property of an object is not the absolute amount of light it reflects, but rather how much light it reflects relative to other objects" (289). In fact, ultimate determination of *relative* reflective values appears to be the goal towards which recovery of SSR is an initial step.

Nevertheless, I think it is wrong to put too much weight on these motivational passages. In order to understand Wandell's considered view, we should attend to what he says in those passages where he is doing color science. And when he considers precise data about surfaces he writes of SSRs, not colors *simpliciter*, and when he considers precise data about sensations, he writes of "color appearances", not colors *simpliciter*. The motivation here is exactly that emphasized by Teller: sensations and surfaces are conceptually distinct, their properties are measured using radically different methods, and it is only by maintaining the distinction between the two that evidence for any structural mapping between them can be evaluated precisely.

In the case of the SSR / color appearance comparison, Wandell first employs computational

models to investigate idealized recovery of SSR using a limited number of basis vectors (295–308). The key idea here is that if we assume color experience evolved in the context of surfaces illuminated by light from a particular illuminant (i.e. the sun), we can dramatically simplify the problem of recovering SSR from the combined $\text{SSR} \times \text{SPD}$ signal incident at the retina. He next evaluates separately the evidence for SSR recovery in color experience provided by asymmetrical matching experiments (309–315). By comparing the success of the visual system at assigning the same color appearance to surfaces with the same SSR across changes in context and illuminant, as measured against the idealized computational models, he can assess the evidence for the hypothesis that the goal of color vision is recovery of SSR. The conclusion is that “asymmetric color matches do not compensate completely for the illumination change” (314). There are two fundamental points here. First, Wandell’s estimation hypothesis is not a conclusion, it is a proposal which motivates a specific research program. Second, the evidence in favor of this hypothesis can only be stated and evaluated if one keeps color sensations and surface properties conceptually distinct. A practice which insists on distinguishing two concepts cannot provide evidence that they are metaphysically equivalent unless it is supplemented with an antecedent commitment to such a reduction.

To conclude: *if* one approaches color science with an antecedent commitment to the one property reading of **CR**, *then* the surface recovery / estimation thesis appears to provide support for physicalism. (Although there are also alternative theses in color science and, as Wright emphasizes, room for debate about the heuristic value of the estimation hypothesis.) Conversely, if one approaches color science without an antecedent commitment to either one or two property readings of **CR**, then even the practice of those who endorse the estimation thesis supports the two property reading. This blocks ontic realism about color properties, but given that there are many reasons to support semantic and epistemic realism, the natural considered view becomes structural realism.

4.3 Ecological Views

In this section, I briefly discuss two ecological theories of color which share many features with structural realism. The ecological approach takes color experiences to represent features of the environment of evolutionary importance, for instance those of functional significance to the organism (e.g. edibility, availability for mating, constituting a threat). These theories may be cashed out as instances of direct realism (Noë, 2004; Matthen, 2010) or of relationalism (Hatfield, 2003). It is perhaps unsurprising that ecological accounts should share so many similarities with structural realism given their origin in the work of J. J. Gibson, who emphasized many of the same features of perception as the present account. To list only one, Gibson discusses heat perception in much the same terms as it is discussed above, including even use of the term “calibration” to describe the effect of context in determining sensations of heat (1966, 131).

4.3.1 Mohan Matthen

Matthen (2005, 2010) defends a view with all the components of the one developed here but a different conclusion. He emphasizes the importance of the structural relationship between colors

for understanding their content, and he argues that successful denotation does not depend upon the denoted object possessing the attributed property. He motivates this claim with the example of a radar operator who successfully refers to a plane depicted on her screen by a red dot as “the red plane”, “although she does not mean or imply that the aircraft is red” (2010, 78). He calls this “projective” denotation, and summarizes his overall view in the thesis that “Color experiences constitute a structured projective denotational system” (2010, 79). Finally, he argues that this is a *semantic* theory of the relationship between color experience and the world, and should be understood as analogous to the relationship between a calibrated measurement scale and properties in the world (2005, 259).

So far, the basic ingredients of Matthen’s system appear very close to those of structural realism. Where we differ is again on the question of **CR**, and in particular whether one can safely identify colors as experienced with colors as they are in the world.

In a semantic theory, color experiences *denote* colors. Just as the word ‘cat’ denotes the property that cats share ... [color experience] is a *symbol* internal to the workings of the mind, a token by which the color-vision system passes to other epistemic faculties ... the message that ... it has determined the color of this visual object to be orange. (Matthen, 2010, 77)

Matthen takes his semantic theory to imply that external correlates, the “denotations” of color sensations, must be identified which stand in the same similarity relations as color sensations. To a first approximation, these colors are surface reflectance properties (2010, 75), but the similarity relation between them is defined in terms of the role they play as a “substrate for conditioning”, i.e. two color properties in the world are similar because they can be used to condition similar responses (2010, 82). A consequence of this definition is “pluralistic realism” about colors (2005, 200–9).

Frankly, I feel Matthen has not followed his own insights to their logical conclusion. I take the analogy with the radar screen to show precisely that it does not follow from the fact that a red experience is doing the representing that the property it represents is redness. Furthermore, this interpretation should be strengthened by the analogy with symbols. In fact, it is this very analogy which motivates Tolliver above (c.f. his talk of “internal code”) in his rejection of the one property reading of **CR**. Despite his claims to a semantic view, Matthen appears to have fallen into the trap of conflating the semantic question with the ontic one.

We can clarify the difference between Matthen’s view and structural realism by looking more closely at the relationship between “cat” and *cats* (the group of objects in the world). We put quotes around “cat” when we consider it *as a symbol* precisely because we know that its properties as a symbol are disjoint from the properties of the category it represents. For instance “cat” is orthographically similar to “car”, yet we do not take that to have any implications for the similarities (or not) between *cats* and *cars* as categories in the world. If we are going to take color experiences as analogous to symbols we should take that analogy seriously, and this means acknowledging that the properties of colors as experienced may be disjoint from the properties they represent. A consequence is that we should deny the one property reading of **CR**.

Now, it may be that Matthen does reject the one property reading of **CR**, he simply considers it helpful shorthand to use the same term for both the symbol and that which it denotes. We safely use *cats* to stand for the category denoted by “cat” because that category stays relatively fixed across contexts. As discussed above, however, we do not find this fixity in the measurement relation between color experience and the world. For Matthen, the calibration of color experience is a one time evolutionary event, establishing a fixed relationship with surface properties. Like the “cat” / *cats* relationship, the red experience / red surface property relationship is fixed. I have argued above, however, that the calibration of color perception changes regularly with context. A better semantic analogy here would be with a comparative like “tall”. We don’t identify “tall” things with some category *tall* in the world, precisely because the veridicality of attributions of tallness changes with context. Just as a mouse may be veridically assessed as tall standing next to his brothers, but veridically assessed as short next to an elephant, the very same surface patch may be veridically assessed as white in one context and black in another. And this context dependence of perceptual calibration defeats not only the one property reading of **CR**, but also the assumption of a denotation relationship between particular *color*₁ sensations and fixed *color*₂ categories.

4.3.2 Gary Hatfield

Hatfield (2003, 2007) defends an ecological theory of color which is both relationalist and objectivist. He rejects the analysis of color content in terms of SSR because of the problem of metamers (2007, 141). Without a straightforward physical analysis of representational content, it appears that “[t]he existence of color as an attribute of objects depends on the normal effects of objects on perceiving subjects” (2003, 195). Since this view depends on reference to “perceiving subjects” in the definition of color content, it is relational. Nevertheless, Hatfield’s view is still objective because color categories “sustain factual claims” and “pertain to publicly available states of affairs” (2003, 199).

A crucial difference between Hatfield and other ecological views is in his analysis of the function of color vision. On the ecological approach, a functional analysis is necessary to understand the content of a representational state, and evidence for this analysis is to be found in the evolutionary history of the organism. Where Noë (2004) and Matthen (2005) cash color function out in terms of elaborate counterfactual patterns of expectation and conditioning, respectively, Hatfield sticks to the straightforward view that the function of color vision is simply discrimination:

The functions of color vision ...are served merely if color vision enables us to better discriminate some objects from other objects, and enables us to reidentify them as those objects when we encounter them again. (2007, 146)

This view, combined with a distinction between phenomenal color experience and its representational content, motivates the conclusion that the information about the world provided by color vision is purely relational.¹⁸

¹⁸i.e. relations between color experiences tell us something about relations between surface properties—again, not to be confused with the view that color is a relational property, *also* endorsed by Hatfield, c.f. footnote 2.

Beyond the implication that, with conditions held constant, surfaces that look different chromatically are different in some way, color qualia of themselves don't contain further content about the properties of surfaces. (2007, 145)

Hatfield's conclusion as stated here is quite close to that of the structural realist. In particular, the position that color experiences merely represent the distinctness of surfaces is consistent with the view that color experience constitutes a nominal scale. As argued above, there is some evidence that the representational structure of the color solid includes more than mere category difference, yet certainly category difference constitutes the weakest interpretation of color experience as measurement. Arguably, Hatfield also endorses the two property reading of **CR**. Certainly, he is careful to distinguish "color as a property of objects" from "color experience" (2007, 135); an interpretation which is further supported by his endorsement of realism about color qualia (133).

So, the basic ingredients of Hatfield's view and structural realism are essentially the same. The main difference is in his relationalism, defining the external correlates of color in terms of their dispositions to produce color sensations in (human) observers. Here is where structural realism can give us traction, however. On both views, the veridicality of color attributions does not depend upon experience assigning the same color to the same surface property in every context. However, once we countenance the artifactual / representational structure distinction, the phenomenon of metameric surfaces no longer defeats any particular analysis of the external properties measured by color experience. Consequently, the structural realist can maintain that these external correlates exist independent of human observers, and are thus not dispositional properties, while agreeing with Hatfield on the need to distinguish them both conceptually and metaphysically from color as experienced.

5 Other Secondary Qualities

In order to demonstrate how structural realism applies to other secondary qualities, I briefly discuss the examples of pitch and odor. The example of pitch perception in a musical context shows that calibration not only establishes a baseline correspondence between experience and the world, it can also change the structure of a perceptual measuring space. The example of odor perception demonstrates how structural realism can illuminate research practices on perceptual topics as yet poorly understood. These examples provide additional support for structural realism by demonstrating its expressive adequacy for explaining the broad diversity of perceptual research on secondary qualities.

5.1 Pitch

We perceive pitch differently in musical and non-musical contexts. More specifically, our assessments of sameness of pitch and of distances between pitches change across these contexts. This empirical phenomenon can be understood on the structural realist view by enriching the concept of calibration. Before, we took calibration to merely establish a baseline correspondence between

a measuring structure and a measured structure. On the richer view, we can take calibration to also “choose” amongst various possible measuring structures. This example is closely analogous to the problem which has motivated structural realism in philosophy of science (multiple theories for a single phenomenon in the world).

The physical property we typically take pitch perception to measure is frequency of vibrations in the air. Both pitch and frequency may be linearly ordered and this ordering constitutes part of the representational structure of pitch sensation. The presence (or not) of additional representational structure in the measuring space, however, is determined by contextual calibration.

If pure tones (sine waves) are used as stimuli and presented to the subject in a random order, we can determine the ratio between the just noticeable difference between two stimuli and the absolute value of a comparison stimulus. This quantity is called the “Weber fraction,” and its measurement is a typical practice in psychophysical research. As with other sensory modalities, the ability to discriminate between frequencies varies with the comparison frequency. While some sensory modalities exhibit ranges of relative fixity in this relationship (i.e. the Weber fraction remains stable), the Weber fraction for pitch varies dramatically across the range of audible frequencies.¹⁹

But this data has a puzzling implication. It is common practice in psychophysics to interpret just noticeable differences as psychologically equal units. This assumption, combined with the data on frequency discrimination, implies that frequencies which are equivalent distances apart from a physical standpoint are perceived in experience as different distances apart. Now, in the context of other sensory modalities (e.g. heat perception), this observation would not be very significant. All it would demonstrate is that the ordering of sensations, but not the distances between them, constitutes representational structure. We veridically perceive the ordering of temperatures, but not distances between temperatures, via our sensations of heat.

This result is puzzling in the case of pitch perception because the equivalent physical distances at issue can in fact be identified as equivalent perceptually. The crucial example here is the octave. If x and y are frequencies an octave apart, then $x = 2y$ (or vice versa). Octaves are important for explaining physical phenomena like sympathetic resonance. They also form the basis for most melodic forms of human music.²⁰ Furthermore, frequencies separated by an octave are easily identified as in some sense “the same” by typical human observers, and the distances between different octave pairs as equivalent. But this well known fact is inconsistent with the psychophysical data summarized by variation in the Weber fraction. For example, if measured in units of just noticeable difference, the perceptual distance between c' and c'' is significantly greater than that between C and c , because a greater number of distinct changes in stimulus frequency can be discriminated between c' and c'' . Nevertheless, in a musical context, we judge the distance between C and c and that between c' and c'' as the same.

These considerations motivated the development of psychophysical techniques for investigating pitch perception in a musical context. For example, rather than presenting sine waves in a random

¹⁹For a recent summary of data on the Weber fraction for pitch see Moore, 2008, 196–204.

²⁰The contrast here is with purely rhythmic forms of music. Although cultures differ in the number of notes and the size of intervals they identify within an octave, they always assign some musical significance to the octave interval itself. (I am omitting here discussion of some subtle questions about how precisely human octave assessments match physical octaves.)

Figure 1: Shepard's toroidal pitch space. The closeness of fifths is not accurately depicted in this image. Since the torus is defined as the cartesian product of two circles of equal size (the circle of fifths (greatly shrunk in this representation) and the chromatic circle), distances can only be accurately recovered if it is considered embedded in four dimensional space. Furthermore, this space is crucially different from the familiar color solid. In particular, it is not a continuous space, since there is no significance to the positions between points around the circle of fifths (nor, in the Western musical tradition, to the positions between steps in the chromatic scale); rather, it is best understood as a discrete graph topologically connected in the shape of torus, and in which equivalent distances are measured by equivalent numbers of discrete steps.²³ This is Figure 5 from Shepard, 1982; see his discussion for full details.

order, one might first play a short musical passage to the subject, then present her with stimuli for comparison. Some of the crucial developments in this project were due to Roger Shepard, including both methods (e.g. the “Shepard tone,” developed as a stimulus for these studies) and theory. On the theoretical side, a major contribution was Shepard's proposal that musical pitch space should be represented by a torus (Figure 1). The toroidal representation incorporates several facts about perception of pitch within a musical context: i) frequencies separated by an octave are perceived as “the same”; ii) frequencies separated by a fifth are perceived as “close together”; iii) steps in a diatonic scale are perceived as equivalent in distance (even though some are whole steps and some are half steps) (Shepard, 1982).²¹

Return again to the realism debate and the status of **CR**. When the one property reading of **CR** is endorsed, realists have felt obligated to find external correlates for color similarities, while antirealists have used their failure to do so as an argument for eliminativism. But how would this debate look transposed into an auditory context? There is not one, but (at least!) two different measuring spaces for pitch perception, each with its own distinct similarity structure. For which

²¹Of course, I elide many details here. To mention just one example, similarity judgments between stimuli which are not “pure” sine waves, but more complex waves exhibiting harmonics, have a physical basis in the agreement of the harmonics across different base frequencies. This consideration provides a physical basis for the assessment of fifths as “similar.”

²³I am indebted to Dmitri Tymoczko for helpful discussion of this example.

of these two spaces should the realist struggle to find external correlates?

By adopting a two property reading of **CR** (for pitch) and employing the artifactual / representational structure distinction, the structural realist has no trouble analyzing this example. Measurement of frequency by pitch sensations is simply calibrated differently in musical and non-musical contexts. Musical calibration establishes a mapping between frequencies and the musical pitch torus, whereas the default calibration for auditory perception merely establishes a mapping from frequency space into a linear pitch space. The two measuring spaces have different structures, the status of which as artifactual or representational can be evaluated separately in each case. There is no meaningful question of whether frequencies separated by an octave are similar *simpliciter*. Rather, there are some physical features they share and some they do not, and the map into musical pitch space veridically represents (some of) the former, while the map into nonmusical pitch space veridically represents (some of) the latter.

5.2 Olfaction

Compared to color and pitch, odor perception is extremely complex and relatively poorly understood. The kind of question about the ontological status of apparent similarities which has dominated the color realism debate cannot even be asked in the context of odor perception. In fact, there is no systematic story about which similarity relations obtain between smells, or even the appropriate terminology for characterizing smell space. Nevertheless, odor science provides support for structural realism. This support follows from the questions asked in olfactory research. These questions correspond to the three fundamental components we have repeatedly emphasized are required for the measurement analogy, and thus for the structural realist account of perception: i) what is the structure of possible odor experience? (the measuring space); ii) what is the structure of the external correlates of smell? (the measured space); and iii) what is the physiological / functional processing which links the odor signal incident at the olfactory bulb to the neural correlates of smell experience? (the calibration procedure). Whereas in the case of color, early progress on (i) and (ii) drove predictions about (iii), the complexity of odor perception frustrated efforts on all three tasks until relatively recently.

Odor is classified with taste as a “chemical sense” since the perceptual response is driven (somehow) by the microstructure of molecules, in the case of odor, those that are volatile. Despite attempts at a categorization of odors and the search for a molecular basis for them in the late 19th and early 20th centuries, in 1942 Boring could assert of the failure to confirm any systematic relationship that “[t]he failure to make the analysis is simply a phase of the failure to make the crucial discovery about smell, to find the essential nature of its stimulus” (449). The simple knowledge that odor receptors respond to molecular structure is not enough to understand the “essential nature” of what is measured because it does not provide a characterization of the systematic variation in molecular structure which drives systematic differences in odor experience, a characterization which is needed to extract quantitative conclusions about odor space from psychophysical methods.

In recent years, much more progress has been made on understanding the nature of the odor receptors and the features of molecules with which they interact, due largely to the development

of increasingly sophisticated techniques for controlling and analyzing both the molecular structure of stimuli and the pattern of neurophysiological responses. For example, Zhao et al. (1998) used an adenovirus to increase expression of a particular olfactory receptor in rat nasal cavities. They found a single compound amongst fifty tested which increased neural firing in the infected tissue, motivating the conclusion that the specific receptor expressed detects that compound. Experiments such as this one have increased understanding of the physiological basis of the odor signal, but without better characterizations of the measuring and measured spaces, they cannot yet give an adequate account of the calibration of odor perception.

A significant advance in the characterization of the measuring space was made by Henning in the early 20th century. He asked subjects to order odor samples with respect to degree of similarity. These experiments motivated his proposal of the Henning odor prism in 1915, the surface of which was offered as an analysis of the space of possible odor experiences (Boring, 1942, 445). Immediate attempts to experimentally confirm the prism concluded that, although its gross features could be recovered, the prism was not a sufficiently close approximation to smell space to support quantitative measurements of distance such as those which had been made in the color solid. For example, Macdonald (1922) both found stimuli which produced sensations which could not be located within the structure of Henning’s prism and failed to find stimuli which could fill out specific regions of it, concluding that “the solution may lie in some other geometrical construction” (551).

Another tradition has attempted to map smell space by asking subjects to describe odor stimuli with a number of smell-related adjectives. Boring discusses the late 19th century efforts along these lines by Dutch physiologist Zwaardemaker (441–4). A mid-century summary of efforts in this direction can be found in Harper et al. (1968) and a large set of data was collected by Dravnieks (1985). Data sets such as these provide characterizations of smell space with a dimension for each adjective used in the experiment (Dravnieks, for example, used 146 adjectives, giving a 146 dimensional smell space). But which of these adjectives describe the psychological dimensions of smell space and which do not? Or what if none of them do? One approach to a high dimensional data set such as this is to use a mathematical technique such as principal components analysis to find a space of reduced dimensionality that preserves relative distances between points in the space. Koulakov et al. (2011) performed such an analysis on the Dravnieks data set, discovering that distances can be preserved on a 2 dimensional curved “potato chip” shaped surface embedded in three space. Although this analysis implies that, despite the large number of receptors, smell space need not be high dimensional, Koulakov has as of yet failed to find any psychologically significant characterization of the dimensions of this reduced space.

The scientific study of olfaction still has a long way to go before it reaches the maturity of color science. Nevertheless, the basic conceptual distinctions which we saw in color science have analogs in the study of odor perception. Since these conceptual distinctions motivate and support the analogy with measurement, olfactory science, even in its present, nascent, state, also supports structural realism.

6 Conclusion

What is the relationship between the world *as we experience it* and the world *as it is*? I have argued for *structural realism*, the claim that the structure of our possible experiences corresponds to the structure of possible ways the world can be. Since this structural correspondence between experience and the world is calibrated differently across different contexts, however, we cannot directly identify particular experiences with particular properties in the world. We cannot identify red with a particular surface reflectance property, warmth with a particular temperature range, pungency with a particular molecular shape, etc. This is why this realism is *structural*: it is not committed to a metaphysical reduction of the properties of the world as experienced to the properties of the world as it is.

Nevertheless, structural realism is still *realism* since the preservation of *relations between properties* across the correspondence between experience and the world ensures that property attributions are in general veridical (semantic realism) and demonstrate knowledge (epistemic realism), so long as they are assessed against a contextually established calibration baseline. The dissociation of ontic from epistemic and semantic realism is motivated by an analogy with measurement. Quantities in the world are not themselves numbers, yet we can use numbers to represent them once we establish a structural correspondence between the real line and possible values of a quantity through an act of calibration. The three basic components of measurement (measuring space, measured space, calibration process) shape the scientific investigation of the perception of secondary qualities. To demonstrate this, we've surveyed the examples of warmth, color, pitch, and odor. Consequently, structural realism is the epistemological analysis of the status of secondary qualities most strongly supported by scientific practice.

Acknowledgements

This paper is dedicated to Patrick Suppes and John Perry, the former for the carrot, the latter for the stick. It has benefitted greatly from the helpful comments of Johanna Wolff, Gary Hatfield, and two anonymous reviewers. This research was supported by NSF grant no. 1028130.

References

- Boring, E. G. (1942). *Sensation and Perception in the History of Experimental Psychology*. Appleton - Century - Crofts, Inc., New York.
- Byrne, A. and Hilbert, D. (2003a). Color realism and color science. *Behavioral and Brain Sciences*, 26:3–21.
- Byrne, A. and Hilbert, D. (2003b). Color realism redux. *Behavioral and Brain Sciences*, 26:52–63.
- Byrne, A. and Hilbert, D. R. (1997). Introduction. In Byrne, A. and Hilbert, D. R., editors, *Readings on Color*, volume 1: The Philosophy of Color. MIT Press, Cambridge, MA.

- Campbell, J. (1993). A simple view of colour. In Haldane, J. and Wright, C., editors, *Reality, Representation, and Projection*, pages 257–268. Oxford UP, New York, NY.
- Cartlidge, E. (2012). Loose cable may unravel faster-than-light result. *Science*, 335(6072):1027–8.
- Chang, H. (2004). *Inventing Temperature: Measurement and Scientific Progress*. Oxford UP, New York, NY.
- Churchland, P. (2007). On the reality (and diversity) of objective colors: How color-qualia space is a map of reflectance-profile space. *Philosophy of Science*, 74:119–149.
- Dravnieks, A. (1985). *Atlas of Odor Character Profiles*. ASTM, Philadelphia, PA.
- Fairchild, M. D. (2005). *Color Appearance Models*. John Wiley & Sons, Ltd., Chichester, West Sussex, 2nd edition.
- Frigerio, A., Giordani, A., and Mari, L. (2010). Outline of a general model of measurement. *Synthese*, 175:123–149.
- Gibson, J. J. (1966). *The Senses Considered as Perceptual Systems*. Houghton Mifflin Co., Boston, MA.
- Hardin, C. L. (1988). *Color for Philosophers: Unweaving the Rainbow*. Hackett, Indianapolis.
- Harper, R., Bate Smith, E. C., and Land, D. G. (1968). *Odour Description and Odour Classification*. J. & A. Churchill, London.
- Hatfield, G. (2003). Objectivity and subjectivity revisited: Color as a psychobiological property. In Mausfeld, R. and Heyer, D., editors, *Colour Perception: Mind and the Physical World*, pages 187–202. Oxford UP, Oxford.
- Hatfield, G. (2007). The reality of qualia. *Erkenntnis*, 66:133–168.
- Isaac, A. M. C. (2013). Quantifying the subjective: Psychophysics and the geometry of color. *Philosophical Psychology*, 26(2):207–233.
- Johnston, M. (1992). How to speak of the colors. *Philosophical Studies*, 68:221–263.
- Köhler, W. (1947). *Gestalt Psychology*. Liveright, New York.
- Koulakov, A. A., Kolterman, B. E., Enikolopov, A. G., and Rinberg, D. (2011). In search of the structure of human olfactory space. *Frontiers in Systems Neuroscience*, 5(65):1–8.
- Krantz, D. H., Luce, R. D., Suppes, P., and Tversky, A. (1971). *Foundations of Measurement*, volume 1. Dover, Mineola, NY.
- Kuehni, R. G. and Hardin, C. L. (2010). Churchland’s metamers. *British Journal of Philosophy of Science*, 61:81–92.

- Kuehni, R. G. and Schwarz, A. (2008). *Color Ordered: A Survey of Color Order Systems from Antiquity to the Present*. Oxford UP, New York, NY.
- Ladyman, J. (1998). What is structural realism? *Studies in History and Philosophy of Science*, 29(3):409–424.
- Luce, R. D., Krantz, D. H., Suppes, P., and Tversky, A. (1990). *Foundations of Measurement*, volume 3. Dover, Mineola, NY.
- Macdonald, M. K. (1922). An experimental study of Henning’s system of olfactory qualities. *The American Journal of Psychology*, 33:535–553.
- MacLeod, D. (2003). Phenomenology and mechanism. In Mausfeld, R. and Heyer, D., editors, *Colour Perception: Mind and the Physical World*, pages 431–434. Oxford UP, Oxford.
- Maloney, L. T. (1986). Evaluation of linear models of surface spectral reflectance with small numbers of parameters. *Journal of the Optical Society of America*, 3(10):1673–1683.
- Matthen, M. (2005). *Seeing, Doing, and Knowing*. Clarendon Press, Oxford.
- Matthen, M. (2010). Color experience: A semantic theory. In Cohen, J. and Matthen, M., editors, *Color Ontology and Color Science*, pages 67–89. MIT Press, Cambridge, MA.
- Mausfeld, R. (2002). The physicalist trap in perception theory. In Heyer, D. and Mausfeld, R., editors, *Perception in the Physical World*, pages 75–112. John Wiley & Sons, West Sussex.
- Mausfeld, R. (2003). ‘Colour’ as part of the format of different perceptual primitives: The dual coding of colour. In Mausfeld, R. and Heyer, D., editors, *Colour Perception: Mind and the Physical World*, pages 381–430. Oxford UP, Oxford.
- Moore, B. C. J. (2008). *An Introduction to the Psychology of Hearing*. Emerald, Bingley, UK, fifth edition.
- Müller, I. and Weiss, W. (2005). *Entropy and Energy*. Springer, Berlin.
- Nassau, K. (2001). *The Physics and Chemistry of Color*. John Wiley & Sons, Ltd., New York, NY, 2nd edition.
- Noë, A. (2004). *Action in Perception*. MIT Press, Cambridge, MA.
- Psillos, S. (1999). *Scientific Realism: How Science Tracks Truth*. Routledge, New York.
- Shepard, R. N. (1982). Geometrical approximations to the structure of musical pitch. *Psychological Review*, 89:305–333.
- Shepard, R. N. (1992). The perceptual organization of colors: An adaptation to regularities of the terrestrial world? In Barkow, J. H., Cosmides, L., and Tooby, J., editors, *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*, pages 492–532. Oxford UP, New York, NY.

- Smith, G. E. (2001). J. J. Thomson and the electron, 1897–1899. In Buchwald, J. Z. and Warwick, A., editors, *Histories of the Electron*, pages 21–76. MIT Press, Cambridge, MA.
- Stevens, S. S. (1946). On the theory of scales of measurement. *Science*, 103:677–680.
- Teller, D. Y. (2003). Color: A vision scientist’s perspective. *Behavioral and Brain Sciences*, 26:48–49.
- Tolliver, J. T. (1994). Interior colors. *Philosophical Topics*, 22:411–441.
- Tye, M. (2006). The puzzle of true blue. *Analysis*, 66.3:173–178.
- van Fraassen, B. (2008). *Scientific Representation*. Oxford UP, New York, NY.
- Wandell, B. A. (1989). Color constancy and the natural image. *Physica Scripta*, 39:187–192.
- Wandell, B. A. (1995). *Foundations of Vision*. Sinauer, Sunderland, MA.
- Worrall, J. (1989). Structural realism: The best of both worlds. *Dialectica*, 43(1–2):99–124.
- Wright, W. (2010). Perception, color, and realism. *Erkenntnis*, 73:19–40.
- Zhao, H., Ivic, L., Otaki, J. M., Hashimoto, M., Mikoshiba, K., and Firestein, S. (1998). Functional expression of a mammalian odorant receptor. *Science*, 279:237–242.