ALESSANDRA BUCCELLA
University of Pittsburgh
alb319@pitt.edu

NATURALIZING QUALIA

abstract

Hill (2014) argues that perceptual qualia, i.e. the ways in which things look from a viewpoint, are physical properties of objects. They are relational in nature, that is, they are functions of objects' intrinsic properties, viewpoints, and observers. Hill also claims that his kind of representationalism is the only view capable of "naturalizing qualia". After discussing a worry with Hill's account, I put forward an alternative, which is just as "naturalization-friendly". I build upon Chirimuuta's color adverbialism (2015), and I argue that we would better serve the "naturalizing project" if we abandoned representationalism and preferred a broadly adverbialist view of perceptual qualia.

keywords

qualia, naturalism, adverbialism, perceptual experience, appearance, reality

Phenomenology and Mind, n. 12 - 2017, pp. 152-161
DOI: 10.13128/Phe_Mi-21114
Web: www.fupress.net/index.php/pam

ISSN 2280-7853 (print) - ISSN 2239-4028 (on line)

1. Introduction

Perceptual experiences are the filter through which the world reveals itself to us. All we *know* about the world, all we can *do* in the world, how it *feels* to be in the world: it's all informed by what experiences "tell us".

Philosophical theories aim at giving an account of what perceptual experiences are, and what their relationship with the world out there is. In particular, one theory seems to provide a quite elegant explanation of the relationship between our experiences of the world and the world itself: representationalism. The main idea is simple: perceptual experiences have *contents* which *represent* objects as having certain properties, which may or may not be the properties objects really have. Representationalists (e.g. Dretske, 1988; Fodor, 1987) think that, in order to *explain behavior,* we must understand perception as "organizing" the world for us by representing it as being in a certain way. The idea is the following. It is intuitively plausible to say that we act on the basis of how we experience the world to be, and that sometimes we fail at what we do precisely because we take the world to be in a way it is not. Representationalists then argue that this fact can only be explained if we postulate the existence of something that "mediates" between us and the environment as it really is, namely the representational content of perceptual experiences.

Representationalism has been found particularly attractive not only by philosophers who are interested in explaining perception-guided behavior, but also by those committed to an attempt to "naturalize" perceptual experience itself and conscious experience more generally.[1] Although a proper definition of 'naturalism' in philosophy of mind is hard to provide,[2] I take both Christopher Hill (my main target in this paper) and myself to share the following definition. To naturalize experience is (1) to explain the fact that we are consciously "in touch" with the world by means of our perceptual (and cognitive) capacities in a way that is maximally coherent with neurosciences, psychology, and other empirical sciences of the mind, and (2) to avoid commitments to entities which don't clearly belong to the physical world. For many naturalists who accept the rough definition above, (accurately/veridically) representing the world is simply the evolved function of the perceptual system, and therefore

1 E.g. Dretske (1995), Tye (1995, 2000), Dennett (1990, 1991), Burge (2010), Hill (2014, chs. 11-12).

2 Starting from the "early" naturalists such as John Dewey, or Roy Wood Sellars, authors who call themselves naturalists disagree with respect to how strong the commitment to coherence with the science should be. Moreover, the distinction between methodological and metaphysical naturalists is often overlooked and not sufficiently spelled out by naturalists themselves.

representationalism is the most straightforward route to the naturalization of experience.[3] On the other hand, philosophers who are less committed to the so-called "naturalizing project" have one main worry. Nagel (1974) puts the point in the following way. Usually, explaining something in naturalistic terms means explaining it in a way that is public, shareable, and accessible by many people. This entails describing the phenomenon we want to explain in increasingly objective terms, so that the description can "go beyond" our idiosyncratic experience of the phenomenon. However, it is hard for a naturalist to capture *every* component of conscious experience – and perceptual experience in particular – in such objective terms. According to Nagel and those who followed his lead,[4] conscious experience has an intrinsically subjective component.

Traditionally, such an intrinsically subjective aspect of conscious experiences (not limited to the perceptual case) is called *phenomenal character*, and it is captured intuitively by the idea of "what-it-is-like" to have a certain experience. Take the example of having a dull pain in your ankle.[5] You want to describe the experience to someone else, perhaps a neuroscientist who studies pain experiences. The scientist may have put electrodes all around your skull, and may be measuring blood pressure, heartbeat, and other stress indicators in order to understand how your body is reacting to pain. Additionally, you may try really hard and describe how it feels like to experience pain by using terms and phrases in a shared language as precisely as you can. Together, all these elements will help the scientist reach a naturalized explanation of pain experiences, that is, an explanation that uses objective and public descriptions in order to explain what pain is. However, Nagel would point out that neither words and phrases in a shared language, nor physiological indicators, seem to be the kind of things that can "disclose" *what this particular pain-experience is like for you*, from a first-person perspective.[6] The notion of a *quale* (plural: *qualia*) can be helpful in order to systematize this idea of "what-it-is-likeness", and characterize more precisely the *explanandum* for which naturalists should find a place in their theories of conscious experience.

The subjective component of perceptual experiences is constituted by *perceptual qualia*. Perceptual qualia are defined by Kim (2006) as "the ways that things look, seem, and appear to conscious observers" (p. 225). The particular brown-ness of the table I am sitting at, the particular softness of my cat's fur when I touch it, the particular sharpness of her meowing, etc. These all count as perceptual qualia: they make up the intrinsically qualitative aspect of *my* perceptual experiences.

Hill (2014, chs. 11-12) takes on board the challenge of giving a plausible account of perceptual qualia in a naturalistic framework, or, in his own words, "to bring qualia into the physical fold" (p. 197).[7] In fact, Hill thinks that representationalism is the *only* theory that can succeed in the hard enterprise of naturalizing qualia,[8] and he tries to articulate how exactly such a

3 See, for example, Burge (2010).

4 Apart from Nagel (1974), people who have highlighted the presence of an intrinsically subjective component of conscious experience are, among others, Jackson (1982, 1986) and Chalmers (1996, 2007).

5 An analogous example was presented by Robert Howton in a lecture. I thank him for putting things this way, which I found very helpful.

6 For a more thorough explanation of the difference between first-person and third-person descriptions of conscious experience, see Chalmers (2001).

7 Other representationalists who chose the "naturalizing route", though in a less explicit way, are Shoemaker (1975, 1981) and Kriegel (2002).

8 "As far as I can tell, there is only one theory of experience that provides a satisfactory way of dealing with this problem [i.e. the problem of placing qualia in the physical world] – representationalism" (2014, p. 201). Hill quickly dismisses naïve realist-type theories in the light of the "mysteriousness" of the notion of direct acquaintance, especially in the context of an attempt to explain qualia and justify something like a genuine appearance-reality

naturalization of qualia could go.
In this paper, I shall motivate my disagreement with Hill. While I grant that Hill does succeed in giving a naturalistic account of qualia, my disagreement regards the claim that his view provides the *only* good naturalistic account of qualia. In particular, I will argue that a kind of *adverbialist* account can do equally well. By questioning the very categories employed in the traditional debate about qualia, the kind of adverbialist[9] position I articulate opens a different, promising path for the naturalizing project while avoiding a problematic commitment of Hill's view. I will first sketch the main points of Hill's view, by at the same time emphasizing what I take to be its most controversial commitment. Secondly, I will sketch an adverbialist solution which avoids Hill's counter-intuitive commitment regarding the nature of qualia while at the same time being "naturalization-friendly".

2. Perceptual Qualia: What Are They?

Earlier on I have given the definition of qualia endorsed by Hill. He understands the notion of *quale* as the notion of *look,* or *appearance*, in a *phenomenal* sense. This notion is spelled out by Hill by collapsing it into the notion of "comparative" look:[10]

> When one says that an object looks small to an observer, using "looks small" in this phenomenological sense, one is not claiming that the observer's perceptual experience supports the judgment that the object really is small. One is not saying that the observer's experience represents the object as small. Rather one is drawing an analogy between the observer's current visual experience and the visual experiences he has when is viewing objects that are reasonably close at hand and really are small (2014, p. 198).

Take the *particular smallness* that we experience while observing a truck on the street very far away from us. The truck, even though it is big, *phenomenally looks small* (or looks$_p$ small) to me. There is a sense in which this phenomenal smallness is the same phenomenal smallness that other objects look as having. For instance, a miniature toy truck looks$_p$ small as well, with the only difference that the latter *is* small, whereas the former *is* big. The objective sizes of the real truck and of the toy truck are very different; however, when the real truck is far away and the toy truck is in our hands, they both look small to us *in the same way.*
If this is plausible enough, however, naturalistic accounts of perceptual experience face what we may call the "problem of qualia": how to explain in a naturalistic framework the fact that the real truck and the toy truck, despite having completely different objective sizes, *look* of the same size *to me* in particular circumstances? *Where* is the smallness the two trucks seem to share? Is it in our head? Is it in the world? Is it nowhere at all? Hill has a proposal.

distinction in a naturalistic framework. He also dismisses the group of so-called doxastic theories, which hold that awareness of qualia depends on the subject having some sort of doxastic attitude (e.g. a belief or a judgment) towards the perceived scene. Against doxastic views, Hill appeals to phenomenology, and claims that doxastic views don't accommodate all the properties of experiential awareness (2014, pp. 207-209).

9 In fact, Hill provides an argument for rejecting adverbialism as a viable alternative in the "naturalizing qualia" project (2014, p. 208). I will briefly address the objection after the form of adverbialism I favor has been introduced.

10 The notion of phenomenal/comparative look is meant to be kept distinct from the epistemic notion of look, where an observer is actually disposed to judge that an object is small and uses the "looks"-claim as evidence for that judgment.

3. Hill on Qualia and The Appearance/ Reality Distinction

Hill considers qualia as the ways things *phenomenally* look to observers. In particular, since he wants to "locate" qualia in the physical world, he suggests that qualia are a special kind of property: they are phenomenal properties objects have in virtue of how they appear to a certain observer, at a certain time, from a certain angle. The small quale that I experience when I look at (a) the real truck from a distance and (b) the toy truck in my hand are nothing but two distinct physical properties, each belonging to one of the two objects. Both are properties of looking small to an observer in circumstances C(a) and C(b). These are *physical* properties because they can be seen as functions with multiple variables, all of which can be assigned precise values empirically. Hill (2014) identifies phenomenal properties with what he calls *Thouless properties:*[11]

> What then is the general character of a Thouless property? An example is the property P that an external object possesses just in case it is producing a retinal image with a height H such that the result of applying f1, . . . , fn to H and other relevant quantities has the numerical value N. P is a Thouless size. Alternatively, a Thouless size can be seen as a property P* that an external object possesses just in case it is subtending a visual angle V such that the result of applying the computable functions f1*, . . . , fn* to V and other relevant quantities has the numerical value N*. P and P* are both legitimate physical properties of external objects, though they aren't properties that a physicist would want to mention in giving an inventory of the characteristics that play a role in physical laws (2014, pp. 233-234).

Qualia, therefore, are characterized by Hill as physical properties of objects. The representationalist theory says that they are represented by the perceptual system in a certain way, just like every other property. Additionally, representationalism claims that how a phenomenal property appears to an observer constitutively depends on the way it is represented. Hill is a representationalist, thus he commits himself to the idea that there is a distinction to make regarding qualia: there are qualia-as-they-appear, that is, *representations of* Thouless properties, and qualia-as-they-really-are, that is, Thouless properties themselves. What we experience when an object $looks_P$ small to us is how the quale is represented, that is, the quale-as-it-appears.

This distinction allows Hill to claim, on one hand, that qualia belong to the physical world, and, on the other hand, that there is a sense in which it is legitimate to say that a truck seen at a distance looks small just in the same way as a miniature toy truck looks small in our hands. The two objects do not share the same Thouless property: we could say that the real truck at a distance *looks* $small_A$, while the toy truck *looks* $small_B$. However, the two Thouless properties are *represented* by our visual system *in the same way*.

The move of distinguishing qualia-as-they-are from qualia-as-they-appear seems to succeed in accommodating qualia in a naturalistic account of perception: qualia are physical properties, genuinely part of the natural world. A real truck and a toy truck do not share the same apparent size, since their apparent sizes are the results of different Thouless functions with different values assigned to the variables. Yet, our visual system represents the two Thouless properties in the same way. In other words, the real truck and the toy truck only *look as if they looked of the same size*. The sense in which that particular smallness is shared by the two appearances is due to the way the two appearances appear.

11 Robert Thouless was a Scottish psychophysicist who in the 1930s modeled the perception of apparent sizes as a "compromise" between objective sizes and angular sizes. His idea became part of the "mainstream" thanks to the work of the famous vision scientist David Marr in the 1970s (see Marr, 1982).

One does not need to be a hard-core representationalist, neither does she have to be a philosopher, in order to see virtues in Hill's account. Hill seems to correctly capture the way in which we commonly characterize the phenomenology of our visual experiences. Take his description of the famous drawing by Mach (figure 1):

> I am in my study, sitting in a chair. [...] my hand looks much larger than the chimney; it looks slightly larger than the books; and it looks about the same size as my feet. In short, my current visual experience is very much like the one that Mach recorded in his famous drawing of how the world looked to him from the vantage point of a chair in his library (2014, pp. 225-226).

Figure 1: Mach's drawing. Hill (2014), p. 226.

Hill argues that, although the distinction between qualia-as-they-appear and qualia-as-they-are is strongly counter-intuitive, we just have to bite the bullet and commit to it, since it is a direct consequence of the representationalist view of perceptual experience, which in turn is the only view that can naturalize qualia, i.e. that can include qualia in the physical world. In short, if we substitute 'qualia' with 'appearance' – a move that Hill probably considers innocent, given his tendency to use the verb 'to appear' interchangeably with 'to look'[12] – Hill accepts that appearances have themselves appearances, or, in other words, that there are *appearances of appearances*: "A representational theory makes it possible to draw an appearance/reality distinction with respect to qualia. On the one hand, there are qualia-as-

12 See Hill (2014), p. 198.

they-are-represented-by-experiential-representations. On the other hand, there are qualia-as-they-are-in-themselves" (2014, p. 210).

Now, it seems to me that the main advantage of Hill's introduction of Thouless properties was to bring qualia back into the physical world, thus dissolving the "mystery" lingering around them since Nagel (1974) and Jackson (1982) wrote their seminal papers on the topic. However, if Thouless properties aren't enough to account for everything that raised the problem of qualia in the first place – Thouless properties still cannot explain the specific way in which they are represented in experience – then it may be argued that the move isn't really worth making. As Hill himself admits, there is still something that is left out by the identification of qualia with Thouless properties: namely, the way those properties look phenomenally to us.

Hill introduced Thouless properties into the picture so that no explanation of how we go from objective, mind-independent properties to properties that are dependent on our perspective in the world and other features of context. However, even after introducing Thouless properties, Hill needs to provide an explanation of *why* those properties are represented by our perceptual system so that they can phenomenally appear differently from occasion to occasion. In other words, he moved the explanatory burden from the "original" distinction between objective properties and perspective-dependent appearances to a distinction within the domain of appearances.

I happen to be one of those people for whom the notion of appearance of an appearance only brings further and unneeded complexity into an already complex picture, without thereby providing a big explanatory advantage. In fact, the notion of appearances of appearances, where appearances *simpliciter* already mark the difference between how things seem to us subjectively and how things really are, is just redundant, and I would rather do without it.

In other words, my objection to Hill is advocating for parsimony: if representationalism needs to draw an appearance/reality distinction also within the domain of phenomenology (i.e. with respect to qualia), then the theory postulates a further level of complexity in the structure of experience. Consequently, it might be wise to explore alternatives that avoid postulating the level of appearances of appearances without losing much in terms of explanatory power. The adverbialist picture I will now sketch offers a compelling way to avoid the redundant notion of appearance of an appearance, and still shows clear commitment towards the naturalizing project.

4. A Broadly Adverbialist Alternative

The kind of naturalization-friendly adverbialism I start from is defended by Chirimuuta (2015) with a specific focus on color. My goal is to show that what Chirimuuta argues about colors (a) can be extended to other features of the environment traditionally understood as properties of objects, such as size or shape, and (b) can provide a somehow radical solution to the "problem" of qualia by promoting a different way to think about what it is for an object to appear in a certain way to a perceiver.

The key thesis of Chirimuuta's color adverbialism is the following: colors are not properties of objects. Alternatively, colors are the *ways* in which the perceptual systems of certain animals, humans included, interact with the external world. According to Chirimuuta, colors are modifications of the *interaction* between an animal and its environment. This idea is spelled out with the help of the contrast between two hypotheses regarding the function(s) of color vision in the life of an animal with the appropriate visual system. The two models are called *coloring-in* and *coloring-for*, and are in turn tied to two different ways of conceiving the function of perception in general (pp. 69-99).

The *coloring-in* model of color perception falls into the broader "correspondence model" of perception, according to which the main function of perception is to *detect* invariant, mind-independent features of the environment. Perceptual systems have the function of making

the perceiver's subjective experience *correspond* to what is "out there" in the most accurate way possible. According to this model, the visual scene is "colored in" as if it were a page of a coloring book.

It is quite easy to see how this model would be consistent with Hill's view of qualia. Perception works by representing the mind-independent world in a certain way. Phenomenal properties like looking yellow, or looking small, are mind-independent properties of the objects out there, that are represented by a visual system that aims at a correspondence between the subjective aspect of experience and the objects and properties out there. Consider again the two trucks. The real truck's phenomenal property is *looking small*$_A$, whereas the toy truck's phenomenal property is *looking small*$_B$. In order to achieve accuracy, the system "depicts" as many details of the scene as possible, just like a photograph (or Mach's drawing above) would do. The more detailed the representation, the more (real) features of the world are picked out, the more adaptive an animal's behavior based on that representation will be.

On the other hand, the rival *coloring-for* model fits in a framework that takes perception to be more concerned with utility than with correspondence/accuracy. In particular, this model suggests that perception aims at providing the organism with information that is *useful* for it in order to interact with its environment in an efficient way, and this may or may not entail presenting the subject with accurate representations corresponding to mind-independent objects and their features. This second model takes colors to be essential to the very construction of the perceived scene. For an animal equipped with cone cells in the retina, seeing a scene in colors isn't just an additional "embellishment" whose role is to improve the accuracy of the representation. 'Coloring for' means coloring *for utility*. Color vision presents features of the environment in a way that makes them immediately salient for the organism, depending on what the needs of that particular organism are.

For creatures with color vision, the scene is constructed *as colored, starting from colors*, and all the other features, such as size, shape, distance, etc. are constructed derivatively. In humans, for example, Chirimuuta points out that V4, i.e. the part of the visual cortex that was once identified as the "color center" due to the sensitivity of that neural population to wavelength information, has recently been found to have a much more diverse range of functions, such as perception of depth (stereopsis), perception of shape, or visual attention (p. 72).

By presenting a great deal of empirical work (pp. 78-98) regarding what she lists as the "functions of color vision" (p. 77), Chirimuuta argues that perception scientists are much more sympathetic with the coloring-for model, that is, the model according to which "color vision is a way of seeing *things* – flowers, tables, ladybirds – not, in the first instance, a way of seeing the *colors*" (p. 69). I take this to be at least a good reason to think that philosophical theories committed to the naturalizing project should be attracted to the utility-based model, too.

In the light of her arguments in favor of the coloring-for model, Chirimuuta then concludes that a certain type of adverbialism is a very plausible view of color ontology. According to "traditional" adverbialism, color properties are *ways* in which we perceive objects, adverbial modifications of our experience of objects. Canaries aren't yellow; they are experienced yellow-*ly*. If color perception is for utility instead of accuracy, the experience of seeing colors is part of what it is to be related to the world in a way that is useful, i.e. that promotes the animal's survival and reproduction.

What if we extended the account beyond colors? Just like we can say that a canary doesn't possess the color property 'yellow', I suggest that we think of all the other ways in which objects can perceptually appear to us as *not* being properties of objects, either. Analogously to colors, the shapes, sizes, etc. that we experience perceptually, that is, the perceptual shape-qualia or size-qualia, are nothing but modifications of the interaction between the subject and the environment, where such interaction is interest-relative, utility-based, and action-guiding.

Think again about the two trucks: we see the distant real truck and the miniature truck small-*ily* not because either of them has a particular property of 'looking small'. Rather, 'small-ily' is just the way in which we experience objects, whenever it is appropriate or useful for creatures like us to do so.

Hill (2014, p. 208) rejects adverbialist theories of experiential awareness – i.e. awareness of qualia – on the basis of the observation that, for the adverbialist, qualia would only be ways in which we are aware of something *else* (e.g. an object), and this is at odds with our actual experience. Clearly, Hill states, our experience is *of* qualia: when we experience pain, our experience is *of* the pain. Therefore, the adverbialist must allow for a non-experiential awareness of qualia, and this in turn forces her to account for this awareness either in terms of direct acquaintance (collapsing the theory into an acquaintance theory, which Hill already rejected), or in terms of judgment (making the theory a doxastic one, which Hill also rejects). However, let me briefly clarify why I think this objection is misguided, especially in the light of the radical "paradigm change" promoted by the kind of adverbialism I endorse.

First, from the fact that, according to adverbialism, qualia aren't themselves properties we can be aware of in the traditional sense, it does not follow that the adverbialist must concede a form of non-experiential awareness of qualia. In fact, an adverbialist like Chirimuuta would probably say the following: we should revise the way we understand the very structure of experience. Experiential awareness *in general* is not meant to be structured as classical representationalists want it to be, i.e. in the form of an attribution of properties to substances.[13] By itself, experience doesn't make us aware *of* anything. Rather, experiential awareness shapes our interaction with the environment so that we can act in it. Experiential awareness is *itself* the way in which a subject and an environment interact, and asking ourselves which properties we are experientially aware *of* while we experience is asking the wrong question. Adverbialism doesn't say that qualia are ways in which we experience something else. It says that qualia are the way in which we experience, period.

At the beginning, I stated that my plan for this paper was to, first, reject Hill's claim that representationalism is the only view capable of naturalizing qualia, and, second, propose a credible alternative. Once we refute a "paradigm" telling us that the perceptual system is a detection device, and we accept that it is more of an *interaction* device, suddenly we will notice that the need to identify the way things look with special kinds of properties possessed by objects loses most of its urgency. If the perceptual system doesn't aim at accuracy/veridicality in itself, but only insofar as accuracy is compatible with what is useful to the creature, why should we think that a representational model of the kind Hill, Burge, and others defend is the best way to explain how perceptual experience has a "grip" on the world? Why think that the structure of perceptual contents is necessarily one in which a property is attributed to a particular?[14]

The naturalistic adverbialism I sketched invites us to abandon the idea of a correspondence between the world and the representations produced by the perceptual system. Moreover, I urge that we stop thinking that, if objects have properties independently of our experience, experience *itself* has to "tell us" that objects have some kind of properties. If we abandon representationalism, we don't need anything that plays the role Thouless properties play in Hill's theory. In the framework I defend, we wouldn't have to say that the distant truck looks small because the property of looking small is a physical attribute the truck possesses

13 See, in this respect, Burge's notion of "typing" (2010, p. 380).

14 I am indebted to Alison Springle for articulating both the implications of a rejection of the accuracy-based model of perceptual content and the possibilities that this rejection opens.

in certain circumstances. Rather, the adverbialist claims that the property of looking small doesn't belong to the truck, but smallness is a modification of a particular "interaction-episode" between us and the world.

I am aware that more needs to be said in order for naturalistic adverbialism to stand as a solid alternative to Hill's proposal in a broad range of cases. Moreover, an adequate defense of this kind of adverbialism would require more detailed explanation of some core-notions, such as the one of "interaction episode". However, I believe that such further work can be done, and it will yield successful results for the naturalizing project more broadly. Naturalistic adverbialism promotes a radical change of paradigm regarding how we should think of the very role of perception, and of how experience relates to the world. In this framework, to naturalize qualia means to think of them as modifiers of the utility-based interaction between a creature and its environment.

REFERENCES

Chalmers, D. (1996). *The Conscious Mind*. Oxford: Oxford University Press.
Chalmers, D. (2001). The First-Person and Third-Person Views (Part I). [Web log post]. http://consc.net/notes/first-third.html.
Chalmers, D. (2007). Phenomenal Concepts and the Explanatory Gap. In T. Alter & W. Walter (Eds.), *Phenomenal Concepts and Phenomenal Knowledge: New Essays on Consciousness and Physicalism.* Oxford: Oxford University Press, 167-194.
Chirimuuta, M. (2015). *Outside Color*. Cambridge, MA: MIT Press.
Burge, T. (2010). *Origins of Objectivity*. New York: Oxford University Press.
Danks, D. (2014). *Unifying the Mind*. Cambridge, MA: MIT Press.
Dennett, D. (1990). Quining Qualia. In W. Lycan (Ed.), *Mind and Cognition*, Oxford: Blackwell, 519-548.
Dennett, D. (1991). *Consciousness Explained*. Boston, MA: Little, Brown and Company.
Dretske, F. (1988). *Explaining Behavior – Reasons in a World of Causes*. Cambridge, MA: MIT Press.
Dretske, F. (1995). *Naturalizing the Mind*. Cambridge, MA: MIT Press.
Fodor, J. (1981). *Representations*. Cambridge, MA: MIT Press.
Fodor, J. (1987). *Psychosemantics*. Cambridge, MA: MIT Press.
Fodor, J. (1994). *The Elm and the Expert*. Cambridge, MA: MIT Press.
Hill, C. (2014). *Meaning, Mind, and Knowledge.* Oxford: Oxford University Press.
Jackson, F. (1982). Epiphenomenal Qualia. *Philosophical Quarterly*, 32 (127), 127-136.
Jackson, F. (1986). What Mary didn't Know. *Journal of Philosophy*, 83 (5), 291-295.
Kim, J. (2006). *Philosophy of Mind* (2nd ed.). Boulder, CO: Westview Press.
Kriegel, U. (2002). Phenomenal Content. *Erkenntnis*, 57(2), 175-218.
Marr, D. (1982). *Vision*. San Francisco, CA: W.H. Freeman and Company.
Millikan, R. (1984). *Language, Thought and other Biological Categories*. Cambridge, MA: MIT Press.
Nagel, T. (1974). What Is it Like to Be a Bat?. *Philosophical Review*, 435-450.
Shoemaker, S. (1975). Functionalism and Qualia. *Philosophical Studies*, 27, 291-315.
Shoemaker, S. (1981). The Inverted Spectrum. *Journal of Philosophy*, 74, 357-381.
Sprevak, M. (2013). Fictionalism about Neural Representations. *The Monist*, 96(4), 539-560.
Tye, M. (1995). *Ten Problems of Consciousness*. Cambridge, MA: MIT Press.
Tye, M. (2000). *Consciousness, Color, and Content*. Cambridge, MA: MIT Press.