

The fragmentation of phenomenal character

A recurring theme of this book is that the problems of perception have a much more complex profile than is usually realized. In this chapter, I begin to make that point by examining debates about *phenomenal character* – debates about *what it's like* to see a mango, hallucinate the sound of a violin, or vividly imagine the scent of vanilla. I argue that philosophers who discuss phenomenal character have had in mind, not one topic, but at least three.

In particular, I argue that the underlying expression “what it’s like” exhibits a particularly deep form of context-sensitivity (§1-§2). When this expression has been used to define phenomenal character, at least three very different contexts have been operating in the background. As a result, the term *phenomenal character* has been used with at least three radically different referents in mind, which I call *hard character*, *strong singular character*, and *conveyed character* (§3-§5). Once these are sharply distinguished, some new and exciting theories of perception come into view. Pluralism is one such theory. Later chapters will offer a pluralist account of each of these kinds of phenomenal character (§6).

1. A core idea introduced

Roughly put, a core idea of this chapter is that different philosophers have used the term *phenomenal character* while having radically different properties in mind. In this section, I give a detailed analogy to characterize this core idea more precisely.

Begin by contrasting two sorts of context-sensitive expressions. On the one side, for a context-sensitive expression to be *shallow* is for it to have a reasonably narrow range of possible referents. A paradigmatic example is the context-sensitive expression *tall*. To be sure, in principle

this expression could refer to just about any height – just imagine a discussion of the dimensions of a mouse. Nevertheless, *tall* will always refer to some height property or other. Similarly for other paradigmatic gradable adjectives: *rich*, *warm*, *funny*, etc. Or, for a very different sort of shallow context-sensitive expression, consider *today*. This expression might in principle refer to any day, but it will always refer to some day or other.

On the other side, if a context-sensitive expression is not shallow, then it is *deep*. The expression *that* is an especially clear example of a deep context-sensitive term: in principle it could be used to refer to nearly anything. I will also count slightly more specific expressions such as *that object* or *that property* as deeply context-sensitive. For these expressions can in principle be used to refer to radically different objects or properties, respectively, so these expressions have a reasonably broad range of possible referents. By contrast, I will count expressions such as *that color* or *that number* as shallowly context-sensitive, since there is some reasonably specific range of possible referents here.

The distinction between shallow and deep context-sensitive expressions is vague.¹ Still, a distinction does not have to be sharp to be real and useful.

Let us make one more distinction. Let *terms* be individuated by their meanings; let *expressions* be individuated by their phonemes and orthographies. Thus there is a single expression, *bank*, corresponding to two different terms, (*river*) *bank* and (*financial*) *bank*. The expression *bank* has two different meanings, we might say.

Now turn to philosophy of mind. By definition, the *phenomenal character* of an experience is what it's like to have that experience. I will soon argue that the defining expression *what it's*

¹ One tricky example is the expression *I*. There is *some* specificity to what the expression *I* can refer to. But a lot is left open, as well: *I* can refer to something that is not a person (imagine a sign that says “I am written in English”). So is the context-sensitivity of this expression deep or shallow? I do not know.

like is deeply context-sensitive, and that philosophers who discuss phenomenal character have had in mind different definitional contexts. But let me first give a less controversial example as a model.

Thomas N. points to a blue circle and says to his friends, “Let us refer to this property as *property X*.” He is in a context in which color properties are salient. What can we say about Thomas’s new term?

Thomas’s definition relies on a deeply context-sensitive expression, *this property*. So, given Thomas’s context, his term *property X* refers to blueness. Surprisingly, Thomas’s term *property X* is not context-sensitive at all, *even though* it was defined using the deeply context-sensitive expression *this property*! The term *property X* simply refers to blueness regardless of context.

Now suppose that John C. is conversing with his own friends across the room and overhears Thomas. John’s interest is piqued; he wants to talk about the same topic as Thomas. John points to the same blue circle and says to his friends, “Let us refer to this property as *property X*.” However, in John’s context, it is not color properties but *shape* properties that are salient. Still, John assumes that his term is the same as Thomas’s and even intends for that to be the case when he gives his definition. Let us add, perhaps unrealistically, that John does not realize that the expression “this property” is deeply context-sensitive. What can we say about John’s term *property X*?

It is difficult to identify the referent of John’s term, for there are two facts that tug in opposite directions. On the one hand, John intends to be using the same term that Thomas is using. This tugs his term towards referring to blueness. On the other hand, John has given an explicit definition of his term via the deeply context-sensitive expression *this property*. Given John’s context, this tugs his term towards referring to circularity.

In light of these opposed pressures, I see three reasonable views about the referent of John's term *property X*. First, that it refers to blueness; second, that it refers to circularity; third, that it is indeterminate whether it refers to circularity or to blueness. (I repeat that on any of these views, John's term *property X* will be context-*insensitive*.) I will not try to adjudicate among these views. Which view is correct may even depend on further facts about the case that have yet to be specified. I simply wish to observe that no matter which view is correct, there is an obvious sense in which circularity is what John *has in mind* when he talks about property X.

For example, when John discusses what he calls "property X," insofar as he is rational, what his evidence will *non-accidentally support* are conclusions about circularity. He might say things like, "Necessarily, anything that has property X cannot be square or triangular." (His evidence might sometimes support conclusions about blueness, but only by accident: "necessarily, anything that has property X is spatially extended.") Moreover, suppose that John were asked to rephrase his claims without using the expression *property X*. He would presumably use some expression referring to circularity. These observations retain their force whether John's term *property X* refers determinately to blueness, determinately to circularity, or indeterminately to both.

What underlies these facts is that when John speaks of property X, circularity is what he *has in mind*. I will not attempt to give a precise characterization of what it is to have something in mind. The notion is an ordinary one, and I take it to be clear enough for our purposes.

It will be useful for us to develop our example a little further. Suppose that John is not the only person who overhears Thomas. Many others do as well, and they too turn to their friends, point to the same blue circle, and define property X as "that property." However, they are speaking in still further contexts: textures are salient in Susanna S.'s context, surface areas are salient in Michael M.'s context, and so on. Then there are Thomas's friends, such as David

C., Frank J., and Michael T., who were conversing with Thomas and understood the original context. When all of these people mingle with one another and speak about “property X,” they will have many different properties in mind. Mass confusion will be sown.

I say that this is just what has happened in contemporary philosophy of mind. Inspired by Thomas Nagel, philosophers have defined the *phenomenal character* of an experience as *what it’s like* for the subject to have the experience.² But the defining expression “what it’s like” is deeply context-sensitive and has been used in the context of radically different philosophical problems. Thus, philosophers who speak about “phenomenal character” have had radically different properties in mind. (That is true whether or not the expression *phenomenal character* has had the same referent in their mouths.) The result has been mass confusion.

The rest of this chapter defends these claims and draws out some important implications.

2. The deep context-sensitivity of “what it’s like”

According to the utterly standard definition, the *phenomenal character* of any experience is what it’s like for the subject to have that experience. Equivalently, it is what having the experience is like for the subject; equivalently again, it is what the experience is like for the subject.³

Regarding that last equivalence: yes, there is not *in general* an identity between what *something* is like for the subject and what *having that thing* is like for the subject. In the special case of experiences, however, it is plausible that there is such an identity: what having a perception of red is like for the subject is presumably the same as what that perception of red is like for the subject.

² See Nagel (1974). His term is actually *subjective character*, and Nagel focuses in the first instance on what it’s like to be a subject. None of this affects my point.

³ This definition is now so dominant that it would be pedestrian to give further citations.

In any case, for us to understand the expression *phenomenal character*, we must first understand the expression *what it's like*. It is widely appreciated that the latter expression is context-sensitive.⁴ In this section, I argue that this is only the tip of the iceberg. What has remained below the water is that the expression is *deeply* context-sensitive.

Suppose that, in making a claim about what something is like, some property *F* is cited. Sometimes, this will entail that the entity in question really is *F*: “I’ll tell you what Sarah is like – she is honest.” But sometimes it will be obvious that there is no such entailment: “I’ll tell you what blue is like – it is like purple.” At other times a property may not be cited at all: “I’ll tell you what Chicago is like – it is like London.”⁵ To simplify the discussion, my proposal will focus only on uses of “what it’s like” that are of the first type: *F-entailing uses*, as I shall say. I argue in Appendix 2A that the proposal generalizes to uses that are not *F-entailing*.

What I propose is:

The *deep context-sensitivity view*: Any *F-entailing* claim of the form “what *x* is like is *F*” is deeply context-sensitive, because it means that *F* is a property of *x* that is relevant in the context at hand (where this context may be further specified with prepositional phrases and subordinate clauses).

Is the deep context-sensitivity view best understood as a view about (i) *just* the literal meaning, or (ii) the literal meaning *together with* the standard implicatures, of claims of the form “what *x* is like is *F*”? This is a delicate question. However, the answer to it will not matter for my purposes. To glide over this question, I stipulate that here and throughout I am using the term *meaning* to

⁴ See for example Snowdon (2010, p. 22); Stoljar (2016, pp. 1174-1175); Gaskin (2019, p. 679 and p. 682).

⁵ I take this example from Gaskin (2019, p. 675).

refer to the disjunction of (i) and (ii). Thus the deep context-sensitivity view entails that an assertion that what x is like is F , where F is a contextually irrelevant property of x , is at least *inappropriate*, but it is neutral on whether such an assertion must also be *false*.

I offer six lines of support for the deep context-sensitivity hypothesis.

The first line of support. Suppose that we are discussing what India is like. On an internet travel forum, it might be appropriate to say, “It is full of temples, palaces, and other historical sites.” In a geology class, it might be appropriate to say, “It has been largely shaped by extensive melting, and subsequent basalt flooding, underneath the Indian Craton.” And in an economics conference, it might be appropriate to say, “It is enjoying rapid expansion – at present, it has the fastest-growing economy in the world.”⁶ Which claim is appropriate evidently depends on the context. It would be inappropriate, on the internet travel forum, to cite geological or economic features of India.

In short, there is radical variation in which properties of India – including cultural properties, geological properties, and economic properties – can be appropriately cited as what India is like in different contexts. This directly supports the deep context-sensitivity view.

The same sort of evidence can be found when F -entailing claims about what something is like are adorned with a prepositional phrase or a subordinate clause. For example, the traveler might discuss what India is like *in the winter*, the geologist might discuss what India was like *after the Deccan traps were created*, and the economist might discuss what India is like *in the manufacturing sector*. The deep context-sensitivity theorist should say that these adornments will serve to specify the context further (in addition to whatever role they play with respect to the literal meaning of the sentence).

⁶ The second and third of these answers are drawn from Wikipedia.

The second line of support. I have often tried to explain what I work on to the uninitiated: to students and friends, for example. I have *never* been able to get *anyone* to understand the expression *phenomenal character* merely by giving the what-it's-like definition, with no prologue or epilogue. This observation is not just about outsiders to our discipline: even some professional philosophers find the term vague and unhelpful.⁷

This is puzzling. English speakers use what-it's-like talk all the time. How then can they fail to understand the term *phenomenal character* when they are told that it is what it's like to have an experience?

You might answer simply that what-it's-like talk is context-sensitive. But this explanation is too weak. By analogy, imagine that someone wonders who the upper class are. You answer that they are the rich, with no prologue or epilogue. Even though the term *rich* is context-sensitive, your audience will gain a good understanding of the term *upper class*. For the context-sensitivity here is shallow: the term *rich* has a reasonably narrow range of possible referents.

However, if what-it's-like talk is *deeply* context-sensitive, then the phenomenon is explained. For, on this view, saying that the phenomenal character of an experience is what it's like for the subject to have that experience, with no prologue or epilogue, is much like saying, "The phenomenal character of an experience is *that property*," or like saying, "The upper class are the individuals with *that property*," with nothing to indicate which property you are referring to. Your audience will be mystified.

The third line of support. I tell my introductory students to read Nagel's classic paper "What is it like to be a bat?," which offers a what-it's-like definition of the term *subjective character*.

⁷ See Snowdon (2010); Hill (2014, pp. 230-231).

They understand this – but they are baffled when the same definition is served neat. How is this possible?

The deep context-sensitivity view provides a satisfying answer: Nagel’s article provides the missing context for the what-it’s-like definition. To return to our original example, it is unhelpful to hear, out of context, that property X is *that property* of a blue circle. But the very same remark would be very informative once you learn that we are discussing color properties.

The fourth line of support. It is often said that there is some deep connection between the what-it’s-like properties of experience and consciousness. It is also often noted that this is semantically mysterious.⁸ Look at the words “what,” “it’s,” and “like.” Individually they have no particular connection to consciousness, nor does putting them together seem to help. What is happening?

Some theorists propose complex answers to this question.⁹ I offer a simple one: *out of context*, there is no particular connection between the expression *what it’s like* and consciousness. The connection arises almost entirely from the context itself. By analogy, out of context, there is no particular connection between the expression *that* and the color blue. But it is unmysterious how such a connection might arise in certain contexts.

The fifth line of support. It is sometimes held that philosophers are using the expression *what it’s like* in a special, technical sense.¹⁰ This would explain why the bare what-it’s-like definition is unhelpful. It might also explain why Nagel’s article is helpful: perhaps students can glean this special, technical sense of the expression *what it’s like* from context. It might also explain why examining the words “what,” “it’s,” and “like” reveals no special connection to consciousness.

⁸ See for example Stoljar (2016, p. 1162).

⁹ For instance, Lormand (2004) and Stoljar (2016, esp. §7).

¹⁰ See Lewis (1995, p. 140); Byrne (2004, p. 215).

Better, though, not to multiply special technical senses without necessity. That is a virtue of the deep context-sensitivity hypothesis: it does justice to the data above while saying that philosophers use the expression *what it's like* in a perfectly ordinary way.¹¹

The sixth line of support. The deep context-sensitivity view predicts that if we talk about what it's like to have an experience in a suitably distant context, then we should be talking about properties of the experience that are not phenomenal characters. For if what-it's-like talk is deeply context-sensitive, then such talk can be used in different contexts to refer to radically different kinds of properties of experiences.

To test this prediction, consider a case in which you ask what it was like for me to see Cirque du Soleil in New York. I might truly answer that my experience was surprising and exciting. But philosophers do not usually think that the properties of being surprising and being exciting are phenomenal characters of my visual experience of Cirque du Soleil.¹²

To be sure, philosophers *do* usually think that the states of surprise and excitement *have* phenomenal characters of their own. What they do not usually think is that the properties of being surprising and being exciting *are themselves* phenomenal characters of my visual experience of Cirque du Soleil. Various philosophers have identified the total phenomenal character of visual experiences with – for example – (i) the property of being a special type of representation of such-and-such perceived entities, (ii) the property of being a primitive, non-representational relation of awareness to such-and-such perceived entities, and (iii) the neural property that underlies my

¹¹ On the point that philosophical “what it's like” talk is not technical, I side with Farrell (2016).

¹² On the basis of such examples, and contrary to philosophical orthodoxy, Hellie argues that emotional states or properties *are* phenomenal characters (2007a, pp. 261-262). But if the expression “what it's like” is deeply context-sensitive, then Hellie's reasoning is too swift: he has not ruled out the possibility of a shift in context.

visual perception.¹³ None of these theories would treat the properties of being surprising and exciting as phenomenal characters of my visual experience of Cirque du Soleil.

Summary. The deep context-sensitivity view enjoys the following lines of support:

- (a) It is directly supported by ordinary examples of what-it's-like talk.
- (b) It correctly predicts that, and explains why, the bare what-it's-like definition of phenomenal character will be mystifying.
- (c) It correctly predicts that, and explains why, the what-it's-like definition *will* be very helpful when given in a clear philosophical context of the sort that Nagel provides.
- (d) It correctly predicts that, and explains why, the words “what,” “it’s,” and “like” have no obvious relationship to consciousness.
- (e) It treats philosophical uses of *what it's like* as entirely ordinary rather than technical.
- (f) It correctly predicts that, and explains why, in non-philosophical contexts, we can talk about what an experience is like without talking about its phenomenal character.

Some readers will wonder how the deep context-sensitivity view relates to standard theories of the meaning of “what it’s like.” I believe that most of these theories strongly support the deep context-sensitivity view. This is a seventh line of support for the view – see Appendix 2A. Meanwhile, some will object that philosophers *must* ordinarily have been speaking in the same context when defining the term *phenomenal character*. After all, the term is usually introduced against a standard backdrop of supplementary remarks – e.g., it is standard to give a

¹³ These are the paradigmatic theories of phenomenal character offered by representationalists (e.g., Tye (1995)), naïve realists (e.g., Campbell (2002)), and type-physicalists, respectively.

few examples of experiences, to connect the notion of what it's like to have an experience to the notion of consciousness, and so on. I rebut this objection in Appendix 2B.

The next stage of my argument is to show that philosophers have, without realizing it, given the what-it's-like definition of *phenomenal character* with at least three very different contexts in the background.

3. The first context: hard character

I observe that in one strand of the philosophical literature, the term *phenomenal character* gets used in the context of the so-called *hard problem(s) of consciousness*. In particular, at least in the paradigmatic cases, these theorists take it that the phenomenal character of an experience plays the following roles:

1. The *explanatory gap role*: We seem unable to understand how it could be nothing more than something that is conceived purely physically and/or functionally. (See Levine (1983).)
2. The *missing knowledge role*: We seem unable to know everything about it just on the basis of having all physically and functionally conceived information. (See Nagel (1974); Jackson (1982).)
3. The *conceivable absence role*: We seem to be able to conceive of it as present in one subject *S* yet absent in some other subject who is identical to *S* in all functional and physical respects. (See Block (1980); Chalmers (1996).)

4. The *conceivable inversion role*: We seem to be able to conceive of it as invertible across subjects who are identical in all functional and physical respects. (See Shoemaker (1982).)¹⁴

You might immediately object that to say that something plays roles 1-4 is to assume that there is some property of experience that *in fact* generates an explanatory gap, missing knowledge, and so on. Such an assumption is very controversial. This objection is mistaken, however. To say that something plays roles 1-4 is only to assume that there is a property of experience that *has seemed* – *in a reasonably robust way, to a reasonably large group of philosophers, from the late twentieth century to the early twenty-first century* – to generate an explanatory gap, to generate missing knowledge, and so on.¹⁵ This assumption is very plausible.

Why think that the context characterized by roles 1-4 pins down one property that philosophers have had in mind when using the expression *phenomenal character*? One compelling reason is sheer influence. Look at discussions of phenomenal character from the final decades of the twentieth century. You will find over and over that it is the problems implicated in roles 1-4 that are discussed, and it is the works listed above that are cited.

In addition, in the more recent literature many philosophers assert that there is some especially close connection between phenomenal character and roles such as 1-4, especially the explanatory gap role. Below are a few representative quotations. Note that although these quotations use terms other than *phenomenal character*, these other terms are always defined with the help of the expression *what it's like*.

¹⁴ There may be other roles related to the hard problems of consciousness. For further discussion, see Chalmers (2018).

¹⁵ Let the italicized phrase be a specification of the meaning of the word “we” in roles 1-4.

“A property is phenomenal only if it contributes to the hard problem of consciousness, and in particular, only if it gives rise to an explanatory gap.” (Carruthers and Veillet 2011, p. 45)

“Perhaps the most striking feature of phenomenal properties is that they appear to present an explanatory gap.” (McClelland 2016, p. 540)

“The traditional explanatory gap consists in the fact that phenomenal states do not seem to be explainable in terms of physical states [see Levine (1983)]. Phenomenal states are states that there is *something it is like to be in*. In what follows, I distinguish four features of phenomenal states that individually resist explanation.” (Saad 2016, p. 2356; the emphasis and reference to Levine are both his.)

“*Consciousness properties* are properties that contribute to making up what things are or can be *like* for subjects....I will often abbreviate the physicalist identity thesis as the claim that C=P.... Many of us have an exceptional resistance to the physicalist identity thesis.... Our exceptional resistance to physicalism is the “explanatory gap” on one understanding of that phrase. (Sundström 2017, pp. 681-682, emphasis his)

But why think that these remarks are intended as *reference-determining* contextual information for the expression *what it's like*, when that expression is being used to define the term *phenomenal character*? Couldn't these simply be further assertions about phenomenal character?

I doubt it. It is very clear that all of these theorists have in mind *some* restrictions when they use the term *phenomenal character*. None of them would count *every* property of an experience as a phenomenal character of that experience. But I just cannot find anything else in these works that would supply these contextual restrictions. The remarks that I have quoted are the only substantial pieces of dialectically neutral contextual information that these theorists give before they launch into their arguments. So – in the interest of charity, in light of the historical

importance of the hard problems of consciousness, and in light of the deep context-sensitivity of the expression “what it’s like” – I interpret the quotations above as providing contextual information that is reference-determining.

In other cases, I would offer the same interpretation on more holistic grounds. A paradigmatic example is Michael Tye’s 1995 book *Ten Problems of Consciousness*. In the first paragraph of chapter 1 of this book, Tye simply lists a few experiences, mentions their “phenomenal or ‘what it is like’ aspects,” and moves on. But it is impossible to miss the imprint of the hard problems on Tye’s way of thinking: the titular ten problems include the explanatory gap problem, the missing knowledge problem, the conceivable absence problem, and the conceivable inversion problem.¹⁶ Similarly for Dretske’s book *Naturalizing the Mind*. Dretske briskly defines the “subjectivity” of an experience of *F* as what it’s like to experience *F*.¹⁷ Immediately afterwards, however, he launches into a 15-page discussion of the missing knowledge role, as it is characterized by Jackson and Nagel.¹⁸

There is a final reason that I am partial to thinking that roles 1-4 help to fix what some philosophers have in mind when using the term *phenomenal character*. Recall my earlier report that I have *never* been able to get the uninitiated to understand the term just by appealing to what it’s like to have an experience. I now add that I have *routinely* been able to get the uninitiated to understand the term, at least in the way that this group of philosophers understands it, just by describing roles 1-4. In fact, describing such roles is the *only* way that I have found to get the

¹⁶ See pp. 16-17, 12-13, 21-25, and 26-29, respectively. One minor complication is that Tye thinks of phenomenal character not as a *property* of an experience, but as the *content* of the experience. If anything, this is more grist for my mill: here is even more fragmentation in uses of “phenomenal character.”

¹⁷ Dretske (1995), p. 81.

¹⁸ Ibid, pp. 81-95.

uninitiated to use the term *phenomenal character* in the way that these philosophers do – regardless of whether I use the expression “what it’s like.”

I would like to offer a particular proposal about how to more precisely characterize what philosophers have in mind when they define the expression *phenomenal character* in the context of the hard problems. The idea is to treat the term *phenomenal character* as a natural kind term like *water* or *gold*.¹⁹ We can start by identifying the *apparent paradigms* of phenomenal characters as those entities that play enough of (but not necessarily all of) roles 1-4. We can then use those apparent paradigms to lock onto the relevant natural kind, thereby pulling in any non-paradigmatic phenomenal characters, and perhaps also pushing out some entities that initially seemed to be paradigmatic phenomenal characters but turned out not to be phenomenal characters at all.²⁰

More precisely, I propose that:

When philosophers define the expression *phenomenal character (of an experience)* in the context of the hard problems of consciousness, they have in mind what I will call *hard character*: any property of that experience that belongs to the natural kind associated with the properties of experiences that play enough of roles 1-4.^{21,22}

¹⁹ On the notion of a natural kind, see Lewis (1983).

²⁰ Here I am taking inspiration from Lewis (1970).

²¹ If there is no such natural kind, then *phenomenal character* might refer simply to any property of that experience that plays enough of roles 1-4.

There is another possibility: perhaps a few distinct natural kind properties play different sets of roles 1-4. For example, some properties might play the missing knowledge role but not the conceivable absence role (see Carruthers and Veillet (2011); McClelland (2016); Bayne (ms)). In this case, perhaps *hard character* refers indeterminately to these, or determinately to their disjunction. If so, then fragmentation is even more widespread, and my central points have even wider application.

²² I argue in my (2019) that one of the hard problems, the explanatory gap problem, arises not only consciousness, but also for moral goodness and the self. Is this a problem for my definition

Incidentally, there are interesting connections between the current proposal and the proposal to define *phenomenal character* with a long list of positive and negative examples. See Appendix 2C.

4. The second context: strong singular character

Turn now to philosophers working, not in the context of the hard problems of consciousness, but in the context of certain problems involving perception – the paradigm being John Campbell in his 2002 book *Reference and Consciousness*. It is a central aim of Campbell's book to account for the phenomenal character of perception. Yet Campbell has little interest in the hard problems of consciousness: his central discussion of phenomenal character occurs in chapters 6 and 7, but these chapters do not contain even a passing mention of explanatory gaps, inverted spectra, or zombies. These phrases do not even appear in the book's index. What is going on?²³

I suggest that Campbell is understanding the term *phenomenal character* in terms of a theoretical role that has little direct connection to the hard problems of consciousness – in particular:

5. The *strong singular role*: It helps to explain why perceptions not only position us to make singular reference to certain entities, but also position us to know, in an especially deep sense, what we are thereby referring to.

Thus I propose that:

of hard character? No: the definition appeals to a particular *natural kind property*, which is presumably not a property of moral goodness or the self.

²³ Campbell does mention inverted spectra in his (1993, p. 268). But this is brief, and more importantly, Campbell is not discussing phenomenal character here.

When the expression *phenomenal character (of an experience)* is defined in the context of certain problems involving perception, it refers to what I will call *strong singular character*: the property of that experience that plays role 5.²⁴

Let me build my case for these claims gradually. My focus will be on chapter 6 of *Reference and Consciousness*, for it is here that Campbell carefully introduces the notion of phenomenal character and offers his theory of it.

At almost the very beginning of that chapter, Campbell makes some remarks that are worth quoting at length:

Suppose that you live in one of a terraced row of houses, and you sometimes hear noises from the house next door. Being of an enquiring turn of mind, you formulate hypotheses about what objects are to be found next door.... And you might test and confirm your hypotheses over a long period, without ever catching sight of those things.

Suppose now that the day finally arrives when you do get a look inside the house. What does this add to your knowledge? Perhaps the hypotheses you formed had been amply confirmed long before your look inside, so the existence of objects with these particular functional roles does not get significant further confirmation from your observation. Nor is it that you can now refer to those particular objects but could not refer to them before. You could have referred to those particulars before....

The contrast between the knowledge you have now, on the basis of a look at the objects, and the knowledge you had before of the existence of objects with particular functional roles, is that when you see the thing, you are confronted by the individual substance itself....

²⁴ Why do I call this *strong* singular character? Because I will later use the term *weak singular character* to refer to the property of merely positioning the subject to make singular reference to something (regardless of whether she can know what she is thereby referring to).

These remarks fill out something of the sense in which conscious attention to an object can provide you with knowledge of the reference of a demonstrative referring to it. (2002, pp. 115-116)

This is Campbell's characterization of what I have called the *strong singular role*. For ease of reference, I have numbered a few crucial claims. Note that the strong singular role is associated distinctively (though, I will argue later, not exclusively) with perception. Seeing a mahogany table will position you not only to make singular reference to it, but also to know, in a special sense, what it is that you are referring to. Merely hallucinating a table, or sensorily imagining a table that you have not seen before, will not even position you to refer to a real table, much less to know what you are thereby referring to.

A few pages later, Campbell makes the following telling remark:

“Since both the Relational and Representational Views of the phenomenal character of experience have to be taken seriously ..., how are we to decide between them? ... I think that the only way to proceed is to ask why we need the notion of the phenomenal character of experience. We have to look at the role that the notion plays in our reflective thinking, we have to ask what the point is of the notion. I have been arguing that experience is what explains our grasp of substantial objects.” (Campbell 2002, p. 120)

Given Campbell's emphasis on the *notion* of phenomenal character, it is clear that this remark is intended to be definitional.²⁵ It is also clear that Campbell's talk of “explaining our grasp of substantial objects” is meant to refer to the strong singular role. For one thing, Campbell goes

²⁵ To be sure, at the end of the quotation Campbell slides from talk about the *phenomenal character* of an experience to talk of *the experience itself*. But the term *phenomenal character* does not appear anywhere else in the paragraph from which this quotation was taken. Thus, the charitable interpretation is that, according to Campbell, the point of the notion of phenomenal character is to explain our grasp of substantial objects.

on to argue that the special sense in which we can know the reference of perceptual demonstratives is partly a matter of our being able to form a certain conception of the objects being demonstrated – for instance, a conception of those objects as having categorical and intrinsic characters. It is in this sense that we have a grasp of objects as *substantial*. In addition, the present quotation is from a section entitled “The Argument from the Explanatory Role of Experience,” and the section from which the earlier quotation was drawn is entitled “The Explanatory Role of Experience.” Evidently, Campbell intends for his argument to draw on the explanatory role that he identified there.

Further support for this interpretation can be found in chapter 7, where Campbell says:

“Why do we need the notion of the phenomenal character of experience? We have to look at the role that the notion plays in our reflective thinking, we have to ask what the point is of the notion ... I have argued that we use the notion of experience of an object – or more precisely, conscious attention to an object – in explaining how it is that we have knowledge of the reference of a demonstrative.” (Campbell 2002, p. 138)²⁶

This is an even clearer reference to role 5. And, again, Campbell’s use of the word *notion* makes it clear that his remark is intended to be definitional.

Thus, in light of the deep context-sensitivity hypothesis, I conclude that in his book *Reference and Consciousness*, Campbell uses the term *phenomenal character* with strong singular character in mind. For reasons explained in §1, this conclusion is silent about the *referent* of Campbell’s expression. If his dominant intention is to use the same term as certain previous thinkers, then his expression may refer to hard character, even while strong singular character is what he has *in mind*.

²⁶ The remarks from the previous footnote apply once more.

One consequence of this is that Campbell ends up talking past some of his opponents. Campbell argues against the representationalist view of “phenomenal character” championed by Dretske and Tye²⁷ – but these thinkers have in mind hard character, not strong singular character! To be sure, Campbell has many real, substantive disputes with Dretske and Tye. But he is talking past them *insofar as he is discussing phenomenal character*.

The expression “what it’s like” does not appear in any of the above passages from Campbell. Is this a problem? I argue in Appendix 2D that it is not.

Conclusion. When Campbell uses the expression *phenomenal character* in his book *Reference and Consciousness*, it is strong singular character that he has in mind. The significance of this conclusion is far-reaching, for Campbell’s book has a foundational status in contemporary philosophy of perception.²⁸ In particular, it is often agreed that perception gives the subject a special kind of knowledge of the reference of her demonstrative terms, and it is sometimes thought that this requires perceived objects and property-instances to be constituents of phenomenal character. *Reference and Consciousness* is the seminal expression of these ideas. I thus conjecture – though I will not argue here – that Campbell has colored the way that some other philosophers understand the expression *phenomenal character*.²⁹

This flows into a broader observation. Consider *naïve realism*, the view that what it is to perceive something is to stand in a direct, non-representational relation of awareness to it. John Campbell is one of the most influential naïve realists alive; Michael Martin is another. What does

²⁷ See pp. 116-117 for Campbell’s characterization of the representationalist theory of phenomenal character. See p. 146 for Campbell’s attribution of representationalism to Dretske and Tye.

²⁸ It might help to note that *Reference and Consciousness* has received more than 1000 citations.

²⁹ These other philosophers need not have in mind the exact context that Campbell does. They may draw on some of roles 5-7 and some other roles, besides. That said, I find it hard to cite specific examples. The problem is that many philosophers say little to define the term *phenomenal character*, except that it is “what it’s like to have an experience.”

Martin have in mind when he speaks of phenomenal character? I suspect that it is not strong singular character, but something like the notion of *what is present in the subject's stream of consciousness*.³⁰ What is crystal clear, however, is that Martin does *not* have hard character in mind. Search a number of his most influential papers, and you will find virtually no talk of explanatory gaps, color-blind scientists, or zombies.³¹

So I wish to make an observation about naïve realism more generally: whether or not naïve realists have in mind strong singular character or something else, *very few of them* think of phenomenal character as hard character.³² Yet naïve realists usually think that they are engaging with the theories of phenomenal character developed by early representationalists such as Dretske and Tye. They are not.

³⁰ As Martin (1997) puts it, “For any conscious state of mind there is something that it is like for the subject to be in that state – each such state of mind contributes to the character of one’s stream of consciousness” (p. 3). Martin continues to make reference to streams of consciousness in his later work. See for example his (2006, p. 371).

³¹ For a rare exception, see Martin (2004, pp. 83-84), which mentions the possibility of philosophical zombies – but even this mention is tangential. Here Martin is considering the objection that his disjunctivist theory cannot allow that a hallucinating subject is phenomenally conscious at all. He refers to zombies only to make this objection vivid:

the disjunctivist’s conception of these cases seems to introduce a form of philosophical zombie: a subject who may have thoughts and possess the ability to make judgements about phenomenal consciousness but who lacks phenomenal consciousness proper. In that case, one’s intuition will be that the subject so described is not properly conscious and lacks genuine subjectivity or point of view on the world.

Of course, ... we have the strong intuition that an unfortunate subject who is subject to a total hallucination must still be conscious (after all this is the force of the objection to the disjunctivist)” (ibid, p. 83).

³² There is one notable exception: naïve realist William Fish is sometimes clearly discussing hard character. Fish even has a 2008 paper entitled “Relationalism and the problems of consciousness,” where the problems in question are the hard problems. Fish has a substantive disagreement about phenomenal character with representationalist Michael Tye, who is also discussing hard character. But at the same time Fish has a merely verbal agreement *about phenomenal character* with his fellow naïve realist John Campbell (though I am sure that Fish and Campbell are genuinely allied on other topics). For, again, Campbell is discussing strong singular character.

5. The third context: conveyed character

That is not all; there are even more contexts in play. Siegel's 2010 book *The Contents of Visual Experience* is a fine example. What Siegel has in mind when she uses the term *phenomenal character* is what I will call *conveyed character*. The conveyed character of an experience is any property of the form: *conveying ... to the subject*.

Siegel clarifies that an experience *conveys* a content "if [that content] would be a content of explicit beliefs that are natural to form on the basis of visual experience," "if [the content] enables the experience to guide bodily actions," and "if it is manifest to introspection that [the content] is a content of experience."³³ Siegel also speaks of entities other than contents that can be conveyed to the subject – properties, for example – so presumably her remarks above can somehow be generalized to explain what it is for these other entities to be conveyed to the subject.

Why do I think that Siegel uses the term *phenomenal character* with conveyed character in mind? Admittedly, Siegel is not explicit on the point. Like most authors who use the term *phenomenal character*, Siegel introduces that term simply by giving the "what it's like" definition, along with an example of an experience. This, by itself, does not give us much information.

However, Siegel does leave behind some interesting clues. To begin with, Siegel presupposes that if there are properties presented in phenomenology, then "those properties are conveyed to the subject."³⁴ Now, from context, it is clear that this presupposition is meant to be acceptable to all parties in the dispute, which suggests that Siegel is thinking of

³³ Siegel (2010, p. 51).

³⁴ Ibid, pp. 51-52.

phenomenology as something like conveyed character. And Siegel makes it clear that, for her, the term *phenomenology* is interchangeable with the term *phenomenal character*.³⁵

In a similar vein, Siegel describes one of her arguments, the *Argument from Appearing*, as “starting with something that is conveyed to the subject.”³⁶ Shortly afterwards, Siegel characterizes this same argument as “proceeds from premises about the phenomenal character of visual perceptual experience.” This further supports the idea that, for Siegel, the phenomenal character of an experience is what the experience conveys to the subject.

There is another suggestive clue: Siegel insists that “it would stretch the notion of the phenomenal too far if there were phenomenally conscious states that were necessarily completely under the first-person radar.”³⁷ Given Siegel’s use of the word “notion,” she seems to be making a definitional point here. And it is plausible that the metaphors of being conveyed to the subject and being within the first-person radar are intended to mean the same thing.

None of this is conclusive. Like many contemporary philosophers, Siegel does not say enough to make it perfectly clear how she is using the term. But the present interpretation is the best one that I can find.

Significance. This is important because Siegel, like Campbell, has a foundational status in contemporary philosophy of perception. In particular, according to the traditionally accepted *conservative view*, the phenomenal character of a visual experience pertains only to *low-level properties and relations* such as shape, size, and color. Nowadays, however, many philosophers accept the *liberal view* that such phenomenal character sometimes pertains to some *high-level properties and relations*: for instance, causation, natural kind properties, emotional properties, and

³⁵ Ibid, p. 4.

³⁶ Ibid, p. 44.

³⁷ Ibid, p. 81, fn. 3.

affordances. Siegel was one of the thinkers who helped to put liberalism on the map,³⁸ so it is plausible that her way of thinking has influenced many others.³⁹

The problem is that Siegel is talking past her main opponents.

She cites Michael Tye as one of her conservative opponents.⁴⁰ Presumably she does so because Tye holds that properties like *being a tiger* do not figure in the “phenomenal character” of a visual experience⁴¹; such contents concern only low-level properties like “distance away, orientation, determinate color, [and] texture.”⁴² But it is hard character, not conveyed character, that Tye has in mind. Siegel’s arguments thus fail to connect with Tye: it is perfectly consistent to hold that, for visual experiences, the hard problems of consciousness are associated only with low-level properties, while also holding that visual experiences convey information about low-level and high-level properties. Loosely put, the idea might be that there is a broad circle of what visual experience conveys and a much narrower circle of what generates the hard problems of consciousness.

For similar reasons, although Siegel thinks that she is engaging with John Campbell about “phenomenal character,” she is actually talking past him.⁴³ Campbell has in mind strong singular character, not conveyed character.

What is even more confusing is that Siegel may be in merely verbal *agreement* with some other liberals. Two examples are liberals Tim Bayne and Tom McClelland, both of whom see a

³⁸ See Siegel (2005), (2006a), (2006b), (2009), and (2010). The idea caught on quickly. See for example Bayne (2009); Masrour (2011).

³⁹ The remarks from fn. 29 apply here, with the obvious modifications.

⁴⁰ See for example Siegel (2005, p. 501, fn. 18).

⁴¹ Tye understands phenomenal character in terms of the contents delivered from sensory modules to cognition (1995, p. 137) and claims that these contents do *not* concern entities such as tigers (p. 141).

⁴² Ibid, p. 140.

⁴³ See for example Siegel (2010, p. 166, fn. 15).

deep connection between phenomenal character and at least some of the hard problems of consciousness.⁴⁴ Thus Bayne and McClelland may well accept liberalism about hard character.

A conjecture. In Appendix 2E, I argue that the prominent contemporary representationalist Susanna Schellenberg uses the expression *phenomenal character* with some fourth referent in mind. Given that there are many philosophical problems about experiences, it would be unsurprising if other philosophers had used the expression *phenomenal character* with still further referents in mind.⁴⁵ I leave this as an open conjecture.

Summing up and moving forward. I have argued that the expression “what it’s like” is deeply context-sensitive. I then observed that the expression “what it’s like” has also been used to define the term *phenomenal character* in at least three very different contexts. As a result, this term has picked out at least three very different referents: hard character, strong singular character, and conveyed character. Some will worry that these terms might co-refer anyway. I take up this challenge in Appendix 2F. Others will worry that Tye, Dretske, Campbell, and Siegel have used the expression *phenomenal character* to refer to exactly the same property, with the contextual information that I have discussed functioning to highlight different *aspects* of this property. I address this concern in Appendix 2G.

My conclusion is that philosophical debates about phenomenal character have been thoroughly infected with a purely verbal component.⁴⁶ This has a number of important applications.

⁴⁴ Bayne (2009, §6) focuses on the relationship between (i) high-level properties and (ii) the knowledge argument and the explanatory gap problem. McClelland (2016) does the same, but with sole emphasis on the explanatory gap problem.

⁴⁵ Such slippage might even have occurred *within a single text*.

⁴⁶ To skim off such verbal disputes, for the remainder of this book I will not use terms defined via the expression “what it’s like.” I will only use terms like *hard character*, *strong singular character*, and *conveyed character* – terms that are explicitly characterized by their theoretical roles.

I leave some of these applications to the appendices. It has been argued on semantic grounds that hard character does not exist and that Nagel's what-it's-like definition of consciousness is trivial; I rebut these arguments in Appendices 2H and 2I, respectively. In addition, John Campbell's central argument against Tye-Dretske representationalism focuses on "phenomenal character." In Appendix 2J, I show that Campbell has seriously misidentified the source of his disagreement with Tye and Dretske.

The most important application, however, is the one that I discuss in the next section.

6. A new approach to "phenomenal character"

There are several pairs of claims about "phenomenal character" that are initially very attractive but apparently inconsistent. With the distinctions among hard character, strong singular character, and conveyed character clearly in view, we can dissolve the appearance of such inconsistencies. This constitutes an exciting new approach to "phenomenal character."

I have already explained how such an approach is possible with respect to the conservatism/liberalism debate: conservatism about *hard* character is perfectly compatible liberalism about *conveyed* character. Now I wish to make a structurally similar point about hard character and strong singular character.

Hard character is, roughly, the natural kind property that generates the hard problems of consciousness. These problems seem to arise equally, and in precisely the same form, for any perception and matching hallucination. Thus, an initially attractive hypothesis is that any perception and matching hallucination have precisely the same hard character.

Is the same true of strong singular character? This is the property that helps to explain why perceptions position us not only to make singular reference to certain entities, but also to know what we are thereby referring to. When I see a mahogany table, my experience positions

me to make knowledgeable reference to the table, its instances of brownness and table-shapedness, and so on. When I suffer a matching hallucination, my experience does not position me to make any sort of reference, knowledgeable or otherwise, to any such particulars. My hallucination might first-personally appear to do so, but this appearance is mistaken. Thus, the most initially plausible view is that perceptions and matching hallucinations have significantly different kinds of strong singular character, if hallucinations even have strong singular character at all.⁴⁷

This result is not only significant. It is surprising. The standard naïve realist view is indeed that my perception and my matching hallucination differ radically in their “phenomenal characters.” However, it has been said over and over, sometimes by naïve realists themselves, that this view is strongly counterintuitive.⁴⁸ I say that this is plainly incorrect if what is at issue is strong singular character. The view that hallucinations have no strong singular character, far from being counterintuitive, is the most initially attractive view. This is perfectly consistent with saying that any perception and any matching hallucination have precisely the same hard character.

I extend this point in Appendix 2K. There I consider a standard objection to naïve realism: that it cannot accommodate the very plausible claim that any perception and any matching hallucination have precisely the same phenomenal character. I reply that there are several *prima facie* reasonable ways for the naïve realist to accommodate this claim.

Conclusion. Consider this pair of claims:

⁴⁷ This does not yet establish that my hallucination has *no* strong singular character: it remains possible that it positions me to make knowledgeable reference to certain *universals*. I will argue in chapter 3 that this is the case.

⁴⁸ See Fish (2008, ch. 4, for example p. 81); Logue (2012, §4-§5) and (2013, p. 106); [cite more].

The *common kind theory of hard character*: Any perception and any matching hallucination have precisely the same hard character.

Disjunctivism about strong singular character: Any perception and any matching hallucination have radically different strong singular characters.

And recall another pair of claims from the previous section:

Conservatism about hard character: Only low-level properties figure in the hard character of any perceptual experience.

Liberalism about conveyed character: Perceptual experiences sometimes convey information about high-level properties in addition to low-level ones.

I find each of these claims to be very initially attractive. But when they are expressed as claims about “phenomenal character,” the first claim appears to contradict the second, and the third claim appears to contradict the fourth. Only when we express them in this more perspicuous form does the appearance of contradiction vanish; only then does it seem possible to embrace all of them. This is precisely what I wish to do.

The ground has now been cleared for pluralism. It is time to construct the theory proper.

Looking forward. There is something of a neat accident in the way that the next three chapters are structured. Each one develops a single core posit of pluralism: sensory awareness in chapter 3, sensory representation in chapter 4, and categorization in chapter 4. Each chapter also characterizes the nature of at least one kind of mental event: conscious experience in chapter 3,

perception in chapter 4, and hallucination and sensation in chapter 5. Finally, each chapter offers an account of a different kind of phenomenal character: first hard character, then strong singular character, and finally conveyed character.

Appendix 2A. Standard theories of the meaning of “what it’s like”

Here I consider how standard theories of the meaning of “what it’s like” would bear on the deep context-sensitivity view. Start with:

The *resemblance theory*: Any claim of the form “What x is like is ...” literally means, “What x resembles is ...,” or equivalently, “What x is similar to is ...”⁴⁹

The resemblance theorist offers a (more or less) unified semantic analysis of sentences such as the following:

- (1) What Chicago is like is London.
- (2) What Chicago is like is exciting.

The analysis of (1) is obvious enough: this sentence means that Chicago *resembles* London, i.e., that Chicago is *similar to* London. The analysis of (2) is what is interesting. This does not mean that Chicago is similar to exciting; that makes no sense. Rather, it means that Chicago is similar to *things* that are exciting, *in respect of being exciting*. Thus, (2) has the same truth-conditions as:

⁴⁹ See Gaskin (2019).

(3) Chicago is exciting.

But (2) and (3) are quite different in semantic structure.⁵⁰

The resemblance theorist should welcome the deep context-sensitivity view with open arms. For similarity-talk – which is how resemblance-talk is being understood – is itself deeply context-sensitive. For example, different claims of the form “India is similar to ...” will be appropriate in different contexts: the relevant similarities might be cultural, geological, economic, etc.

Indeed, I would encourage the resemblance theorist – or for that matter anyone else – to accept this broader view, which also covers uses of “what it’s like” that are not *F*-entailing:

The *generalized deep context-sensitivity view*: The expression “what x is like is ...” is deeply context-sensitive, where this context may be further specified with prepositional phrases and subordinate clauses.

For consider claims about what dogs are like. In some contexts, it is sensible to say, “I’ll tell you what dogs are like: they are like pigs.” Dogs and pigs are both intelligent, highly social animals that make excellent pets. But suppose that dogs bring great happiness to their owners. Then, in a positive psychology conference, it would be perfectly sensible to say, “I’ll tell you what dogs are like. They are like exercise: they strongly support long-term mental health.” Since pigs and exercise are radically different types of entities, even “what it’s like” talk that is not *F*-entailing is deeply context-sensitive. To streamline our discussion, however, I will primarily focus on *F*-entailing uses of what-it’s-like talk.

⁵⁰ Ibid, p. 21.

Now consider this theory:

The *predicative theory*: Any *F*-entailing claim of the form “what *x* is like is *F*” literally means that *x* has property *F*.

Some predicative theorists hold that the expression “what *x* is like is *F*” gets its literal meaning idiomatically.⁵¹ Others say that it gets its literal meaning compositionally – for example, that it means, “Experience *E* is similar to experiences that are *F* in respect of being *F* – i.e., *E* is *F*.” (Notice that this is a way of combining the predicative theory with the resemblance theory.)⁵² Still other predicative theorists have remained neutral between the idiomatic and compositional approaches.⁵³

There is also a close cousin of the predicative theory, namely:

The *restricted predicative theory*: Any *F*-entailing claim of the form “What *x* is like is *F*” literally means that *F* is a way that *x* is.⁵⁴

Ways are to be understood as properties of a special type: those that can be cited in response to questions about *how* something is, as opposed to *where* it is or *why* it is.⁵⁵

I observe that the predicative and restricted predicative theories would both support the deep context-sensitivity view. For, as I hope is evident by now, there is deep context-sensitivity in claims about (i) what properties India has and (ii) how India is.

⁵¹ See Hellie (2007b, p. 452).

⁵² See Gaskin (2019, p. 693).

⁵³ See Snowdon (2010, pp. 21-22).

⁵⁴ See Snowdon (2010, p. 22); Stoljar (2016, pp. 1169-1172, including fn. 14).

⁵⁵ See Stoljar (2016, fn. 14).

The deep context-sensitivity view would be equally supported by the following theory, which is an extension of the restricted predicative theory:

The *affective theory*: Any F -entailing claim of the form “What it is like to x for y to φ is F ,” when uttered in context c , literally means that F is a way that y ’s φ -ing affects x in c .⁵⁶

Consider claims about what is it like, to you, for you to have a dog. The affective theorist would understand these as claims about the ways in which your having a dog affects you. If we are discussing your physical condition, then you might say, “Having a dog affects me by being helpful – in particular, encouraging me to get a lot of exercise”; if we are discussing your mental condition, then you might say, “Having a dog affects me by being deeply meaningful”; if we are discussing your financial condition, then you might say, “Having a dog affects me by draining my bank account!” There is *some* broad ontological similarity to the properties of being helpful, being deeply meaningful, and draining one’s bank account, insofar as these are all ways that one entity can affect another. Still, I take it that these are properties of radically different types.

Finally, consider:

The *perceptual appearance theory*: Any claim of the form “what experience E is like for subject S is F ,” as used in philosophy of mind, literally means that F is a way that E perceptually appears to S .⁵⁷

⁵⁶ See Stoljar (2016, p. 1173). Stoljar’s actual proposal is about the expression “there is something it’s like ...” rather than the expression “what it is like,” but presumably this is how Stoljar’s proposal would be extended.

⁵⁷ See Lormand (2004, p. 315).

This theory is supposed to entail an inner-sense theory of phenomenal character – so the idea is that an inner-sense theory is *analytically* true.⁵⁸

If the perceptual appearance theory is true, then the deep context-sensitivity view, even if it is true, is irrelevant to philosophy of mind. For the deep context-sensitivity view is a view about *F*-entailing claims of a certain form, and the perceptual appearance theory says that there are no such *F*-entailing claims in philosophy of mind: the fact that an experience *perceptually appears* to be *F* does not entail that it *really is F*.

That said, I believe that there are serious problems with the perceptual appearance theory. For one thing, this theory has been criticized in detail by others, and I am sympathetic to these criticisms.⁵⁹ For another, it does not enjoy lines of support (a)-(d) and (f) described in §2. Finally, it is very hard for me to see how the theory can adequately account for the radically different uses of the term *phenomenal character* that I pointed out in the main text of this chapter.

Conclusion. The deep context-sensitivity view receives strong support from almost every theory of the literal meaning of what-it's-like talk. The major exception is the perceptual appearance theory. While this theory does not undermine the deep context-sensitivity view, it does render it trivial. However, I believe that there are strong reasons to reject the perceptual appearance theory anyway.

Appendix 2B. Some standard supplementary remarks

There are certain standard supplementary remarks that philosophers tend to make when they introduce the term *phenomenal character*. You might hope that these remarks suffice to pin down

⁵⁸ Ibid, pp. 316-317.

⁵⁹ See Hellie (2007b, pp. 443-444 and 449-459); Stoljar (2016, pp. 1186-1188); Gaskin (2019, pp. 676-680).

the relevant philosophical context(s) – but in this section I will briefly argue that this is not the case.

One standard tactic is to mention a few examples: to say that there is something it's like to see purple, to hear a cello, and so on.⁶⁰ However, this popular supplement provides very little of the missing information about the context(s) that philosophers have in mind. For notice: if you ask someone what it's like to see purple, she may truly answer, "It's simply *wonderful*." But, again, it is widely agreed that emotional states or properties are not themselves phenomenal characters of visual experiences. (But see Appendix 2C for a discussion of a much richer definition along these lines.)

Another common way to fill in some missing contextual information is to connect the term *phenomenal character* to other terms in the neighborhood: perhaps *subjectivity*, *awareness*, or *consciousness*.⁶¹ But this is not sufficient. For one thing, these terms are wildly polysemous – for instance, it has been said that the term *consciousness* has more than half a dozen distinct meanings.⁶² And in any case these terms still do not supply enough of the missing contextual information. For, again, if you ask someone what it's like to have a particular episode of *subjective conscious awareness*, then she can truly answer by reporting her emotions.

Thus the standard supplementary remarks do not suffice to pin down the philosophical context(s) that determines the reference of "what it's like talk." Consequently, they leave out some crucial information that is required to settle the meaning of the term *phenomenal character*.

⁶⁰ This practice is so ubiquitous that I trust that citations are unnecessary.

⁶¹ For example, Chalmers uses the "what it's like" definition of consciousness and says that "[t]he subject matter [of a discussion of consciousness] is 'the *subjective* quality of experience'" (1996, p. 4, emphasis mine); Tye equates phenomenal aspects with what-it-is-like aspects and adds that "a mental state is phenomenally conscious just in case there is some immediate *subjective* 'feel' to the state" (1995, p. 4, emphasis mine); Sturgeon (2008) explains that "the conscious character of Good cases" is "what it is *consciously* like to enjoy Good visual experience" (p. 116, emphasis mine).

⁶² Lycan (1996, ch. 1) identifies eight.

Appendix 2C. A definition by example

In Appendix 2B, I argued that giving a few examples does not suffice to pin down any typical philosophical meaning of the term *phenomenal character*. However, it has recently been suggested that extensive examples might do the trick.

The idea is to start with a long list of positive examples: seeing a computer screen; visually experiencing phosphenes; hearing the hum of a fan; feeling the pain of a headache and the position of your limbs; vividly imagining the Eiffel Tower and the tune of “Happy Birthday”; feeling joy, surprise, and sadness; thinking “what a jerk!”; craving a slice of chocolate cake; and dreaming. Add to this a list of negative examples: being disposed to say that $6 \times 4 = 24$; having a standing intention to stop for lunch at 11:45; sleeping dreamlessly; and so on. Then we can define *phenomenal character* as the thunderingly obvious feature that the positive examples have and the negative examples lack.⁶³

I think that this definition is a perfectly legitimate one. However, it does not have any overt connection with “what-it’s-like” talk, so there is a real worry that to work with this definition is to change the subject. In the end, I suspect that this worry can be rebutted: I suspect that *phenomenal character*, as defined by example, co-refers with my term *hard character*. But it would take substantial work to show this.⁶⁴

There is no such worry for my definition: I have argued from the ground up that *hard character* has the same meaning as one type of usage of *phenomenal character*.

⁶³ See Schwitzgebel (2016, esp. §2). Schwitzgebel is defining the term *phenomenal consciousness*, but it is worth seeing whether his definition would work just as well for *phenomenal character*.

⁶⁴ For what it is worth, there is a further complication: *phenomenal character* as defined by example arguably differs in *meaning* (even if not reference) from *phenomenal character* as defined in the context of the hard problems.

Appendix 2D. A concern about my interpretation of Campbell

A perceptive reader might have noticed that the expression “what it’s like” does not appear in any of the passages from Campbell that I quoted in §4. Perhaps, then, Campbell is simply defining the term *phenomenal character* directly, without appeal to what it’s like to have an experience. Would this affect any of my central conclusions?

It would not. To begin with, the quotations would still show that when Campbell uses the term *phenomenal character*, it is strong singular character that he has in mind. We should continue to conclude that Campbell is talking past philosophers like Dretske and Tye.

In addition, the present concern still allows the deep context-sensitivity hypothesis to play a crucial explanatory role. To see what this role is, note that it is somewhat puzzling that Campbell ends up talking past Dretske and Tye. Campbell is plainly familiar with their work. Why then would he fail to understand how they are using the term *phenomenal character*? The deep-context sensitivity hypothesis suggests an answer. Campbell was immersed in certain philosophical problems regarding perception before he encountered the term *phenomenal character*. Thus, when Campbell encountered the standard “what it’s like” definition of phenomenal character, he assumed, mistakenly but quite reasonably, that his interlocutors were speaking in a context in which such problems were salient.

Appendix 2F. Are hard character, strong singular character, and conveyed character really different?

I have argued that some philosophers have used the expression *phenomenal character* with hard character in mind; others have had in mind strong singular character or conveyed character. But are these really three different properties?

The pluralist theory will turn out to entail that they are. But even near the start of inquiry, there is some reason to think that there are three distinct properties here. For there will be a strong tendency for users of a deeply context-sensitive expression, such as *this property*, to have different things in mind in different contexts.⁶⁵ This general observation is reinforced when we notice the stark differences among the three contexts in which the expression “what it’s like” has been used. The notions of hard character, strong singular character, and conveyed character were introduced to do very different kinds of theoretical work. So there is no particular reason to think that any pair of these will be identical. They *might* be, but this would be a fluke.

But what if these three characters are all identical – what consequences would this have? I say that there would still be an important rational defect in philosophical disputes about phenomenal character.

Suppose by analogy that you and I disagree about who “that person” is. Neither of us realizes that we have in mind different contexts of demonstration: I assume that we are both thinking of the person at the bar wearing a trench coat, while you assume that we are both thinking of the person who gave a philosophy talk earlier today. Now, as it happens, the person in the trench coat *just is* the person who gave a philosophy talk earlier today – but neither of us knows this.

In such a situation, even though you and I will in fact be talking about the same person, our discussion will suffer from an important defect: we will fail to give *dialectically effective*

⁶⁵ There will be a weaker tendency of the same sort even for shallow context-sensitive expressions like *here*, *rich*, or *this person*.

arguments to support our claims. For instance, I might show you a picture of the person at the bar, and I would be right to treat this as evidence about the referent of my expression “that person.” But, absent evidence that the person at the bar is the one who gave the philosophy talk, you would rationally conclude that I have not yet given you evidence about the referent of your expression “that person.”

I conclude that, at a minimum, debates about phenomenal character have been rationally defective. More than this, philosophers who use the expression *phenomenal character* have most likely had three very different properties in mind.

Appendix 2E. A fourth use of “phenomenal character”: Susanna Schellenberg

Consider the prominent contemporary representationalist Susanna Schellenberg, whose 2018 book includes a detailed theory of “phenomenal character.” What does she have in mind by this?

Schellenberg barely discusses the hard problems of consciousness. Thus she does not have hard character in mind. Nor does she seem to have strong singular character in mind. She discusses in detail why perceptions position the subject to *make* singular reference to what is perceived, but she never mentions a special kind of *knowledge* of what the subject is thereby referring to. In addition, Schellenberg eventually concludes that phenomenal character is the same across perceptions and matching hallucinations, but, as I argued in §6, this is a very implausible claim to make about strong singular character. Finally, it seems not to be conveyed character that Schellenberg has in mind. For Schellenberg’s theory is that the phenomenal character of a perceptual experience consists of perceptual capacities to discriminate and single out particulars. But surely what a perceptual experience conveys to its subject is not a *perceptual*

capacity – or at least not *just* that. Surely the experience (also?) conveys its content, or perhaps the entities that are perceived.

Thus Schellenberg must have some fourth thing in mind. What might that be?

Although Schellenberg’s systematicity suggests that she has *something* firmly in mind, I do not quite know what that is. For Schellenberg gives *no* definition, not even the what-it’s-like definition, of phenomenal character.⁶⁶ The closest that she comes is to suggest in passing that phenomenal character is identical to consciousness (p. 7) – but given that the expression *consciousness* has many meanings⁶⁷, this is not especially helpful.

Still, we can at least conclude that Schellenberg uses the expression *phenomenal character* with yet a fourth referent in mind.

Appendix 2G. An alternative interpretation

There is another interpretation of the situation. Perhaps philosophers like Nagel, Dretske, Tye, Campbell, and Siegel are all using the expression *phenomenal character* in exactly the same way. Perhaps the differences in background context that I have identified are meant only to highlight the different *aspects* of phenomenal character that are relevant for their purposes. If so, then these philosophers are not talking at cross-purposes after all.

I believe that this interpretation is not sustainable.

Nagel was the seminal figure for popularizing “what it’s like” talk in discussions of consciousness, and I have argued that what he had in mind was hard character. Thus, on the alternative interpretation, Campbell is emphasizing a certain *aspect* of hard character: the aspect

⁶⁶ Are my arguments undermined by the fact that Schellenberg does not use the “what it’s like” definition? No. The points from Appendix 2D apply with the obvious modifications here.

⁶⁷ See Lycan (1996, pp. 2-7).

that helps to explain why perceptions position us to recognize perceived objects as mind-independent, to know the reference of perceptually-based demonstratives, and to recognize perceived objects as the categorical grounds for various dispositions. Similarly, the alternative interpretation has it that Siegel is emphasizing yet another aspect of hard character: the aspect of what is conveyed to the subject.

We should then expect to find *some tangential evidence* that the properties of experience that Campbell and Siegel have in mind are related to the hard problems of consciousness. Somewhere, at least in passing, Campbell should indicate that strong singular character is an aspect of hard character, that it is deeply entangled with the hard problems. But look at his lengthy discussion of phenomenal character in chs. 6-7 of *Reference and Consciousness*. Pay special attention to the quotations I cited in §4 above, where Campbell is describing the very purpose of the concept. You will find nothing of the sort.

Likewise for Siegel: *somewhere* across the five works of hers cited in the footnotes of §5, she should drop some hint that the notion of phenomenal character is deeply bound up with the hard problems of consciousness. She does not.

I conclude that Campbell and Siegel really are speaking at cross-purposes with Nagel, Dretske, and Tye.

Appendix 2H. On the existence of hard character

Some theorists have argued that phenomenal character, taken as hard character, does not exist, because there are no properties of experiences that play enough of roles 1-4. We can now understand why this argument is unsound.

The argument. According to the resemblance theory, any claim of the form “What x is like is ...” literally means, “What x resembles is ...,” or equivalently, “What x is similar to is ...” But

consider a sentence like, “What Chicago is like is exciting.” This does not mean that Chicago is similar to exciting; that makes no sense. Rather, it means that Chicago is similar to *things* that are exciting, *in respect of being exciting*. Thus, this sentence has the same truth-conditions as the sentence, “Chicago is exciting.” But these sentences are quite different in semantic structure.⁶⁸

Assume that this theory is true. Then talk about what an experience is like is just talk about what the experience is similar to. This talk is entirely unmysterious; it does not in any way support the idea that there are any hard problems of consciousness. Thus, there is no particular reason to think that there are any such problems – or so the argument goes.⁶⁹

The problem. For the sake of discussion, grant the truth of the resemblance theory.⁷⁰ Then, yes, there is nothing *in the meaning of what-it’s-like talk* to suggest that there are any hard problems of consciousness. But from this, we should not infer that *nothing at all* supports the existence of the hard problems. There is plenty of support!

It comes from claims such as these: first, that there are certain facts about (for example) seeing red that initially seem impossible to fully explain in physical/functional terms; second, that there are certain properties of a bat’s echolocatory experiences that initially seem impossible for us to know about; third, that we initially seem to be able to conceive of a subject who is physically and functionally just like us when we see red, but whose experience of red lacks a certain crucial property of our own experience; and fourth, that we initially seem to be able to conceive of a subject who is physically and functionally just like us when we see red, but whose

⁶⁸ See Gaskin (2019, p. 693).

⁶⁹ Ibid, p. 695.

⁷⁰ A similar argument might be run with the predicative theory, which, along with the resemblance theory, is discussed in Appendix 2A.

experience of red is similar, in a certain crucial respect, to our experience of green. To put it another way, it is just plausible *from the start* that there are hard problems of consciousness.⁷¹

The mistake of such a resemblance theorist is to ignore the deep context-sensitivity of the expression *what it's like* – in particular, to look for information about the hard problems of consciousness by scrutinizing the expression *what it's like*. This is as wrongheaded as looking for information about blueness by scrutinizing the expression *that property*. Such information can only be found in the context: it is only when (say) Thomas points to a blue circle *in a context in which colors are salient* that the expression *that property* comes to refer to blueness. Likewise for *what it's like* and hard character.

The resemblance theorist's argument against the existence of hard character might be restored if she had any special reasons to resist the deep context-sensitivity hypothesis. In Appendix 2A, I argue that quite the opposite is true: resemblance talk is deeply context-sensitive, so resemblance theorists have even more reason to accept the deep context-sensitivity hypothesis.

Appendix 2I. Nagel's characterization of consciousness

Here is a mildly revised version of Thomas Nagel's famous characterization of consciousness:

⁷¹ My opponent might respond that when she considers Nagel's bat, Jackson's Mary, etc., she sees nothing that seems even initially difficult to explain in physical and functional terms. This response strikes me as denying the obvious. But I can offer one parting observation. I have taught several courses on the hard problems of consciousness, and on the very first day of class, I introduce the topic by mentioning the standard examples (of Nagel's bat, of what Mary seemingly could not know, etc.). I then take a poll to see whether the hard problems really do feel hard to my students. Very close to 100% of them agree that they do. I can confirm on the basis of many other polls that students are not reluctant to disagree with me or with each other. This suggests that there really is something that *strongly seems* to be difficult to explain.

(C) For a state to be *conscious* is for there to be something it's like for a subject to be in that state.⁷²

According to the *restricted predicative theory*, any *F*-entailing claim of the form “What *x* is like is *F*” literally means that *F* is a way that *x* is. *Ways* are to be understood as properties of a special type: those that can be cited in response to questions about *how* something is, as opposed to *where* it is or *why* it is.

Some restricted predicative theorists have objected that (C) cannot be right. For it is trivial to satisfy the right-hand side: what-it's-like talk refers to the ways that the entity under discussion is, and every entity is some way or other.^{73,74} We are now in a position to appreciate why this objection is not a good one.

For argument's sake, grant the truth of the restricted predicative theory. Even so, given the deep context-sensitivity view, this reasoning is clearly mistaken. Yes, *in principle* what-it's-like talk can be used to refer to any way that an entity is. But *in context* that is not so. Nagel was speaking in the context of a discussion of a hard problem of consciousness: he was discussing the

⁷² Nagel's original characterization is that “an organism has conscious mental states if and only if there is something that it is like to be that organism” (1979, p. 166). I think that the reference to organisms is misplaced. We should leave it conceptually open that subjects other than organisms (robots, planets) might have conscious states.

⁷³ As Snowdon puts it: “My hypothesis, then, is that when people talk of what an experience was like they are using ‘like’ in this non-comparative way, where it simply means being some way.... If this is what [the what-it's-like characterization of conscious states] means, then it does not characterize experiences in any way that would not apply to *absolutely everything*. Everything is some way or other” (2010, pp. 22-23, emphasis his). A similar concern is raised by Hellie (2007, p. 461, fn. 13).

⁷⁴ You might think that the expression “for a subject” renders Nagel's definition non-trivial. After all, not every property of a state characterizes what it's like *for a subject* to be in that state. Fair enough! Still, Nagel's opponent might justly reply that it is at least too easy to satisfy the right-hand side of the definition. There is something it's like for a subject to be in a state of poverty (“degrading”), but poverty is not a conscious state.

properties of bat experience that seem to elude explanation in purely physical and functional terms. So we can rephrase (C) roughly as follows:

(C*) Let F be the property of a bat's experience that makes it seemingly impossible to fully explain the experience purely in physical and functional terms. For a state to be *conscious* is for it to have F .

Whether or not this characterization is correct, it is certainly not *trivial*. And, for what it is worth, I find the characterization very initially plausible. Conscious experiences quite generally appear to elude explanation in purely physical and functional terms. Isn't it initially plausible that there is a single property F that explains this across the board?

The restricted predicative theorist might be able to restore her argument if she could resist the deep context-sensitivity hypothesis. I argue in Appendix 2A that such resistance would be unwise: talk of ways is quite generally deeply context-sensitive, so she in fact has an extra reason to accept the deep context-sensitivity hypothesis.

Appendix 2J. A re-evaluation of an argument from Campbell

Consider a veridical perception of a mahogany table and a matching hallucination – call these the *table-perception* and the *table-hallucination*, respectively. Dretske and Tye say that the table-perception and the table-hallucination have precisely the same “phenomenal characters.”⁷⁵

⁷⁵ See Tye (1995, p. 152); Dretske (1995, p. 101). Similar utterances may be found in Alston (1999); Chalmers (2006); Pautz (2007); Schellenberg (2018); D'Ambrosio (2019). Qualified support is also expressed in Siegel (2010, ch. 6); Brewer (2011, p. 101).

Campbell says that their “phenomenal characters” are radically different.⁷⁶ As I have said before, Campbell makes it explicit that he takes himself to be disagreeing with Dretske and Tye on this point.^{77,78} I say that this is a merely verbal dispute, one that has obscured the shape of the real dispute. In this appendix, I further substantiate this charge: I argue that Campbell’s misunderstanding causes him to mislocate his disagreement with these opponents.

According to the *representationalist* theory of perception defended by Dretske and Tye, perceptions are nothing more than representations of a certain distinctive kind; non-representational relations are not involved in the analysis of what a perception is.⁷⁹ According to Campbell’s *naïve realist* theory of perception, perceptions are nothing more than non-representational relations of awareness of a certain distinctive kind; representations are not involved in the analysis of what a perception is.⁸⁰ This is a real, substantive dispute between Dretske and Tye, on the one hand, and Campbell, on the other.

One of Campbell’s most central arguments against Tye-Dretske representationalism may be regimented as follows:

C1. The phenomenal character of the table-perception helps to explain why it positions the subject to make knowledgeable reference to the table.⁸¹

⁷⁶ See Campbell (2002, p. 117). Similar utterances may be found in Martin (2004) and (2006); Soteriou (2005); Fish (2009); Logue (2012); Allen (2015); Steenhagen (2019).

⁷⁷ Again, see for example Campbell (2002, p. 146).

⁷⁸ Alternative views are possible. For instance, there is the view that their phenomenal characters are the same in some significant respects, but also different in some significant respects. See Beck (2019) for a view *somewhere* in this neighborhood.

⁷⁹ Tye (1995, ch. 5); Dretske (1995, ch. 6).

⁸⁰ See Campbell (2002, esp. chs. 6-7).

⁸¹ See the quotations that I cited in §4.

C2. If C1 is true, then the phenomenal character of the table-perception is constituted partly by a table.⁸²

C3. The phenomenal character of the table-hallucination is not constituted partly by a table.⁸³

C4. If the phenomenal character of the table-perception is constituted partly by a table objects and the phenomenal character of the table-hallucination is not constituted partly by a table, then the table-perception and the table-hallucination differ in phenomenal character.⁸⁴

C5. If Tye-Dretske representationalism is true, then the table-perception and the table-hallucination do not differ in phenomenal character.⁸⁵

C6. Tye-Dretske representationalism is not true.

The argument is valid, and Campbell treats almost all of the premises as obviously true. The exception is C2: Campbell offers a detailed defense of this premise that spans much of chapters 6-7. For our purposes, however, the crucial premises turn out to be C1 and C5.

Suppose that we interpret *phenomenal character* as hard character. Then C5 would indeed be indisputable: it would express a central commitment of Tye-Dretske representationalism. However, C1 would enjoy very little support! The hard character of an experience is roughly the natural kind property that generates the hard problems of consciousness. *But the fact that a property of experience plays roles 5-7 has no evident bearing on whether it generates the hard problems.*

Suppose that we instead interpret *phenomenal character* strong singular character. Then C1 is true by definition – but C5 no longer expresses a commitment of Tye-Dretske

⁸² See for example Campbell (2002, pp. 121, 130, 147-150).

⁸³ This is granted by nearly all parties – even sense-datum theorists, if the objects in question are understood to be ordinary, mind-independent ones.

⁸⁴ See Campbell (2002, p. 117).

⁸⁵ Ibid, p. 117. See p. 146 for Campbell's citation of Dretske and Tye.

representationalism. Tye and Dretske *are* committed to the claim that, for any perception and matching hallucination, there is a common property that explains why both experiences generate the hard problems of consciousness. That is, they are committed to the claim that the table-perception and the table-hallucination have the same *hard* character. But this has no obvious bearing on whether these experiences have the same *strong singular character*. Perhaps Tye and Dretske will say that the strong singular character of the table-perception is to be explained in terms of a special kind of causal link to the perceived table and its instances of brownness and table-shapedness. In the case of the table-hallucination there is no such table, no such instances of brownness and table-shapedness, and no such causal link. Hence this experience might have a radically different kind of strong singular character from the perception.

On either interpretation, then, Campbell's argument is unsound.

That is not the real lesson, however. For there is an obvious way to repair Campbell's argument: Campbell should argue that no representationalist account of strong singular character will be adequate. Many of Campbell's arguments against representationalism can be charitably reconstructed along these lines, and I engage with such reconstructions of Campbell's arguments in later chapters. The real lesson is simply that we must skim off the purely verbal dispute before we can even identify where Campbell disagrees with Tye and Dretske, much less assess who is in the right.

Appendix 2K. An objection to naïve realism reconsidered

It has frequently been said that naïve realists have great difficulty accommodating the very plausible claim that any perception and any hallucination have the same phenomenal character. Here I argue that this is not a good objection as it stands. I illustrate the point by focusing on Campbell's version of naïve realism.

Campbell thinks that the strong singular character of a perception consists partly of a relation to the ordinary objects and properties that are perceived. But the subject is not related to any such objects or properties in a matching hallucination. This commits Campbell to the conclusion that the perception and the hallucination have different *strong singular characters*.

But, as I have argued, this is not the sense in which it is plausible that these experiences have the same “phenomenal characters.” What is plausible is that these experiences have the same *hard characters*. And it is perfectly possible for Campbell to accommodate this view. To pick a pair of views arbitrarily, perhaps the hard problems of consciousness are generated by the neural properties that underpin a given experience,⁸⁶ or alternatively by perceptual capacities to single out and discriminate particulars.⁸⁷ These neural properties or perceptual capacities could be precisely the same across a perception and a matching hallucination, resulting in identical hard characters. None of these claims are in any obvious tension with Campbell’s theory of strong singular character, nor with his naïve realism more generally.⁸⁸

Ultimately, I believe that there are further problems for these alternative naïve realist approaches. My point is simply that much more would be required to press the objection home.

⁸⁶ See Block (1996). This approach might pair well with the *neurocomputational naïve realism* developed in Beck (2019a) and (2019b).

⁸⁷ This view is inspired by Schellenberg (2018).

⁸⁸ The point does not generalize to every version of naïve realism: Fish (2008) offers a naïve realist theory of the hard character of perception that cannot be generalized to matching hallucinations. Still – and quite remarkably – the point applies to *many* versions of naïve realism.