Sensation, Introspection, and the Phenomenal

Jonathan Ellis

[This is a late but not final draft of my chapter in *Wittgenstein and the Philosophy of Mind*, eds. J. Ellis and D. Guevara (New York: Oxford University Press, 2012).]

1. Introduction

There is perhaps no notion more central to debates in contemporary philosophy of mind than that of phenomenal character. I use the expression "phenomenal character," but that is merely one expression of many that I might have chosen to indicate the notion I aim to introduce. I could have just as well used "phenomenal qualities," "phenomenal properties," "qualitative feels," "qualitative character," "raw feels," "what-it's-like properties," "qualitative properties," "immediate subjective qualities," "subjective experiential properties," "the overall feel," "the very feel," "sensational qualities," "phenomenal consciousness," "sensational components," or even "qualia," all of which are used in the debates I wish to discuss. These terms are sometimes used interchangeably, sometimes not. And while each of them brings its own idiosyncrasies, the thoughts behind them are similar. I shall use "phenomenal character" and "phenomenal qualities" primarily, but much of what I say will apply equally to philosophical lines of thought expressed in the other terms.

What is phenomenal character? I don't mean to ask what theory we should give of it, but rather, what are we talking about when we talk about phenomenal character? This question is in fact part of what I shall be discussing, but for the moment, note that philosophers often introduce what they have in mind by giving examples—what it's like to bite into a chocolate bar and taste

the chocolate, what it's like to experience a sunset, or to have an itch on your forearm, and so on: what it's *like* to have these experiences. Sometimes, in order to excuse their not saying more to identify their topic, philosophers then follow Ned Block in appealing—only half-jokingly, it seems—to something Louis Armstrong said. When asked "What is Jazz?" Armstrong said "If you got to ask, you ain't never gonna get to know" (Block 1978). I don't fault these philosophers for employing this method in trying to identify their topic. When reading these introductions, I myself typically think I know exactly what they're talking about—as they clearly think they do too—even if we are not yet sure what philosophical account to give of it. But I think identifying the topic may be more problematic for philosophers than we realize.

Let me begin by introducing the two most prominent philosophical accounts of phenomenal character today. The first account, sometimes called "the standard view" (Loar 2003), holds that the phenomenal character of a sensation (or experience, etc.) is constituted by (or explained by, or otherwise involves) qualities of the sensation that are both "intrinsic" to the sensation and introspectively accessible (see Peacocke 1983; Boghossian and Velleman 1989; Block 1990, Searle 1992; and Loar 1997). The phenomenal character of my headache, for instance, is constituted by properties of the headache that are intrinsic to the headache as well as accessible to me in introspection. To say that a property of something is intrinsic to it, in this debate, is typically to say that the thing has that property "in itself," apart from its relations to other things. Being a brother is not an intrinsic property; having a particular mass arguably is.[1] The chief motivation for the standard view can be traced to the quite natural thought that when we introspect our sensations (and experiences, etc.), we are able to inspect or observe them. Our awareness of them involves their presence, their presence to our awareness.[2]

The standard view, however, is thought to provide substantial problems for physicalist conceptions of the world. In order for a physicalist to account for an intrinsic property of something, it is said, that property must be identifiable with—"nothing but"—one or another intrinsic *physical* property (a particular neurophysiological property, say). And many philosophers find it difficult to understand how phenomenal qualities could be identical to, nothing but, intrinsic physical properties.

As a result, many physicalists reject the standard view. The opposing view that has received by far the most attention in the last two decades—and here is the second of the two views I shall mention—is "representationalism." According to representationalism, the phenomenal character of a sensation is identical to (or exhausted by) particular representational properties of the sensation. The throbbing aspect of my headache, for instance, is said to be nothing but, or nothing over and above, the property of the headache whereby it represents my head to be throbbing. These representational properties are not intrinsic to the mental phenomena of which they are properties; rather, those phenomena represent what they do in virtue of the relation those phenomena stand in to other things. According to representationalists, there are no properties of experience that are both intrinsic and introspectively accessible.[3]

Representationalism too, however, is thought to face significant difficulties. Many philosophers highlight mental phenomena that are allegedly identical in representational content but different in phenomenal character (Peacocke 1983; Block 1996), or that have phenomenal character but no representational content at all (moods, for instance, such as depression or elation). Another influential problem for representationalism is that, conjoined with a standard "externalism" about representational content, representationalism appears to lead to externalism about phenomenal character: physically identical individuals could have sensations with different

phenomenal qualities. But this view strikes many philosophers as extremely implausible. And there are further difficulties.[4]

While these two views—the standard view and representationalism—do not exhaust the views about phenomenal character in currency, they certainly garner the most attention and support today. They are the sorts of views that Block had in mind, when he said—in another passage that is quoted frequently—"The greatest chasm in the philosophy of mind —maybe even all of philosophy— divides two perspectives on consciousness. The two perspectives differ on whether there is anything in the phenomenal character of conscious experience that goes beyond the intentional, the cognitive and the functional" (Block 1996, 19). Representationalists think that there is not. Proponents of the standard view think that there is.

What I want to do in this chapter is to introduce a third perspective—or really, to reintroduce to contemporary philosophy of mind a third perspective. It is an approach that, in many ways, I find in Wittgenstein. But this is not an interpretive paper, and there will certainly be many interesting questions as to the extent to which Wittgenstein would approve of what I say, and the way in which I say it.

There is an assumption guiding so much philosophy of mind today that in order for a broadly "physicalist" conception of the world to be correct (more on physicalism in a moment) every sensation, experience, emotion, and so on, that has phenomenal character—every "bearer" of phenomenal character—must be identical to, nothing but, a particular physical phenomenon (a physical process, a physical state, property, event, etc.), that is, that every bearer of phenomenal character must be identical to a phenomenon that is fully specifiable in physical terms. The attendant idea is that if bearers of phenomenal character are not identical to physical phenomena, we will be saddled with one or another form of dualism, and one that threatens our contemporary

physicalist conception of the world. Thus, physicalist proponents of what I called "the standard view" identify sensations and the like with physical processes or states of the body (mostly or entirely of the brain), processes specifiable in the terms of (say) neurophysiology. And they identify the phenomenal qualities of those phenomena with intrinsic physical properties, also specifiable in such terms. Representationalists too identify bearers of phenomenal character (e.g., sensations) with physical processes or states of the body. This is true even of representationalists who are externalist about representation, as many representationalists are, and thereby hold that many or all bodily states have the representational properties they do in virtue of the relations those bodily states stand in to things outside the body.[5]

I want to suggest that this guiding assumption—that physicalism depends on every bearer of phenomenal character being identical with a physical phenomenon, one fully specifiable (at least in principle) in physical terms—too often goes unscrutinized, and that there is considerably more reason to doubt it than is typically appreciated. Indeed, I think the assumption is false, and that it is based in a fundamental misconception of those mental phenomena that are taken to have phenomenal character.

I should make clear, my interest in arguing for these conclusions stems not so much from an interest in defending physicalism. In the philosophy of mind, physicalism is often conceived as a view about supervenience: no two possible worlds can be identical in their physical properties but differ in their, for example, mental, social, or biological properties.[6] However, the difficulties in specifying in an adequate way what is meant by the expression "physical" are substantial and sometimes dismissed too quickly as details that will ultimately be finessed. What proves difficult is to define "physical" in a way that makes the thesis of physicalism at once plausible and not vacuous (see, e.g., Hempel 1970 and Crane and Mellor 1990).

At any rate, my primary interest in this chapter is to expose what I see as the fundamental, albeit very seductive misconception of sensation (and experience, and so forth) responsible for the idea that the thesis of physicalism (if it can be made out) is threatened by such mental phenomena. Because so many debates in the philosophy of mind are motivated by concerns about physicalism, in the interest of making this point I will place to the side concerns about the notion of the physical and proceed on the basis of the assumption that an adequate specification of the view (one that by and large captures the spirit of the view as it is discussed in contemporary debates in the philosophy of mind) is in the offing. My primary goal in the chapter, then, is to sharpen and make plausible a particular form of what I shall call "Non-Identificatory Physicalism" (NIP). According to this view, bearers of phenomenal character are not identical to physical phenomena, yet physicalism is true: most relevantly for our purposes, any world physically identical to ours would not differ in its sensations, experiences, and the like. I will take sensations as my primary example, but much of what I say, it should become clear, will go for other bearers of phenomenal character (experiences, emotions, and so forth).[7]

How could physicalism be true, and people have sensations, yet sensations not be identical to physical phenomena? What *are* sensations then? What is it "in virtue of" that we have them? And what is it in virtue of that those sensations have particular phenomenal qualities? One goal of the chapter is to sketch a way of approaching these questions that makes plausible a NIP account of sensation. There are a variety of ways in which a NIP account of sensation might handle these questions, many of which are not promising. Adverbialism, expressivism, eliminativism, quietism—all of these are views which, on at least some understandings of them, could be construed as compatible with NIP. I do not endorse any of these views—or at least, I will not characterize the approach that I entertain in any of these ways. One of my chief goals in

the chapter is to bring one version of a NIP view into clearer focus, and to do so in a way that preempts the kinds of misunderstandings and caricatures it can elicit.

I will spend considerable space on questions of the sort: "What is it *in virtue of* that we have sensations?" "In what do sensations (or phenomenal character) consist?" These are standard questions in the philosophy of mind. In the introduction to his influential volume on pain, for instance, Murat Aydede writes, "The question is: in what does the painfulness, the hurting quality of pains consist?" (2005, 28). To anticipate, I shall be questioning the expectations and requirements with which philosophers pose such questions. A rejection of these expectations and requirements is not incompatible with physicalism either.

The general idea is of course not new. It is the sort of broad strategy that is often associated with Wittgenstein—rightly or wrongly: Such-and-such question or philosophical project is misconceived; it is misguided to try to pursue it, no less answer or fulfill it; one needs to understand why one was led to pursue it in the first place. That general approach is met with such resistance nowadays, especially when it is applied to phenomena such as sensations or phenomenal character, which it rarely is.

Of course, one can certainly find these, or at least related, ideas in Wittgenstein scholarship. My point is that they're almost entirely absent from mainstream philosophy of mind. That is to say: they're absent from the sorts of discussions that appear in our most reputable journals and at major conferences like those of the American Philosophical Association. The view is virtually not even on the table.

2. Road Map

In order to sharpen and make plausible the particular kind of NIP view that I favor, I must first introduce and analyze what I take to be the most powerful sources of the conception, or conceptions, of sensation that would seem to be in blatant tension with a NIP treatment. This will be the focus of sections 3 through 5, in which I aim to identify, and then provoke suspicion about, some very influential lines of thinking.

In section 6, I will take stock and make explicit some pressing questions. Then, in section 7, I will underscore some dialectical features of the larger debate, and in particular, the relationship between representationalist theories of phenomenal character and the lines of thinking that ground resistance to the NIP approach. Another important theme of the chapter is that a good deal of what has historically motivated the representationalist approach to sensation and phenomenal character is no less supportive of a NIP approach. And the NIP proposal I will advertise does not face the substantial difficulties thought to beset representationalism.

In sections 8 and 9, I will elucidate the NIP approach I am proposing. I finally return in section 10 directly to the notion of phenomenal character. Once we see why sensations might receive a NIP treatment, it will be a short step to see why phenomenal character might as well. Another possibility I will entertain, though, is abandoning that notion (and its cognates) altogether.

3. Sources of Resistance and the Price of Milk

One way to begin to uncover the source(s) of the formidable resistance that any NIP approach to sensation will surely face is to give an example of something else to which a NIP view would

apply. Doing so will also make more explicit how I am understanding what a NIP account says of a particular phenomenon and what it does not say.

The example I will give here is an example of something that is, by all means, very different from sensations and experiences, but it will be instructive nonetheless. I beg my reader's patience in making assumptions as to which aspects of the following example I wish to claim are also aspects of the case of sensations and which are not. Once I have introduced and elucidated the example, I will make explicit precisely what I hope to do with it, and what I do not intend.

Suppose I am at the market with a gallon of milk in my hand. In particular, consider the price of the gallon of milk. The price tag on it says "$2.79." And it is correct; this is what the grocer is charging today for a gallon of milk. Suppose that, standing in the market, holding this gallon of milk, I begin to wonder to myself about the metaphysics of price. "What is price?" I ask myself. What does it consist in? What is it identifiable with?

My friend, who overhears me muttering to myself, might perhaps say, "You're wondering what price is? A thing's price is the amount of money (or goods, etc.) one has to give in order for it not to be illegal to take it [or something like that]." "Yes," I say, "I know what the word 'price' *means*, I'm wondering what the price *is*—you know, what it is identifiable with, or what it consists in." My friend's initial answer has satisfied my initial question no more than a seventeenth-century chemist would have been satisfied to be told, in response to her question "What is water?": "You know, water is that liquid that is in the oceans and lakes and that comes out of faucets." The chemist knows that water is that stuff; she wants to know what that stuff, as it were, *is*. And at any rate, the price couldn't be identical to (or consist of, or be nothing but) a required amount of money, for the milk *has* its price but it does not have that amount.[8]

"In fact," I go on, "I even know what 'makes it true that' something has the price it does, or, if you prefer, 'in virtue of what' it has its price. Or at least I know roughly the sorts of things that do. Take the gallon of milk in my hand. What makes it the case that its price is $2.79—indeed, what makes it the case that it has any price at all—is a very complicated set of facts or state of affairs involving social customs, human attitudes, behavioral dispositions, and so on (e.g., the grocer's disposition to chase me or call the police if I were to take the milk without paying for it). But still [I continue] those facts are not the price. The price is not identical to that collection of attitudes, dispositions, and customs (nor to facts about them), even if those things are in virtue of what the milk has the price that it has. The price, after all, is $2.79. And it would be incorrect to say that that entire collection of customs, attitudes, or facts is itself $2.79."

The correct thing to say in this case, I believe, is that price is not identical with any phenomenon (object, property, state, etc.) fully specifiable in physical terms. Price is something we ascribe to things, and correctly, but prices cannot be identified with—said to be nothing but—any particular physical phenomenon. But for all that, prices provide no problem for physicalism.[9] It is quite compatible with this nonidentificatory approach to price to think that no two worlds could be the same in all physical respects yet diverge in its distribution of prices or economic properties.[10]

Some might elaborate this view of price by saying that price is not a *thing* at all, or not an "object," but I think we should be wary of characterizing matters in this way. I will return to this issue in section 9. One of my intentions here is simply to provide an example of something for which it is correct to give a NIP account.In fact, there are many such things for which it is correct to give a NIP account. The same sorts of things could be said, I think, of some or all of the following: obligations, future appointments, debts, deadlines, absences, scarcities, sakes (as

in, for Paul's sake),[11] prospects, errands to run, best interests (as in Paul's best interests), schedules, and free time. It would be misguided, I suspect, to try to identify any of these with phenomena specified fully in physical terms. We might quibble about some of these cases; nothing hangs on whether a NIP account is correct for all or even many of these. I am simply trying to motivate the idea that there are some phenomena of which those committed to physicalism should give a NIP treatment.[12]

What does this have to with the source of resistance to a NIP account of *sensation*? The juxtaposition of the case of price and that of sensation will immediately worry many readers. Even if I am correct that we should give a NIP account of price, sensations are extremely different from prices! More to the point, the aspects of price that make a NIP account of it plausible are decidedly not present in the case of sensation, one might insist. One can imagine an objector responding initially as follows:

> "Consider the gallon of milk in your hand. In an important sense, there's nothing really *there* that is the price. Prices are, you might say, "abstract" (or "abstracta"). But sensations certainly are not! On the contrary, sensations are concrete phenomena which we can identify and reidentify, and can inspect and observe, in introspection. Indeed, we might say that the price of the gallon of milk is not really a thing at all, or at least that it is not an object. But a sensation certainly is.
>
> Moreover, there are other differences between prices and sensations that are relevant to sensations' not being amenable to a NIP account: Unlike sensations, something has a price only in virtue of particular social customs. Prices are also "relational"; something's having a price is constitutively dependent on that thing's standing in particular relations to other things. Unlike sensations, prices are not located, they do not exist anywhere.

> They are also not occurrences or processes. And finally, being abstract, prices do not enter into causal relations, but sensations surely do."

These are some of the natural ways of responding. I count roughly nine or ten ideas here, many intimately related. Some of the ideas expressed in this reply are difficult to sharpen, such as the idea that prices are not "there." And notions such as those of the abstract and of an object are notoriously difficult to specify. But I will not hold this at all against my objector. At least not at this juncture. I think it is extremely natural to harbor these albeit rough ideas about prices and sensations; indeed, I feel their force myself.

I hope by the end of the chapter to give my reader a sense, for each of these ideas concerning differences between prices and sensations, as to why I think that appealing to it is ineffective in arguing against the plausibility of a NIP account of sensation. My treatments of them will come in section 9 and will fall roughly into three categories. For about half of them—many of the most trenchant—I will argue (to a first approximation) that there is a sense (or way) in which sensations are as the objector claims, and a sense in which they are not, and that the particular sense in which they *are* as the objector claims does not threaten the prospects of a NIP view of sensation. For one of the remaining ideas, I will claim that what the objector says about price is incorrect. And for the remaining ideas (such as that prices are social), I argue that the differences, which I grant at that juncture for the sake of argument, are irrelevant to the plausibility of a NIP account.

4. A Powerful Line of Thinking

From where does the objector acquire the ideas about sensations that inform his objection?—that sensations are, as it were, "there," that they are observable, that they are objects, and so on? The tremendous persistence of this way of conceiving sensation has its ultimate source, I believe, in a line of thinking—or really, an activity—that philosophers find extremely natural to engage in. Indeed, the line of thinking is one I find myself engaging in all the time when thinking philosophically about the mind—or at least, it's one I find myself thinking I'm engaging in. In thinking about the philosophical problem of sensation or phenomenal character, I decide to inspect my own sensations or phenomenal character. This might be because a philosopher whose writing I am reading requests that I do so (e.g., Harman 1990, 667; Shoemaker 1994, 202) or because I am simply thinking about one or another philosophical question about sensation (as I was about price in the market) and decide to inspect some of my sensations in order to better understand the phenomenon of sensation (or phenomenal character). What *are* sensations, I might ask myself. I might then attend to an itch I notice I have on my arm, or to the feeling of the collar of my shirt against my neck. Then, while attending to this sensation, or to some property of it, I contemplate what that thing *is*. Not what kind of sensation it is—that it's dull as opposed to sharp, say—but more something about what it, I want to say, *is*, or consists in. This is why I attended to a sensation after all: to investigate what these things called "sensations" *are*, and what account we should give of them. Most often in this situation, I then find myself very soon concluding that what I'm attending to is something that—and I don't know quite how else to put this—something that is "there." There seems to be something there that I'm inspecting. I might not know how best to describe what that thing is—whether it's an object, a property, something in an inner space—but there is something there that I'm inspecting, I think to myself.

This line of thinking comes in different varieties of course. I depicted how it often happens to me: at the end, I form the thought that there is something "there" that I'm inspecting.[13] Others find themselves concluding that they're attending to an object that exists in some inner space. Or that what they're attending to is a "something."[14] Or that it is an intrinsic property.[15]

In the next section, I will urge that this line of thinking—one that I myself and many others engage in—is problematic, but it cannot be denied that it is an overwhelmingly common and influential train of thought. Of course, many in the history of philosophy have inveighed against its legitimacy, in one way or another. There is the standard concern, expressed in different ways, that for all we know the very act of introspection changes or alters or obscures what one wants or intends to introspect.[16] There are Sydney Shoemaker's influential Royce Lectures in which Shoemaker argues against "inner-sense" models of introspection, according to which introspection is or involves a form of perception or observation. I will not assume that the line of thinking depends on this model of introspection, but it does seem most natural on such a model.[17] And rife in contemporary philosophy of mind is the idea that experience is "transparent." When we turn our attention inward in order to inspect our experiences, all that we find are the intentional contents or objects of our experiences; we do not find the experiences themselves or their intrinsic features.[18] The same goes for sensations, according to many philosophers.[19]

But for all that, the line of thinking still has enormous appeal and continues to impact a great deal of philosophical thinking (perhaps most often implicitly). In the next section, I shall scrutinize the line of thinking by attending carefully to a variety of critical issues, ones I believe have not received sufficient attention. I want to underscore certain crucial steps typically taken in the line of thinking that ought not to be as readily accepted as they are. I will not argue that the conclusions drawn in the various steps are incorrect; I will argue that we should not be nearly as

confident as many philosophers typically are in the legitimacy of the inferences involved. I will offer four considerations, each concerning a different step in the line of thinking. Considered together, the four points I raise should give pause to anyone who accepts the line of thinking uncritically, and should thereby call into question a primary source of some central projects in the philosophy of mind.[20]

5. Analysis of the Line of Thinking

5.1. Everyday versus Philosophical Introspection

Let us begin, then, by noting that the line of thinking differs in some significant ways from the sort of thinking we do when we introspect in everyday contexts. Introspection is not always done for philosophical purposes. When the doctor asks you what the pain in your elbow is like, you introspect and report that it's sharp, say, not dull. That is not the context in which we, as philosophers, draw the conclusions we draw at the end of the line of thinking—that what we're attending to is something that's there, or that has introspectible intrinsic properties, and so forth. I think that appreciating the differences between introspection as it is done for philosophical purposes and introspection as it is done in everyday contexts might be revealing. And for a variety of reasons.

To mention just one: Why are we so inclined to trust the line of thinking and the introspection involved in it? I suspect that one reason we do is that we think that in engaging in it we are not doing anything we do not normally do. We introspect all the time, we think. But if the two sorts of introspection differ in important ways, then this way of justifying the philosophical methodology is questionable.

One way in which the philosophical case differs from the everyday case concerns the *motive* or goal of the introspection. In the doctor's office, your goal is to tell the doctor the kind of sensation it is. You're trying to classify the sensation as being a particular kind of sensation—a sharp one, a throbbing one. In the philosophical case, your goal is not to investigate the kind of sensation you're having, or how the sensation you're having at that moment feels—whether it feels sharp, or dull, or throbbing—but rather it is to ascertain something more like what it, or sensations more generally, "are," or consist in, or what their "nature" is.

5.2. First Consideration

A number of questions arise now concerning the cognitive capacities and judgments that must be employed in the philosophical case in order to form the intention that you do—that is, in order to intend to ascertain what sensations are, by attending to a sensation. That is an early step of the line of thinking: we form an intention to investigate what sensations are, by turning our attention inward, as it were, and focusing on one. That is a complicated intention. And in order simply to form that intention, we need to be exercising a variety of cognitive capacities.

First, let us ask this: What *judgments* must one make in order to form this intention? For many other mental phenomena Wittgenstein persuades us that the intention to observe or "watch" the phenomenon, for the purpose of understanding what it is, is a mistake. For instance, about thinking, he writes:

> In order to get clear about the meaning of the word "think", we watch ourselves thinking; what we observe will be what the word means! — But that's just *not* how this concept is used. (It would be as if without knowing how to play chess, I were to try and make out what the word "checkmate" meant by close observation of the last move of some game of

chess.) (PI §316)

What we would be attending to in these cases, Wittgenstein says, are merely some "characteristic accompaniments" (PI §152 [3rd ed.]) of the thinking, such as perhaps mental images, or inner vocalizations. But thinking does not "consist in" the presence of any such accompaniments.

Of course, sensations may be among the sorts of things Wittgenstein sometimes has in mind when he speaks of "accompaniments." And I admit that it certainly does not follow, from the fact that "watching ourselves" while we think would not get us clear about what thinking is, that watching ourselves while we have a sensation would not get us clear about what sensation is. Nonetheless, Wittgenstein's point here should remind us that our intention to attend to our sensations, for the purpose of understanding what sensations are, does not come without significant assumptions. We wouldn't form this sort of intention in many other cases. Consider my philosophical question about price in the market. In order to ascertain what price "is," or what prices are, I do not form the intention to inspect the price of the gallon of milk in my hand. And if I did try to inspect the price (again, not just to determine that it is, e.g., $2.79, but to determine what the price "consists in," etc.) I would not be successful.

And so it is difficult to see how, when we do intend to attend to our sensations, in order to investigate what they "are," we are not prejudging the issue, at least to some extent. We are assuming at least that sensations are things that we can attend to for the purpose of investigating what kind of things they are.[21] That prejudgment may well be correct, but it's worth bearing in mind that we make it in this very standard line of thinking. And of course, in order for us to have good reason to make the prejudgment at this juncture of the line of thinking, that reason could not come from the line of thinking itself.[22]

5.3. Second Consideration

Let me continue, though. After that step—that is, after intending to attend to some sensation—we then attend to one. There are two steps here, in fact, which I didn't explicitly distinguish before: identifying a sensation to inspect—locating one, as it were—and then inspecting it. The sensation we identify is the one we inspect—or at least, that's the idea.

So let me just recapitulate the line of thinking for my reader: (a) something gets me contemplating a particular philosophical question about sensation; (b) I form an intention to inspect my own sensations, in order to ascertain what I think the correct thing to say about sensations is; (c) I identify a sensation to inspect; (d) I inspect this sensation, with the goal of ascertaining what it "is"; and (e) I draw the conclusion that what I'm inspecting is something that is "there," or that it's a something, and so forth. (Or, if it is phenomenal character that I intended to inspect, then I conclude that what I'm inspecting is, e.g., an intrinsic property of the sensation.) I now want to discuss steps (c) and (d): identifying a sensation to inspect and then inspecting it, with the goal of ascertaining what it is.

What is involved in our identifying a sensation to inspect? Precisely how do we identify this sensation? There are many interesting questions to think about here. One question I am interested in is: What are the cognitive capacities we call upon in the identification? What sorts of concepts do we employ in identifying the sensation? What judgments do we rely on in identifying the sensation?

Here it is important to bear in mind some of the lessons from Wittgenstein's discussions about ostension. Consider a case in which I ostend, or point to, something nonmental—the greenness of a lime, say. The considerations are familiar: For me to point to the greenness of a lime on the table across the room, it is not enough that I raise my hand and extend my finger

such that the tip of my finger, its three knuckles, and the lime form a straight line. I might make just the same movements whether I intended to point to the greenness of the lime, or to the lime's elliptical shape, or to the part of the table on which the lime is resting, and so forth. Such movements alone are not sufficient. What must be added, it seems, is that when I make such movements, I at the same time concentrate on, or otherwise have in mind, the particular feature or thing to which I am attempting to point. Of course, what *having it in mind* involves, or requires, is a substantive issue. But, in any case, what I manage to ostend or to identify is determined in part by what concepts I employ, and by my intentions.[23]

So, what are the relevant concepts and intentions at play, in the philosophical context, when we identify our sensations? And are they different from the concepts and intentions one employs in the case of everyday introspection—in the doctor's office, for instance? And if so, does that difference affect how the two cases of introspection ultimately unfold? For instance, does it affect what it is we end up identifying and attending to? I want to say: it is difficult to judge, and we make these questions seem easier than they really are when we simply assume that what we are doing, or think we are doing, in the philosophical case is similar to, or the same as, what we are doing in the everyday case.

I have indicated that the *motives* according to which we typically proceed in the two cases are different, but this of course does not entail that the concepts that play a role in the identification are different. However, it would not be surprising if they were. Consider that how one sees things—*what* one sees—can be affected by what one is expecting, or thinking about, or doing, or *wanting to see*.[24] Indeed, we ought to be wary of taking it for granted even that the same notion of sensation is employed in the two cases. The concept of sensation that I employ in the philosophical case may be tainted by my historical engagement with debates in the

philosophy of mind, and may be different from the one I employ in the doctor's office. And even if the concepts *are* the same, it is not just the particular concept of sensation that determines what it is I ultimately identify and attend to; my present concerns, and intentions, and judgments might affect it as well.[25]

And even if we did identify the very same kind of thing in the two cases—a sensation—there is still the next step, in which we inspect the sensation. Inspection too requires cognitive capacities. And which capacities are brought to bear will certainly affect how the inspection unfolds. In the everyday case, in the doctor's office, the purpose of our inspection is to report to the doctor what *kind* of sensation it is—whether it's sharp or dull, and so on; in the philosophical case, the purpose of our inspection is to determine what it "is." The cognitive capacities are bound to be substantially different—as is the resulting inspection. Moreover, what cognitive capacities we employ in the philosophical case depends on what it is we do when we attempt to inspect something for the purpose of ascertaining what it "is," or what it "consists in," and so on. And that—that is, what we do when we attempt to do this—is hardly patent.

Now, in questioning whether we can trust the line of thinking, I am not questioning whether we can trust our introspection of our sensations in everyday contexts. After all, we also introspect our beliefs in everyday contexts. But introspection of belief does not involve finding, or coming upon, a "something" (in the strong sense above, whatever that is) or something that is *there*. There is perhaps a more innocent, distinct sense in which when introspecting one's beliefs one might sometimes say that one "finds something" or comes upon something that is "there" (e.g., "I now see a connection there between my career choice and my childhood hobbies"). But it is not the sense that is involved in the sorts of strong conclusions drawn at the end of the line of thinking. Nor do we infer from the fact that we introspect beliefs in everyday contexts that we

can ascertain what beliefs "are" or "consist in," and so on, by inspecting them in introspection. What the introspection of belief informs one of (when it does) is *what one believes*.

Indeed, this is connected up with the first of the four considerations I offered, and further supports it. Even in the case of the mental phenomenon of belief, it would be incorrect to suppose that introspective inspection of the phenomenon will yield an understanding of its nature; unlike prices (the example I gave in offering the first consideration) beliefs are things we do introspect in everyday life. So to be sure, I grant that we sometimes introspect our sensations in everyday contexts, but that does not give us reason to think that sensations are the sort of thing for which introspective inspection is a promising means of metaphysical understanding.

Now this point—that philosophical introspection may unfold differently than everyday introspection, on account of the concepts, intentions, goals, and judgments at play in the two cases—would not on its own, or without further exploration, convince anyone that the line of thinking is not trustworthy. The same is true of the point I made earlier about the prejudgment we make at the beginning of the line of thinking. But these points are worth bearing in mind as we proceed, and I will return to them. Indeed, many of the considerations I offer in this section involve underscoring assumptions on which the line of thinking depends and to which attention is rarely given. In many cases, I do not advance an argument that the assumption is false but rather argue that it may well be false, or at the very least, that one should not simply assume that it is true. In some philosophical contexts—certainly, the one we find ourselves in here—successful arguments of this sort are of significance. If we can move the proponent of the line of thinking to reduce his confidence in an assumption from 100 percent to even just 80 percent, say—and do that for four independent assumptions—he will have reason to believe that the line of thinking is more likely to be unsound than sound.[26]

Some readers may worry that the critical eye with which I am looking at the line of thinking is more critical (and suspicious) than that with which it would be fair to look at, say, ordinary perception (i.e., of the "external world"). I am about to take this sort of objection up below.

5.4. Third Consideration

The particularly puzzling step in the line of thinking, though, is the final one. We have just been discussing my inspection of my sensation. That inspection is what is ultimately supposed to lead to my having justification for my judgment about what that sensation is—that it is a "something," or that it's something that is *there*, or that it has intrinsic properties of which I am aware. But precisely how should we understand the relation between the inspection, on the one hand, and my having justification for forming that judgment, or having that belief, on the other? What justifies this judgment? Is the judgment the result of an inference from another judgment? If so, what is that other judgment, and how did I form *it*? If it is not the result of an inference, then if it is justified, what is the nature of that justification?

In raising this last question—If the judgment is not the result of an inference, then if it is justified, what is the nature of the justification?—I do not intend to intimate that one cannot have justification for forming a judgment, or believing something, unless that judgment or belief is the result of an inference one has made. At this juncture, I am merely raising some questions as to whether I am indeed justified in forming the judgment about what the sensation "is." And one possibility that we cannot yet rule out is that the judgment is not formed on the basis of an inference but is justified nonetheless. The point of this third consideration is, first, to pose these questions and note that they are worth careful reflection, and second, to explain why one tempting approach to answering the questions, to which I now turn, is inadequate.

One way of explaining how I am justified in forming this judgment is to appeal to the idea that my attention to my sensation is a *perceptual* sort of attention, that is, to appeal to an "inner-sense" model of introspection, according to which introspection is or involves a form of perception or observation. In section 4, I noted that the line of thinking would certainly seem to be most natural on such a model, but that I will not assume that the proponent of the line of thinking is committed to it. For one thing, adjudicating whether the line of thinking does depend on such a model is a substantial task, involving (among other things) establishing precisely what is required of a model to count as an inner-sense model. Many models of introspection liken introspection to perception in some way or other. But also, the vast majority of philosophers of mind today claim to reject inner-sense models, often citing Shoemaker's arguments with apparent approval.[27] Of course, the fact that a philosopher *says* she doesn't endorse an inner-sense view is one thing; that she does not implicitly rely on one is another. Indeed, I suspect that inner-sense models of introspection ultimately play a very significant, if implicit, role in a great deal of contemporary thinking about sensation and phenomenal character, including the line of thinking itself. But I will not pursue that suspicion here.

Regardless, I here want to address the inner-sense theorist in particular because an inner-sense theorist might be thought to have an effective answer to the question of what justifies the judgment I form at the end of the line of thinking. An inner-sense theorist might argue as follows: In order to perceive something, what I perceive must have perceivable intrinsic properties, or must be something that is *there* (in the robust sense), and so on; otherwise, I would not be able to perceive it. And since introspection is itself a form of perception, it follows from the fact that I do introspectively attend to my sensation that my sensation is something that's "there," or that it has intrinsic properties that I am perceiving, and so forth.[28]

However, one problem with making this argument is that it is difficult to see how a philosopher could be justified, at this juncture of the line of thinking, in appealing to an inner-sense model of introspection. One's reasons for holding an inner-sense theory would presumably be tied up with one's ideas about the sorts of things we find in introspection—for instance, that they're the kind of thing that can be perceived. But if one has not yet completed, or does not yet have reason to trust, the line of thinking, what justification would one have for such ideas about the sorts of things we find in introspection? That is not to claim that the inner-sense theory is false; it is to doubt whether one would be justified in assuming it is true for the purpose of justifying the formation of the final judgment in the line of thinking.

It is bound to be objected here that this sort of reasoning, which I myself have just engaged in against the inner-sense theorist, would lead to skepticism if similarly applied to perception of nonmental objects, to objects in the "external world." And, one might argue, since skepticism is clearly false, or at least to be rejected at all costs, the mode of reasoning is itself to be rejected. I have never much understood the insistence that if a premise or line of reasoning leads to skepticism, so much the worse for the premise or line of reasoning, but let us grant it. I will make four points against this concern.

First, the reasoning I have just employed against the inner-sense theorist's proposed way of justifying the conclusion drawn at the end of the line of thinking would not be applicable to, or similarly forceful against, the justification we have for our perceptual beliefs about external objects (or beliefs of the sort that, say, the chair is "there" or a "something"). Our justification for our perceptual beliefs is grounded in a rich variety of factors that are not present in a similar way in the case of introspection, such as the fact that there is not one but many mutually supporting perceptual modalities: vision, smell, taste, and so forth; that our predictions based on our

perceptual beliefs are extremely reliable; and the more general fact that our perceptions and perceptual judgments are themselves supported by our entire "web of belief," which is itself grounded in such perceptions and judgments. In the absence of such facts, reasoning leading to skepticism would have even more force.

Second, and closely related, while skepticism of the external world is often rejected, in spite of the compelling nature of the reasoning leading to it, skepticism about particular domains of beliefs is much more frequently countenanced. Consider, for instance, the fact that many philosophers and scientists argue that beliefs in which we ascribe colors to objects in the external world (such as my belief that the strawberry on the table is red) are systematically erroneous; the strawberry is not red, nor is it any other color (see, e.g., Boghossian and Velleman 1989; Hardin 1993). This view is a central contender among theories of color.

Third, even if ordinary perception *would be* threatened by the sort of reasoning I have offered, and even if the correct response to that threat would be simply to reject that reasoning, the correctness of that response would be a result of the absurdity of skepticism, or of our alleged inability to believe that skepticism is true (or something of that sort). However, the idea that the philosophical line of thinking is itself false (or that an inner-sense model of introspection is false) is hardly absurd or something we are unable to believe.[29]

Fourth, it is worth noting that on the most celebrated antiskeptical tack in recent decades—contextualism (e.g., DeRose 1995; Lewis 1996)—the reasoning that leads to skepticism is *not* rejected, but remains correct, intact, not erroneous, when engaged in in philosophical contexts. The reasoning is faulty only in, or as concerning, "ordinary" contexts.

Let us return, then, to the third of the four considerations. In offering this third consideration, which concerns the justification for forming the judgment that is formed at the end

of the line of thinking, I have focused only on the inner-sense theorist. The formation of the judgment is even more puzzling on views of introspection that do not liken it to perception. At the very least, we would certainly need plausible answers to—indeed, very convincing answers to—questions about the justification for making this judgment, in order for us to continue to give the line of thinking as much credence and influence as we do.

5.5. Fourth Consideration

Allow me to press on, and turn to my fourth consideration. I invite you now to attend for the moment to the feeling of your shirt (or some other article of clothing) against a part of your skin. Now, while you're attending to it, allow me to pose a question. Before I mentioned any such feeling—and thus before you attended to it—did you have it? I'm not asking whether the feeling of your shirt against your skin, before I asked you to attend to it, was the *same* as the feeling you have now on your skin, now that you're attending to it—that is, whether attending to the feeling *changed* the feeling. The question I want to ask is whether you had any feeling at all on that part of your skin, before I invited you to introspect. Most likely you had not focused once while reading this chapter on the feeling of your clothing against that part of your skin. And if you had, it should be easy for you to focus on the feeling of your clothing against some other part of your skin on which you had not focused.

This is a question quite familiar to philosophers.[30] And it is remarkable the difference in intuitions we have about it. My intention here, though, is not to argue or suggest that you did *not* have a sensation or feeling on that part of your skin before I invited you to introspect. Rather, I want to claim that you do not have strong reason to *believe* that you did, that you have at least as much reason to believe that your act of introspecting (or perhaps your act of intending to

introspect) created the sensation—or at the *very* least, that this is a very live possibility. On this possibility, agents of introspection are not always passive, unobtrusive voyeurs; but rather their acts sometimes affect, change, even create what is found in introspection.[31] Indeed, if this possibility is actual, we might expect the creations to occur most often when there is not otherwise anything to be found.

This is important, first, because it reminds us that unlike most external perception, perhaps, the very act of introspection could contaminate or otherwise affect the purported objects of introspection. And so the possibility is made vivid, not far-fetched, that in the line of thinking our introspection is in some way creating an illusion. Second, this is most likely to be the case when there is nothing there to be found, when, one might say, one's investigation is idling. It is in precisely such situations in which one is spinning one's wheels that introspection would be most likely to create something. The idea we must take seriously is that what it creates in the line of thinking is the illusion of our coming upon a "something," or upon something that is "there."[32]

One might protest: But this is quite a far-fetched possibility! Perhaps it is logically possible that our introspection in the line of thinking creates an illusion, but is there any reason to think that it is so? Let us recall three aspects of our situation that, in my view, together make it a real possibility. The first is that the line of thinking we are discussing is an activity that human beings engage in extremely rarely, only in certain philosophical contexts. This fact on its own is not very significant. If through an equally rare activity—a novel geological procedure, say—we discovered a kind of rock human beings had never discovered before, we would not, after adequate testing, be particularly dubious that we had really found what we thought we had found. But there is a difference between this case and the line of thinking. In the rock case, we already have good reason to trust our scientific and perceptual methodology. One point I've been

emphasizing in this chapter is that we do not yet have similar good reason to trust our method of investigation in the philosophical line of thinking. That is not to say that the method is a faulty one, it is to say something about what we can assume at this juncture.

But there is something else. When, in any investigation, we come upon something that would seem to provide problems for our larger conception of the world—as inner "somethings" have struck many philosophers as doing (or as intrinsic, introspectively accessible properties have)—this typically makes us, and should make us, less confident that we have truly found what we think we have found. When we are in a dark attic, and out of the corner of our eye we seem to see a white figure shimmering beside us, we might first think with fright that there is a ghost! But then we think again—or at least we should. We don't believe in ghosts; it must have been something else. Perhaps it was the light of the moon reflecting off the window, or an idiosyncrasy of our aging eyes. So it is all of these factors—the rareness of what we are doing, the status of the methodology with which we are doing it, and the relation of what we would be finding to our current scientific conception of the world—all of these factors, taken together, I think should make one less confident in the conclusions we form at the end of the line of thinking. They should make us more open to the possibility that philosophical introspection creates a context that makes us form false judgments about what we are attending to.[33]

5.6. Conclusion

I have focused in this section primarily on sensations, for the sake of example, but it should now be clear that the considerations I have raised apply equally well to the variant of the line of thinking on which what one intends to attend to is phenomenal character: the prejudgments we make, the comparison between the philosophical case and the everyday case (and the questions

raised by the comparison), the justification for the judgment we form at the end, and so on, all apply.

As I say, it may be that no one of the four considerations I have raised would itself have substantial impact on one's subscription to the line of thinking. For me, it is especially when the points are considered together that they have a real effect. What they do is make me doubt that the line of thinking is on solid ground. And that line of thinking was, in the first place, the primary source of a conception of sensation that would make it seem as if a NIP approach to sensation is clearly misguided.

6. Looking Forward

I have done the best that I can, for now, with the line of thinking. What I want to argue in the rest of the chapter is this: *If* you're willing to suspend judgment about, or at least confidence in, the soundness of that line of thinking, much of the resistance to the NIP account of sensation I will sketch in sections 8 and 9 ultimately falls away. Recall my objector's initial insistence that, even if a NIP view of price is correct, such an account could not be correct for sensation. The chief source of that thought too is the line of thinking. We will continue to see that this is so.

I have three central goals in the remainder of the chapter. One is to return to the differences that my initial objector appealed to between price and sensation. Another is to sketch the particular kind of NIP account of sensation that I think should be taken more seriously in the philosophy of mind. To this end, I will address many questions that my readers are bound already to have in mind: What, on my account, *should* we say about sensations and phenomenal character? Am I suggesting that we don't in fact inspect or attend to sensations in introspection? Or, is the point rather that sensations are not objects? What is it *in virtue of* that people have

sensations? In the case of price, after all, even if we cannot identify the price with any physical property or phenomenon, we can still say what it is in virtue of that something has a price, we can still point to what the milk's property of having such and such a price "supervenes" on. What is the parallel explanation in the case of sensations? And where does all this leave us with phenomenal character?

It will be the point of sections 8 and 9 to achieve these two goals. But there is a third task that it will be helpful to tackle before either of these two. And that is to discuss the relationship of what I have said so far in the chapter to the representationalist position about phenomenal character. Understanding this relationship is instructive.

7. Representationalism

Many representationalists will agree with the particular conclusions of section 5 that are most important for my purposes. Which conclusions do I consider most important? Recall the primary role that my treatment of the line of thinking is meant to play in the overall argument of the chapter. I am proposing that we adopt a NIP account of sensation (and thus, ultimately, phenomenal character as well). In section 3, in order to provide an example of a phenomenon to which a NIP account applies, as well as to generate immediate resistance to what I am proposing in the chapter, I introduced the case of price. The juxtaposition of price and sensation is one way of provoking and illuminating the substantial resistance many readers will have to a NIP account of sensation. Even if I am correct that we should give a NIP account of price, it is thought, sensations are extremely different from prices. And, more to the point, the aspects of price that make a NIP account applicable to it are decidedly absent in the case of sensation.

I then imagined some of the purported differences between prices and sensations that an objector might suppose are relevant here. Unlike prices, sensations are "there," they are "observable," they are "objects," and so forth. Many of these ideas about sensation have their primary source, I claimed (though I have yet to complete my argument for this claim), in the line of thinking. The goal of section 5 was to add to the significant considerations philosophers have already offered against the soundness of that line of thinking—and, in particular, to question the truth of the judgments we form at the end of the line of thinking, such as that in introspection we are coming upon a "something" that is the sensation, or that we are coming upon intrinsic properties of the sensation, or that the sensation or its phenomenal character is something that is "there."

Most representationalists will agree with (at the very least) the spirit of these conclusions. Recall that many representationalists argue that sensations (e.g., pains) are transparent (Harman 1990; Tye 1995; Crane 2003; Dretske 2003). As with experiences, when we introspect our sensations, they claim, all we come upon, or "find," are the *intentional contents* or *intentional objects* of the sensation. We do not observe the sensation itself or any intrinsic property of it. Insofar as we *do* observe anything, or something that is "there," it is the intentional object.[34] In the case in which sensations are veridical (a notion I will return to below), what the subject finds is a physical state or process of the body. This is a physical state or process distinct from the sensation; it is what the sensation is *of*, its intentional object. In the case in which sensations are nonveridical, what the subject observes does not exist. Thus, the representationalist will agree that we should be suspicious that introspection reveals to us something that both (1) is "there," and (2) is the sensation or a (nonintentional or intrinsic) feature of the sensation.

Representationalists will thus welcome the conclusions of section 5 and consider the arguments I provide in their service, if correct, to yield yet further support for representationalism.

But I want to offer an alternative to representationalism. Let us, for the moment, grant the representationalist that what one "finds" when one introspects one's sensations is the intentional contents of the sensations. We then face the subsequent question: what should we say about sensations and phenomenal character? The representationalist proceeds to identify sensations with physical phenomena. He identifies sensations with those physical states or processes that he claims represent what the sensations represent (i.e., he identifies the sensations with the physical states and processes that he claims *have* the corresponding representational contents). And he identifies the phenomenal character with the representational properties of these physical phenomena. I, on the other hand, propose that sensations are not identical to any physical phenomenon (i.e., any phenomenon fully specifiable in physical terms), and that this does not threaten physicalism.

What favors one approach over the other? Part of this will depend ultimately upon what we take ourselves to be trying or needing to explain. This issue is important and is intimately connected to the concerns I raised at the very beginning (and to which I will return) concerning what it is we are speaking about when we speak of phenomenal character. For now, two facts about representationalism are especially important to mark: one concerning the substantial difficulties that have been raised in the literature for representationalism, the other concerning the negative form of the most influential arguments advanced in support of representationalism. I will discuss these in turn.

As I noted in section 1, many philosophers argue that representationalism faces substantial difficulties. First, a wide variety of alleged counterexamples have been offered to the

representationalist's thesis: for instance, cases in which there is phenomenal character but no representational content (moods, say, and even sensations themselves (McGinn 1997; Rey 1998; Burge 2010)); cases of pairs of mental phenomena that have identical representational contents yet different phenomenal qualities (Peacocke 1983; Block 1996); and cases of pairs of mental phenomena that have different representational contents yet the same phenomenal character (Block 1990). Second, assuming "externalism" about representation, representationalism appears to lead to externalism about phenomenal character, a view many philosophers find extremely unpalatable. If an experience's phenomenal character is exhausted by its representational content, and if representational content is externally individuated, then phenomenal character itself is externally individuated, it is argued. However, even if internally identical individuals could have different thoughts, it is argued, if one of them has a headache, or a tingly sensation, so must the other.[35] Third, some philosophers argue that mental phenomena presumably have the representational properties they do partly *because of* the phenomenal qualities they have, and thus the latter cannot simply consist in the phenomenon's representational capacity.[36] Fourth, on the representationalist picture, on which what it is to be a sensation is to represent (in a particular way) something as being a particular way, every sensation is either veridical or nonveridical, correct or incorrect. But speaking of correctness and incorrectness as applied to sensations is at best rather foreign; it is rare, for instance, that we conceive of a particular sensation as being *illusory*, say, or, for that matter, as being *correct*. Fifth, and not least, there is substantial controversy concerning whether representational content can itself be explained in physical terms.[37] The assumption that it can constitutes much of the motivation to identify phenomenal character with representational content in the first place.[38]

No less significant than the problems besetting representationalism, for our purposes, is the negative form of the arguments that have historically motivated representationalism. There are various sorts of arguments that have been advanced in favor of representationalism, but it is fair to say that two are paramount and have been most influential. The first concerns the considerations about transparency I have already introduced: in introspection, one finds neither the sensation itself nor intrinsic features of the sensation. The second stems from the idea that all *other* physicalist approaches in currency to questions about phenomenal character face insuperable difficulties.[39] Neither of these arguments, however, suggest anything like the positive thesis of representationalism—that phenomenal properties are identical to (or exhausted by) particular representational properties. The argument from transparency concerns only what we do not find in introspection; and the argument concerning alternative views appeals only to their inadequacy. It is true that the argument from transparency often emphasizes that what we do come upon in introspection is the intentional (or representational) contents of a particular mental phenomenon. But even if it is *granted* that all sensations that have phenomenal character also have intentional content, it is a substantial, subsequent step to the conclusion that sensations are identical to, or nothing but, the physical states or processes that purportedly have that content, or to the conclusion that the phenomenal character of a sensation is identical to, or exhausted by, that content.[40] Most importantly, neither of these arguments for representationalism itself intimates the truth of the representationalist thesis any more than it does the truth of NIP.

8. Room for a View

What is it *in virtue of*, then, that we have sensations? Many readers, I presume, will have been waiting a long time for me to address this question. What *makes it true* of someone that she

is having a sensation? What do sensations "supervene on"? In granting physicalism for the sake of this chapter, I have granted that no two worlds can be physically identical yet differ in any mental respect. My interlocutor wants to know precisely what facts guarantee or necessitate the existence of particular sensations. Moreover, in order to satisfy the interlocutor, the answers need to be of a particular sort. The sorts of answers my questioner is asking for are not to involve more talk of sensations, for instance. It would not answer the question in a way satisfactory to the questioner for me to reply that what makes it true of someone that she is having a sensation is that she is having a sensation; or that someone has a sensation in virtue of having a sensation. A satisfying answer will also be sufficiently detailed. It would not satisfy the questioner to be told that one has a sensation in virtue of some set of physical facts obtaining. Which set of physical facts?

The representationalist, after all, provides an answer to such questions, one that appeals to facts that are expressed at what he understands to be an intermediate level of explanation. According to the representationalist, one has a sensation in virtue of one's being in a particular kind of representational state (or, in terms of facts: in virtue of particular facts about one's representational state), in virtue of a particular physical state of one's brain or body having particular representational properties. The sensation *is* that physical state of the body, and that physical state of the body *is* the sensation, they claim, in virtue of the representational properties the state has. That is a noncausal, *synchronic* sense of "in virtue of." It is the sense of "in virtue of" on which most philosophers of mind (physicalists and nonphysicalists alike) require physicalists to specify (in physical terms) what it is in virtue of that someone has a particular sensation. The representationalist's answer is satisfactory (to him) only because (1) he finds it plausible to suppose that some representational states are sufficient for a person's having a

sensation (as well as to suppose that being in some such state is necessary for a person's having sensations); and (2) the representationalist (at least he who is physicalist) finds it plausible to suppose that facts about representational states can be satisfactorily explained in terms of facts expressed in physical terms (whether by way of facts expressed at yet other intermediate levels, or not).

And indeed, one might say (and as I granted in section 3 during my reflections in the market) we can also say what it is "in virtue of" that things have prices. Things have prices in virtue of a very complicated set of facts or states of affairs involving social customs, human attitudes, behavioral dispositions, and so on (e.g., the grocer's disposition to chase after me or call the police were I to take the milk without paying for it). Or at least, so one might argue.

How then does an advocate of a NIP approach to sensation respond to the question "What is it in virtue of that a person has a sensation?" There are in fact a variety of distinct ways in which an advocate of NIP might answer or handle the question. Some linguistic behaviorists, for instance, would propose that one has a sensation in virtue of one's having dispositions to behave in particular ways, and of one's behaving in particular ways. Some adverbialists would propose that one has a sensation in virtue of one's sensing in a particular *way*. Both of these views are notoriously problematic though, and I wish to underscore an alternative.[41]

What I want to encourage is that we take seriously the possibility that (i) answers of the sort that the questioner is asking for are unavailable; one cannot say, in the way the questioner desires, what it is in virtue of that someone has a sensation (this "unavailability" could take different forms, have different sources, as I will discuss); and (ii) the unavailability of such answers does not provide a threat to physicalism. In the limited space I have here, I will not be

able to make a compelling case that (i) and (ii) are both true; rather, my goal is to argue that they may be, and that the thesis that they are should be clearly situated on the table of positions.

Let us call the thesis that (i) and (ii) are both true, conjoined with physicalism: "Non-Explanatory Physicalism" (NEP). (NIP) does not entail (NEP), nor does (NEP) entail (NIP). But on the view I am entertaining in this chapter, both are true.

Most philosophers of mind proceed on the assumption that a NEP view of sensation is false; a physicalist about sensation must be able to provide a satisfactory answer to the "in virtue of" question. Their reasons are deep and complex and require careful scrutiny, which I cannot provide here. In the limited space I have left, I wish only to make explicit several possibilities.

Note first that the availability of an answer to the "in virtue of question that would satisfy the interlocutor is not something that is in any way entailed by physicalism. The mere thesis of physicalism (as we are understanding it in this chapter) does not entail that such an answer can be given. That is a crucial point that is too often overlooked. Of course, one source of the assumption that a NEP view of sensation is false is perhaps the desire or hope for the sort of answer in question. But from the fact that we have this hope, it does not follow that it is fulfillable. Nothing about physicalism, as we are understanding it in this chapter, itself entails that it is fulfillable.

The first point, then, is simply a logical one, but an important one. It is especially important given the current state of debate concerning phenomenal consciousness, and in particular the fact that all views of phenomenal consciousness in currency are thought to face extreme difficulties. As a result of these difficulties, central figures in the philosophy of mind have argued for such counterintuitive views as panprotopsychism, eliminative materialism, and epiphenomenalism. It is difficult to see how the reasons to resist a NEP view—whatever those are—could be that much

more compelling than the reasons to resist these other positions. Given this, a NEP view ought at least to have a place on the table of going options.

The second point is that a nonexplanatory physicalist view is in currency as applied to other phenomena discussed in philosophy of mind and language. Here is Barry Loewer who entertains this view as applied to semantics:

> Of course, the failure of naturalization proposals to date does not mean that a successful naturalization will not be produced tomorrow. But another possibility, and one that philosophers have recently begun to take seriously (such as McGinn, 1993), is that while semantic naturalism is true, we may not be able to discover naturalistic conditions that we can *know* are sufficient for semantic properties; that is, perspicuous semantic naturalism may be false. It may be that the naturalistic conditions that are sufficient for semantic properties are too complicated or too unsystematic for us to be able to see that they are sufficient. Or, it may be that there is something about the nature of semantic concepts that blocks a clear view of how the properties they express can be instantiated in virtue of the instantiation of natural properties. This position, though it may be correct, is not by itself intellectually satisfying. (1997, 121–122)

Indeed, the unavailability of a certain kind of perspicuity connects very closely with Wittgenstein's work on meaning, understanding, and thought more generally. Barry Stroud characterizes the upshot of some of Wittgenstein's writing as follows:

> [I]n giving descriptions of the practices we and others engage in we must employ and rely on the very concepts and practices and capacities that we are trying to describe and understand. One reason it is difficult to describe them correctly is that we see right through them, as it were; they are too close to us to be seen for what they are. This

> inability to command a clear view of our concepts, and the apparently natural tendency in philosophical reflection to wrongly assimilate the use of one kind of expression to that of another, is for Wittgenstein a continuous source of philosophical problems.
>
> The conclusion can be philosophically dissatisfying or disappointing in another, and deeper, way. If facts of what expressions mean, of the correctness of certain ways of understanding them, or of the rules by which speakers and hearers proceed, can in general be expressed only in semantical or intentional statements which make use of the very concepts that they attribute to those they describe, then they would seem not to be the kinds of facts that could ever explain how language or meaning in general is possible, or what facts or rules human beings rely on, as it were, to get into language in the first place, from outside it. That can seem to leave the phenomena of meaning, understanding, and thinking as philosophically mysterious as they would be on the hypothesis of an occult mental medium. No explanation of thought or meaning in non-semantical or non-intentional terms would be available. That can be felt as deeply dissatisfying.
>
> Much of Wittgenstein's later philosophy deals in one way or another with the aspiration or demand for a different and potentially more illuminating kind of explanation of meaning. (1996, 317–318)

This is not the venue in which to explain and assess the substantial arguments Stroud advances and himself finds in Wittgenstein.[42] To be sure, many contemporary analytic philosophers of mind and language would disagree with Stroud's assessment of the possibility of explanation of meaning, understanding, and thinking. Moreover, Stroud himself does not discuss these issues in service of physicalism or naturalism. But they do remind us of the possibility that perspicuity—perhaps even philosophical satisfaction—is not to be had.[43]

The view I am entertaining in this chapter is that while physicalism is true, "perspicuous" physicalism about *sensations* is not. Now, Loewer mentions two—but just two—ways in which, or reasons for which, perspicuity might be unavailable in the case of semantic properties. The first is that "the naturalistic conditions that are sufficient for semantic properties are too complicated or too unsystematic for us to be able to see that they are sufficient." The other is that "there is something about the nature of semantic concepts that blocks a clear view of how the properties they express can be instantiated in virtue of the instantiation of natural properties." While I think we should remain open to the first of these two possibilities with respect to sensation, I entertain more seriously the second possibility, which itself could take different forms.

One thing that the interlocutor, in the case of sensations, demands of a satisfactory answer to the "in virtue of" question is that the answer not, in the end, leave the questioner with a feeling of puzzlement, puzzlement that might get expressed by the thought "How could *that* sort of thing (the explanans) ever be sufficient for, or 'necessitate,' or on its own 'give rise to,' a *sensation*?" Precisely what is required to preclude such puzzlement is an extremely complicated and difficult question. One thing that is clear, however, is that, whatever is required, the answer and its availability depend in part on the nature of the concepts expressed in the explanandum. That is because any understanding that is required to be produced by the answer to the "in virtue of" question must concern the facts to be explained, which are described using particular concepts, in this case concepts such as *sensation*, *pain*, and so forth. The answer itself must thus be expressed or grasped partly in such concepts. If the answer is to take the form "X in virtue of Y," the "X" needs to be stated or expressed in sensation concepts.

One of the possibilities that I propose we pursue further is that the special features of our concepts of sensations (more on what those might be in a moment) are such that facts expressed in those concepts cannot be satisfactorily explained by facts expressed entirely in the concepts of physics (even by way of facts expressed in intermediate concepts)—and not simply on account of our (perhaps "contingent") cognitive limitations. The idea is that the former facts are such that it could not be understood how or why facts expressed in physical concepts (or in appropriately intermediate concepts) "necessitate" facts of the former sort, or "are in virtue of what" the former facts obtain—and yet it still be true that no two worlds could be identical in their physical properties but differ in their distributions of sensations. Nothing about physicalism entails that there are not some concepts such that some facts expressed in terms of them cannot be explained in this way. Indeed, this is a point that has been made, in different ways, in areas beyond philosophy of mind and language (see, e.g., McDowell 1983; Williams 1985).

Before going a bit further, let me submit that what I have said so far in this section should be of import on its own. I have introduced one general way in which it might be open to an advocate of a NIP approach to sensation to treat the sort of "in virtue of" questions that will inevitably be asked. That way is to resist the expectation and requirement that a physicalist must be able to provide satisfactory, positive answers to such questions. Physicalism does not entail the availability of such answers. And the idea that such answers are not available has currency elsewhere in philosophy. I then introduced one—but just one—way in which, or reason for which, such answers would not be available. On this more specific possibility, it is something about the nature of sensation concepts that makes them unavailable. It is especially once we have suspended belief in the line of thinking discussed in section 5 that these possibilities gain plausibility.

Allow me, though, to go one step further, and to advertise a perspective on sensation on which this more specific possibility might be actual. It is suggested by some of Wittgenstein's later writings in which he emphasizes the importance of investigating the roles that sensation (and other) concepts play in our lives, and the purposes for which we engage in discourse about sensational (and other) phenomena. Wittgenstein admonishes us, for instance, about construing ascriptions of psychological phenomena on the model of *descriptions* (PI §§180, 585). The idea is not that we do not sometimes describe mental states, but that the phrase "description of a state of mind" characterizes a particular "language-game" (or perhaps more than one), the structure and role of which is very different from what we typically suppose when doing philosophy (LW I §50).

Consider, for instance, how Ian Hacking describes what he takes to be Wittgenstein's approach to the particular phenomena of seeing and imaging in *Remarks on the Philosophy of Psychology*:

> We are not to think of seeing and imaging as being different phenomena in themselves but as verbs distinguished by the ways in which they "relate to a host of important kinds of human behavior, to the phenomena of life." The phenomena are not the seeing and the imaging but the practices in which they are embedded. (Hacking 1982, 224)

This is a provocative, perhaps ultimately helpful, way of characterizing the general perspective, even if it would need to be qualified and sharpened for it not also to be misleading.[44] We might explore these ideas as applied to sensations. About pain Wittgenstein writes:

> The concept of pain is characterized by its particular function in our life.
> Pain has *this* position in our life; has *these* connexions; (That is to say: we only call "pain" what has *this* position, *these* connexions). (Z §§532-533)

One of the many virtues of Wittgenstein's idiosyncratic, "criss-cross" (PI, Preface) journeys across the tremendous variety of psychological phenomena and concepts consists in its bringing to light some of these surprising "positions" and "connections." So often in Wittgenstein's philosophy one

finds unexpected truth in the idea that "It is necessary to descend to the application, and then the concept finds a different place —— one which, so to speak, one never dreamed of" (PPF §165).

One might balk at the idea that the concept of sensation, or our thoughts involving that concept, or our sentences involving sensations words, are much different from how contemporary debates construe them. What is more obvious or quintessential a case of "object and designation" (PI §293) than concepts of sensation and their corresponding linguistic expressions? The force of such a reply, however, should be mitigated considerably once we have lost our confidence in the line of thinking explored in section 5. Once we are out from under the grip of that line of thinking, we are free to look anew at our sensational concepts and discourse.

What then, one might ask, *is* one doing, or accomplishing, when one says, for instance, that she is in pain? Is the view I am entertaining here a form of "expressivism," according to which sensation reports are nothing but expressions, akin to exclamations such as "Ow!" and neither descriptions nor assertions? No. That is the sort of caricature view that is sometimes attributed to Wittgenstein and precisely the sort of neat, all-encompassing account (in this case, of "sensation reports") that Wittgenstein's philosophy would suggest should not be expected.[45]

Our practices involving these concepts and words might be extremely difficult, even impossible, to delineate in a satisfying way. For one thing, it may be "that we engage in endless loosely overlapping language games" (Hacking 1982, 224). We cannot know at this juncture how perspicuous a view could be had, on this perspective of sensation. Perhaps though we *could*, on this perspective, one day understand our practices in enough detail to answer the "in virtue of" question in a way that would satisfy our interlocutor, or in a way that would itself dissolve the force of his question. But it is also possible that such an understanding of our practices could never be had.

These are bold ideas, and mentioning them does nothing to support them. What I wish to urge is that we recall that it *is* a possible approach to philosophical inquiry about sensation, and one that has more to speak for it than is typically acknowledged. Wittgenstein's discussions of the concepts of thinking, imagining, expecting, fearing, hoping, noticing, grieving, wishing, pretending, seeing, imaging, feeling, and so on, are rich and illuminating and make vivid the possibility that this is a promising direction to pursue. Indeed, investigations into them—their interrelations (RPP II §454), the "kinships and differences" (RPP I §1054) among them, the contexts of their application, conditions for their mastery, and so on—and into the *limits* to which such investigations can ultimately be pursued, preoccupied Wittgenstein substantially in the last years of his life, following his completion of *Philosophical Investigations*. Many of these discussions have yet to be fully appreciated or explored.

9. On Characterizing Sensations: Returning to Price

It is time now to return to my juxtaposition in section 3 of price and sensation. Recall that I introduced the case of price for two reasons. The first was to make more explicit how I am understanding what a NIP account says of a particular phenomenon and what it does not say. The second was to uncover the source(s) of the substantial resistance that any NIP treatment of sensation will inevitably face. My objector claimed that the aspects of price that make a NIP account of it plausible are decidedly absent in the case of sensation:

> Consider the gallon of milk in your hand. In an important sense, there's nothing really *there* that is the price. Prices are, one might say, "abstract" (or "abstracta"). But sensations certainly are not! On the contrary, sensations are concrete phenomena which we can identify and reidentify, and can inspect and observe, in introspection. Indeed, we might

say that the price of the gallon of milk is not really a thing at all, or at least that it is not an object. But a sensation certainly is.

Moreover, there are other differences between prices and sensations that are relevant to sensations' not being amenable to a NIP account: Unlike sensations, something has a price only in virtue of particular social customs. Prices are also "relational"; something's having a price is constitutively dependent on that thing's standing in particular relations to other things. Unlike sensations, prices are not located, they do not exist anywhere. They are also not occurrences or processes. And finally, being abstract, prices do not enter into causal relations, but sensations surely do.

It may come as little surprise that, from my vantage point, many of the aspects of the objector's conception of sensation are rooted primarily in the line of thinking I discussed in sections 4 and 5. To be sure, sensations and prices are extremely different. What I must deflate, though, is the sense that the aspects of price that make a NIP account of it plausible are absent in the case of sensation.

The objector points to nine or ten differences between prices and sensations that he intimates may be relevant for the applicability of a NIP account. I can indicate of course only briefly how I would approach each one. For instance, the objector claims that prices are not "things," but that sensations clearly are. The implicit assumption he relies on here is that anything that is a "thing" must either be identical to a physical phenomenon or provide a threat to physicalism. Am I saying then that sensations are not things? Some philosophers do argue as much. For instance, Daniel Dennett sometimes puts it that way: "When you have a pain, it isn't like having a penny; the pain *isn't a thing* that is in there" (2007, 75; his italics). But I am wary of characterizing things in this way. I suspect that the idea Dennett is attempting to convey here is one with which

I would ultimately be sympathetic. However, we must proceed with caution when it comes to categorizing phenomena by metaphysical groups such as "thing," "property," "object," and so forth. In our everyday discourse, we employ these expressions to draw different distinctions in different contexts. One and the same can count as a thing in one context and not in another.[46] If we do ask whether a particular phenomenon belongs to some such group, it is of utmost importance that we be very clear as to the criteria for inclusion in that group. Let me be clear: I am not insisting that in order to use the word "thing" without problem we must always be explicit about the criteria for something's being a thing. But in murky and dangerous philosophical waters such as we are in we should avoid asking for such categorizations without such clarification.

What then are the intended criteria for being a thing? If the criteria one intends in employing the expression "thing" includes being identical to a physical phenomenon, then no, on my account, sensations are not things. If it is sufficient for being a thing that what is in question bears properties, then certainly, sensations are things. But on that notion, so are prices.

What, though, does the *objector* here mean by "thing"? I doubt that he has a clear notion in mind. But that is beside the point. What I wish to argue is that, if he does have a clear notion in mind (or once such a notion has been specified) that notion either will be such that the primary reason for thinking that the notion applies to sensations stems from the line of thinking we are suspending, or will be a notion that does apply to sensations but the application of which does not threaten a NIP account. (Recall my sentence in section 3, after discussing price: "In fact, there are many such things for which it is correct to give a NIP account."

I would say similar things for the objector's claims that sensations are "objects," and that they are "concrete" and not "abstract." Am I claiming that sensations are not concrete? Again,

that depends on the criteria for being concrete. Certainly, sensations are such that we can have the same one, or same type of one, more than once. There is that sensation again, we say. But the same goes for prices. "There's that price again, $29.99; it seems as if everything in this store is $29.99." Of course, the ways in which we identify and reidentify prices and sensations are different (one's own sensations at least). To identify the price of something, we might look at a price tag or ask someone in the store. Our own sensations, however, often bring themselves to our attention. In some cases, less often perhaps, we become aware of sensations through introspection or by turning our attention "inward." But to say that in some cases we become aware of our sensations through introspection or inner attention is not yet to say anything about what such an activity amounts to or involves. As I argued in section 5, we also introspect our beliefs, but we do not suppose that that involves a process of observing the beliefs. Nor would we conclude that beliefs are something "there."

But sensations *hurt*!—or some do at least. I do not deny that. The crucial question is: What makes my objector think that the fact that some sensations hurt threatens a NIP account of sensation? Why does the objector take that fact to suggest that either sensations are identical with physical phenomena or physicalism is false? I believe that the inference the objector makes here—if some sensations hurt, then a NIP account of sensations cannot be correct—involves implicit reliance on a conception of sensation that has its primary source in the line of thinking. What else would support such a conception if not what we conclude from trying to inspect our sensations in introspection? We must be careful, now that we have suspended judgment about the soundness of the line of thinking, not to rely upon an assumption about, or conception of, sensation that is based in the line of thinking. I invite readers who are inclined to make this inference to consider the source of the conception of sensation, or of hurting, upon which their

inclination is based. I suspect that there will be little support for that conception other than that which appeals in one way or another to something like the line of thinking.[47]

About the features I just addressed—being a thing, being an object, not being abstract, and so forth—I argued that whether such predicates apply to sensations depends on what is meant by those predicates. On some senses, the predicates apply; on others they don't. And on the ones on which they apply, their application does not threaten a NIP account.

Unlike these features, however, there are others that, for the sake of argument, I will grant that prices have but that sensations do not. For instance, one might argue that prices are not occurrences or things that *happen* but that sensations are. If this is a genuine difference, it is a red herring. First of all, things also come to have a price, and prices go up and down. These are things that happen, and NIP accounts are no less applicable to them. But even if such things did not occur, my objector would need to provide reason for thinking that NIP accounts are not applicable to occurrences or happenings. Why should we think that? I would say the same about the social aspect of price, that is, the fact that having a price appears to depend on the existence of social customs. This too is a red herring. What makes a NIP approach applicable to price is not that it depends on social customs. Nor is it the fact that prices are "relational"—that is, the fact that a thing has a price only in virtue of standing in certain relations to other things—that makes the NIP approach applicable. There are many things that are relational (sunburns, for instance) that are not amenable to a NIP approach.[48]

Finally, let us turn to the objector's appeal to the causal efficacy of sensations. On the basis of the rough idea that prices are not "there," or that they are "abstract," the objector claims that prices themselves do not stand in causal relations, yet that sensations do, the idea being that NIP accounts do not allow for causal efficacy. But prices do stand in causal relations. As surely

sensations do too. The high price of gas has caused more people to fly instead of drive. Or, some philosophers claim that only events stand in causal relations. We could grant that and focus on events involving prices. The Federal Reserve's announcement caused a rise in prices; the rise in prices caused a run on the banks. NIP accounts are applicable to the activities of prices no less than they are to prices. NIP accounts do not preclude causal efficacy.

Before concluding, let me mention two other issues. First, are sensations physical on my view? Well, are prices physical? We would not want to say that prices are *immaterial*, at least not in any way that would threaten physicalism. Nor do I about sensations. But that does not mean that they are identical with physical phenomena. Here again, it would be crucial to delineate precisely what is being asked.

Second, it may seem as if, on my view, sensations could not be a subject that concerns scientists. On the contrary. The NIP approach is compatible with facts such as that what sensations one has depends crucially on the physical processes transpiring in one's body, and the fact that we can often specify some of the physical events that cause a particular sensation. Itches are often the result of physical impingements on the skin, headaches of certain physical processes in the head. I do not deny such facts. What I deny is that sensations are to be identified with any such causes, or with any physical phenomena that are the effects of such causes. Nor do I deny the importance and promise of scientific investigation into, say, the brain's role in these matters. To be sure, there is a great deal to be learned from the science of the brain as to what makes it possible that we have sensations at all. But we will not learn from such a science what sensation *is* (i.e., is identical to), just as we would not learn what price is (i.e., is identical to) by studying physical phenomena such as those in the market, or in banks, or anywhere. What we can learn about, though, is what we can do to the body, or skin, or external environment, in order

to produce or remove particular sensations, just as we might be able to learn a lot from economics about what we can do to effect particular changes in price.

10. Phenomenal Character

Let me summarize what I have argued in this chapter and then return to the notion with which I began, that of phenomenal character. One of my central goals has been to make a case for a Non-Identificatory Physicalist approach to sensation, according to which sensations are not identical to physical phenomena yet physicalism is true. I traced the paramount source of resistance to such an approach to a very common and influential line of thinking, and argued at length that we should be suspicious of that line of thinking. Another goal of the chapter was to make explicit the fact that many of the considerations that lead philosophers to endorse a representationalist approach to sensation and phenomenal character support a NIP theory of sensation no less than they do representationalism. Representationalism is not the only option once one has rejected the line of thinking. Representationalism proceeds to identify sensations with physical phenomena, and phenomenal qualities with some of the representational properties those phenomena have. But this step is not required; and by not taking it, we avoid the problems that beset representationalism. What, on my view, is it in virtue of that we have sensations? A third goal of the chapter has been to argue that we take seriously the possibility that answers to questions of this sort are unavailable, and perhaps not merely on account of our (perhaps "contingent") cognitive limitations. And indeed, physicalism does not entail its availability. This treatment of "in virtue of" questions, conjoined with my initial replies to questions concerning how we should thus characterize sensations (whether they are thus not "objects," or "concrete," or "physical," and so forth) was intended to sharpen the particular form of NIP approach that I am advertising.

I have focused primarily on sensations. I would urge that, for similar reasons, we give a NIP & NEP treatment of other, related phenomena—mental images, sense-perception, kinesthetic experiences, and so on. What about phenomenal character, the notion with which I began? At the outset, I suggested that the very topic of phenomenal character—or the very content of the notion of phenomenal character—may be more problematic to identify than we realize. I asked, "What are we talking about when we talk about phenomenal character?" Nothing in what I've argued in the paper depends on the thought that identifying this is more problematic than we realize; I want to emphasize this point. But some of the considerations that I've offered along the way would suggest that it would behoove us to pay more attention to this fundamental question.

The notion of phenomenal character is a philosopher's artifact, and so its content depends on the way in which it is, or was, formed by philosophers. How is it formed? There are a variety of ways in which it is introduced. Whether they result in the same notion, or the notion philosophers assume they are forming, is difficult to say. Our discussion of the line of thinking in section 5 rears its head once again. Many of the ways in which the notion is formed rely, either directly or ultimately, on it. But if we are suspending judgment about that line of thinking, we would similarly need to suspend judgment about the legitimacy of a notion formed on its basis.

There are other ways in which it might be proposed that we form the notion of phenomenal character or phenomenal quality, other than by way of the line of thinking. Sometimes it is formed by employing related notions that we already possess, such as that of "what it's like," as in "what it's *like* to have" such-and-such experience, an expression I used at the beginning of the chapter, or the notion of "how something *feels*." But the content and origin of these notions are themselves difficult to specify as well.[49] These notions are used in myriad ways in everyday discourse–they are "widely ramified" (RPP II §§218ff.)—and, at the outset at least, there's no

telling whether when philosophers attempt to form the notion of phenomenal character by way of them, they succeed in identifying a particular set of things or properties (or whatever) to which the notion applies. The mere fact that philosophers *say* "you know, the what-it's-like aspect of sensations, experiences, and the like" does not guarantee that they do.[50]

As I say though, the central arguments and conclusions of this chapter do not depend on there being anything problematic at all about the notion of the phenomenal. If we do continue to employ the notion, then for the same reasons for which I urged we should embrace a NIP account of sensation, we should also embrace a NIP account of phenomenal character. We should be no more confident about the line of thinking as it applies to phenomenal character than we should about the line of thinking as applied to sensations (that is, where what one intends to attend to in introspection is conceived of as phenomenal character, and not as a sensation, or where what one allegedly finds is conceived of as phenomenal character). Nor should we be more trustful of those questions, which I introduced in section 1, that are the starting place for much philosophy of mind—questions such as, "In what does the painfulness, the hurting quality of pains consist?". On the view I am entertaining in this paper, these questions, and the expectations and requirements with which philosophers standardly ask them, may be misguided from the beginning. Physicalism may well be a problematic thesis. If it is, though, it is not on account of the absence of satisfying answers to questions such as these.[51]

Notes

1. This way of characterizing intrinsic properties is not without its difficulties (see Langton and Lewis 1998 for a helpful discussion), but it is standard in the debates in question (see, for instance, Harman 1990, 664).

2. The inference from the idea that we observe or inspect phenomenal qualities to the idea that those properties are "intrinsic" properties is typically unspecified. Presumably, some nonintrinsic properties are observable as well, such as the similarity one shade of green bears to another. I will discuss this inference at length in section 5.

3. Prominent representationalists include Fred Dretske (1995), Michael Tye (1995, 2000), and Bill Lycan (1996). Representationalists differ on a variety of issues. They can be divided, for instance, according to whether they believe that the representational properties of mental phenomena can be entirely explained or accounted for in physical terms. The three pioneers I list above argue that they can. A rarer breed of representationalist believes that they cannot (e.g., Chalmers 2004). Hereafter, when I speak of representationalists, I mean the former kind. Another division concerns precisely what representational properties are sufficient for phenomenal character; representationalists acknowledge that not every representational property constitutes (or is sufficient for) phenomenal character, but they disagree as to what more is needed. (Lycan, for instance, points to the functional role of the representation.) Another difference is that some representationalists identify the phenomenal character not with representational properties of the experience but with properties that are represent*ed*, by the experience, i.e., with properties of the intentional content or object of the experience. In either case, experiences with phenomenal character have representational properties whereby the experience represents something as having particular properties, and there is no more to phenomenal character than all of that.

4. A tremendous amount has been written, on both sides of the debate, about these difficulties for representationalism, especially the alleged counterexamples. I will return in section 7 to the significance of these and several other problems representationalists face.

5. Similarly, what makes the condition of the skin on my arm a *sun*burn, as opposed to another kind of burn (or to a condition that has come about in a way other than through heat), is that the skin stands in a particular causal relation to the sun. But the sunburn itself is on (or perhaps in) my arm and does not

extend beyond my body. Also note that, if some philosophers do identify bearers of phenomenal character with physical phenomena that include phenomena outside of the body (as, for instance, Chalmers and Clark 1998 identify particular forms of cognition), these accounts are no less in concert with the guiding assumption. (Noë 2004 may be an example.)

Not all physicalist views identify bearers of phenomenal character with physical phenomena. Eliminative materialism, for instance, does not, but that is because, according to that view, there are no sensations or experiences at all. Nor perhaps does logical behaviorism or adverbialism, but these latter views are largely remnants of the past, and I do not wish to resurrect them; I want to provide a more plausible way of flouting the guiding assumption. I will return briefly to these views later in the chapter.

Standard functionalist views, on the other hand, which are certainly alive and well today (sometimes conjoined with representationalism, as in Lycan 1996) do identify bearers of phenomenal character with physical phenomena. Many functionalists identify sensations and the like with whatever physical states or processes play the relevant functional roles (or stand in the relevant causal relations). Other functionalists (sometimes called "role functionalists") do not identify the bearers with the physical states or processes themselves (i.e., with those that play the relevant roles) but rather identify them with higher-level properties, such as the property of *being in a state that* plays the relevant functional role (or that stands in the relevant causal relations). This higher-level property is no less specifiable in the terms of physics.

Daniel Dennett's interesting and rich approach to consciousness may not endorse the guiding assumption, though it is difficult to say. Dennett's approach is complicated and has undergone various iterations. I will unfortunately be unable in this chapter to give his views the attention they deserve. For some representative writings, see Dennett 1978, 1988, 1991, and 2005.

6. This is often taken to be a "minimal" physicalism (e.g., Lewis 1983), and it is what I shall mean by "physicalism" hereafter.

7. I choose sensations because in many ways they would appear to provide the most difficulty for the view I am proposing.

8. Nor would it answer the question I intended to ask to reply, "In this case, the price is $2.79." At any rate, the milk has its price; it does not have $2.79.

9. Or at least, they do not for anyone who supposes that the psychological attitudes alluded to above, on which economic customs depend, do not themselves provide a problem for physicalism.

10. Four observations are in order here. First, I suppose that one could dig in one's heels at this point and insist that the price is identical to—nothing but—a relational property that the milk has, the property of standing in all the relations it does to the complicated set of customs, attitudes, and so on I mention above. This strikes me as incorrect. For one thing, it would not seem correct to say that the complicated relational property of standing in all the relations that the milk does to the customs, attitudes, etc. is itself $2.79. But that is what the milk's price is. Moreover, it is unclear to me what the motivation for insisting as much is. Is it a worry that if we do not identify the price with something like this, we will be threatened with a worrisome "dualism," or with the prospect that physicalism is false?

Second, one might wonder, now that one has seen in application what I intend by a "NIP" approach, whether representationalists do not themselves give NIP accounts of phenomenal character. Some of them may; this is a difficult interpretive question. Regardless, the view that I claimed in section 1 is virtually not on the table—and which representationalists clearly do not hold—is a NIP account of *sensations*. What I will argue in the latter part of the chapter is that once we are open to giving a NIP account of sensations, we have little need for representationalist theories of phenomenal character.

Third, in saying that we should give a NIP account of price I am not saying we should give a NIP account of the property of *having a price* (though nor am I saying we should not). What I am ruminating about in the market, in my confused line of thinking, is what the milk *has* when it has a price.

Fourth, the example of price I hope gives the reader more of a picture of the sort of question I am claiming is misguided, as applied to sensations, and the sense in which it is misguided. Asking what

price *is*, in the sense of what it is identical to, or consists in (at least in the sense in which I am asking myself that in the store) is misguided: it is based on a mistake, a misconception of the phenomenon of price. The goal of the chapter is to make the case that we should take seriously the idea that this is true for similar questions asked about sensation (or at least, those that require answers that are not expressed in the very concepts of sensation; I will say much more about this requirement, and the questions that have it, in section 8).

11. Thanks to Paul Dunn for this example.

12. Note that at the beginning of the paragraph I said, "In fact, there are many such things." I trust my reader did not object to this use of "things." This brings out one of the potential dangers of saying that on a NIP account of price "price is not a thing." One can even see how one might be inclined—if only to indicate some of the issues that must be navigated here and of which we might be careful—to say about price: that it is not a something, but not a nothing either.

13. What I've described might be thought of as more than a line of thinking, as it's supposed to involve, in addition to thinking or reasoning, the introspection of the phenomenon, and on some conceptions of thinking and introspection, introspection is not a form of thinking. I will continue to refer it as "the line of thinking"; however, my doing so does not involve a prior commitment to a particular conception of thinking or introspection. (Of course, a prior subscription to the soundness of the line of thinking might involve such a commitment.)

14. This is how Brian O'Shaughnessy 2000 puts it, for instance (ch. 18).

15. In fact, the line of thinking is ultimately also the primary source of what I called at the outset "the standard view" of phenomenal character, according to which phenomenal character involves properties that are both intrinsic and introspectively accessible.

The line of thinking can vary with respect to precisely what one *intends* to attend to in introspection. That is, one might intend to inspect a sensation or an experience (i.e., to inspect a bearer of "phenomenal character"), or one might intend to inspect what one conceives of as the phenomenal character of one's

sensation or experience. Accordingly, one might conclude that "it" (the sensation) is what is "there"—or that it has introspectively accessible, intrinsic properties—or, one might conclude that "it" (the phenomenal character) is what is "there," and that it is an accessible, intrinsic property. I'll take as my example the former case, where the intention and attention concern sensations. But the considerations I'll offer will apply similarly to the case of phenomenal character.

16. For instance, see Comte 1974. The fourth of the four points I will make in section 5 connects up with this general idea.

17. The extent to which the line of thinking does rely on an inner-sense model of introspection is a difficult issue. Placing it to the side will not affect the points I wish to make. Midway through section 5 I will have something to say about philosophers who embrace both the line of thinking and an inner-sense model of introspection.

18. Gil Harman's articulation of this idea is perhaps the chief source of the recent rise of representationalism (Harman 1990, though as he acknowledges, he was not the first to make the point); subsequent appeal to the idea has proven one of the most influential ways to resist the standard view of phenomenal character.

19. E.g., Harman 1990; Shoemaker 1994; Tye 1995, Crane 2003; Dretske 2003. There are many other arguments that would tell against the soundness of the line of thinking. Many considerations in Wittgenstein's work speak against it, only a small portion of which are to be found in the celebrated and overworked passages concerning private language and the beetle in the box. For just one explicitdiscussion, see *Philosophical Investigations* §412 and the surrounding passages. For another compelling explanation of the inadequacy of crucial elements of the line of thinking, see David Finkelstein's treatment of what he calls "detectivism" (2003, ch. 1).

20. In saying that the line of thinking is the primary source of the fundamental misconception I am arguing against, I am not suggesting that every philosopher who endorses the conception must have at some point engaged entirely in the line of thinking. I do think that many philosophers do acquire the

conception in this way; however, my argument does not rely on this empirical fact. Even if some philosophers do not acquire the conception from having engaged in the line of thinking, it is difficult to see what compelling support for the conception can be provided that itself does not ultimately depend on the soundness of the line of thinking. I will return to this point in section 9.

21. Wittgenstein: "It indicates a fundamental misunderstanding, if I'm inclined to study my current headache in order to get clear about the philosophical problem of sensation" (PI §314). See also PI §§294, 370.

22. Strictly speaking, the line of thinking does not *require* that we assume that sensations are things that we can attend to, for the purpose of investigating what they are. We need merely to assume that sensations *may be* such things. But as a matter of fact, we *do* make the former assumption, and thus we approach our inquiry about sensation fully engaged with a prior and substantial conception of sensation.

23. The link between, on the one hand, the cognitive and conceptual capacities and mechanisms involved in the introspection of sensations, experiences, and their phenomenal qualities, and on the other hand, the content of the introspection—i.e., *what* is introspected—receives relatively little attention in discussions of phenomenal character and its introspection. Frequently lacking, in particular, is discussion of the dependence of the latter on the fulfillment of the formidable conditions on the former. (Papineau 2002, ch. 4, is one exception.) For more on this issue and its significance, see Ellis 2010.

24. See Balcetis and Dunning 2006, 2007, 2010; Changizi and Hall 2001; Li and Warren 2004.

25. Notice another significant difference between everyday and philosophical introspection. I have been discussing here the step of identifying a sensation to attend to. However, in everyday contexts, intentions to turn one's attention to a sensation are in fact quite rare. Not that we are not often aware of our sensations in everyday contexts, or that we do not sometimes introspect or attend to them. But the way in which we typically become aware of them, or begin to attend to them, is by their making themselves present to our awareness, by their intruding on our attention, as it were. Only in quite rare

cases, as in the doctor's office, do we form an intention to turn our attention to a sensation that we are not already attending to. (And even in the doctor's case, it is a sensation that we have presumably mentioned to the doctor and thus already been thinking about.) This is no doubt related to another point, also significant for our purposes, that the expression "introspection" is considerably less natural as applied to sensation than it is to attitudes such as belief and desire.

26. $0.8^4 = 0.41$. I am assuming that if any of the assumptions is incorrect then the line of thinking is unsound. Introducing percentages may strike the reader as unnecessary. I do so for emphasis. Especially in cases in which a line of thinking is so dominant and influential, the exponential effect of doubt is crucial to bear in mind. Of course, for many of the assumptions I have identified, I believe the probability assigned should be quite a bit lower than 0.8.

27. This is true even of many proponents of the standard view of phenomenal character. There are some exceptions, of course, even among those who resist the standard view. Bill Lycan, for instance, endorses a version of the inner-sense model (Lycan 1996).

28. Note that not all inner-sense theories would support such an explanation. The explanation would seem to depend on a theory according to which what we "perceive" in introspection are "objects," and not, say, facts. See Shoemaker 1994 for the difference between these two kinds of inner-sense theory.

29. One might reply that, if the reasoning is faulty in the skeptical case, then it is faulty full stop. But how could such a generalization be justified? What makes the reasoning faulty in the skeptical case could be something having to do with that case; we cannot assume that whatever antiskeptical argument is ultimately provided to undermine the reasoning leading to skepticism is similarly applicable in the introspective case.

30. For a rather rare and quite helpful, sustained discussion of the question, see Schwitzgebel 2007.

31. This is a distinct possibility from, and does not entail, the possibility that we have sensations only when we are engaging in introspection, i.e., that we never have unintrospected sensations. That you do not have strong reason to believe that you had a sensation on that part of your skin before I invited you

to introspect the feeling of your clothes against it does not entail that you do not have reason to believe that you sometimes have sensations that you are not introspecting.

32. Recall, we cannot at this point appeal to everyday introspection in order to motivate the idea that introspection is rarely if ever illusory, for the relation between the two forms of introspection is still in question.

33. There is also to be taken into account the substantial amount of empirical work that purports to show that introspection even in everyday contexts is not as reliable as we might have assumed. See, for instance, Schwitzgebel 2011.

34. Harman: "Look at a tree and try to turn your attention to intrinsic features of your visual experience. I predict you will find that the only features there to turn your attention to will be features of the presented tree, including relational features of the tree 'from here'" (1990, 667).

35. Fred Dretske (1995), for instance, sees this as the most significant obstacle for representationalism. Dretske, himself a representationalist, proceeds to argue that externalism about phenomenal character may in fact be correct. Also, a note on the expression "internal": In debates about phenomenal character, to say that two individuals are "internally" identical is to say that they are *physically* identical. The distinction intended is thus quite different from the distinction Putnam has in mind in his influential arguments concerning meaning (Putnam 1975).

36. See Lowe 2000, 58, for a suggestion along these lines.

37. See, e.g., Loewer 1997. I will return to this particular issue in section 8.

38. Of course, representationalists have argued ardently against the objections I have outlined here, and I cannot here adjudicate the prospects of their significant replies. It is also worth noting that some of the objections I have mentioned are sometimes supported in a way that may themselves ultimately depend on the soundness of the line of thinking I discussed in section 5. However, the force of, at the very least, the first, fourth, and fifth kinds of objection does not require such dependence.

39. In his general survey of representationalism, for instance, Bill Lycan, himself one of the pioneers of the view, catalogs four arguments in favor of representationalism. He discusses at greatest length the two I mention here. The other two, which he calls "the argument from veridicality" and "the argument from seeming" (Lycan 2008), have had much less influence.

40. Many philosophers have made the point that considerations about transparency are consistent with a variety of views of the mental. See, for instance, Martin 2002, who subsequently employs them in service of disjunctivism about perception.

41. I have in mind especially the traditional form of adverbialism, as in Chisholm 1957; see Jackson 1977 for an influential critique. As I noted above, whether adverbialism and linguistic behaviorismare forms of NIP is a matter of interpretation, andmy own proposal does not depend on their so being. Indeed, a variety of views have been called "adverbialism." See Butchvarov 1980 for a discussion of some of them and how they differ. See Kriegel 2007 for an interesting, nonstandard form of it.

42. For another exposition and defense of some of these ideas, see Stroud's contribution to this volume, as well as Stroud 1990, 1996, and 1998. See McDowell 1984, 1993 for some similar themes.

43. Here we return to the fifth of the five kinds of objections I mentioned have been raised for representationalism. Stroud's conclusion stands opposed to the standard assumption, made by all physicalist representationalists, that the representational properties of mental phenomena can be satisfactorily explained entirely in physical terms. That assumption is most often nowadays grounded in one or another form of "informational" (e.g., Fodor 1990) or "teleological" (e.g., Dretske 1995) semantics. However, these approaches to representational content face difficulties, many of which have their ultimate source in Wittgenstein's considerations about rule-following.

Some of these difficulties, it is worth pointing out, might ultimately render unavailable particular sorts of answers to "in virtue of" questions concerning price (e.g., "What is it in virtue of that the milk has a price?"). In my rumination at the grocery, I claimed to be able to say roughly what it is in virtue of that the gallon of milk has its price; it is a very complicated set of facts or states of affairs involving social

customs, psychological attitudes, behavioral dispositions, and so on (e.g., the grocer's disposition to chase me or call the police if I were to take the milk without paying for it). The difficulties Stroud discusses may not threaten this claim. But they may threaten the further idea that the question "What is it in virtue of that the milk has a price?" can be adequately answered by citing facts entirely expressed in physical terms.

44. It would be incorrect, for instance, to conclude from Hacking's depiction that Wittgenstein is suggesting that what we are ultimately interested in, and should focus upon, are primarily or entirely words and language.

> One ought to ask, not what images are or what happens when one imagines something, but how the word "imagination" is used. But that does not mean that I want to talk only about words. For the question of what imagination essentially is, is as much about the word "imagination" as my question. And I am only saying that this question is not to be clarified — neither for the person who does the imagining, nor for anyone else — by pointing; nor yet by a description of some process. The first question also asks for the clarification of a word; but it makes us expect a wrong kind of answer. (PI §370)

It would also be incorrect to conclude that we do not have images after all, or that our images are identical to the practices in which they are embedded.

45. For a more nuanced and plausible account of what can be learned from Wittgenstein's writings about the relation between expressions of sensations and first-person ascriptions of sensations, see Finkelstein 2003.

46. See PI §§10, 13 and their surrounding passages for the significance of the fact that one word can be used to draw different distinctions.

47. This is why, as I explained in note 20, my argument does not depend on claiming that the only way in which the conception of sensation I want to reject is formed is by first engaging in the line of thinking. What is most important is that compelling support would appear not to be available from other sources.

48. Moreover, nothing in my approach to sensations entails that sensations are not relational in this sense either; nor that people's having them does not depend on social customs.

49. See Hellie 2007 for the difficult enterprise of ascertaining the syntactic and semantic properties of the expression "there's something it's like".

50. Yet another way in which the notion of phenomenal character is sometimes introduced or taught is by appeal to similarities and differences among particular mental phenomena, such as that "the experience of seeing purple is more like, in respect of phenomenal character, the experience of seeing blue than it is like the experience of smelling vanilla" (Byrne 2001, 200). How though does one identify those "respects"? One way perhaps would be to underscore what the experiences are *of*, what they represent; but that might be to presuppose, or to stipulate, a conception of phenomenal character as something exhausted by its representational content. If one does not appeal primarily to the representational features of such experiences though, then simply to grasp the ways in which, or the mere fact that, some of those experiences are more similar to one another than to others would seem to require a prior conception of phenomenal character. But this is a difficult issue to which I cannot do justice here.

51. I am grateful to my audiences at the "Wittgenstein and Philosophy of Mind" conference at UC Santa Cruz (June 2007) and the Wittgenstein Workshop at the University of Chicago (February 2008). I am especially indebted to Brendan O'Sullivan, my commentator at the conference, with whom I have been discussing issues in the philosophy of mind for almost twenty years now, albeit with little success of persuading of him of anything. Conversations with Brendan are invariably helpful. I also benefited a great deal from comments I received on earlier drafts of this chapter from Jason Bridges, William Child, James Conant, Alice Crary, Janette Dinishak, Christoph Durt, Todd Ganson, Daniel Guevara, Miles Hatfield, Sam Levey, Nicoletta Orlandi, Jay Peters, Robert Schroer, Abraham Stone, and Julie Tannenbaum. My largest gratitude is to Daniel Guevara, whose wisdom never fails to impress me. I could not ask for a better colleague, friend, or philosophical interlocutor. Directing this conference with him was a pleasure from start to finish.

References

Aydede, Murat. "Introduction: A Critical and Quasi-Historical Essay on Theories of Pain." In *Pain: New Essays on Its Nature and the Methodology of Its Study*, ed. Murat Aydede. Cambridge, Mass.: MIT Press, 2005.

Balcetis, Emily, and David Dunning. "See What You Want to See: Motivational Influences on Visual Perception." *Journal of Personality and Social Psychology* 91 (2006): 612-625.

Balcetis, Emily, and David Dunning. "Cognitive Dissonance and the Perception of Natural Environments." *Psychological Science* 18 (2007): 917-921.

Balcetis, Emily, and David Dunning. "Wishful Seeing: More Desired Objects Are Seen as Closer." *Psychological Science* 21 (2010): 147-152.

Block, Ned. "Troubles with Functionalism." *Minnesota Studies in the Philosophy of Science* 9 (1978): 261–325.

———. "Inverted Earth." In *The Nature of Consciousness: Philosophical Debates*, ed. Ned Joel Block, Owen J. Flanagan, and Güven Güzeldere. Cambridge, Mass.: MIT Press, 1997. First published in *Action Theory and Philosophy of Mind*, (Atascadero, CA: Ridgeview, 1990).

———. "Mental Paint and Mental Latex." *Philosophical Issues* 7 (1996): 19-49.

Boghossian, Paul, and David Velleman. "Colour as a Secondary Quality." *Mind* 98 (1989): 81–103.

Burge, Tyler. *Origins of Objectivity*. Oxford: Oxford University Press, 2010.

Butchvarov, Panayot. "Adverbial Theories of Consciousness." In *Studies in Epistemology*, ed. Peter A. French, Theodore Edward Uehling, and Howard K. Wettstein. Minneapolis: University of Minnesota Press, 1980.

Byrne, Alex. “Intentionalism Defended.” *Philosophical Review* 110 (2) (2001): 199–240.

Chalmers, David. “The Representational Character of Experience.” In *The Future for Philosophy*, ed. Brian Leiter. Oxford: Oxford University Press, 2004.

Chalmers, David, and Andy Clark. “The Extended Mind.” *Analysis* 58 (1) (1998): 7–19.

Changizi, Mark, and Warren Hall. “Thirst Modulates a Perception.” *Perception* 30 (2001): 1489-1497.

Chisholm, Roderick. *Perceiving: A Philosophical Study*. Ithaca, N.Y.: Cornell University Press, 1957.

Comte, Auguste. *The Essential Comte*. Trans. Margaret Clarke. New York: Barnes and Noble, 1974.

Crane, Tim. “The Intentional Structure of Consciousness.” In *Consciousness: New Philosophical Perspectives*, ed. Quentin Smith and Aleksandar Jokić. Oxford: Clarendon Press, 2003.

Crane, Tim, and D. H. Mellor. “There is No Question of Physicalism.” *Mind* 99 (1990): 185–206.

Dennett, Daniel. “Toward a Cognitive Theory of Consciousness.” In *Brainstorms: Philosophical Essays on Mind and Psychology*. Cambridge, Mass.: MIT Press, 1978.

———. “Quining Qualia.” In *Consciousness in Modern Science*, ed. A. Marcel and E. Bisiach. Oxford: Oxford University Press, 1988.

———. *Consciousness Explained*. Boston: Little, Brown, 1991.

———. *Sweet Dreams: Philosophical Obstacles to a Science of Consciousness*. Cambridge, Mass.: MIT Press, 2005.

———. “Philosophy as Naïve Anthropology: Comment on Bennett and Hacker.” In Maxwell Bennett, Daniel Dennett, Peter Hacker, and John Searle, *Neuroscience and Philosophy: Brain, Mind, and Language*. New York: Columbia University Press, 2007.

DeRose, Keith. “Solving the Skeptical Problem. *Philosophical Review* 104 (1) (1995): 1–52.

Dretske, Fred. *Naturalizing the Mind*. Cambridge, Mass.: MIT Press, 1995.

———. “How Do You Know You Are Not a Zombie?” In *Privileged Access: Philosophical Accounts of Self-Knowledge*, ed. Brie Gertler. Aldershot: Ashgate, 2003.

Ellis, Jonathan. “Phenomenal Character, Phenomenal Concepts, and Externalism.” *Philosophical Studies* 147 (2) (2010): 273–299.

Finkelstein, David. *Expression and the Inner*. Cambridge, Mass.: Harvard University Press, 2003.

Fodor, Jerry. *A Theory of Content and Other Essays*. Cambridge, Mass.: MIT Press, 1990.

Hacking, Ian. “Wittgenstein as Philosophical Psychologist.” In *Historical Ontology*. Cambridge, Mass.: Harvard University Press, 2002. First published in *The New York Review of Books*, 1 April, 1982.

Hardin, C. L. *Color for Philosophers: Unweaving the Rainbow*. Indianapolis: Hackett, 1993.

Harman, Gilbert. “The Intrinsic Nature of Experience.” In *The Nature of Consciousness: Philosophical Debates*, ed. Ned Joel Block, Owen J. Flanagan, and Güven Güzeldere Cambridge, Mass.: MIT Press, 1997. First published in *Philosophical Perspectives* 4 (1990): 31–52.

Hellie, Benj. “‘There’s something it’s like’ and the Structure of Consciousness.” *Philosophical Review* 116 (2007): 441-463.

Hempel, Carl. “Reduction: Ontological and Linguistics Facts.” In *Philosophy, Science, and Method: Essays in Honor of Ernest Nagel*, ed. E. Nagel, S. Morgenbesser, P. Suppes, and M. G. White. New York: St. Martin’s, 1970.

Jackson, Frank. *Perception*. Cambridge: Cambridge University Press, 1977.

Kriegel, Uriah. “Intentional Inexistence and Phenomenal Intentionality.” *Philosophical Perspectives* 21 (2007): 307–340.

Langton, Rae, and David Lewis. “Defining ‘Intrinsic.’” *Philosophy and Phenomenological Research* 58 (1998): 333–345.

Lewis, David. “New Work for a Theory of Universals.” *Australasian Journal of Philosophy* 61 (1983): 343–377.

———. “Elusive Knowledge.” *Australasian Journal of Philosophy* 74 (1996): 549–567.

Li, Li, and William Warren Jr. “Path Perception during Rotation: Influence of Instructions, Depth Range, and Dot Density.” *Vision Research* 44 (2004): 1879-1889.

Loar, Brian. “Phenomenal States.” In *The Nature of Consciousness: Philosophical Debates*, ed. Ned Joel Block, Owen J. Flanagan, and Güven Güzeldere. Cambridge, Mass.: MIT Press, 1997.

———. “Transparent Experience and the Availability of Qualia.” In *Consciousness: New Philosophical Perspectives*, ed. Quentin Smith and Aleksander Jokić. Oxford: Oxford University Press, 2003.

Loewer, Barry. “A Guide to Naturalizing Semantics.” In *A Companion to the Philosophy of Language*, ed. B. Hale and C. Wright. Malden, Mass.: Blackwell, 1997.

Lowe, E. J. *An Introduction to the Philosophy of Mind*. Cambridge: Cambridge University Press, 2000.

Lycan, William. *Consciousness and Experience*. Cambridge, Mass.: MIT Press, 1996.

———. "Representational Theories of Consciousness." In *The Stanford Encyclopedia of Philosophy*. Fall 2008. Online. Available: http://plato.stanford.edu/archives/fall2008/entries/consciousness-representational/.

Martin, Michael. "The Transparency of Experience." *Mind & Language* 17 (2002): 376–425.

McDowell, John. "Aesthetic Value, Objectivity, and the Fabric of the World." In *Pleasure, Preference and Value*, ed. Eva Schaper. Cambridge: Cambridge University Press, 1983.

———. "Wittgenstein on Following a Rule." *Synthese* 58 (1984): 325–363.

———. "Meaning and Intentionality in Wittgenstein's Later Philosophy." In *Midwest Studies in Philosophy*, ed. Peter A. French, Theodore Edward Uehling, and Howard K. Wettstein. Notre Dame: University of Notre Dame Press, 1993.

McGinn, Colin. *Problems in Philosophy: The Limits of Inquiry*. Oxford: Blackwell, 1993.

———. *The Character of Mind: An Introduction to the Philosophy of Mind*. Oxford: Oxford University Press, 1997.

Noë, Alva. *Action in Perception*. Cambridge, Mass.: MIT Press, 2004.

O'Shaughnessy, Brian. *Consciousness and the World.* Oxford: Oxford University Press, 2000.

Papineau, David. *Thinking about Consciousness*. Oxford: Oxford University Press, 2002.

Peacocke, Christopher. *Sense and Content: Experience, Thought, and their Relations*. Oxford: Clarendon Press, 1983.

Putnam, Hilary. "The Meaning of 'Meaning.'" *Minnesota Studies in the Philosophy of Science* 7 (1975): 131–193.

Rey, Georges. "A Narrow Representationalist Account of Qualitative Experience." In *Language, Mind, and Ontology*, ed. James E. Tomberlin. Oxford: Blackwell, 1998.

Schwitzgebel, Eric. “Do You Have Constant Tactile Experience of Your Feet in Your Shoes? Or Is Experience Limited to What’s in Attention?” *Journal of Consciousness Studies* 14 (2007): 5–35.

Schwitzgebel, Eric. *Perplexities of Consciousness*. Cambridge, Mass.: MIT Press, 2011.

Searle, John. *Rediscovery of the Mind*. Cambridge, Mass.: MIT Press, 1992.

Shoemaker, Sydney. “The Royce Lectures: Self-Knowledge and ‘Inner Sense.’” In *The First-Person Perspective and Other Essays*. Cambridge: Cambridge University Press, 1996. First published in *Philosophy and Phenomenological Research*, 54 (1994): 249-314.

Stroud, Barry. “Wittgenstein on Meaning, Understanding, and Community.” In *Wittgenstein—Eine Neuberwertung: Wittgenstein—Towards a Re-evaluation*, ed. R. Haller and J. Brandl. Vienna: Verlag Hölder-Pichler-Tempsky, 1990.

———. “Mind, Meaning, and Practice.” In *The Cambridge Companion to Wittgenstein*, ed. H. Sluga and D. Stern. Cambridge: Cambridge University Press, 1996.

———. “The Theory of Meaning and the Practice of Communication.” *Critica* 30 (88) (1998): 3–28.

Tye, Michael. *Ten Problems of Consciousness*. Cambridge, Mass.: MIT Press, 1995.

———. *Consciousness, Color, and Content*. Cambridge, Mass.: MIT Press, 2000.

Williams, Bernard. *Ethics and the Limits of Philosophy*. London: Fontana, 1985.

Wittgenstein, Ludwig. .*Zettel*, ed. G. E. M Anscombe and G. H. von Wright, trans. G. E. M. Anscombe. Oxford: Blackwell, 1967.

———. *Remarks on the Philosophy of Psychology*, vol. I, ed. G. E. M. Anscombe and G. H. von Wright, trans. G. E. M. Anscombe (Oxford: Blackwell, 1980).

———. *Remarks on the Philosophy of Psychology*, vol. II, ed. G. H. von Wright and H. Nyman, trans. C. G. Luckhardt and M. A. E. Aue (London: Blackwell, 1980).

———. *Last Writings on the Philosophy of Psychology Vol. 1: Preliminary Studies for Part II of the* Philosophical Investigations, ed. G. H. von Wright and Heikki Nyman, trans. C. G. Luckhardt and Maximilian A. E. Aue. Chicago: University of Chicago Press, 1982.

———. *Last Writings on the Philosophy of Psychology Vol. 2: The Inner and the Outer*, ed. G. H. von Wright and Heikki Nyman, trans. C. G. Luckhardt and Maximilian A. E. Aue. Chicago: University of Chicago Press, Chicago. 1992.

———. *Philosophical Investigations*, 3rd ed. Trans. G. E. M. Anscombe. Oxford: Basil Blackwell and Mott, 1958.

———. *Philosophical Investigations*, ed. P. M. S. Hacker and Joachim Schulte, trans. G. E. M. Anscombe, P. M. S. Hacker, and Joachim Schulte, 4th edn (Oxford: Wiley-Blackwell, 2009).

———. *Philosophy of Psychology—A Fragment*, in L. Wittgenstein, *Philosophical Investigations*, ed. P. M. S. Hacker and Joachim Schulte, trans. G. E. M. Anscombe, P. M. S. Hacker, and Joachim Schulte, 4th edn (Oxford: Wiley-Blackwell, 2009).