

Phenomenal Character and Philosophical Introspection¹

Jonathan Ellis

1. Introduction

I want to talk today about a notion that's central to a great deal of work going on in the philosophy of mind. The notion of phenomenal character. Before the expression "phenomenal character," the term of choice was "qualia"; and before that, "what it's like" qualities. Other closely related expressions are "phenomenal properties," "raw feels," and "phenomenal consciousness." These terms are sometimes used interchangeably, sometimes not. And while each of them brings its own idiosyncrasies, the thoughts behind their discussions are similar. I'll take phenomenal character as my example, but much of what I say will go for all the rest.

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'll be discussing today, but for the moment, let me say 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. 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."² I don't fault these philosophers for taking this strategy in identifying their topic. When reading

¹ This is a draft; please do not quote from it without author's permission (jellis@ucsc.edu).

² Block, N. "Troubles with Functionalism".

these introductions, I myself typically think I know exactly what they're talking about—as clearly they think too—even if we're 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 first remind you of two very prominent accounts of phenomenal character today. The first view, sometimes called “the standard view,”³ holds that phenomenal character—for instance, the phenomenal character of my headache—is constituted by properties of the headache that are both intrinsic and introspectively accessible. To say that a property of something is intrinsic, according to these philosophers, is to say that the thing has the property independently of the relations that thing stands in to other things. Having a sister is not an intrinsic property; having a certain shape is.

This view, however, is thought to provide serious problems for so-called “materialist” approaches to the mind. In order for a materialist to account for an intrinsic property, it's said, that property must be identifiable with some intrinsic material or physical property. And many theorists find it difficult to see how phenomenal character could be identical with an intrinsic material property.

The second view, which is receiving a tremendous amount of attention these days, is called “representationalism.” According to this view, the phenomenal character of my headache is to be identified with some set of my headache's representational properties. The throbbing aspect of my headache is nothing but the property of the headache whereby it represents my head to be throbbing. According to representationalists, there are no properties that are both intrinsic and introspectively accessible.

³ Loar, B. “Transparent Experience and the Availability of Qualia” (2003).

While these two accounts of phenomenal character do not exhaust the views, they certainly garner the most attention and support.⁴ They're the sorts of views that Ned Block had in mind, when he said—in another line of his that's 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.”

What I want to do today is to introduce a third perspective—or really, to re-introduce to contemporary philosophy of mind a third perspective. It's an approach that, in many ways, I find in Wittgenstein. But this isn't an interpretive paper, and there'll 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. I'm sure that he'd resist some of it, if not all of it.

There's a current assumption guiding so much philosophy of mind that we need to say or ascertain what phenomenal character *is*, or what things that have phenomenal character, such as sensations, are, and that until we do, we will have on our hands a significant mystery, or problem, or something we don't understand—that we will have the so-called “Problem of Consciousness.” And so the views in currency attempt to do just that. Sensations are brain processes; phenomenal character is identical to a set of representational properties; and so on. I want to suggest that we take seriously the idea that this guiding assumption is false, that it is

⁴ Shoemaker (1994), for instance, holds a hybrid view according to which phenomenal qualities are intrinsic but *not* introspectively accessible. This view is much less common. It is also worth noting another view, “intentionalism,” which holds that phenomenal character supervenes on the representational properties of experience: there can be no difference in phenomenal character without a difference in content (e.g., Byrne 2001). Intentionalism is weaker than representationalism. Many intentionalists do deny that anything has intrinsic, introspectively accessible properties (and are representationalists); however, they need not. Intentionalism and the standard view are compatible.

a misguided enterprise to try to say what phenomenal character or sensations are, or what they consist in, or what they are identifiable with. Not because we don't have sensations, but because the question about what these things "are" is, in an important sense, misguided or ill-formed. The same goes for such questions as applied to particular sensations. Questions like, "What is my headache?" It's not just questions about sensations in general.

Not that the words "What are sensations?" are never appropriate to utter. A child might ask his mother "Mommy, what are sensations?" And she might say, "Well, you know how when you bang your toe and it hurts? That's a sensation." Or the doctor might ask "Now, what is this sensation you keep complaining about in your elbow?" The idea I want to entertain is that it's the philosopher's question that is problematic.

This isn't a new proposal of course. It's the sort of general strategy that is often associated with Wittgenstein—rightly or wrongly: Such-and-such question is ill-formed; it's misguided to try to answer it; you need to see why you're led to ask that question in the first place. That strategy is met with such scorn and resistance nowadays, especially when it's applied to phenomena such as sensations or phenomenal character, which it rarely is. But, in my view, the strategy deserves much more attention than it gets.

Now, one can certainly find these, or at least related, ideas in Wittgenstein scholarship. My point is that they're remarkably 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 the APA. The view isn't even on the table.

My chief goal in this paper, then, is to bring one version of this view into clearer focus, and to do so in a way that preempts the kind of misunderstandings and caricatures it can elicit. I want to remind those who work in this field that this is a possible approach to take to these

questions—and indeed that it’s plausible. I don’t hope to establish here that it’s the correct approach. But I do hope to convince you that it should at least be taken much more seriously than it is.

As I say, what I’m proposing will strike many as decidedly implausible. *Why* is it misguided to ask what phenomenal character is, or what sensations are? Those certainly seem to be perfectly appropriate questions, indeed urgent ones. After all, aren’t there many things about which it does make good sense to ask what it is, or what it consists in? Water, for instance. And oxygen. We now have answers to what those things are. It’s easy to say just don’t ask the question about sensations; you need to give me reason for thinking I shouldn’t.

In the first half of the talk, then, I want to identify and scrutinize what I take to be the paramount source of the conviction that it is appropriate in philosophical settings to ask what sensations and phenomenal character are, or what they consist in. In the second half, I’ll discuss what, on the approach I have in mind, we should or can say about these phenomena. To this end, I shall discuss some of Wittgenstein’s passages about signification, early in the *Investigations*. Some of Wittgenstein’s discussions about language can be so helpful when inquiring about the mind philosophically—and in a strikingly wide variety of ways—even if those discussions are not explicitly about mental phenomena at all.

Before proceeding, though, let me say something about what I mean when I say these questions may be misguided or ill-formed. Am I saying they’re nonsensical, or meaningless? Well, that would depend on what we mean by “nonsensical” and “meaningless.” Most helpful at this point would probably be to draw an analogy.

I think the sense in which it may be misguided to ask these questions is similar to the sense in which it is misguided to ask what, for instance, *numbers* are—or what particular

numbers are. I take the following questions to be misguided when asked by someone who already understands the meaning of the word “number” and who is able to add, subtract, and so forth: What does the number 2 *consist in*? What constitutes the number 2? What kind of thing or substance is the number 2? What is the number 2 made up of? What *is* the number 2? Is the number 2 physical or something over and above the physical? Is the number 2 identifiable with some natural object or property or is it a supernatural object or property? What are numbers?

Now to say that certain philosophical questions about sensations and phenomenal character are misguided in a similar way is *not* to say that sensations and phenomenal character are at all like numbers. They are of course very different, and in countless ways. For now, I am only trying to convey the very rough sense of misguidance or ill-formation I have in mind. It’s not clear, for instance, whether we should say that the questions above about numbers are nonsensical or meaningless, but the questions are (in my view) misguided nonetheless. I’ll return later to what I think we should make of the question “What does having a sensation consist in?”

2. Philosophical Introspection

So, on to the first part. The most influential source of the conviction that these questions are appropriate is also the source of the first view of phenomenal character I mentioned, the standard view. The tremendous persistence of both ideas have their primary source in a line of thinking—or really, an activity—which philosophers find extremely natural to engage in. It involves the philosopher’s introspection (as it is said) of his or her own sensations or phenomenal character. 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. First, something gets me contemplating a particular philosophical question about the mind—about sensation, or consciousness, phenomenal character. Perhaps I’m reading J. J. Smart’s essay in which he argues that sensations are brain processes. Or I’m listening to a talk in which the speaker claims that the what-it’s-like aspect of experience is identical to some of the experience’s representational properties. At some point, in order to assess the view I’m thinking about—to ascertain what I myself think is the correct thing to say—I decide to inspect my own sensations or phenomenal character. I might 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 sensations *are*, and what philosophical 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—but 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 introspecting, I think.

Now, I have no resistance to the idea that this line of thinking that I and certainly many others engage may be problematic—in fact, that’s what I’m going to try to urge—but it can’t be denied that it’s an overwhelmingly common and influential train of thought. Lines of thinking such as these constitute the chief source both of the standard view of phenomenal

character and of the idea that we can appropriately ask and theorize about what sensations and phenomenal character are.

What I've described is really more than a line of thinking, as it's supposed to involve, in addition to thinking or reasoning, the introspection of the phenomenon in question. But I'm going to continue to refer to it simply as "the line of thinking."

It 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. 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".⁵ Or that it is or has intrinsic properties.

The line of thinking can also vary with respect to precisely what one intends to attend to in introspection. That is, one might intend to inspect a sensation, or to inspect the phenomenal character of some sensation. Accordingly, one might conclude that "it" (the sensation) is what is "there"—or that it has accessible, intrinsic properties—or, one might conclude that "it" (the phenomenal character) is what is "there" (whatever that would mean), 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.

Of course, many philosophers have already inveighed against this sort of thinking. But it still has remarkable force. In the first part of the talk, I'm going to provide four considerations about it. My goal is not to show that the line of thinking is false. But I do want to make a case that we should be much less confident that it is true. I want to give those who take it for granted pause. This would be important if the line of thinking is truly the primary source of some central projects in the philosophy of mind.

⁵ O'Shaughnessy, B., *Consciousness and the World*, ch. 18.

Let me 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 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 philosophical line of thinking, and the introspection involved in it? Why do we give it so much credence? I suspect that one reason we do is that we think we're not doing anything we don't normally do. We introspect all the time, we think. But if the two sorts of introspection differed in important ways, this way of justifying the philosophical methodology would be 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 to ascertain something more like what those sensations are, or consist in, or amount to, or what their “nature” is.

A number of questions arise, then, concerning the cognitive capacities 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 one or more sensation. That's 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's a complicated intention. And in order simply to form that intention, we need to be exercising a variety of cognitive capacities.

First let's 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 while we think; what we observe will be what the word means!—But this concept is not used like that. (It would be as if without knowing how to play chess, I were to try and make out what the word "mate" 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) of the thinking, such as perhaps mental images, or inner vocalizations. And thinking does not "consist in" the presence of any such accompaniments.

Of course, sensations are among the sorts of things Wittgenstein has in mind when he speaks of "accompaniments." And at any rate, it certainly doesn't follow from the fact that "watching ourselves" while we think wouldn't get us clear about what thinking is that watching ourselves while we have a sensation wouldn't 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 substantial assumptions.⁶ We wouldn't form this sort of intention in many other cases. Consider the phenomenon of price—as in the price of a particular gallon of milk you've just taken out of the refrigerator at the store. Suppose you start thinking about what price is. What *is* price? What are *prices*? The gallon of milk you have in your hand has a price. But in order to ascertain what price is, or what prices are, you wouldn't turn your attention to the gallon of milk. And if you did try to inspect the price, you wouldn't find it. Suppose there's no price tag on the bottle in your hand—and even if there were, the price tag isn't the price, it's just the price tag. Nor would you think, upon not finding the price on the milk, well, the gallon of milk does have a price, and since I can't observe the price anywhere on the bottle, it must be something inside the bottle, in some part of the milk. Or, if it's not in the milk, perhaps it's some mysterious immaterial entity that I don't quite understand.

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 they are things that we can attend to for the purpose of investigating what kind of things they are. Of course, 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, for us to have good reason to make this prejudgment at this juncture of the line of thinking, that reason could not come from the line of thinking itself.

Let me continue. 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

⁶ Wittgenstein: "It shows a fundamental misunderstanding, if I am inclined to study the headache I have now in order to get clear about the philosophical problem of sensation." (PI 314.) See also PI 294.

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 recap the line of thinking for you: a) something gets me contemplating a particular philosophical question about sensation; b) I form the intention to inspect my own sensations, in order to ascertain what I think is the correct thing to say about sensation; c) I identify a sensation to inspect; d) I inspect that 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. I want to discuss steps (c) and (d) now: 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’m 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 make in identifying the sensation?

Here it’s helpful to bear in mind some of the things we learn from Wittgenstein’s discussions about ostension. Consider a case in which I ostend, or point to, something non-mental—the greenness of a lime, say. The considerations are familiar. For me to point to the greenness of a lime across the room, it’s 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’d make just the same movements whether I intended to point to the greenness of the lime, or to the lime’s elliptical shape, and so forth. Those 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 to which I’m attempting to point. What I

manage to ostend or to identify is determined in part by what concepts I employ, and by my intentions.

So, what are the relevant concepts and intentions at play in the philosophical context when we identify our sensations? And whatever they are, 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 they are different, 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?

You might think that the sorts of concepts we employ in the two cases are the same: in both cases we use something like the concept *sensation*. But this isn't obvious. First, the concept of sensation we employ in the doctor's office may well be different from the concept of sensation we employ in the philosophical case. And even if they're the same, it's not just the concept *sensation* that determines what it is we identify and attend to; our present concerns, and intentions, and thoughts might affect it as well.

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 cognitive 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're employing here depends on what it is we're doing when we

try to inspect something for the purpose of ascertaining what it is, or what it consists in, or amounts to. And that—that is, what we’re doing when we try to do this—is hardly patent.

Now this point—that philosophical introspection may unfold differently from 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, do much to move anyone to suppose that the line of thinking may not be 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. Those are the first and second of the four considerations—the point about the prejudgment we make, and the point about the appeal to everyday introspection to support the appeal to philosophical introspection.

The particularly puzzling step in the line of thinking is the final one. We’ve just been discussing my inspection of my sensation. That inspection is what’s supposed to allow me to draw a conclusion about what that sensation is—that it is a “something,” or that it’s something that is *there*, or that it has intrinsic properties. But precisely what is the inference or move that I make here? That is, the move that leads me to such a conclusion? And is it justified? Precisely what do I do here that leads me to conclude something about what the sensation is? The result of the move is a judgment. It’s something like: there’s something there that I’m inspecting, and so forth. On the basis of what do I form this judgment? If I infer it from another judgment, what is that judgment? And how did I form it?

One way of supporting the inference I make is to appeal to the idea that my attention to the sensation is a perceptual sort of attention: to appeal to a model of introspection on which

introspection is a form of perception—sometimes called an “inner-sense” model of introspection.

Now, I haven’t said anything yet as to the relation between the line of thinking I’m exploring and inner-sense models of introspection. The line of thinking would certainly seem to be most natural on such a model. But I’d prefer not simply to assume that the proponent of the line of thinking must endorse such a model. For one thing, the vast majority of contemporary philosophers of mind claim to reject inner-sense models, even proponents of the standard view of phenomenal character. They often cite Sydney Shoemaker’s Royce Lectures against inner-sense with apparent approval. Of course, the fact that these philosophers *say* they don’t endorse inner-sense views is one thing; that they don’t implicitly rely on them is another. Indeed, I suspect that such a model of introspection ultimately does play a large role in the line of thinking, whether proponents of the line acknowledge it or not. But I won’t pursue that suspicion here. Regardless, I won’t make any assumptions today as to whether the line of thinking ultimately depends on an inner-sense model, nor as to whether such a model is correct.⁷

So, back to the questions I was just raising. What justification do I have for making the inference or move I do when I conclude that what I’m attending to is something there? An inner-sense theorist might argue as follows: In order to perceive something, what I perceive must have intrinsic properties, or must be something that’s there, etc.; otherwise, I wouldn’t be able to perceive it. And since introspection is 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.

⁷ There are some philosophers who still explicitly endorse it, e.g., Lycan (1996).

One problem with making this argument here, though, is that it's difficult to see how one could be justified at this juncture of the line of thinking in appealing to an inner-sense theory 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 hasn't yet completed, or doesn't yet have reason to trust, the line of thinking, what justification would one have for such ideas—that is, about the sorts of things we find in introspection? That's not to claim that the inner-sense theory is false; it's to doubt whether one would be justified in assuming it's true for the purpose of supporting the final inference of the line of thinking.

Indeed, the final inference of the line of thinking is even more puzzling on views of introspection that don't liken it to perception. At the very least, we would certainly need plausible answers to—indeed, very convincing answers to—questions about the basis for this inference in order to give the line of thinking as much credence and influence as we do. So that's the third consideration.

Allow me to press on, though, and turn to my fourth consideration. If you wouldn't mind, attend for a moment to the feeling of your back against the back of your chair. If you're sitting forward, you can attend to the feeling of your feet in your shoes. The question is, before I even mentioned this feeling—and thus before you even attended to it—did you have it? I'm not asking whether the feeling you had in your back before I asked you to attend to it was the same as the feeling you have in your back right now, now that you're attending to it—that is, whether attending to the feeling changed the feeling. The question I'm asking is whether you had any feeling of your back against the chair at all. Most likely you had not focused once while I've been talking today on the feeling of your back against the chair. Do

you nonetheless think that you did have a feeling in your back during that time? Or, do you think that in attending to what your back felt like against the chair you thereby created the sensation you attended to—out of thin air, as it were?

This is a question I'm sure we've all thought about before, and it's remarkable the difference in intuitions people have about it. My intention here is not to suggest that you didn't have a sensation in your back before you introspected. Rather, I want to ask whether you have good reason to believe that you did. I don't think you do.⁸ If you do have good reason to believe that, it would seem the reason couldn't solely be that you had a sensation in your back when you introspected. Any reason, it seems, would have to be an inferential one, as it were.

But first, why do I introduce this point? As I'll explain, I think this point leads to two ideas which, together, suggest that we shouldn't be as confident as we are in the line of thinking. So I'm going to offer a two-step argument. The first idea is that: even if we did find intrinsic, accessible properties in introspection, that wouldn't give us reason to infer that anything we're not introspecting has such properties. I'll explain why I say this in a moment. The second idea says that if this first idea is true—that is, if finding intrinsic, accessible properties in introspection really wouldn't give us reason to infer that anything we're not introspecting has such properties—then we shouldn't be as confident as we are that we truly find such properties even in introspection. So let me try to support these two ideas. This fourth consideration is a little circuitous, so bear with me.

Let me start by making explicit one way in which I do *not* want to support the first idea. And that is by saying that the case of the back suggests that we don't have reason to think we

⁸ For an excellent discussion of this question, see Eric Schwitzgebel's "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).

ever have sensations except when we are introspecting. Even if it's true that you don't have reason to think you had a sensation in your back before you introspected just now, that wouldn't mean that you don't have reason to think you ever have sensations except when you are introspecting them.

Someone might then say to me: If you're open to the idea that we do sometimes have sensations when we're not introspecting, then wouldn't a simple process of induction show that the first of your two ideas is wrong? Wouldn't your finding that your introspected sensations have intrinsic, accessible properties would give you justification for inferring, on the basis of induction, that unintrospected sensations have such properties too.

Let me explain why I don't think I'd be justified in making this inductive inference. What the back example emphasizes is that we don't have good reason to think that introspection does not sometimes bring about or create things that didn't exist before. Or at least, I think this is something it is reasonable to take from the back example, and I'm going to take it for granted here. Now, even if I did come upon intrinsic properties of my sensations during philosophical introspection, how should I regard the possibility that the reason my introspected sensations have these properties—intrinsic or even “what its like” properties—is that my introspection somehow causes them to have such properties? One reason to wonder whether I would have justification for making the inductive inference from introspected sensations to unintrospected sensations is that it's not easy to see what justification I would have for ruling out that possibility.

But that's only one way to question the inductive inference to unintrospected sensations. There's another way. Even if one did find, during philosophical introspection, something with intrinsic, accessible properties, does one have reason to think that that thing is a

sensation? I know, what else could it be? But keep in mind Wittgenstein's admonishment that it would be foolish to watch ourselves while we think in order to get clear on what thinking is. I don't think that at this juncture one can simply assume that what one has attended to that has the intrinsic properties is a sensation. And if we don't yet have good reason to think it's a sensation that has these properties, then we're not committed to inferring, from the fact that we sometimes have unintrospected sensations, that unintrospected sensations have such properties. And so perhaps the subject's conclusion that he's introspectively attended to intrinsic, accessible properties would only provide him reason to think that there exist things with such properties in contexts of introspection.

So that's the first idea: even if we did find intrinsic, accessible properties in introspection, that wouldn't give us good reason to infer that anything has such properties when we're not introspecting them. The second idea is that appreciating this first idea should make us less confident that we really find such objects in introspection after all. If we really agree that the line of thinking would only provide us reason to think that there exist things with such properties when we're engaging in this philosophical activity—an activity, it's worth noting, human beings engage in extremely rarely, only in certain philosophical contexts—I think we become at least more open to the possibility that something is going awry in the line of thinking.

Of course it's not merely that the activity is one we engage in very rarely. If we found something rare in a non-philosophical context—a very rare rock, say, a kind of rock humans 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's 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 thing I've been suggesting today is that we do not yet have similar good reason to trust our method of investigation in the philosophical case. That is not to say the method is a faulty one, it's to say something about what we can assume at this juncture.

But there's something else. When, in any investigation, we come upon something that would seem to provide problems for our larger conception of the world—as intrinsic, introspectively accessible properties have struck many philosophers as doing—this typically makes us, and should make us, less confident that we have truly found what we think we have found. When we're in a dark attic, and out of the corner of our eye seem to see a white figure shimmering beside us, we might first think with fright that there's a ghost. But then we think again—or at least we should. We don't believe in ghosts; it must have been something else. Maybe it was the light of the moon reflecting off the window. So it is all of these factors—the rareness of what we would be finding, the status of the methodology with which we would be finding 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 line of thinking. They should make us more open to the idea that we're not truly attending to such objects even during philosophical introspection; more open to the possibility that philosophical introspection creates a context that makes us form false judgments about what we are attending to; or at the least, more open to the idea that we should not automatically trust this philosophical activity.

One might protest, “Wouldn't that be strange if philosophical introspection ‘creates a context,’ as you say, that makes us form false judgments about the nature of what we

introspect? What would be the explanation for that? After all, there's nothing wrong with or misleading about everyday introspection."

I would want to say three things in response to this objection. First, we shouldn't yet make assumptions as to the extent to which introspection in this philosophical case is like introspection in everyday contexts. As we've seen, the goals and context in the philosophical case are substantially different from the goals and context in the everyday case. And the philosophical case comes with prejudgments that the everyday case does not. And the concepts, intentions, goals, and judgments at play in the two cases, we saw, may substantially affect the way the two cases of introspection unfold. So it would be unwise, it seems, to assume at this juncture that we should trust introspection in philosophical contexts on account of introspection's being something we do all the time in ordinary life.

Second, it's not clear that even everyday introspection is as trustworthy as we think. A variety of philosophers and psychologists have recently argued that everyday introspection may be mistaken much more frequently than we might have thought.⁹

Third, we must keep in mind some of the difficulties that the view that there are intrinsic, introspectively accessible properties is thought to generate. Many proponents of the Standard View hold that such properties provide us with an intractable problem about the relation between these sorts of properties and material properties. That's a problem, of course, only for philosophers who wish to advance a materialist conception of the mind, which I do not here need to endorse. However, most philosophers who believe in accessible, intrinsic properties do wish to advance such a view, and so they, by their own lights, face

⁹ See, e.g., Schwitzgebel and Gordon "How Well Do We Know Our Own Conscious Experience? The Case of Human Echolocation" *Philosophical Topics*, 28 (2000), 235-246

some difficult explanatory questions. And so the fact that some questions would remain about the possibility I'm entertaining—namely, the possibility that the line of thinking creates a context that leads us to make false judgments about what we attend to in introspection—does not show that we should not take that possibility seriously.

I must say, though, I feel a bit self-conscious raising the questions I've raised about the line of thinking—that is, without exploring them further or anticipating more responses. It's easy to raise questions—or make suggestive points—and then move on. For me, it's only once the points I have made are considered together, in tandem, that they have a real effect. What they do is make me wonder why on Earth I've been so confident about this line of thinking I engage in when trying to ascertain the nature of sensation. And that sort of line of thinking is what primarily led think there are deep philosophical problems about sensations in the first place.

I feel a little more justified taking this kind of approach at a workshop on Wittgenstein. What's so unbelievably amazing about Wittgenstein's later work is how his sections tie together and weave back on one another, and support one another. How often one sees depth and plausibility in what Wittgenstein is saying in a particular section because of what he says elsewhere (and not necessarily before). So often in his writing, no particular section would move anyone off of whatever view Wittgenstein is calling into question at that point. There is often, in any particular passage, something apparently quite plausible that Wittgenstein's interlocutor could say. That is part of what makes some people balk at the significance of Wittgenstein's philosophy. It can appear—and I say appear—that there is never any sustained argument in his writing. What gives it all such depth, though—and what sustains all the arguments (or many of them)—is the “criss-cross” nature of them all (PI, v). A particular

view is attacked in so many ways that the proponent of the view starts to lose a grip on the reasons he ever had for favoring the view in the first place.

I don't pretend to have accomplished even a smidgeon of what Wittgenstein accomplishes in any one train of passages. But this general strategy was the only way I could see how to begin to try to disarm the line of thinking. And perhaps once the considerations I've offered are supplanted with many other considerations, and cultivated in more depth, they could start to have some real effect.

3. Sensations and “Phenomenal Character”

Now, if what I've said so far is correct, then what *should* we say about sensations and phenomenal character? What can we say? Well, two things to note right away. First, you will have certainly noticed that I haven't said much at all so far about phenomenal character, only about sensations. I focused on sensations for the sake of example, but I do think that many of the considerations I've raised, perhaps all of them, apply equally as well to the variant of the line of thinking on which what we intend to attend to (and what we do attend to) is phenomenal character: the prejudgments we make, the conceptual capacities employed in the philosophical case and the everyday case, the basis of the inference we make at the end, and so on. But I don't have room to make good on that claim here, and you will be right to be suspicious of it. I will continue to focus on sensation as my example, but will return to phenomenal character in the end.

Second, even if everything I've said so far were true, even that wouldn't entail that the conclusion of the line of thinking is false. All I've done so far is to try to show that we should not take the line of thinking for granted. But, if I'm right, that is the paramount source of the

conviction that it's appropriate to ask and say what sensations and phenomenal character are—as well as of the standard view of phenomenal character.

What about representationalism, the second view of phenomenal character I introduced? I of course think representationalists are right to resist the standard view. But I think they go too far. And yet that, in another way, they don't go far enough. They go too far in continuing on, after resisting the standard view, then to identify phenomenal character with representational properties—or really, to identify it with anything. This is related to the sense in which they don't go far enough: and that is that they adopt many of the standardists' presuppositions. They too think it is quite legitimate to ask and say what phenomenal character is. In fact, I suspect that the source of their thinking this ultimately involves the very same line of thinking, even if representationalists explicitly reject that line. But that is merely a speculation, and I won't pursue it here.

Instead, I want to discuss what one might say about sensations if one does resist the theorist's questions about what they are. Am I saying that we don't find sensations in introspection? Do we not have sensations? Or is the point rather that sensations are not objects? What are we ascribing to people when we ascribe sensations to them? And what is it in virtue of which our sensation reports are true?

There are a variety of different ways in which one who resists questions about what sensations are might try to answer these questions. And each set of answers amounts to a distinct approach to sensation. What I want to do in the remainder of the paper is to bring into clearer focus the approach to these questions that I find the most plausible, the one that I think should at least be on the table.

In order to do this, it will actually be helpful first to spend a little time on what Wittgenstein says early in the *Investigations* about the notion of signification. Wittgenstein's philosophy of language naturally gets even less attention from philosophers of mind than his philosophy of mind. But the considerations Wittgenstein provides about language contribute substantially to the plausibility, sometimes even accessibility, of his treatment of mental phenomena.

Consider, then, sections 10 and 13 of the *Investigations*. In section 8, you recall, Wittgenstein had introduced an expansion of the familiar language-game he had introduced in section 2. In this expanded language-game, A will give an order to B such as “d—slab—there.” And at the same time, A will show B a color sample, as well as point to a place on the building site when he says the word “there.” B will then take one slab for each letter of the alphabet up to “d”, of the same color as the color sample, and bring them to the place to which A pointed. If A says “c—pillar—there” and holds up a blue sample, B will bring three blue pillars to the place to which A points.

In section 10, Wittgenstein says about this language-game:

10. Now what do the words of this language signify?—What is supposed to shew what they signify, if not the kind of use they have? And we have already described that. So we are asking for the expression “This word signifies this” to be made a part of the description. In other words the description ought to take the form: “The word. . . . signifies. . . .”.

Of course one may for short describe the use of the word “slab” by saying that this word signifies this object. This will be done when, for example, it is merely a matter of removing the mistaken idea that the word

“slab” refers to the shape of building-stone that we in fact call a “block”—but apart from this point it is already known what sort of ‘referring’ this is—that is, how these words are used.

Equally one may say that the signs “a”, “b”, etc. signify numbers; when for example this removes the mistaken idea that “a”, “b”, “c”, play the part actually played in language by “block”, “slab”, “pillar”. And one may also say that “c” means this number and not that one; when for example this serves to explain that the letters are to be used in the order of a, b, c, d, etc. and not in the order a, b, d, c.

But assimilating the descriptions of the uses of words in this way cannot make the uses themselves any more like one another. For, as we see, they are absolutely unlike.

That’s section 10, and then in section 13, he says:

13. When we say: “Every word in language signifies something” we have so far said nothing whatever; unless we have explained exactly what distinction we wish to make. (It might be, of course, that we wanted to distinguish the words of language (8) from words ‘without meaning’ such as occur in Lewis Carroll’s poems, or words like “Lilliburlero” in songs.)

I think these sections contain the germs of a number of important and interrelated ideas. One thing section 10 suggests is that in order for an utterance of a sentence of the form “the word ‘X’ signifies Y” to be effective, the person for whom the sentence is uttered must already understand how certain kinds of words are used. For instance, in the case in which someone says about the language-game of 8, “the signs ‘a’, ‘b’, ‘c’ signify numbers,” the

recipient already has to know what it would be for a word to signify a number; one would already have to have a conception of the sort of role that number words play. And once we appreciate that point, the possibility presents itself that the function of the word “signify”—and the function of sentences of the form “the word ‘X’ signifies Y”—may be quite different from what we might have thought. The possibility presents itself that what we’re primarily doing when we say “the word ‘X’ signifies Y” is saying something like “You know that kind of way of using a word? That’s how this word is used.”

Now, there’s a very tempting way of characterizing what this reveals about signification which can be quite misleading and dangerous. And that is to say this. What these sections bring out is that the standard picture of signification is misguided. According to the standard picture, when we say that the word “X” signifies Y, what we are saying is that there are two things, the word “X” and the thing Y. And that the two things are related to one another by the relation of signification. And that that relation is what the word “signify” in these sentences refers to (or itself signifies). Of course, there are different theories of what that relation is, but it’s something like a pointing relation, or a reaching-out relation. But—and this is still the dangerous characterization—what sections 10 and 13 show is that this just isn’t how the word “signify” works: there is no relation—that of signification—that the word “signify” in these sentences refers to. Rather, once we focus more on the role that sentences involving the word “signify” play in our lives, what they achieve (as Wittgenstein often says)—and not so much what they “describe” or refer to, we won’t be inclined to posit such a relation, nor thus to ask or say what that relation is or consists in. Now, while this way of characterizing the significance of these sections can be helpful, it can also be very misleading

and dangerous. And we would need to be extremely careful in what we take from it. Let me explain.

The potential hazard is related to something else that sections 10 and 13 bring out. For some sentences—for instance, the sentence “Every word in language signifies something”—in order for an utterance of it to “say anything at all,” as Wittgenstein says, it must be clear what distinction is being made. Absent any such distinction, the sentence wouldn’t say anything at all, so to speak. What sometimes happens in philosophy is that the required distinction is absent—or, if you prefer: insofar as there is a distinction provided, it is not suitable for one’s philosophical purposes.

Wittgenstein’s treatment of understanding, later in the *Investigations*, is illuminating in this regard, so let me say something about that. This will help me convey the dangers of that characterization of the lessons of sections 10 and 13. It is sometimes said that one of the conclusions Wittgenstein draws later in the *Investigations* about understanding is that understanding is not a process. What he says, though, is that “*In the sense in which there are processes (including mental processes) which are characteristic of understanding, understanding is not a mental process.* [New para.] (A pain’s growing more and less; the hearing of a tune or a sentence: these are mental processes.)” [PI, 154, my emphasis] And earlier in the passage, he says “Try not to think of understanding as a ‘mental process’ at all.—For that is the expression which confuses you.” “Try not to think of,” he says, not “It is not.” To say “understanding is not a mental process” might suggest that, isolated from any context, and independent of a particular distinction one wants to draw, there is a category—the category of processes—which some things fit into absolutely, once and for all, and others don’t. Instead of saying understanding is not a process and then trying to articulate precisely

what he means by “process,” Wittgenstein gives examples of the sorts of things that, on one sense of the word “process,” we would call processes: a pain’s growing more and less, the hearing of a tune or a sentence.

It’s natural for someone to protest here, “Surely this is not a very helpful philosophical point about understanding Wittgenstein is making! Yes of course understanding is not a mental process *in the sense that the hearing of a tune is a mental process*. Understanding is different from the hearing of a tune!” But we must keep in mind what Wittgenstein’s trying to do at this juncture. He’s simply trying to move us off of a particular, very tempting way of seeing understanding. He doesn’t intend here to be mounting a decisive argument to prove or establish that that conception of understanding is false. Rather, he’s trying to alter, or get us to question, our presuppositions about the sort of account of understanding we should be looking for, or expect. It is only a move at an early or ‘primitive’ stage of a philosophical treatment of questions about understanding. And to do that, one need not be able to specify in some fully explanatory way precisely what understanding is not.

That is not to say that it would be wrong, or false, to say that “understanding is not a process”—or even that to utter that sentence would be to say nothing at all. Indeed, saying that could be effective; it could have the same effect as saying “try not to think of understanding as a process” does. But it could also be dangerous, and misleading, for the reasons I’ve mentioned.

All of the same issues apply to how we might be tempted to characterize the significance of sections 10 and 13. One way of characterizing it is to say that when we say “the word “X” signifies Y,” the word “signify” does not refer to some relation that obtains between the word “X” and the thing Y, or that the word “signify” doesn’t refer to anything; it’s just a tool we

use. But on some distinctions, it would be quite fine to say that the word “signify” refers to something, as Wittgenstein’s example of “Lilliburlero” brings out.

In the case of the price of milk, for instance, we might have concluded “Look, there’s just no thing or entity that the expression ‘the price of the gallon of milk’ refers to; that’s where you’re going wrong in looking for the price.” But one might well have said to that, “Sure there’s something the expression ‘the price of the gallon of milk’ refers to: it refers to its price! What do you mean by ‘thing’ or ‘entity’?”

Still, if Wittgenstein is correct about the way sentences involving the word “signify” work (and of course we’ve only looked at one kind of such sentence), then the concept of signification—and, if you like, the phenomenon of signification itself—is very different from what one might have thought. This is what we often find in Wittgenstein’s work. As he says elsewhere, about another concept, “It is necessary to get down to the application, and then the concept finds a different place, one which, so to speak, one never dreamed of.” (PI II, xi, 208.)

Now, what does this all have to do with sensations? I think the themes these sections about signification bring out have quite a lot to do with philosophical questions about sensations. One of the ways in which they’re related concerns the way in which one might be tempted to characterize the approach to sensations I want to entertain. And that is to say that on this view sensations are not objects. Or that sensations are not things or entities. But once again, this is a dangerous way of speaking. What distinction is being made here? What is an object? What is a thing?

What about the question: Do we have sensations, or not? Here there are two ways to respond, and I’m actually not sure which I prefer. The first is to say something like, “Yes, of course we have sensations; it’s just that what it is to say that we have sensations is different

from what you might think.” The second response would be to say, “I certainly don’t want to say we don’t have sensations, but in the context we’re in right now—the philosophical, non-everyday context—the sentence ‘We have sensations,’ in fact has no application.” I worry about this latter response, but am also at times quite drawn to it. My worry concerns the implications this response may have for the application of some of the sentences I’ve uttered in the rest of my talk. Do they have no application either? But this is a large topic, and one I can’t pursue here. Let me continue to sketch how we might begin to approach some of the questions about sensation I said might naturally be posed to someone who resists the line of thinking.

So what are we doing then, one might ask, when we ascribe sensations to ourselves and others—when we make sensation reports? I mean, we do make such reports and often think of them as being true or false. One who has resisted the line of thinking might be inclined to try to answer this question in some neat, all-encompassing way. And certainly some philosophers have tried to do just that. Take first-person sensation ascriptions, for instance—such as the ascription “I am in pain.” One view—sometimes attributed to Wittgenstein in fact—holds that first-person sensation ascriptions are expressions, and not descriptions or reports. I suspect, though, that to say, once and for all, that this is what first-person sensation ascriptions are—namely, expressions—would be to give in to precisely that urge for simplicity or explanation that Wittgenstein works so hard in his writing to show us is so dangerous. It would amount to precisely the sort of caricature view that is sometimes attributed to Wittgenstein and then dismissed as being implausible. Which is not to say that it is not in some situations helpful to think of first-person ascriptions as being merely expressions.

Likewise, one might try to give some version of an “adverbialist” account. Something like: to say of a person that she’s having a sensation of red is to say that she is sensing redly, or doing something redly. On the view I want to entertain, there is no guarantee that such a general, all-encompassing description of sensation reports is available.

What if in response to the question “What are we doing when we make sensation ascriptions,” one were to say “well, we’re ascribing sensations!” That certainly wouldn’t satisfy our interlocutor. The interlocutor’s question would seem to request an answer that does not involve more concepts like *sensation*. It’s not clear, though, why the proponent of the view I’m entertaining should be required to be able to say what it is we’re doing when we make sensation reports, or what we’re doing when we make any particular one sensation report, in a way that does not use more sensation discourse. Why must that be required for the view to be plausible? Must it even be possible to do that?

And indeed, isn’t that what we learn from Wittgenstein is the case for a wide variety of other mental phenomena—understanding and thought, for instance? In fact, I think that once we become open to the possibility that the line of thinking is false, we should be open to the possibility that our philosophical treatment of sensations should be broadly similar to Wittgenstein’s treatment of understanding and thought—perhaps all intentional phenomena: that it is misguided to ask what such phenomena consist in, that in our philosophical treatment of these phenomena we should focus more upon the role that our sentences about these phenomena play in our lives and what utterances of these sentences achieve—and not so much on what these sentences refer to or “describe”; that we may not be able to explain ascriptions of these phenomena without using the sorts of terminology and concepts we employ in making those ascriptions. For instance, about imagination Wittgenstein writes:

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

I of course can't do any justice here to Wittgenstein's many important discussions of thought, understanding, and the like. But I think it's a good place to start in trying to see the plausibility of the approach to sensation I have in mind. Even those who are sympathetic to Wittgenstein's treatment of understanding and thought will have a harder time believing that that sort of treatment is even broadly applicable to sensations. That's because the line of thinking is so powerful as it concerns sensations and the like. It's much easier to disarm the line of thinking as it is applied to understanding, for instance. The phenomenon of thinking, in particular—occurrent thought—is a good intermediate case to pursue, I believe. Some philosophers who are willing to go along with Wittgenstein's treatment of understanding will be less willing to endorse those ideas as they apply to thinking. Wittgenstein's many discussions of the concept of thinking and the role in our life of the words “thinking” and “think” in our lives are rich and promising, and still have not been fully appreciated or explored.

But one still feels a temptation to say, “but still, what is it *in virtue of* which we have sensations?” What we should say about this question depends on precisely what kind of

answer the questioner is asking for, and in what circumstance. When I say that the question “What are sensations?” is misguided, I’m not saying that it is never appropriate to utter the words “What are sensations?” Recall the mother explaining to her child what sensations are. In that vein, it might be fine to answer the question, “What is it in virtue of which we have sensations?” by saying something like “In virtue of our having certain feelings in our body.” Or, if a causal sense was intended, I might say “I had the painful sensation in my hand in virtue of my putting my hand too close to the flame.” But these are not answers to the question, expressed in the same words, that the interlocutor is here asking. That question seems to be a constitutive one, and one that asks for an answer that does not simply employ more sensation discourse—words or concepts like “sensations” or “feelings”. *That* is the question we should be suspicious about. The big issue—and what should occupy our attention more—is why we are so inclined to think that in order for the philosophical treatment of sensations I’m entertaining to be correct, an answer of the sort the interlocutor is asking for would need to be available, even in principle.

Now, I’ve focused today primarily on sensations. But broadly similar treatments might be given to other, related phenomena—mental images, sense-perception, kinesthetic experiences, and so on—though the details will always depend on the particular case. What though of phenomenal character, the notion with which I began? On the view I’m entertaining, perhaps philosophers should just abandon the expression “phenomenal character” altogether. The notion is a philosopher’s artifact, and, in many ways, a misleading one. The same goes for “qualia,” “raw feels,” and the rest of them. But at the very least, if we do continue to employ the term, we must take seriously the view that asking and attempting to say what these things are is a misguided enterprise. In contrast, expressions and

concepts such as “sensations” and “feelings” and “mental images” have their origin in everyday life and are importantly not terms of art. These are expressions and concepts we *should* continue to focus on—even if we should not ask what such phenomena are. Once we really get down to the application of these concepts and expressions, they might well find a different place, one which, so to speak, we never dreamed of.