

Eliminativism, First-Person Knowledge and Phenomenal Intentionality

A Reply to Levine

Charles Siewert
Department of Philosophy
University of Miami
PO Box 248054
Coral Gables, FL 33124
U.S.A.

csiewert@miami.edu

Copyright (c) Charles Siewert 2003

PSYCHE, 9(03), January 2003
<http://psyche.cs.monash.edu.au/v9/psyche-9-03-siewert.html>

KEYWORDS: consciousness; eliminativism; first-person knowledge; phenomenal color experience; conscious thought.

REPLY TO: Levine, J. (2001) Phenomenal Consciousness and the First-Person: Comments on *The Significance of Consciousness*, *Psyche*, 7(10).
<http://psyche.cs.monash.edu.au/v7/psyche-7-10-levine.html>

ABSTRACT: Levine suggests the following criticisms of my book. First, the absence of a positive account of first-person knowledge in it makes it vulnerable to eliminativist refutation. Second, it is a relative strength of the higher order representation accounts of consciousness I reject that they offer (as my account does not) explanations of the subjectivity of conscious states and their special availability to first-person knowledge. Further, the close connection I draw between the phenomenal character of experience and intentionality is unwarranted in the case of both color perception and conceptual thought. In response to Levine's critique, I argue that the eliminativist can be rebutted and higher-order representation theories found wanting, even without offering a positive account of first-person knowledge. Also, I note that I actually have begun to offer an account of this based on my conception of phenomenal consciousness. Finally, it will be seen that Levine's concerns do not undermine my views on color experience, conscious thought, and intentionality, once their justification and character are made clear.

I thank Joseph Levine for his thoughtful comments on *The Significance of Consciousness* (Siewert, 1998). (All pages references will be to this book.) His criticisms raise important issues, challenging me to develop my views in directions I agree I need to explore. Levine's main areas of concern are these. He says that until I provide a positive account of what first-person knowledge or warrant consists in my account of consciousness is either vulnerable to eliminativist objections, or (at best) seriously incomplete. Further, he doubts whether the close connection between phenomenal character and intentionality that I (correctly) say holds in the case of visual spatial experience can be extended, on the basis of first-person reflection, as far as I claim with respect to color experience, and non-imagistic, conceptual thought. I will address these issues in turn.

1. Against Eliminativism

The first concern Levine raises seems the most threatening to my views. The fact that I leave it open just what the "distinctiveness" of first-person warrant consists in creates, he suggests, a deep vulnerability opponents could readily exploit. Where I am silent, they might offer their own positive account, and use this to justify an eliminativist position on phenomenal consciousness. More specifically, they might hold that the specialness of first-person knowledge consists in our possession of an "internal monitoring" mechanism that is reliable in its pronouncements -- but reliable only within certain bounds, bounds that do not include the self-attribution of what I call "phenomenal consciousness." Then they might say: "Set to the side any attributions of phenomenally conscious experience, and ask: can we adequately explain what is left (behavior, information-processing) without re-introducing any conscious experiences into the picture? The answer is: yes, we can. Thus the inner monitor's reliability does not extend to whatever self-attributions of conscious experience it might generate." So experiences with phenomenal character in my sense are idle, unnecessary explanatory posits, and my book purports to discuss the significance of a phenomenon we have no right to say ever occurs.

The critic here assumes that we have first-person warrant for the application of this class of psychological predicates -- those expressing phenomenal features -- only if that can be justified by claiming it is needed to account for explananda conceived of in some way that thoroughly excludes the members of this class. In my book, I identify that underlying eliminativist assumption, and question why we should accept it. This begins in Chapter 2 (pp. 49-61), where I argue against a certain aspect of Cartesian epistemology that has, ironically, been adapted to eliminativist ends -- namely, the view that the third-person application of psychological terms must be justified by appeal to their indispensability in accounting for explananda that have been thoroughly "de-mentalized." And I argue that, even if first-person applications are not warranted without warranted third-person applications, this does not show that first-person knowledge is dependent on a third-person knowledge of mind conceived of in this way. Therefore, we cannot fairly impugn our first-person knowledge of mind by arguing that it is dependent on a third-person source of justification that regards all mentalistic talk as a system of explanatory posits that fails to explain as it should. Then in Chapter 5 (p. 152), I propose that this sort of

response be given, not just, as before, to an eliminativist who wants to throw out all everyday mentalistic talk on such a basis, but also to one who would mount such an argument, more specifically, against all attributions of phenomenal experience. That is, I challenge the grounds for assuming that the warrant for attributions of conscious experiences is to be assessed by asking whether such attribution better explains data conceived of in a way that excludes any phenomenal experience, either directly or by entailment.

In conversation it has been put to me by Georges Rey, who advocates a view along the lines Levine describes, that the eliminativist epistemology is supported by the following analogy. Suppose we met someone who is firmly convinced that his dreams give him knowledge of a separate world. He says that in his dreams he visits a peculiar place that exists even when he does not dream. Though he has no idea where this place is in relation to places encountered in waking life, nor how he is transported there, still, he is adamant that he knows it is real. Of course, we should not believe him. For we should think his warrant for his beliefs about this "dream place" depends on how well his dreaming experience is explained by the truth of these beliefs. And that is to say: not well at all. Now if this is right, we should accept something similar regarding anyone's assertions that his or her experience has phenomenal character, understood as I propose. Unless these outputs of their internal monitors can be justified by saying their truth better explains data described in such a way that does not assume any phenomenal features have instances, we should regard whatever beliefs we hold that they do -- however firmly we hold them -- to be just as groundless as beliefs about a mythical place one visits in one's dreams.

Levine himself does not offer Rey's analogy, but I would like to get to the bottom of this style of objection, so let's examine it. First, I do not find initially plausible the idea that my warrant for judging that I have conscious experience should be assessed in a way similar to that in which the dreamer's judgments about his dream place should be assessed: by asking whether accepting their truth yields the best explanation of something else. Why should I accept that the two cases are analogous?

Perhaps it will seem that the presumption is in favor of the analogy; unless we can produce some reason to say the situations are not epistemically analogous, we should believe that they are. However, I don't think the analogy carries any such presumptive weight. Consider first: it seems that an eliminativist who draws such an analogy should recognize some domain to which it does not extend. For ask yourself: Does it extend, for example, to beliefs about the real places we visit when awake? Are we to say that just as we should assess the warrant for beliefs about dream space and the goings-on therein by asking how well they explain other goings-on, so too we should assess the warrant for beliefs about waking space, and the goings-on in it: by asking how well they explain something else entirely? We might want to say this, for we might consider facts about the phenomenal character of one's experience to constitute the "neutral" non-question-begging data, the explanation of which determines the fate of one's beliefs about things in (real) space. This would be the traditional Cartesian way of viewing one's epistemic situation. But note that the would-be eliminator of phenomenal consciousness cannot

share it. For then beliefs about space would give him no standpoint from which he could criticize and reject beliefs about phenomenal experience.

So eliminativists should recognize some domain where the analogy with "dream space" beliefs does not hold. And this invites us to oppose the eliminativist's analogy with another. The class of beliefs about the phenomenal character of one's experience is, epistemically, not so much like the class of hypothetical beliefs about dream space, as it is like the class of beliefs about the space of waking life, at least in this respect -- neither is in need of justification by viewing it as comprising a candidate theory to explain evidence that entails the truth of no one of its members. So at this point we face a choice of analogies: is our warrant for beliefs about phenomenal experience similar (in the just mentioned respect) to our warrant for beliefs about the real space perceived and inhabited in waking life, or, is it (as the eliminativist says) more like one's warrant for beliefs about the behavior of things in the purported "space of dreams"? We have no right, it seems to me, to suppose that the latter analogy enjoys some presumption of correctness, as against the former.

The eliminativist needs some positive justification for his favored analogy, if it is to support the epistemological assumptions essential to his argument. No plausible justification for accepting the analogy occurs to me. However, I can think of some reasons to reject it. For consider: just what is supposed to serve as the point of analogy with the dreamer's experience; what constitutes the data to be explained by the "postulation" in me of phenomenal experiences? Not yet another level of experiences -- some experiences of the phenomenally conscious experience, that would stand to beliefs about phenomenal character much as the dreamer's dream experiences stand to his beliefs that there are people and things moving about in dream space. No. Well, what then? Perhaps the data are my judgments that I have this or that sort of conscious experience, at one or another time. But this seems unsatisfactory for another reason. My understanding of the utterances by which I express such judgments depends on my success in referring to illustrative instances of the kind of phenomenal character involved. So, if I don't actually manage to identify for myself real instances of these ways of seeming (ways in which it looks to me, feels to me, etc.), then I quite literally do not understand what I mean by the terms I use for them. But in that case, I do not even make the judgments that were our most recent candidates for the data to be explained by the postulation of consciousness. Those "data" do not exist. In that case, eliminativism winds up doing away with the data to be explained as well as the "posits" to explain them. Or else the only data to explain now are bits of unintelligible verbal behavior.

It now emerges that the eliminativist position is not that my reports of consciousness are false -- rather they are unintelligible or somehow at bottom incoherent, like beliefs in magic perhaps. Of course the eliminativist, like the logical positivist of yore, might relish the prospect of accusing opponents of talking nonsense. But now, as then, we should recognize that the evidential standards for such an accusation ought to be high. And we should notice that the earlier argument against me has, in effect, been abandoned for a new one, claiming something like: "It is, on examination, unintelligible to suppose that, to someone in Belinda's situation -- to someone with Connie-like powers to discriminate

left field visual stimuli -- it might look no way at all on her left." But just what is the argument for that?

The eliminativist may be tempted to retort here that asserting the reality of phenomenal experience in my sense would leave us with metaphysical puzzles that admit of no truly satisfactory solution. We would be condemned to a dismal choice between, on the one hand, some form of property dualism (which inevitably brings epiphenomenalism), and on the other, acceptance of some brute materialist necessity (leaving us with an unclosable explanatory gap). And where some putative property gives rise to such metaphysical nightmares, we are justified in refusing to accept that it has instances.

But even if accepting consciousness does leave us with only these choices (and is our grasp on cause and necessity so firm that we cannot entertain a doubt about this?), to react by denying the reality of consciousness would be a cure much worse than the disease. Anyway, one wonders what entitles us to assume reality should leave us with pleasing metaphysical options and whether consistently applying the principle, "metaphysically problematic, therefore not real," will leave us with much reality at the end of the day.

I want to emphasize that the challenge I pose here is not merely to provide a reason to believe I have wandered into some kind of incoherence in my blindsight thought experiment. The challenge I address to the critic is this. Engage in the process of reflection that starts with the argument for first-person warrant and with it gathers together broad classes of examples that I say share the feature of being conscious (experiences differing in phenomenal character). Proceed thence to conceiving of various forms of blindsight from the point of view of one who has them. See then whether you yourself are capable of deploying (or at least seeming to yourself to deploy) the notion you want to argue contains some secret incoherence. If you refuse to engage in this exercise of reflection at some point or another, then I beg you to consider whether your refusal is justified, or whether it manifests only an aversion to sources of doubt about your philosophical position. Of course you may argue for the incoherence of my understanding of phenomenal consciousness. But I ask that you first give it a fair hearing, and that means that you yourself at least try to engage in the reflection that I maintain manifests such understanding. If you do, you may find your intellectual battle against this idea is as much a struggle with yourself as it is with me. Must we really strain against our own minds this way?

There are other objections worth noting to the line of criticism Levine mentions. We might argue that if the reliable inner mechanism account of self-knowledge invoked by Levine's eliminativist will not support our claim to warranted judgments about the phenomenal character of our experience, well: so much the worse for that reliable inner mechanism account of self-knowledge. The question this raises is a serious one. Just when is it justified to use a theory of first-person knowledge to argue against our warrant for a broad class of first-person judgments? That is, when can we fairly argue that, since a theory of first-person knowledge does not show us as having a distinctive warrant for a given class of judgments (though some are disposed to claim it), the members of that

class are all warrantless? We should treat this form of argument with suspicion. For we need to ask: where does the warrant for accepting a theory of first-person knowledge come from anyway? It seems to me that if a given account of first-person knowledge fails to grant us knowledge of such and such occurrences that, independently of epistemological theorizing, we presume ourselves to know to be true in a distinctively first-person way, then that is, *prima facie*, a reason to reject the account. Those arguing that we don't know we have conscious experience because their theory of first-person knowledge doesn't grant us that knowledge need to explain why this is not just an indication that they are offering a defective theory of first-person knowledge. Perhaps the fault lies not in our claims to knowledge, but in their epistemology.

Finally, there is this objection to the eliminativists. Suppose I play along with their premises, and ask myself: Just what would be the best explanation of why I and others are disposed to attribute to ourselves experiences whose character, on reflection, enables us to engage in the kind of conceptual exercises embodied in my blindsight thought-experiment -- phenomenally conscious experience? One explanation that at least has the merit of being straightforward is this: we do indeed have these conscious experiences. That's why we think we do. It is far from clear there is any credible alternative "error theory" that does not attribute to us confusions of which we are not guilty or assumptions we do not make.

I conclude that appeal to a positive account of first-person knowledge does not bolster an eliminativist rejection of phenomenal consciousness. However, I also do not suppose that the friend of consciousness can happily stay silent about the distinctive nature of first-person knowledge. So let me now turn to the next of Levine's concerns that I mentioned at the start.

2. Subjectivity and Self-Knowledge

I agree with Levine's suggestion that there is an important connection between an experience's being conscious and its availability to a distinctively first-person knowledge. And I definitely believe that a satisfactory philosophical discussion of consciousness needs to account for this. I also agree that one needs to do justice to a certain aspect of experience that is difficult to get a hold of, even in a provisional way -- that which Levine expresses by saying a person's conscious experiences are experience for that person. My own experiences are experiences "for me," in some special sense that leads us to want to speak of their being peculiarly subjective, and of a distinctive "first-person perspective" on them. However, I find less surface appeal than does Levine in the way higher order representation theories of consciousness speak to these concerns, and I find much less alluring than many do the perennially resurgent image of consciousness as a theatre, an inner space where the mind's antics "present" themselves. Nonetheless, I do agree that part of what gives higher order representation theories their appeal, and what makes the theatre image so difficult to lay to rest, is that they promise to say something readily

accessible to the intellect or imagination about these fundamental but elusive aspects of consciousness -- aspects which, regrettably, I was not able to address in my book.

Still, I steadfastly maintain: neither a higher order representation theory of consciousness nor the theatre of consciousness metaphor provides a good way of understanding either of these important aspects. They illuminate neither the role of consciousness in first-person knowledge, nor the subjectivity of consciousness. And I make no concession to the idea that bad theories are better than no theories at all.

My reasons for rejecting higher order theories are given in Chapters 6 and 7 (and also I elaborate on them in my Psyche responses to Ludwig and Lycan). I won't recapitulate those arguments here, though I might just comment on Levine's complaint that I don't answer the higher order thought theorist's argument that we are justified in positing the occurrence of nonconscious higher order thoughts shadowing our every experience, because "[w]hat we think of as conscious experience is thereby explained."

My response to this (expressed -- admittedly rather tersely -- on p. 201) is that the claim is either question-begging or false. For what is explained exactly by the positing of non-conscious higher order thoughts? The fact that the states they target are conscious? I don't see why one would think they explained this, unless one just assumed higher order thought was necessary for consciousness. But that is (part of) what is at issue. Maybe they are supposed to explain the reportability of conscious states? But to explain this, we need not postulate the occurrence of nonconscious thinking. We could instead say that we generally form dispositional belief states about conscious states, where such a belief is the capacity to think a conscious higher order thought. To explain the capacity to report, one need posit nothing more than the capacity to think the thoughts that would be expressed in a report -- no actual occurrence of thought is needed.

But now: what do I have against the theatre metaphor? First let us ask why the image of the mind's theatre is attractive. It seems we want some image or analogy to help us explain why conscious states are available to the reflection of the person whose states they are, in a way that are not available to others. And so it is tempting to think of consciousness as a special place, whose unity and boundaries help explain why things in it are "observable" only from "inside." But if, like me, one admits no "inner observation" but certain forms of thought or judgment about one's own experience, what is wanted is an account of how first-person judgment is related to experience that helps us understand the extent and limits of its reach. But it is totally unclear how the proposal that consciousness occurs in a certain literal or figurative place is supposed to accomplish that. The theatre is, at best, a misleading metaphor for what we want explained.

The question then arises: do I think there is a better way to explain what first-person warrant and the subjective (or first-person) "perspective" consist in? I think there is. But doing so is a big project in its own right, one on which I have only begun, and of which I can here give only a rough idea.

The key challenge to be met, I believe, is to explicate the special relationship one's conscious experiences have to one another in virtue of being phenomenally conscious -- what I would call the "phenomenal unity of experience." Though it is, I believe, the phenomenal unity of experience that one is tempted to represent as the unity of a place, as in the theatre metaphor, the unity in question is not a unity of place. And I would distinguish phenomenal unity from the notion of just being a conscious experience attributable to a particular person. My approach involves saying that there is a kind of unity among conscious experiences that helps account for their special availability to first-person reflection and for the distinctiveness of the first-person point of view. But what kind of unity this is, is explicable (like the notion of phenomenal consciousness itself) only through a use of example and contrast that appeals to a first-person approach.

It is a challenge to clarify the relevant notion of unity sufficiently. I have tried to make a start on that job (in "Self-Knowledge and Phenomenal Unity," *Nous* 35:4 (2001) 542-568). I will not try to recapitulate those thoughts here. But supposing I eventually manage this task of clarification well enough, the next question is: how does that sort of unity help us to understand the special scope and limits of first-person judgment? I look for an answer in the following idea. Underlying our first-person judgments of experience (e.g., 'I feel...', 'It looks to me...') is a capacity to form conscious demonstrative thoughts both about and phenomenally unified with conscious experience (e.g., "This is a feeling...", "This is its looking to me as if...'). What gives these thoughts their special demonstrative character, and makes them distinct from first-person attribution of experience, is the form of attention on which they are based -- a form that can link experience only with thoughts that are phenomenally unified with the experience. Because this way in which one's experience is available to one's thought is made possible and delimited by the unity of phenomenal consciousness, and because the special warrant enjoyed by first-person judgment is explained with reference to that kind of availability, ultimately first-person warrant and knowledge does get explained by phenomenal consciousness. This style of account does not depend on, and in fact repudiates any notion of consciousness as inner observation or inner stage. On my view, talk of inner observation is a distorted way of speaking of the kind of attention and judgment made possible by the unity of consciousness. And talk of the mind's theatre is an ultimately unhelpful, if imaginatively arresting, metaphor for that unity. On my account this gives us a way of understanding, not only the distinctiveness of first-person warrant, but also the subjectivity, the "for me" quality of consciousness, the specialness of the first-person point of view.

Let me say a little about this. A conscious experience is "for me" because its character depends on its belonging to a certain phenomenal unity of experience, the belonging to which is sufficient for being mine. But I find it misleading to think of the special access of first-person judgment to experience that is based on this unity of consciousness as reflecting a special perspective on the experience. Since we do not, in any substantive sense, perceive our own experiences, they are not presented to us from anything like a certain "angle" or "point of view." What makes the link between judgment and experience special is not the perspective from which the former speaks of the latter. Rather, this specialness of first-person reflection is to be sought in the way judgment and experience together constitute a single perspective or point of view, on account of their

phenomenal unity. These are only hints at how I can go about addressing the challenge of accounting for first-person knowledge and the subjectivity of experience. At least I want to have raised hopes that there are some alternatives for us to explore, consistent with an embrace of phenomenal consciousness, and a rejection of higher order theories and theatre metaphors.

3. Color Experience and Intentionality

Now I will respond to the worries Levine raises regarding how thoroughly and richly intentional the phenomenal character of experience is. I claim that visual phenomenal features are intentional features of a kind, and Levine finds this acceptable for the phenomenal character of one's experience of space, partly since it is indeed hard to see how its seeming as it does for it to look as if something is X-shaped could, through some kind of reinterpretation, be made accurate or inaccurate, through the presence or absence of, say, O-shaped things. But he worries about my claim that the experience of color is inherently intentional. One might, he points out, say that, e.g., the color green could look to someone else the way red looks to me (and red could look to him the way green looks to me), and yet both of us have "accurate" color experience. It's just that the chromatic "look" that for me "stands for" green, for this other person "stands for" red. We use, so to speak, a different phenomenal vocabulary to visually represent the same objective colors. And the regular causal connections that determine what color our color experience stands for do indeed have the status of interpreting conditions, imposing conditions of accuracy on an experience's phenomenal character that it cannot furnish of itself.

My response? First, I want to concede that I do not pretend my discussion of the intentionality of color experience (in section 7.7) settles questions concerning whether the view of color and experience just sketched or some other view of the matter is correct. My point was simply this. Even if there is something about the intentionality of color experience that its phenomenal character leaves undetermined, so that this has to be added through the environment's "interpreting" the experience, it is important to recognize that phenomenal character of normal color experience cannot be broken down into a phenomenal spatial component that is intentional, and a phenomenal chromatic component that is not intentional. The reason is this. That different colors appear differently distributed in space is inseparable from the way they look -- from the way it seems to have such visual experience of color as one has. The phenomenal character of the experience makes that experience one of a color, which looks a certain way, occupying a certain place. It is not, then, just an experience of a certain place and at the same time an experience of a certain chromatic look. What this shows, I think, is that the differences in the phenomenal character between, e.g., my experience of a red equilateral triangle on a white background, and that of a green equilateral triangle (with same size, orientation, location) on a white background, constitute not only differences in the phenomenal features I have, but in the intentional features I have. I experience the space differently "chromatically filled" in the two instances, on account of the phenomenal character of the two experiences. Now this will, if the theory to which Levine alludes is

correct, leave open the matter of which color these phenomenally different experiences are to be judged accurate or inaccurate of. On that theory, this would have to be determined by the environment. But what is not left open is whether the experiences are of different colors in certain places. That aspect of the intentionality of color experience is already fixed by phenomenal character. And, whatever one goes on to say about the nature of color experience and how its accuracy conditions are determined, it needs to be compatible with this. For if color A and color B (generally) look different to x, x's experiences cannot "represent" things that "look A" as having the same color as things that "look B."

4. The Phenomenal Character and Intentionality of Thought

Now the other domain where Levine worries that reliance of first-person knowledge will not yield the kind of thesis I want relating phenomenal character and intentionality is the domain of non-iconic, conceptual thought. He objects that it is unclear to him, when consulting his own experience, that there is some particular way of seeming distinctive of, e.g., thoughts about Mondays, distinguishing them from thoughts about Tuesdays.

I think that what I say about the phenomenal character of thought probably does not commit me to asserting about Levine's experience what he finds himself powerless to discover in it. Let me clarify what my basic claim here is. Often we think, out loud or silently, and our thinking does not consist (at least, not entirely) in the forming of mental images (that is, what we are thinking about is nothing we either visualize or otherwise form an image of). Now there are variations in the way it seems to one to do such thinking that are not rightly identified just with differences in the way it seems to us to sense or image things, for they can be identified only by reference to what one is non-iconically thinking, the "content" of one's thought, if you like. To support this claim, I first invite one to recognize, through first-person reflection, that the way it seems to one to understand and follow what one (or another) is saying differs from the way it seems to one to perceive what is said uncomprehendingly, or without following it: this sort of difference is evidently not just a matter of a difference in the phenomenal character of sense or imagery experience identifiable without reference to one's understanding or meaning something by an utterance. Also I invite you to acknowledge the occurrence of episodes of sudden un verbalized (though verbalizable) non-imagistic thought. I take both of these to illustrate the point that there are differences in the phenomenal character of thinking that are not (or at least not entirely) differences in the phenomenal character of imagery or sensing. The next main point is that this non-imagistic, non-sensory difference is not a single difference. There are distinguishable differences of this sort in the phenomenal character of thought. The phenomenal character of thought is continually changing, in a way not identifiable other than relative to the expression of thought. These differences in phenomenal character are manifest in the "experience of interpretive switch" discussed on pp. 278-283. Contrary to Levine's suggestion, these differences are

not subject to "wholesale reinterpretation, of the sort that Quine, say, imagines." For the differences are only identified relative to some "interpretation" -- some thought. We have no way of conceiving of these shifts in phenomenal character as uninterpreted utterances, merely formal items, or bits of symbolic behavior, to which varying "contents" might be attached in radical interpretation, since our only way of considering them is relative to such contents.

Notice, the view is not that you know what thought you're thinking by recognizing its distinctive non-iconic phenomenal character. Certainly no inference is involved. You identify the phenomenal character by the thought you are thinking, in a sense. You conceptualize the phenomenal character by thinking of it as the way it seems to think a certain thought -- it is in such terms that you know what that way of seeming is. But in my view there is no knowledge of either, independent of the other, which serves as the basis of your knowledge of it.

Perhaps it will be clearer whether this conflicts with Levine's convictions if we apply it to his "Monday/Tuesday" case. Suppose I am writing a letter, or preparing to leave a voice message, and I need to say what the day is. The thought occurs to me -- suppose I express it to myself, silently -- "It's Monday." But then in the next moment I realize my mistake, and correct myself, thinking, "No, it's Tuesday." So my claim would be, first, that the way it seems to me to utter these phrases to myself comprehendingly differs from the way it seems to me to do so meaninglessly (as I would, for example, if I repeated them to myself again and again, until they became to me "mere sound"). Second, the way it seems to me to utter these phrases with understanding changes: it is not as though there is some unvarying phenomenal character of understanding that changes not at all with what is understood.

I do not say this is adequately appreciated just by focusing on this case, when it occurs. For it may not be evident that the phenomenal change is anything more than might be distinguished by reference to the different imaged sounds 'Monday' and 'Tuesday.' So the indirect strategy I employ in the book, applied to this case, would involve appeal to the occurrence of un verbalized non-iconic thought, and to what I call the experience of interpretive switch. Here I suppose we could consider an instance in which, after saying to oneself, 'It's Monday,' the sudden corrective realization that it is Tuesday is not verbalized. This can happen, I think. And if it happens, there is a positive change in the phenomenal character of experience that can't be pinned on any changes in imagery or sense experience. We also could adapt the "interpretive switch" point to this example. One might, it seems use the sound of 'Tuesday' to mean, not a day of the week, but something like, 'the day of Tou' -- where 'Tou' is a proper name. (So maybe it is Mr. Tou's turn today to take out the garbage: it's Tou's day.) Then when someone says, 'It's Monday -- no, it's Tou's day,' we can imagine an experience interpretive switch, thinking first of the third day of the week, then of the day on which it is our friend Tou's turn. Assuming you grant all this, I would argue that such change in phenomenal character is not confined to these cases where its distinctness from merely sensory or imagistic differences in phenomenal character are most directly evident. For that, I believe, can be shown to have unacceptable consequences (pp. 281-2).

I am not sure what in this Levine would reject or balk at. I suspect that what he hesitates to accept is not the phenomenology just indicated, but something else. I suspect what bothers him is the idea that we can, through first-person reflection, justifiably claim to discern a unique phenomenal character specific to all and only "Monday" thoughts. But I think I can honor this skepticism. For I think I can say what I want to say, that episodes of thought differ in phenomenal character in ways distinguishable only by reference to what one is thinking, while allowing also that these episodes cannot be broken up into phenomenal bits, corresponding to the words with which they are expressed, such as 'Monday,' each of which have a separately re-identifiable phenomenal character. The phenomenal character of thought does not resolve into recognizable phenomenally distinguishable units corresponding to the words with which the thoughts are expressed, something on the order of Lockean "simple ideas." But that is consistent with what I have to say in Chapter 8.

Also I would like to accept that there are considerable disanalogies between the phenomenology of thought I am now discussing, and the kind of variation in phenomenal character we are likely to fasten on initially in sense experience -- such as differences in the way colors look, tones sound, or bodily sensations feel. These latter differences admit of certain sorts of stable subjective comparative judgments. We are able to answer questions such as: 'Does red look more like orange than blue?', 'Does this tone sound lower than this one that the previous did?', 'Does this current pain feel more or less intense than the previous one?' By contrast, it seems absurd to ask whether Monday thoughts seem more like Tuesday thoughts than Wednesday thoughts, and so on. But that is fine with me; I don't think that differences in phenomenal character need to lend themselves to this sort of question -- even if many do. After all, even readily appreciable differences in phenomenal character of visual experience don't seem amenable to these kinds of judgments. Think of Gestalt shifts in visual experience. Can we sensibly ask whether seeing the duck/rabbit as a duck seems more like seeing the vase/faces as a vase or more like seeing it as faces in profile?

So I want to suggest that resistance to what I have to say about the phenomenology of thought may come from sources which, properly exposed, can be seen to provide no real objection to it. I believe that, once we clearly detach the notion of differences in phenomenal character from the empiricist tradition's conception of phenomenal simples - - ideas or impressions -- and if we avoid assuming that the kinds of comparisons to which some differences in phenomenal character are amenable must be applicable to all, my claims about the phenomenology of thought will seem much less objectionable.

5. Summary

Let me now summarize my response to the challenges Levine raises. First, there is the question of whether I can get away with introducing and relying on a distinctive first-person warrant for my judgments about attitudes and experience, while leaving open just what positive account of the mind's self-knowledge is to be preferred. I think it is actually

desirable to avoid presupposing any such epistemological view at the start of inquiry, and I deny that this attitude makes my account vulnerable to eliminativist attack. The eliminativist asks us to assess our warrant for attributing phenomenal experiences to ourselves ultimately on the model of our assessment of theoretical postulates or explanatory posits in scientific theorizing. But why should we adopt this epistemology? Support may be sought in something like Rey's "dream place" analogy, but it would be a mistake to suppose such an analogy enjoys some presumptive favor. We can draw analogies differently, and in ways counter to the eliminativist's case, with at least as much, if not more, plausibility. And, the analogy seems to break down when one considers the importance, in understanding the notion of phenomenal character, on identification of illustrative instances in the first-person case, without adverting to other putative data to be explained by them. So it seems eliminativists are driven to seek another line of attack. They must argue that, because I determinedly leave open the possibility of a person such as my blindsighted Belinda, my conception of phenomenal features suffers from some kind of incoherence, and that ultimately neither I nor anyone who thinks along my lines knows what we are talking about. The prospects of a case for this harsher line are unpromising.

Furthermore, we need to consider whether an epistemological theory's failure to supply us with knowledge of the phenomenal character of experience should cast doubt, not on the reality of phenomenal consciousness, but on the adequacy of that theory of knowledge. And we should ask whether our tendency to think we have phenomenal experience is not rather better explained by saying we actually do have it, than by a rival "error theory."

I do acknowledge the need to build on what I have said in the book about phenomenal consciousness to provide an account of how it is involved in the special availability of experience to first-person reflection and the special warrant enjoyed by first-person judgments. Here I can just say: I have made a start on this at least. I'll keep working on it.

As to Levine's worries about my treatment of intentionality-phenomenality connections, I think my claims here will seem less objectionable once their character and scope are properly understood. My remarks on the way chromatic experience is infused into the phenomenal character of spatial vision were not meant to settle the vexed issues to which Levine alludes, regarding the accuracy conditions and objectivity of color experience. But I believe they set important constraints on what will constitute an acceptable answer - - constraints that may prove crucial. Finally, I would say, if one doubts my remarks about the phenomenology of thought, one needs to take account of the actual route of reflection that led me to them and attempt to retrace it oneself, before declaring one cannot find the way to my destination. I ask that one take care not to confuse my claims about thought with certain others, however much they might be suggested by famous philosophers of the past.

References

Siewert, C. (1998). *The Significance of Consciousness*. Princeton: Princeton University Press.