

Phenomenal Thought

Charles Siewert

1 Introduction

Does consciousness—that is, *phenomenal* consciousness—embrace both sensory experience and conceptual thought? We may feel embarrassed to ask such a question. Shouldn't it be introspectively obvious what's in consciousness? But surprisingly perhaps, simple reflection on our own mental lives leaves some stubborn disagreements. That will seem less astonishing if we find we lack a clear, shared interpretation of the issue, and approach it against a complicated background of varying, unarticulated influences. What underlies the lack of agreement here may, in fact, be so obscurely complex, or so entrenched, that a consensus will never be reached. But we should not avoid the question on that account—unless we simply want to avoid philosophical disputes.

To face the issue we need some initial grip on the relevant—phenomenal—notion of consciousness. We may start with Ned Block's formulation, which expresses the now common idea that the phrase 'something it's like' aptly conveys what we're after here: "What makes a state phenomenally conscious is that there is something it's like to be in it" (Block 2002). And just *what* that "something" is—what it is like for you to be in a given state—is said to be its "phenomenal character."

The notion of phenomenal consciousness is also commonly explained by appeal to sensory examples. Its feeling painful to you when you have a toothache, something's looking somehow to you when you see its color or shape—these are prototypical phenomenally conscious states. And the *phenomenal character* of pain consists in its feeling to you just as it does to be in pain; that of vision, in its looking to you as it does when you see colors and shapes.

But what is the significance of using *sensory* examples? Notoriously, it can be very unclear just how far to extend the use of a concept beyond paradigms. Should we conclude that phenomenal consciousness is *exclusively* a sensory affair? When you doubt that something is so, or wonder whether it is, when you think about what to do, or consider reasons for or against a claim, when you try to remember why you've come

into the kitchen and recall you were looking for the pliers—you engage in conceptual thought. Do such episodes—and not *merely* their sensed expression or some concurrent imagery—belong to phenomenal consciousness? Are they as much a part of this as the experiences of hearing a thud, seeing a spiral, tasting some chocolate, or feeling an ache? Is there something it's like for you to think just as there is something it's like for you to sense? This is the gist of the issue.

Perhaps reliance on sensory paradigms would not leave this issue unsettled, if we had a relatively settled understanding of what it means to speak of there being “something it's like” to be in a state. But we don't. I will return to this crucial point.

Another problem here is that it is not immediately obvious how to get the right conceptions of the *sensory* and the *conceptual* to set up a discussion. We do not want to start by assuming too much—too theoretically loaded or ambitious an understanding of these notions. And we do not want simply to stipulate that the sensory can never *also* be conceptual. I propose we initially frame the issue as follows. Let's classify as “conceptual” such cognitive activity we enjoy that is (or can be) expressed in language, and requires capacities for voluntarily making inferences, classifications, and analogies. Next let's call “sensory features” those features whose possession is found in the activity of various standardly recognized perceptual modalities (vision, hearing, etc.) along with bodily feelings of pain and pleasure, cold and warmth, and kindred sensations, together with whatever analogs of these there might be in imagery (visualization, hearing words or music “in one's head,” etc.). Finally, let's call “*merely*” sensory features those whose possession by a subject during a time is insufficient for the occurrence at that time of some conceptual activity.

Plausibly, there are merely sensory features. For it is plausible that, for example, colors, shapes, positions, movements, distances, and sizes, can *appear* somehow to a being who simply cannot voluntarily *classify* colors, shapes, etc., in a sense that requires the capacity to make *inferences* and *analogies* relevant to understanding these classifications, or who at least is not *then and there* so classifying them. But if this is mistaken, and there are in fact *no* “merely sensory features” in this sense, or if there are, but we (normal human adults) at least typically don't have them (because even our sensory experience is—not just linked—but “permeated” with conceptual activity¹), then all (or most) instances of phenomenal sensory features with which we are familiar from our own experience will already involve the exercise of conceptual powers, and so it will be unavoidable to include conceptual activity in phenomenal consciousness. I do not wish to deny this. But the specific way of including conceptual activity that I will discuss and defend here is neutral with respect to this idea that the sensory is conceptually permeated or “loaded.” On the view I'll present, whether or not our sensory features are *merely* sensory, phenomenality should be granted to the occurrence of conceptual activity that is found in *thinking* about something, even when what

¹ I take this way of putting it from McDowell (1994).

we are thinking about is not sensorily apparent or imagined. That, roughly speaking, is what I will call an “inclusive” view of phenomenal consciousness.² It seems fair to say that such a view is advocated in, for example, Goldman (1993), Horgan and Tienson (2002), Pitt (2004), Siewert (1998), and Strawson (1994). The opposing view, which *excludes* conceptual thought from phenomenal consciousness, I will call “exclusivism.” This can be found, for instance, in Georgalis (2003), Jackendoff (1987), Nichols and Stich (2003), Robinson (2005), Tye (1995, 2000), and Wilson (2003).

The issue needs much more refinement. But why should we even care about it? What’s at stake? Considerations of length do not permit me to go into detail. Suffice it to say that whether or not phenomenal consciousness is understood inclusively has consequences for what one will take to be a viable candidate *explanation* of it (since some explanatory frameworks assume an exclusively sensory view of consciousness). The issue also bears on what resources one has for understanding the role of phenomenal consciousness in *self-knowledge*. Further, it affects one’s view of the kind of *value* consciousness has, and how that is related to the value and respect we should accord human beings. Finally, it provides a crucial test of our ideas about how to address apparent disagreements in which divergent “introspective” judgments about pervasively shared aspects of experience seem to be involved.³ Does further first-person reflection have a role in forming a reasoned response to such controversies? Does philosophy?

2 A plea for historical self-consciousness

Though an answer to the question I am raising has serious implications, it does not always get a serious hearing. Since I believe it should be taken seriously, I want to start with some remarks about historical background that may help with this. Thirty years ago or so, many philosophers concerned with what is now often called “phenomenal consciousness” would likely have just assumed that the issues here had to do entirely with “sensory qualia” of some sort—sensations of pain and after-images of color being standard examples. From their standpoint, to be “inclusive” would be to suggest—bizarrely—that there is, for instance, a “*belief quale*”—where this means that much as it feels somehow to be in pain, it *feels* somehow to hold a belief—that there is perhaps an “intrinsic,” “ineffable,” “felt quality” to belief. And they would see this as pretty obviously absurd. They would admit that activities of *thought* and *understanding* are not guaranteed by sensory experience and imagery alone. But they would say there is a special “qualitative” sense of “consciousness” confined to the sensory domain.

² Note that my issue leaves to the side the question of whether there might be some way to reduce conceptual thought to the activation of complex *dispositions* for, or *operations upon*, merely sensory features. Prinz (2002) proposes and defends such a theory.

³ Schwitzgebel (2008).

However, we should wonder a little at the fact that, a *hundred* years ago or so, philosophers concerned with consciousness also would likely have found *this* perspective puzzling and strange. What happened in the interim? Was there some philosophical discovery that consciousness, in a certain sense, is inherently sensory?

Let's go back a little over a century to see how "consciousness" is disambiguated in Husserl's *Logical Investigations*. Husserl there ([1900/01] 2001, V § 2) canvasses three concepts. One concept of consciousness, he says, is "inner perception." Another identifies consciousness with intentional mental acts generally. The remaining sense he introduces pertains to the "stream of consciousness." This includes whatever is *experienced* by someone—in the (internal accusative) sense in which an experience can itself be experienced—e.g. one experiences a sensation. Husserl rejects the idea that all experiences that are conscious in the "stream" sense are objects of some kind of "unbroken activity" of ongoing reflective judgment, such as is included in what Brentano calls "inner perception." And he thinks that non-intentional ("uninterpreted") sensations (of color, for example) belong to consciousness in this same ("stream") sense. Now then: if you purport to identify a concept of consciousness that distinguishes *being conscious* from *being the object of higher-order judgment*, and say that *non-intentional sensations* would be conscious in this sense, and characterize what is conscious as *occurrent* and *experienced* at a personal level, it may well seem that what you're talking about is none other than what some of us today would call phenomenal consciousness.

But now notice: if Husserl would be happy to include sense experience and imagery in this "stream of consciousness," he is far from *confining* it to this. He says: "In this sense, percepts, imaginative and pictorial representations, acts of conceptual thinking, surmises and doubts, joys and griefs, hopes and fears, wishes and acts of will etc., are . . . 'experiences' or 'contents of consciousness.'" Thus the sort of consciousness recognized by Husserl, most plausibly regarded as phenomenal, is supposed to be found in a rich variety of mental activities—spanning sensory experience, imagery, emotions, volition, and all sorts of conceptual thought. There is no hint at a sense of consciousness that would cast all conceptual thought into *unconsciousness*. A little earlier than Husserl, William James (1890) similarly employs a notion of consciousness embracing all sorts of thought: James' famous "stream of consciousness" is a "stream of *thought*" in a very liberal, inclusive sense of "thought." And James' and Husserl's near contemporary, G.E. Moore (1962: 72), does not hesitate to include in consciousness the event of understanding what we say, as we say it: ". . . [S]omething happens in your minds—some act of consciousness—*over and above* the hearing of the words, some act of consciousness which may be called understanding their meaning." And Moore regards it as a matter of "Common Sense" that "acts of consciousness" include not only seeing, hearing, feeling, and imagining, but also *thinking* and *believing* (the latter being for Moore an event of judging some proposition) (*ibid.*: 16, 281–3).

The contrast between this late nineteenth-/early twentieth-century way of thinking about consciousness and that which came to inform much philosophy in the latter half

of the twentieth century becomes evident, if we consider the position of a philosopher who, directly and indirectly, has had a considerable impact on this, Gilbert Ryle. In his (mid-century) *The Concept of Mind*, Ryle remarks he can find no use for the phrase “stream of consciousness,” unless perhaps it refers to a “series of *sensations*,” sensations which, by their very nature, are incapable of being correct or incorrect, and manifest no quality of intellect or character (1949: 203–5).

Ryle, working relentlessly to dispel what he saw as the “Cartesian Myth” of a mind as the “ghostly” locus of hidden occurrences, continually strove to *exteriorize* the mind by interpreting mental concepts in terms of public, observable performances and dispositions to these, while shrinking down as much as possible the class of mental events eluding this treatment. If there are some it is hard to *behaviorize* in this way (e.g. sensations, and words and tunes “heard in one’s head”), they represent only the fading shadow of the Cartesian ghost. If Ryle cannot completely dry up the “stream of consciousness,” he is determined to reduce it to a trickle. For only then, Ryle seemed to think, can we hope to neutralize the threats the dualist legend would otherwise invite: a hopelessly mysterious nexus of mind and matter, and an insoluble problem of other minds. Thus the private mental life that previous philosophers, novelists, and psychologists had imagined so rich and significant becomes, in Ryle, an impoverished sensory residue, to be contrasted with what really matters—publicly displayed qualities of intellect and character. And the sense of “consciousness” at work in James’ and Husserl’s “stream” talk is attached to the former, belittled side of this contrast. Ryle’s “logical geography” crucially influenced seminal writings in the materialist philosophical tradition of the 1950s that emerged in his wake (those of U.T. Place and J.J.C. Smart), as the subject that came to be known in university curricula as “Philosophy of Mind” was established.⁴

From roughly this time on—due in no small measure to the impact of Ryle, as well as to influential interpretations of the later Wittgenstein—it became commonplace for philosophers to contrast the sensory with the conceptual, or sentience with sapience, or sensed qualities with propositional attitudes, while conceiving of the first member of these pairs to comprise some special “qualitative” domain, perhaps marked with

⁴ Place (1956), citing Ryle with approval, starts from the idea that both “cognitive” concepts (e.g. knowing, believing, understanding) and “volitional” concepts (wanting, intending) are suitably analysed in terms of behavioral dispositions. But he then worries that this still leaves what he calls an “intractable residue of concepts clustering around the notions of consciousness, experience, sensation and mental imagery” that must somehow be fit into the materialist worldview in *another* fashion. (And here “the identify theory,” with its consciousness/lightning analogy, comes to the rescue.) Smart (1959) announces that he takes Place as his point of departure. And, with a sympathetic nod to what he sees as Wittgenstein’s and Ryle’s noble (if failed) efforts to subject talk of pain to “behavioristic analysis,” he sets himself to warding off any form of mind/body dualism, and completing a view of ourselves as nothing but “physico-chemical mechanisms.” This means vowing not to accept the reality of anything “irreducibly psychological,” even in the face of those few remaining, stubborn metaphysical hold-outs: “states of consciousness”—that is to say, “sensations.” Place’s and Smart’s articles, with their Rylean “sensory residue” conception of consciousness, sustained by its role in the fundamental project of vindicating a materialist vision, provided the pervasive background against which subsequent functional and causal accounts of mind (and criticisms of them) emerged.

a special sense of “consciousness” and strongly associated with ill-articulated doctrines of “ineffability” and “intrinsic-ness,” a domain that potentially poses a distinctive challenge for behavioral or functional accounts supposedly well suited to intellect, belief, and desire.⁵ The surface plausibility of this framework was also perhaps bolstered by schooling in empiricism’s doctrines of sensible qualities and copy-like ideas, together with standard critiques of these—critiques that fault them *not for impoverishing consciousness*, but for failing to account for linguistic *meaning, belief, and concepts*.

It is true that, as the “cognitivist revolution” consolidated, it came to seem intellectually respectable, even obligatory, to embrace a very unRylean (but now physicalized, “ghost” free) picture of the mind, buzzing with “inner processes” awaiting theoretical discovery. To that extent, the behaviorist mentality waned. Still, these “inner states” were to be thought of “functionally” or “computationally.” And construing them thus seemed to preserve a radically limited conception of experience similar to Ryle’s—but with the difference that the old behavior-dispositional construal of the intellection lying outside experience gave way to a functional or computational one. Thus the idea survived that something distinctively *sensory* exemplifies some special *sensory notion of consciousness*, to be set over against the arena of intelligence—or as it came to be called, the information-processing aspect of mind.

This history has led many who have been educated in Anglophone philosophy since the 1950s to find familiar an account on which the mind divides into some “qualitative” aspect that is purely sensory, and some behavioral/functional aspect, home to intelligence and the use of concepts. One may reject this bifurcated picture in various ways. But one lives still largely in thrall to it, if one rejects it on the grounds that the second behavioral/functional aspect of mind is all there really is to it, since the first “qualitative” aspect turns out to be either reducible to some special province of the second, or illusory and eliminable.⁶ Insofar as one’s grasp of “phenomenal” talk is rooted in the first of these purported aspects, in associated use of the jargon of “qualia,” and in certain ways of engaging with classical empiricism, one may be inclined to think that “consciousness” in one sense just *means sentience*, or, in any case, an exclusivist view is likely to seem “natural.” Introductory lectures and textbooks that take for granted some version of the framework I’ve described, and the canonical status of writings that have instituted it, promote a philosophical formation that can make an inclusivist alternative seem bafflingly eccentric, the kind of idea an ambitious graduate student would be well-advised to treat with casual scorn.

⁵ This can be seen in Wilfrid Sellars’ (1956) stark contrast between “sentience” and “sapience,” which places intelligence, understanding, and reason all on the latter side, while leaving the former to mysterious sensory “qualia.” Richard Rorty (1979), having come of age under Sellars’ and Quine’s tutelage, embraces this sort of dualism, and pronounces the notion of mind an arbitrary amalgam of sensation and intentionality, whose spurious sense of unity is to be blamed on (who else?) Descartes—whereupon the “Cartesian” sensory remnant is rhetorically bludgeoned into insignificance before being eliminated. Putnam (1981), though with more indifference than hostility towards sensations, perpetuates a similar dualism.

⁶ Tye (1995, 2000) I would regard as illustrating the former attitude, Dennett (1991) the latter.

However, it is not inevitable to mix rejection of classical empiricist views about meaning and concepts with retention of a Humean view of consciousness as comprised of cognitively primitive sense experiences (of shape, color, flavor, etc.) and the images formed of them. Hume may have been wrong *both* in trying as he did to extract all thought and meaning from a meager, flattened-out experiential basis, and in viewing consciousness in narrowly sensory terms. And the fact that James, Husserl, and Moore would have agreed on this suggests such a stance does not depend on holding either eliminativist or reductionist attitudes toward consciousness (since they did not). Further, we should consider that the tendency to recognize there is more to *mental* life than sensory activity and imagery, but to reserve *consciousness* in some sense (flagged with “phenomenal”) exclusively for what is merely sensory, comes as the relatively recent product of a certain intellectual history. And it’s hardly clear that what *explains* that history provides *justification* for this picture. Did Ryle really show that the stream of consciousness is as meager as he makes it? If some now find something in the neighborhood of a Rylean view natural, that may only reflect their inheritance of an under-examined mindset dominated by theoretical agendas that emerged, mid-twentieth century, in philosophies steeped in a particular reception of British Empiricism, and driven by a certain strategy of militant anti-Cartesianism.

But if I am right about how the exclusivist’s picture of consciousness can hold one captive, we should not expect to avoid it simply by prompting “intuitive” reactions to questions like: “Does *your* conceptual thought have phenomenal character?” If your intellectual upbringing primes you to expect any phenomenally conscious state to be highly similar to sensations or after-images, such general and direct appeals to introspection will only confirm exclusivism. On the other hand, if you are suspicious of this heritage, and your interpretation of “phenomenal consciousness” does *not* hinge on a close resemblance to its sensory paradigms, you may find James, Husserl, and Moore to have had a properly inclusive attitude, and see the dubious legacy of Ryle and Hume in shrunken conceptions of experience. For an inclusivist like me, the variations in sensory appearance that the exclusivist admits into phenomenal consciousness are but species of a broader genus of *ways of seeming*, including also its seeming to one as it does to *think* and *understand*—in a univocal, phenomenal sense of “seems.”

My challenge, then, is to recreate an alternative to the maps that sustain exclusivism. Setting aside talk of “intrinsic, ineffable qualities” as obscure and prejudicial, I will explain what I mean by “phenomenal consciousness” so that this applies appropriately to sensory paradigms, though in doing so I neither preclude nor assume, but facilitate a recognition, of phenomenal thought—while holding at bay the tyrannizing anxieties and ambitions of mind-body metaphysics. As background to my efforts, I urge a reconsideration of—and a wary self-consciousness about—the historical situation from which our question arises. Nothing dictates that we see an exclusivist conception of consciousness as natural, sensibly cautious, or commonsensical. It is just a perspective that has only lately become entrenched in certain segments of philosophical culture, entitled to no default authority. We should approach it skeptically, and open our minds

to understanding phenomenal consciousness in a way that does not tie it inextricably to the merely sensory, a way we may find as broadly inclusive as conceptions of consciousness in James, Husserl, and Moore, without ignoring developments since their day.

3 How to interpret “What it’s like”

One development we must not ignore concerns the way in which the phrase invoked at the outset, “what it’s like for someone,” has gained prominence as a marker of *phenomenal* consciousness. In the past, I have not relied on this to introduce the notion of phenomenal consciousness, explaining it instead by means of paradigms (its looking somehow to someone), and contrast cases (blindsight).⁷ But while that certainly does not rule out an inclusive view, some may take it to tie phenomenal consciousness to sensory exemplars so tightly that they cannot include thought within it. For this and other reasons, I want in this section to come to terms with the use of “what it’s like” talk to convey a grasp of the notion of phenomenal consciousness.

To tell my story, I want to show how it emerges in response to a difficulty with Block’s formulation, quoted earlier: “What makes a state phenomenally conscious is that there is something it’s like to be in it.” That formulation might be taken to imply that, if there is something it’s like to be in a given type of state, then any instance of that type is guaranteed to be a phenomenally conscious state. But this seems dubious. Suppose that I tell you I ate a durian fruit. You may intelligibly respond, “Hmm, I’ve never eaten durian—I wonder what it’s like.” Similarly, you might say: “I wonder what it’s like to weigh over a 1000 lbs.” It seems there is something answering to such expressions of curiosity. There was indeed something it was like for me to eat durian. And there is something that it is like to weigh over 1000 lbs—something it’s like *for* a 1000 lb-plus person. But if that’s right, then it is problematic to say that the fact that there is something it is like for one to be in a given state makes it phenomenally conscious. For I don’t think we should say that *eating a durian* and *weighing over 1000 lbs* are themselves phenomenally conscious states—in the way in which *tasting* a durian or *feeling like* one weighs over 1000 lbs are. Or rather, that should not be a result so cheaply won. On the face of it, Block’s dictum seems to spread phenomenal consciousness around too readily.

Now you may say that there is something it’s like for one to *eat* a durian, but only in a *different sense* from that in which there is something it’s like to *taste* durian. And it’s only this second sense that makes something a genuine phenomenally conscious state. That seems reasonable. But it remains unclear how the senses differ, and they do not seem *unrelated*. For there seems to be “something that it’s like” to be in a type of state—something it’s like for the one who is in it—in the sense *somehow* relevant to

⁷ Siewert (1998).

phenomenal consciousness, if it is a type of state, about which it is sensible to express what we might call a *subjective curiosity*. By this I mean: a type of curiosity that can be satisfied only by occupying, either actually or imaginatively, the position of a subject of that state, that is, by “adopting the subject’s point of view.” For example: you wonder what it’s like to *taste* durian, and the satisfaction of your curiosity seems to require that you either taste durian yourself, or somehow be able to imagine tasting durian. But the problem is, much the same could be said of *eating* durian. You wonder what it’s like to eat durian, and the satisfaction of your curiosity seems to require you either eat some durian yourself, or can somehow imagine doing so. So again: the application of this idiom seems too unrestricted to elucidate phenomenal consciousness, and to be of much use in addressing controversies about it, where its *extent* is precisely what’s at issue.

You may think a remedy is ready to hand. You might want to alter Block’s dictum to something like this:

What makes a state phenomenally conscious is that it is a *mental* state, and there is something it’s like to be in that state.

Thus, eating a durian is disqualified from being a phenomenally conscious state, because (intuitively) eating is not a *mental* state. But for several reasons this solution is unsatisfying. For brevity I will speak only of the reason most directly pertinent to the current issue, the controversy over phenomenal thought.

The suggested revision does not equip us to understand that controversy. For even those who say that phenomenal consciousness is ultimately a merely *sensory* matter can accept that there is something it is like for one to be in the mental state of engaging in evidently conceptual thought. They may say: “When I think about the financial crisis (say) *I hear myself speaking*—either out loud or in imagination—and *that* is sensory experience of a sort. And perhaps this is also accompanied by various other sorts of *sensory imagery*. And maybe this is coincident with certain bodily feelings, which I interpret as anxiety. Inasmuch as there is something it is like for me to be in this accompanying complex sensory state, there is something it is like for me to think about the financial crisis.”

But inclusivists will not be content to say that conceptual thought is phenomenally conscious just insofar as *something else accompanies* it, and there is something *that* is like. For an inclusivist, conceptual thought does not merely borrow its “what it’s likeness” from some concomitant state of mind, distinguishable from it in kind, any more than does sense-experience. So it is not enough to say that what makes a state phenomenally conscious is that it is a mental state, and there is something it’s like to be in that state. That would fail to distinguish mental states whose “what it’s likeness” is merely *derived from* those that *do not need to derive it* in the same way.

Now I think we *can* get our understanding of “what it’s like” into decent shape, so as to use it to address the question of phenomenal thought. We can start by rendering this critical (derived/non-derived) contrast a little more precise with reference to the durian

example. We may observe: if eating durian is something with respect to which we can correctly claim, and sensibly desire, knowledge of what it's like, this is only insofar as we tacitly assume this *comes with something else that confers that status*. I can know what it's like for me to eat durian, and you can wonder what that's like, just because, besides eating a durian, there is something distinguishable from that going on, something to which eating durian is not strictly essential, which I subjectively know and you are curious about—for example, something's *tasting* a certain way to me, or its consistency and texture *feeling* somehow to me, as I eat it.

The point can be generalized. Suppose you know what it's like to have some feature Φ , because you know what it's like to have certain *other* features to which Φ is inessential, and because certain *further* conditions obtain, which don't themselves require any additional knowledge of what it's like for one to have some feature. If you know what it's like to have Φ only in this way, you know what it's like *only derivatively*. And if there is no other way to know what it's like to have it, there is only derivatively something it's like for one to have it.

But what does it mean to say there's "something it's like for one" to have such features? This needn't be a primitive notion. We can say: there is something it's like for one to have Φ just when it is a feature suited for a certain kind of *subjective knowledge or curiosity*. That is to say, it is a feature of which one can either correctly *claim to have*, or sensibly *want to have* a knowledge of *what feature it is*, which requires one have or can imagine having that very feature.

Now this is *not* just to say it is the sort of feature whose *instances are introspectively knowable* by the person who has them. The "subjective knowledge" of features invoked is specifically a knowledge of *what features they are*—where this "what feature" knowledge is not just the same as knowledge *that a feature has this or that particular instance*. Let me illustrate. Suppose you say that some car has that "new car smell" (knowing well what it's like for something to smell that way). I, for my part, wonder *how that smells*. In this case, the knowledge that I *want* and that you *claim* is not simply the knowledge *that some car smells that way to you*. In some sense I already know that it smells that way to you. If there is a sense in which I *don't* know this, it is because there's a sense in which I'm still wanting a knowledge of *what that way of smelling is*. This missing knowledge is the sort of "what feature" knowledge that subjective curiosity longs to possess.

Now notice, this "what feature" knowledge also should not be equated simply with *introspectively* knowing that the feature has a particular instance. It is true that *by* introspectively knowing this—e.g. that some car smells that way *to me* (i.e. that way new cars smell)—I can know *what way of smelling that is*. But this introspective knowledge is then a *means* to satisfying one's desire for knowledge of *what feature this is*, a desire that one might seek to satisfy by other means: namely, by *imagining* having the feature (e.g. imagining tasting durian, or imagining the smell of a new car).

Note further that this sort of knowledge is not *theoretical* knowledge. It does not consist in the ability to *explain*, e.g. what that way of smelling to someone consists in. And note, finally, that this knowledge is not just an ability to respond to the presence of

the feature. The sheer fact that a creature can respond to an itchy feeling (for instance) by (e.g.) scratching, is not enough for it to *know what that feeling is*. What one wants when one is curious about what it is to feel a certain way is not just the ability to respond somehow to that feeling. All this yields the following proposal:

There is something it's like for one to have a feature Φ , just when one can either correctly *claim to have*, or sensibly *want to have*, a certain sort of knowledge *of what feature Φ is*—i.e. a sort that is both:

- *non-theoretical* (i.e. it doesn't require one be able to explain what having Φ consists in), and
- *subjective* (i.e. it *does* require one have or be able to imagine having Φ).

There is non-derivatively something it's like for one to have Φ just when Φ is not suitable only for *derivative* subjective knowledge or curiosity of this kind. That is: it's not the case that one has this sort of subjective “what feature” knowledge of Φ , if and only if there are some *other* features to which Φ is inessential, of which one has such subjective knowledge (and additional conditions hold, which require no subjective “what feature” knowledge).

This gives us a grip on “what it's like” talk that will allow us to frame the issue of phenomenal thought. For now we can see this in terms of the question of whether there is non-derivatively something it's like to engage in conceptual thinking, something it's like that isn't entirely derivative from what it's like to have concomitant sensory features.

But we may now wonder: does this help us zero in a bit more on what a phenomenally conscious state is, and what phenomenal character is? And how should we connect our conception of this with the distinction between derivative and non-derivative “what it's like”? I would propose the following:

A state is phenomenally conscious just when it is an instance of a phenomenal feature.

A feature Φ is phenomenal, just when *necessarily*, there is non-derivatively something it's like for one to have Φ . That is: Φ is inherently or essentially suitable for non-derivative subjective knowledge or curiosity regarding what feature it is.

Instances of phenomenal features Φ and Ψ —phenomenally conscious states, i.e. experiences—*differ in phenomenal character* (and these *features* differ *phenomenally*), just when what it's like for one to have Φ differs from what it's like for one to have Ψ . That is to say: these features differ in a manner suitable for non-derivative subjective knowledge or curiosity. (E.g. what it's like for you to feel dizzy differs from what it's like for you to feel thirsty just in case these two (phenomenal) features differ in such a way that you can correctly claim to have, or sensibly want to have, a subjective, non-theoretical knowledge of what that manner of difference is, which is not derivative from other such knowledge.)

By this standard, I take it we can interpret various appearance words—“look,” “smell,” “taste,” “feel,” “sound,”—so as to count its looking, smelling, etc., various ways to people as phenomenal features. At the same time we can allow that there can be something it's like to have certain *other* features (and one may know or wonder what

it's like to have them)—features such as eating durian, weighing over 1000 lbs, flying through the air, being near the epicenter of a massive earthquake—without just assuming illegitimately that these are themselves phenomenal features. For we may want to say that there is something it's like for one to have these features only *derivatively*, in virtue of *other* features that occur with them, of which one can have or desire subjective knowledge. Or: what it's like for one to have the first is entirely derived from what it's like for one to have the others. And this would disqualify the first set from being phenomenal features.

Further, on my account, we can allow that there may be features there is non-derivatively something it's like for one to have, which—because this status is not *essential* to them—still aren't bona-fide *phenomenal* features. This bears significantly on the present topic. Suppose we grant that some type of (conceptual) thinking can go on in either a phenomenally conscious form, or a non-phenomenal form. We might then say that *in some cases* there is something it's like for one to think in this sense, and in some cases not. But we might suppose that where there is, what it's like to think cannot be derived from what it's like to have some further feature, to which thinking is inessential. For we may suppose that, when there is something that it's like for us to think, *what* it's like is not entirely *derivative* from the merely sensory, and actually thinking is essential to this phenomenal character of thought. And yet, we may allow that thinking *in the generic sense* is not itself a phenomenal feature, because it can occur, even when there is nothing that it is like for it to occur.

In the present context, my interpretation of “what it's like” talk, and my associated conceptions of phenomenal consciousness and character display the following advantages. First, I make use of this “what it's like” talk widely held to be key to getting at phenomenal consciousness, but without simply relying on a brutally intuitive or ostensive grasp of this locution. Second, my interpretation of it allows me to count as “phenomenal” cases that *should* be so counted, if *any* should be, without illicitly assuming a verdict on more controversial cases. I leave open that (at least some) conceptual thought might be phenomenally *unconscious*, and, at the same time, that there might be (at least some) form of conceptual thought that is genuinely phenomenal. Third, the conception of phenomenality I deploy makes no assumptions about whether or not phenomenal character is “exhausted by” some proposed physicalist or functionalist characterization of its nature. This initial *neutrality* is surely desirable, when we set out to address the question of phenomenal thought. Finally, I think it counts in favor of my approach that it takes into account something we have seen is essential to framing that question adequately—namely, the difference between there being something it's like to be in a state *derivatively* and *non-derivatively*.

4 What the issue is

To apply the preceding conception to the issue at hand, we need to see a bit better where inclusivists and exclusivists can find common ground. We can agree that, not

just generally, but on this or that *occasion*, we take or understand what we hear or read, say or write, in this way or that. That is, there is such a thing as *occurrent understanding*. To make this stand out: consider the difference between, on the one hand, reading some text (either to yourself or out loud) while attentive to its meaning, following what is being said, understanding it in a certain way, and, on the other, reading it *without* following it, without ongoing understanding, so that you need to restart at the beginning, concentrating on the meaning.⁸ And if you can in this way either follow or fail to follow the meaning of your own *recitation*, you can occurrently understand your own *spontaneous* speech. (There's no reason to think that the *spontaneity of your speech* should destroy the *occurrentness of your understanding*.) And the occurrent understanding of speech can be said to be a way of thinking, in the sense that requires *conceptual* abilities, as glossed earlier. So, *occurrent understanding is a form of episodic conceptual thinking*. We may also point out: we are often engaged in episodic conceptual thought such as is involved in (for example) *trying to think of the answer* to a problem or in *supposing* that or *wondering* whether something is so, or in *doubting* that it is so, or in *considering reasons* for doubting it, or *concluding* it is so—whether or not such thought is audibly or visibly expressed. Finally, we can agree that ongoing differences in *the ways we think or understand* are subjectively discernible to us. Not only can we judge with warrant, in a distinctively first-person way, how things look, sound, feel, and smell to us, it is also subjectively discernible to us how we're thinking or understanding from one time to the next. (I intend here to leave quite open just *what* distinguishes this “first-person way.”) All this we can accept without yet speaking explicitly of the phenomenal character of such thinking/understanding, or of what it's like for one to understand occurrently.

Against this background, we can see the issue of phenomenal thought as turning primarily on two questions. First there is what I'll call the question of *Reducibility*. Suppose there is only derivatively something it's like for us to think and understand as we do, because knowledge of what it's like is *entirely derivative from* knowledge of what it's like to possess concurrent, separable sensory features—i.e. ones whose possession at that time is insufficient to guarantee anything about our thinking and understanding as we then do. Then knowledge of what it's like to think and understand in the manner we do is *reducible* to this knowledge of separable sensory features. This is the thesis of *Reducibility*. If you say yes to *Reducibility*, you are *exclusivist*; if you are *inclusivist*, you say no.

Second, there is what I'll call the question of *Variation*. Do significantly many subjectively discernible differences in ways of thinking and understanding constitute

⁸ When I speak of an occurrent understanding of an utterance, of course I do not intend to imply that this is always a *correct* understanding of the utterance as meant by the speaker (or of its meaning in the language). I am just saying that the utterance is then *taken* in a certain way, *understood in a certain way*, whether or not this coincides with speaker or sentence meaning. As it would be too cumbersome to make this qualification explicit every time, I will often omit it, trusting the reader to supply it as needed.

differences in *what it is like for us* to have the experience we do? If you are exclusivist, you say no to Variation; if you say yes, you are inclusivist.

To affirm Variation is to deny Reducibility. Now you could coherently deny both Reducibility and Variation. And admittedly, I have left it undefined whether you would then be inclusive or exclusive. But it does not matter much to me how we extend my labels to such a position; what I am most interested in here is the sort of inclusivism captured by the combination of Irreducibility and Variation.

Framing the issues as I propose allows us to understand questions about the phenomenal character of thought as follows. If Reducibility is false, then the phenomenal character of one's occurrent thought and understanding is not "exhausted by"—there is "more to it than"—the phenomenal character of sensory experience separable from it. And since there is (non-derivatively) something it's like to *think and understand*, there is *inherently* something it's like for one. Then there is, let us say, some *cognitive* phenomenal feature. And if Variation is true, then there are significantly *many* phenomenally differing cognitive phenomenal features. In that case, the phenomenal character of cognition, of conceptual thinking, is *rich*. Questions will then arise about just how this relates to *what* is understood or thought—questions about *content*. I will come to these. But right now let's focus on the questions of Reducibility and Variation.

Seeing the issue of "inclusion" in these terms allows us to avoid certain prejudicial ways of framing it. Occurrent conceptual thought is typically (maybe sometimes even necessarily) *unified* with either sensory or imagery experience. Occurrent differences in ways of thinking are often, in some sense, *intimately bound to* experienced differences in verbal expression, whether silently imaged or publically perceptible. As I *think* differently, I *speak* differently—and I experience this speech. But someone might suppose that to confirm inclusivism, we need to find, *alongside* both our sense experience of speech and verbal imagery, *and* hitherto subjectively discerned differences in ways of thinking/understanding expressed in it, something *additional*, which one may also always find humming along all on its own, utterly disentangled from utterances sensed or imaged, and phenomenally just the same whether or not it "accompanies" experienced or imaged speech. Perhaps one will imagine that we need to find sensory and cognitive phenomenal features occurring as distinctly and autonomously as, say, visual and auditory appearances do, if there really are cognitive phenomenal features to be discerned.

But these would be false assumptions. For we can now see that the inclusivist does *not* require that, *in addition* to the ongoing differences in ways of thinking and understanding that you already discriminate in first-person reflection, there is a phenomenal *something more* we should look for. Rather what you're asked to consider is that some such reflectively recognizable variations in one's manner of thought and understanding *are themselves phenomenal differences*, whether or not you think of them as *such*. Furthermore, it forms no essential part of the inclusivist picture to suppose that words are merely the "clothing of thought," needed for public appearances, which might just as well be shed in the mind's privacy, where phenomenal cognition can

nakedly frolic. No, as far as the inclusivist is concerned, it may very well be that *thought needs talk*, in various ways. To approach the issue fairly, all this needs to be clear.

5 A case for being inclusive

Now I will give three distinct though closely related arguments for being inclusive. I do not think such matters should be decided in a single blow. We should consider a variety of relevant issues from a number of angles. Indeed, I think further arguments for inclusiveness need to be examined beyond those I will express here. However, for reasons of length I will confine myself to these three—linked by their common focus on the *experience of understanding language*.

Let me first say something about the burden of proof. Struck by the broader agreement that at least sensory appearances and imagery count as phenomenal, one may be tempted to think the burden rests on those who also somehow include conceptual thinking. And (one might think) if this burden is not met, we should default to exclusivism. I would disagree. First, generally: the mere fact that group A is inclined to extend a classification more generously than group B does not, by itself, show that A owes us more of an argument or is on shakier ground than B. (Maybe group B is unreasonably *stingy*.) Second, more specifically to the case at hand: a hesitation to be inclusive does not necessarily reflect higher epistemic standards. It may, for example, come from assuming (mistakenly, as we now see) that we can apply the notion of phenomenality only in virtue of close similarity to sensory paradigms.⁹ Or one may decide the matter based on misconceived, misleading, or gratuitous tests—such as whether one is inclined to say there is a “qualitative feel” to belief,¹⁰ or whether, beyond the subjectively discernible variations in how one understands one’s words, one can detect, alongside this, some further, phenomenal “extra.”¹¹ Or one may mistakenly think inclusivism assumes that our conceptual thinking can or does have no need at all of verbal expression.¹² Or that it requires we have a certain view about *concepts*—that they are introspectible objects, rather than abilities.¹³ Or (relatedly) one may take an inclusive view to hold that we have some appearance-like “acquaintance” with the *contents* of our conceptual thoughts, that these are thus “presented” to us, yielding a “cognitive phenomenology” that is “pure” of any sensory presentation.¹⁴ However, we can now see that the version of inclusivism I ask us to consider assumes *none* of this, and is not threatened by their doubtfulness, though it has been assumed otherwise. There is this additional reason why the occurrence of cognitive phenomenal

⁹ Georgalis (2003).

¹⁰ Nichols and Stich (2003: 198).

¹¹ See Robinson (2005: 253): “. . . a further phenomenological something.”

¹² Jackendoff (1987: 289–91). (I respond to Jackendoff’s arguments here in Siewert (1998: 361–2).)

¹³ Putnam (1981: 17, 20–1). Note: the inclusivist view is fully compatible with the idea that to have a concept is to have an *ability*.

¹⁴ See Levine, this volume.

features is more controversial than the occurrence of merely sensory ones. For if we admit that there is something it's like to have sensory features, the idea that there might only *derivatively* be something it's like to have them can get no traction. For whether these be “merely” sensory features, or “conceptually loaded” ones, they will pretty clearly sometimes occur in the absence of any *candidates to consider* as that from which their “what it's likeness” can be derived or “borrowed.” By contrast, in the case of what it's like to think and understand, it's not clear that we're ever lacking such candidates—namely, concomitant sensory features.

The overall point here is: all the just mentioned factors might explain the greater ease with which the phenomenality of sense experience is admitted. But they do not *justify* exclusivism. And if occurrent thought and understanding *should* be segregated from phenomenal character, that is not a result we get by default. With this in mind, I would like you to consider some arguments for being inclusive.

5.1 Argument (i): *The elusive duplicate*

Here is a passage from Thomas Jefferson:

Every difference of opinion is not a difference of principle. We have called by different names brethren of the same principle. We are all Republicans, we are all Federalists. If there be any among us who would wish to dissolve this Union or to change its republican form, let them stand undisturbed as monuments of the safety with which error of opinion may be tolerated where reason is left free to combat it.

Reading this in the midst of an article, and with a contemporary reader's diet of flat contemporary prose, one might well not quite follow at first what Jefferson is saying in the last sentence, and find one needs to come at it again with more attention to its meaning before one gets it. (I confess I had to re-read it when I came across it.) If this example does not work for you, others should be easy to encounter, especially if you read philosophy—or if you just crack open your Shakespeare.

In such cases we find the sort of contrasts earlier introduced—contrasts between reading *while following what one is reading* (understanding it to mean something as one is reading it), and reading *without following the meaning* (or only partly “getting it”) because one is not concentrating adequately. Now ask yourself: is what it's like for you to read a passage without understanding it different from what it's like to then re-read the same passage and understand it in a certain way? I will assume you can answer this question affirmatively, on the interpretation of “what it's like” talk earlier explained. You can correctly claim to have (and someone could sensibly want to have) a subjective, non-theoretical knowledge of what the manner of difference in question is, non-derivative from other such knowledge. The question will be whether *that* difference is a difference in separable *sensory features*—whether what it's like for you to read with understanding is reducible to what it's like to have some concurrent merely sensory features that could occur in the absence of understanding.

What it's like to understand should not be conceived of merely negatively. The phenomenal difference with which we're concerned does not consist entirely in the fact that one *lacks* in the understanding case some *feeling of confusion* one has in the uncomprehending case, so that what it's like to read comprehendingly is just what it's like to read *without* following the meaning—except in the semantically deficient case, there's *phenomenally* an addition: a certain feeling. The experience of re-reading a passage, this time with understanding—the experience of “getting it”—is a *positive phenomenon*; it is not just the experience of saying words to yourself without simultaneously feeling confused or frustrated.¹⁵

If this is not already obvious, consider a bit more closely what we might mean by speaking here of “feelings of confusion.” First, if such feelings are supposed to be *merely sensory* experiences of some sort (perhaps the feeling of knitting one's brow, or perhaps something similar to somatic flutterings of anxiety), they do not *invariably* or even *usually* coincide with the experience of not following what one is reading. Often I'm reading along and I simply realize that I didn't get what I just read, and ceased to follow the meaning, and so am prompted to re-read it—but without having experienced any discernible sensory disturbance. Even so, what it's like to have the experience of not “getting it” is different from what it is like to get the meaning, when I do, upon re-reading. Now if *this* sort of “not getting it” experience counts as a “feeling of confusion,” then it includes some sudden (generally unverballed) *realization* that you didn't understand what you just read. And if you admit this, then you admit there is something it's like to be struck by a thought, where no sensory reduction is to be had. (A willingness to call such realization a “feeling” in no way shows its merely sensory status.) But, if one admits this, then it would be arbitrary to exclude from phenomenal consciousness the way one understands the sentence when one *does* get it, so as to maintain a purely negative view of what it's like to understand.

But suppose you don't admit that you know what it's like for it to occur to you that you didn't get what you just read in such circumstances. Still, you shouldn't maintain the purely negative view. For notice that when understanding a passage requires intense concentration, one can read it while *refraining from any such effort*, and thereby experience a lack of comprehension. But in such cases, one doesn't in any sense *feel confused or frustrated* (nor, for that matter does one always concomitantly experience some distinctive somatic sensation). The experience of “not following” when one *refrains from even trying* is not an experience of *confusion or frustration* (which presupposes some effort or interest in comprehension). But still, what it's like to read without following the meaning differs from what it's like to read while following the meaning. And where there is no *feeling of confusion* involved in the former (nor any supposedly

¹⁵ This is partly in response to Robinson (2005). He suggests, in effect, that we sometimes should reduce the phenomenal character of understanding to the absence of a “feeling of frustration”—though he also adds a positive element: a “feeling of relief” when the frustration goes away. In this connection, see also Pitt's (2004) remarks about the experience of “bewilderment,” which my points here are meant to supplement.

special somatic feeling), and what it's like not to understand is held to involve no concurrent *realization that one isn't following*, then no *positive* experience of incomprehension remains, about which we can propose: "there's nothing more to what it's like for me to understand, beyond the absence of *that*." On the contrary, what it's like to read *without* understanding will need to be seen in terms of the absence of what it's like to read *with understanding*.

Let's now return to the question of Reducibility, approaching this via the case of durian. It is plausible that what it's like to eat durian derives entirely from what it's like to have other features, having to do with the way durian tastes and smells, and how it feels in one's mouth. This is because it seems to us that actual durian eating could be absent, though *what that was like would remain the same* (because the right taste, smell, and feel were still there). This could happen if we ate cleverly engineered *fake* durian, or (perhaps) if we vividly *hallucinated* or *dreamed* we had a meal of durian. But suppose this weren't the case. Suppose there were just no conceivable way that, in the absence of eating durian, what it's like for one to have what remained would be just the same as what it's like for one to eat durian. Then what it's like to eat durian would *not* be entirely derivative from what something else was like. For there wouldn't be a "something else" that preserved what this was like for one, whose presence could serve to endow durian eating with its phenomenal character.

Apply this to the case of thought and understanding. It is not now disputed that there is positively something it's like for one to think and to understand what one is saying, hearing, or reading as this happens—any more than it's disputed there is positively something it's like to eat durian. What's disputed is whether this derives entirely from what it's like to have concomitant, separable sensory features, as Reducibility requires. If it does, then in every case we should be able to find some sensory features that exactly duplicate what it's like for us to read comprehendingly, but in the absence of any way of understanding what is read, much as we could duplicate what it's like to eat durian without actually eating genuine durian.

Can you find such features? You could try, by comparing *actual* cases where you read a given text without comprehension and then you re-read it *with* understanding. But this won't support Reducibility, unless you sincerely judge that what two such actual experiences of that text are like for you is just the same. And I doubt you will find this. I think you will admit that in actual cases, when you read a passage you do not understand or follow, what this is like for you differs from what it is like to read it again, this time following it, understanding it. Since comparisons of *actual* experiences of reading do not support the judgments of phenomenal sameness needed to warrant Reducibility, for this you need to conceive of a *hypothetical* case, constructed by first isolating some sensory features present in the case where you understood, and then *subtracting in imagination your occurrent understanding of what's written*. And once you do that, then it needs to be clear to you that what it would be like for you to have this *hypothetical experience of uncomprehended text* would be just the same as what it is like for you to have your *actual experience with understanding*. Then you could reason: since what

it's like would be just the same if you simply duplicated the appropriate sensory concomitants, the fact that there is something it's like to understand is derived entirely from its association with these.

But by appeal to what separable sensory features could we hope to construct in thought *totally semantically clueless* phenomenal duplicates of our actual comprehending selves? These sensory features would need to differ from those found in the actual cases of reading without comprehension, where what it's like for one to read admittedly differs from what it's like to read the same text comprehendingly. But they must be only *subtly* different, for they would also need to be found in the ordinary case of reading that very same text *with* understanding.

Will they perhaps involve some *sensory imagery*? But what would this be—perhaps: my aurally imagining (saying “to myself”) the words I'm reading? That suggestion won't work for those who say they don't utter words to themselves when they read silently. But while I'm *not* such a person, and do say the words to myself, I also do this when I don't follow what I'm reading. Is the separable sensory element then perhaps *extra-verbal* imagery that supposedly always coincides with what it's like to read with understanding? But when I don't understand what I read, and then go on to a second, more successful try, I do not find this does always (or really, even usually) coincide with the sudden upswell of imagined shapes, smells, colors, or noises. (And even if it just happens to do so, I don't get the requisite sameness judgments, if I try to imagine an analogue of myself, whose uncomprehending mind is filled with sensory, extra-verbal images stripped of all relevant conceptual interpretation, when the passage is re-read.)

Does reading with understanding perhaps then always involve an enhanced appearance of the details of shape and position of the letters or words? But equally detailed experience of shape and position will readily be found in a *phenomenally different* case without understanding. (The words of Jefferson's text and their placement on the page appeared to me quite clearly when I didn't follow the meaning.)

Is there then perhaps some distinctive way of experientially *grouping* the words, with a certain rhythm, or saying them to yourself with a certain intonation, which is notable in the actual experience of reading comprehendingly, but could also be found in reading without understanding, such that when I imagine what *that* would be like for me, I find it to be just the same as what the first, actual case is like for me? I find I can, in actuality, read along, quite *fluidly grouping words as a competent speaker/reader of the language*, without just then *following the meaning* of the text—and what *that* is like for me is, again, *different* from what it's like when I re-read it, this time following the meaning. (This is especially clear where I'm reading text in a foreign language, in which my competence is *minimal*—when it is all too easy for me to read along fluidly with little or no understanding.) And while re-reading a passage with comprehension may sometimes involve saying the words to oneself with a certain emphasis or intonation, I can also do this without understanding, and often enough I re-read such passages without any such exaggerated stress when understanding them. The upshot is, when I try to imagine reading a text as *fluidly* as I do with understanding, but now *without*

understanding it, it does *not* seem to me that what this would be like for me is just the same. For again, what I am asked to imagine is not clearly discernibly different from actual cases of reading without understanding—where what *that's* like is *not* phenomenally the same as reading with comprehension.

Perhaps I should consider: one can experience what one reads as *syntactically well formed* in a certain manner—as grammatically “parsed”—even if one doesn't understand it. Would that give us the desired segregation of conceptual activity from phenomenal character? The suggestion would be this: “Take the sentence you experience with ongoing understanding, and in imagination subtract any occurrent understanding of what it means, leaving only a grasp of grammar. Now do you judge that what *that* would be like for you is just the same as what it is like for you to read the same passage comprehendingly?”

What I am here asked to imagine as hypothetical is again something I already *approximate* in actual experience—because sometimes when I read along without (entirely) following the meaning, this does not (as near as I can tell) involve any snag in getting the mere *syntax* of the passage. I am reading a text aloud, and suddenly it occurs to me I haven't been following; still I recognized the sentences as well formed. But when I re-read with understanding, again, *what that is like for me is different*. (And surely it wouldn't change things—I wouldn't get a judgment of phenomenal *sameness*—if somehow I managed to imagine away every *trace* of semantic understanding from the case in question.)

Perhaps I should try to make the comparison where there is, on neither side, a deficit in attention to the meaning of the text. Maybe I should consider actual cases where I experience a sentence as well formed grammatically, but do not understand it as I read it, though not through lack of semantic attention—because in this instance I simply have *no understanding of many of the terms involved*. Then by analogy with this I could try to construct a relevant hypothetical case. A familiar example to work with would be the opening of Lewis Carroll's “Jabberwocky”: “Twas brillig, and the slithy toves/ Did gyre and gimble in the wabe;/ All mimsy were the borogroves,/ And the mome raths outgrabe.” So the hypothetical I am being asked to imagine is that what I actually read with understanding becomes for me “Jabberized”—I experience it *merely phonologically and syntactically*, as I at least *come close* to doing with the opening lines of Carroll's poem.

But to the extent I can imagine what that would be like, I certainly do *not* get a clear judgment of phenomenal sameness. If anything, I imagine experiencing a disturbing alienation from the meaning of what I speak and hear. And this takes on a nightmarish quality, as I realize that the exclusivist view implies that if my reading were to become *routinely* Jabberized, *what that would be like for me would be the same as my actual experience of reading*. And that doesn't seem right at all. Of course, it is fun to read a poem like “Jabberwocky.” But the thought of what it would be like for my experience of reading *generally* to become Jabberwocky-like is horrifying. Or I might at least say: all reading would become horrifically *boring*—even if it were still accompanied by imagining lots

of uninterpreted shapes and sounds.¹⁶ In any case, I find again I cannot carry out the relevant hypothetical construction in a way that secures the judgments of phenomenal sameness required for Reducibility. Are you confident you *can* do this?

We must not hastily conclude that we have done this, based on a tendentious or cursory consideration (as, for example, in Putnam 1981). And we must take care not to confuse conceiving of a case where understanding is defective or partial, with conceiving of a case where understanding is *totally absent*.¹⁷ The issue here is certainly *not* whether we might or do sometimes “feel like we understand” something better than we do. The issue is: do we find, once we contrast the experiences of reading with and without understanding, that we can construct a hypothetical case of reading, utterly devoid of understanding, which yields the right judgments of phenomenal sameness? If not, then we should reject the Reducibility thesis.

But if this is rejected, can't we still reasonably deny Variation? One may propose that the additional aspect of what it is like to think and understand is *not* constituted by the subjectively discernible ongoing variation in our manner of thinking and understanding, but by something else. However, we should accept this proposal, only if there is some other reasonable candidate for this something else. *Is there?* One might suggest: here there is, not a separable *sensory* experience, but a *non-sensory* experience of understanding—a perhaps more or less homogeneous “understanding-feel”—separable from the varying ways of understanding one's words as one speaks. And that non-sensory “understanding feel” exhausts the more-than-sensory aspect of phenomenal character found in the cases under discussion.¹⁸

But is it really clear to you that there is, subjectively distinguishable both from your occurrently understanding a phrase *in a certain way* as you read it, and from the *experience of words* comprising it, always *something else* that occurs, and that this additional component, together with some separable sensory one, exhausts what it's like to read comprehendingly? Suppose that is *not* clear to you. Then at least: there clearly *are* occurrent variations in ways of thinking and understanding. And since that is all that clearly *remains* to serve as the aspect of what it is like for you to read comprehendingly that isn't exhausted by separable sensory features, you should just default to view that this does constitute part of the phenomenal character of experience.

¹⁶ Galen Strawson has suggested in conversation that something like this point could be converted into an argument of its own: we can see that understanding is included in phenomenal character, because otherwise reading would be intolerably boring.

¹⁷ Putnam (1981) invites us to imagine being hypnotized into having a “feeling of understanding” as we utter words in a language we don't in fact understand at all. But this approach to the challenge is flawed in several ways, as I argue in Siewert (1998: 293–6), where I also discuss the importance here of the distinction between partial understanding and total incomprehension.

¹⁸ Somewhat similarly, Robinson (2005) confines the phenomenal character of understanding to a generic “feeling of familiarity.” But “familiarity” won't do here, since you can cease or fail to follow the meaning of a passage whose words nonetheless are quite familiar to you. And then when you do understand, the phenomenal difference does not involve a sudden injection of some hitherto absent sense of familiarity.

Maybe you are tempted to endorse this idea of a separate, non-sensory “understanding-feel.” Then consider again what it’s like for you to read a particular passage comprehendingly. Now suppose what it was like for you to read it then were just what it was like for you to read *without* understanding. In that case, wouldn’t you *think* you were *not* understanding what you were reading? And so you would *not know* what you were understanding what you were reading to say? This seems right. For otherwise, why are you prompted to *re-read* the passage, when what it’s like for you is what it’s like to read without following what’s being said?

But consider: if how you varyingly understood what you were reading lay strictly outside of what it was like for you to read, this would seem unintelligible. For first: in that case, in order to know how you are understanding what you are reading, you’d need to know *much more* than what it’s like for you to read comprehendingly on that occasion. (Especially, if we’re assuming the “understanding feel” is uniform: for that at most could yield the mere knowledge you are *understanding something somehow*.) But now, anything additional, which you would need to know and which would be available for you to know—about how you understand a given phrase—would *already be enough to do the job by itself*. For what is available for you to know is what specifically you are taking utterance “u” to mean. But if you already know *that*, you already know you are understanding “u.” There is no gap in your semantic knowledge that would need to be filled by *also* knowing you were having some accompanying phenomenal tinnitus-like “buzz of understanding.” That would tell you nothing you didn’t already otherwise know; it would just be *epistemically superfluous*. But we granted earlier that knowing what it was like was not epistemically superfluous. For we said we wouldn’t know we were reading with understanding in a given case, if what it was like for us to read then were no different than what it is like for us to read without comprehension: we’d think we didn’t understand.

By contrast, phenomenal character *can* be relevant to knowing how one is understanding utterances, provided we admit that one’s way of understanding is inseparable from what it’s like for one to read with understanding. For, in that case, in the absence of phenomenal character, one wouldn’t know how one was understanding what one was reading, for this reason: the manner of occurrent understanding itself would then be absent.

Argument (i) can now be summed up like this. What it’s like to read a passage *with* understanding positively differs from what it’s like to read the same passage *without* understanding. This difference will be reducible to a merely sensory phenomenal difference, only if we can isolate separable sensory features in each such case of understanding, and imagine having them in a hypothetical case *without* understanding, in which we find what that would be like is *just the same*. But efforts to identify these separable sensory features and secure the relevant judgment of phenomenal sameness do not pan out. So *Reducibility is rejected*. The acceptance of Variation will follow, if we can subjectively discern no good candidate to serve as the occurrent additional aspect to phenomenal character, other than variable ways of understanding. Some might suggest

there is, separable from this, an “understanding experience” that serves as the sole locus of phenomenality. But it’s hardly clear that such a separate experience *is* subjectively discernible, whenever there is something it’s like to read comprehendingly. And if we insisted on confining phenomenal character to some such factor, we would render unaccountable the fact that we take ourselves not to understand what we are reading, when what it is like for us to read is what it is like to read uncomprehendingly. Thus *Variation*.

5.2 Argument (ii): Delayed understanding

Now consider this sort of case. I meet a friend, and she asks me, “Did you bring the book?” For a moment I am at a loss as to what book she’s talking about—and then I realize in an instant what book it is. (It is: *A Supposedly Fun Thing I’ll Never Do Again*, by David Foster Wallace, which I had said I would lend to her.) But I didn’t *say* the title, even to myself, even partially, in the moment of realization, nor did I visualize the cover of the book. Though I needn’t, I *might* have said, a moment after realizing what she meant, “Oh, *that* book.” But that comes after the fact. Now I know what it’s like for me to realize suddenly what was meant on such occasions. So there *is* something it’s like. And this is not merely what it’s like for a separable sensory “confusion experience” to *go away*. For even if we construe the experience of being momentarily “at a loss” in terms of some such sensation, once it dissipates I might be struck by no positive realization, and what *that’s* like for me differs from what it’s like to *have* the realization. But now: there’s also no separable sensory feature I have, concurrent with the realization, of which I have subjective knowledge that might reasonably serve as a candidate from which to derive the knowledge of what it is like to realize what was meant. Thus, there is a change in phenomenal character, but no relevant coincident change in sensory features separable from the advent of understanding. And so Reducibility is rejected.

This argument cannot be evaded by saying: “What it’s like here is merely what it’s like to experience the *onset of a disposition* to make and experience certain *sounds* in response to *others*”. (For instance, to someone’s ‘What book?’ I am disposed to utter, ‘*A Supposedly Fun Thing*, etc.’) But what sounds I am disposed to utter depends on what I would take them to *mean*. Without that link to understanding, I would lack such phonetic dispositions. And so what it’s like to understand here would not in this manner be confined to sensory features, severed from understanding.

This sort of case differs from that in argument (i). Here my realization of what was meant did not come through some sort of repetition (e.g. re-reading) with greater attention to the meaning (as in the Jefferson example). Notice too that here, something like the earlier “parsing” suggestion for how we are to construe the phenomenal character does not even get a foothold. One cannot suggest that the phenomenal difference between not getting what was meant and getting it is limited to some purely syntactical grasp, peeled away from one’s semantic understanding—for there is just no new grammatical grasp to get.

However, the same basic argument could be reformulated in cases where the delayed understanding does occur after reiteration of the sentence, and there are “parsing” problems—the sort of examples Pitt (2004) and Horgan and Tienson (2002) use to support inclusivism—of “machine gun” and “garden path” sentences. (An example of the former (in Horgan and Tienson): “Dogs dogs dog dog dogs.” An illustration of the latter (in Pitt): “The boy the man the girl saw chased fled.”) Here one usually does need re-readings or re-hearings before one can make sense of these sentences (and there is something that it’s like for one to do so). The experience is something akin to recognizing a visible pattern or form to which one had been oblivious. But these are not just cases of sensory figure/ground organization. Nor is it merely a matter of experiencing a certain rhythm or pattern of stress or intonation. For while saying the sentences with a certain rhythm, stress, or intonation can *help* one to recognize the meaning, one can also say or hear them in that way and still not “get it.” But then when one *does* get it, there is a difference in what that is like. And that is not just the departure of some positive, separable sensory, “confusion feeling,” which anyhow is not present in every case (if indeed it ever occurs at all). Further, if someone should here raise the suggestion that the phenomenal difference noted is in some sense *merely* syntactical, the response would partly parallel what I said in connection with argument (i). That proposal implies that if our grasp of what we read were to become “Jabberized,” what it’s like for us to read it would be just the same. But that is not a consequence we should find acceptable.

Another point to bear in mind: the fact that one can or does *sometimes* image something relevant at the time of understanding (the book cover, the face of David Foster Wallace, a bunch of dogs) should not detain us. What’s crucial is simply that *sometimes* this sort of delayed understanding occurs *without* the advent of such imagery (just as *undelayed* understanding does). And yet I know what it’s like for this understanding to occur, even at such time as no reasonable candidate for a reduction of what it’s like to a detachable sensory element presents itself. (Note: beware here of what we might call “the Hume effect.”¹⁹ Once you *ask yourself* if you imagined dogs, you find yourself *now* imagining a bunch of dogs. And you take this to show that you *were* imagining dogs before. This kind of slip should not prevent us from recognizing that, often enough, there is a phenomenal difference between (first) not understanding what’s meant by an utterance and (then) understanding it in a certain manner, even though *at that time* no image pops up in one’s mind.)

At this juncture similar issues arise as before about the move from rejecting Reducibility to accepting Variation, and similar arguments are available. Will someone say that what it is like to “get it” is phenomenally just the same regardless of what is “got”

¹⁹ You ask yourself whether you form an image of a horse when you speak and think of a horse, and by so asking *prompt yourself to picture a horse*—forgetting that commonly enough when you speak or hear of horses you don’t *then* visualize any horses. I call this the Hume effect because I think it accounts for the appeal a Humean image theory of meaning can have to those coming to it (and the issues) initially.

(or even, perhaps, whether *anything* is got)? I myself do not discern any experience meeting that description, always occurring in addition to the *specific way of understanding* I take myself to have on this or that occasion. In the moment when I realize what book she meant, it is not as though I am aware of two distinct things: (a) my occurrently understanding her to mean a specific book by “the book,” and (b) a (more-than-sensory) “feeling” of understanding her to mean *something* by “the book.” And if I were to insist otherwise, and confine phenomenal character to (b), then I would have the epistemic problem again. Consider: I wouldn’t have known what book I took her to mean, had there been nothing it was like for me to realize she meant *A Supposedly Fun Thing I’ll Never Do Again*. That is, if what it was like for me at that moment were no different from what it would have been like for me to be just then *utterly devoid of thought*, I wouldn’t have believed I was getting what she meant at all. But this seems unaccountable, on the hypothesis that what it was like for me would have been just the same regardless of what I thought she meant. For then the most the presence of that phenomenal character could get me would be knowledge that *I am understanding her to mean something or other by that phrase*. But what more is subjectively discernible that I would need to know to get more specific knowledge, or to confirm the accuracy of the “understanding signal”? It is simply: the specific way I took her to mean the phrase. But this additional thing would be enough *on its own* to do the job of giving me knowledge that I understood her to mean something by her utterance, without the need of any separate “phenomenal indicator.” By contrast, if the way I understand what she’s talking about is itself *part of* what it’s like for me to understand it, the relevance of the latter to knowledge of the former is restored.

Argument (ii) comes to this. Sometimes there is a delay in one’s understanding a given utterance, and when one does understand, no relevant, coincident, separable sensory change occurs. If you know what it is like for you to experience such a delayed understanding, then there is something it’s like to understand occurrently, even when there is no sensory feature detachable from understanding to which that is reducible. Moreover, the specific way one understands on such occasions will be a phenomenal feature, and different ways of understanding will differ phenomenally, unless phenomenal character can be confined to some “understanding experience” that remains invariant regardless of what is understood. But the claim that there always occurs, at the relevant time, some such distinct, subjectively discernible experience, separable from one’s specific way of occurrently understanding, not only is highly dubious in itself, but renders unaccountable how the phenomenal character of understanding could matter to one’s knowledge that one understood what was said in a certain way. Finally, once *some* subjectively discernible differences in ways of occurrently understanding are admitted into the phenomenal realm, it would be arbitrary not to make the admission *general*. It would be absurd to suppose that having been *delayed* is essential to making different ways of understanding phenomenally different. Variation then ensues.

5.3 Argument (iii): Interpretive switch

My last argument concerns cases—not where you go from *not understanding* a phrase to *understanding* it, but—where you switch from understanding it in *one* way to understanding it in *another*. So, consider this conversation overheard, on a sweltering day in Miami.

A says: “I am *so hot*.”

B responds: “Well, okay, but you don’t have to brag about it.”

Many punning attempts at humor illustrate this experience of alternatively resolved ambiguity—where getting the joke involves some sudden interpretive switch. Of course, efforts at humor need not be involved. You may hear someone say, “I think I’ve got a virus,” and switch from thinking he says he’s coming down with something to taking him to mean *his computer* has a problem—or vice versa. The question is, in these cases, do you know what it’s like to switch from one interpretation to the other? If so, then is there always some difference in how the words sound or look to you, or some difference in accompanying imagery or bodily feel, which you can conceive of independently of the difference in meaning or interpretation involved, and to which you might reduce the phenomenal difference in question? If not, then again, we have phenomenal differences involved with conceptual understanding that can’t be segregated to some separable sensory features.

Again, here we may consider the potential gap between denying Reducibility and affirming Variation. Maybe someone will say that in these cases there is just a homogeneous (though admittedly non-sensory) “switch” experience—but no phenomenal difference between understanding “hot” or “virus” in one way, and understanding them in another. Again, I myself find in my experience no separately distinguishable phenomenal “switch” indicator, apart from the very change between interpretations itself. And there is a problem with such a suggestion in any case. If that were true, then there would be no difference in what it is like to switch in one direction rather than another. For example, what it is like first to understand “virus” as *biological infection*, then as *computer problem*, would be exactly what it is like first to take it to mean *computer problem*, then *biological infection*. But I can tell the difference between what it’s like to resolve these ambiguities in one order rather than another. So the premise is mistaken.

If you balk at that, then consider again the epistemic issue. Had it not been like what it actually *was* like for me to switch interpretations, had my experience not had that phenomenal character, I would not have thought the switches (e.g. from biological to computer virus, from temperature-hot to sexy-hot) occurred—and I wouldn’t have known they did. But if all that were going on *phenomenally* here was some uniform non-sensory “Switch!” experience, the most I would get from this was that *some switch or other* is going on. What *more* is there then available for me to know that I would need to add to get knowledge of *what switch it was*? There seems to be nothing available to appeal to here but the *varying occurrent understanding itself*. But that would leave *what it’s*

like to switch interpretations epistemically otiose. But it isn't. Best then to include phenomenal character in the varying ways of understanding comprising the switch. But if we include *these* ways of understanding, then it would be arbitrary to exclude others. It would be absurd to suppose an occurrent way of understanding is phenomenal only if it's involved in a "switch." And now we've gone all the way, from rejecting Reducibility to accepting Variation.

6 Phenomenal variation and thought content

Suppose, in response to my arguments, one says,

"I grant that there are *phenomenal* differences that occur as we think and understand that are not reducible to *sensory* differences. And I grant that some more-than-sensory differences in the ways we think and understand occurrently involve *conceptual capacities*. But I deny that the two classes of difference overlap. For the former (phenomenal) differences are not, even in part, differences in *what* one thinks or in *what* one understands an utterance to mean—they never constitute differences in *thought-content*."

This would be to try to open a gap between differences in how we think and understand, recognized as phenomenal, and any difference in the *content*, and by doing so, to open a gap between the phenomenal character of thought and thought itself. The suggestion is: while *something* more than sensory is included in the phenomenal, my inclusivism mistakenly identifies this with differences in ways of thinking and understanding essentially tied to conceptual activity, defined in terms of a grasp of semantics, of thought content.

But I think this proposed gap is not real. For I still do not find subjectively discernible differences available, identifiable with the phenomenal differences that occur as we think and understand, but distinguishable from differences in ways of thinking and understanding, where these are characterized in a manner essentially tied to conceptual capacities—and thus to differences in what is thought/understood (differences in content). *For I have no satisfactory way to identify these phenomenal differences without speaking in ways that imply there is something thought or understood.* When I think of what, specifically, it was like for me to understand the passage from Jefferson, I have no recourse but to speak of *what I understood it to mean*, repeating the very words I read, or others I take to mean the same. When I think of what it was like for me to realize what book my friend meant, I speak in those very terms—of what it was like for me to realize *what I understood her to mean*. When I conceive of the phenomenal character of my interpretive switch, I speak of what it was like for me to switch from understanding "hot" to mean *high in temperature* to *sexy*, i.e. in terms of a difference in *what I understood* the expression to mean. What I don't seem to have then, is a grasp of these specific phenomenal differences in terms that allow me to detach them entirely from genuine thought or understanding, or to attach them will-nilly to quite different ways of thinking/understanding.

Here one might object: even where there are no sensed or imaged differences in familiar natural language utterances to which I can reduce phenomenal differences in ways of understanding, still there are perhaps hidden, coinciding differences in *brain language* “utterances”—the tokens of the “language of thought,” which are typed purely syntactically, independently of any thought-content they might express or encode. And perhaps then I can think of the phenomenal differences in “ways of understanding” I’ve been discussing “contentlessly” as nothing but covert differences in brain writing.²⁰

But first, even granting this notion of correlated “brain writing,” I still don’t know how to think of *which* neural inscriptions I am allegedly discerning in given cases, if not (for example) as: *the brain writing that coincides with what it’s like for me to take “hot” to mean sexy*. So I actually still *haven’t* been given what is wanted here: a manner of identifying what it’s like for me to think the thoughts I do that separates them from having thought or understood something. Moreover, even if the imagined correlations were made, that would not show that I can conceive of someone for whom what it is like is just what it is like for me, *but who understands nothing by the words he hears or imagines*. And that is still what is called for, if we really want to justify segregating phenomenality from semantics.

Now I may perhaps conceive of a phenomenal duplicate whose head contains the same “writing” as mine. But that isn’t enough. What I would still need to do is to identify what it’s like for me to think and understand in a way that enables me to project the *very same phenomenal character* as I subjectively discern in my own case into the life of a hypothetical thoughtless subject whose head is filled with the same writing. That is, I would again be seeking the “semantically clueless phenomenal duplicate” of Section 5 above. I cannot find him, merely by *understanding the proposition* put to me that my phenomenal character consists in something such a thoughtless subject might have (e.g. the detection of meaningless brain writing), or by simply *stipulating* this to be so. That would do nothing to alter the fact that once I have made myself aware (as in Section 5) of the contrast between what it’s like for me to speak with and without understanding, I don’t see how what it is *actually* like for me to speak, and what it *would* be like for some *imagined mindless other* to speak, could always be *just the same*. For I lose my grip on the very phenomenal contrast I am being asked to conceptually export to another, and I am no longer clear that this is indeed “the *same* phenomenal character” I am exporting, when told that it must be drained of all difference in thought and understanding, leaving nothing but Jabberized chatter.

But what about the “externalist” idea that subjects must have the right connections to their environments to enjoy thought content? We may think that gives us what we need to radically divorce phenomenal character from genuine thought. However, this

²⁰ In discussion (and later in a draft of his paper for this volume he has kindly shown me), Joseph Levine has proposed something like this strategy for giving me the “content-less” grasp of what it’s like for me to think/understand that I say I lack here.

would be unwarranted. If I have difficulty conceiving of a literally thoughtless phenomenal duplicate of myself, then just adding that he is to lack my environmental embedding does not help matters. If I *assume* the truth of content externalism, my conceptual difficulty here might lead me to infer that having the right environment is essential for the phenomenal character of thought. But the radical separability of thought and phenomenal character certainly doesn't follow. And the thought experiments crucial to arguing for externalism do not warrant such a separation. *At most* they show only that some ways of identifying thoughts would distinguish two subjects content-wise who were phenomenally the same. They do *not* show that my phenomenal twin might conceivably be *utterly bereft* of thought, or *utterly dissimilar* to me in what he was thinking.

Take, for example, the standard (H₂O) externalist thought experiment Putnam made famous: a (chemically uninformed) subject on Earth is thinking that *there is water in a glass*, while another, similarly ignorant and phenomenally type-identical subject (on "Twin Earth") is thinking (not this, but) that *there is twater in a glass*. ('Twater'—we assume—names a Twin Earth substance superficially similar to, but microstructurally quite different from, Earth's H₂O.) Even if we accept the attribution of distinct thought contents, this does not show that one of the subjects might be devoid of thought. Furthermore, attribution of distinct thought contents leaves it open that two such subjects may still be significantly cognitively similar. We might say that the conception one has of *water* is no different from that the other has of *twater*. And having what they share would be enough to make someone a thinker of thoughts. So acceptance of this Twin Earth-style scenario does not show that phenomenal character is insufficient for some kind of thought content. Similar considerations hold, I believe, for the kind of social externalism defended by Tyler Burge.

These points could also be developed so as to allow that where there is phenomenal invariance through a difference in "singular thoughts," still there is genuine phenomenal *thought*. The phenomenal character of my thought, actual and dispositional, cannot on its own make it true that I am thinking of the one and only *Brasilia* (whereas my twin is thinking of the equally charming *Twin Brasilia*). Still, it is enough to make it true that I am in some way *thinking*, and enough to *contribute* significantly to making my thought a thought of *Brasilia*. And the essential contribution of the (extra-phenomenal) *context* to fixing the reference and truth conditions of my thoughts still leaves the role of phenomenal character in meaning dramatically different from the role that belongs to those merely *formal* characteristics of a symbol or image that constrain its *interpretation* in a way compatible with radical *re*interpretation, or with the total absence of meaning. The way the phenomenal character of thought constrains what it can be about is not like *that*. Further, I see no reason not to speak of the part of what is thought that is fully determined by phenomenal character as a kind of thought *content*, even if the truth of attributions of content (e.g. "I think that *Brasilia* is the capital of Brazil") are not determined by phenomenal character alone, independently of extra-phenomenal context.

None of this runs afoul of what I said earlier about phenomenal character and self-knowledge. What it's like for me to realize what book my friend was speaking of when she said "the book" may not by itself make it the case that I was then thinking of a book by (the one and only) David Foster Wallace. But it can still play a crucial part in my knowing what I took her to mean. And it's not an objection if we say something extra-phenomenal is needed to yield the full content of the thought—my being in the right *context*, living on the David Foster Wallace planet, not his twin's. For it's not the case that I would then need to have knowledge of extra-phenomenal facts that would turn out to render my knowledge of what it's like for me epistemically otiose. To know that I'm thinking of David Foster Wallace on a given occasion, I just need to know what it's like for me then to think as I do, and actually to be in that context in which the expression I am prone to give of my thoughts would be correctly interpreted to refer to David Foster Wallace.

Let me be quite explicit about this: inclusivists are not necessarily anti-externalists about content. We could consistently hold that having the "right" brain-environment links is strongly necessary (and not just contingently causally required) if a brain is to belong to a subject of genuine conceptual thought. We can even hold that there are differences in thought content due to environmental differences that do not supervene on phenomenal differences. We may, however, wish to speak of a kind of thought content that does supervene on phenomenal differences in any case: "phenomenal content." Nonetheless, acceptance of inclusivism in the sense and for the reasons discussed here is compatible with various forms of content externalism. So they should not be taken to refute it.

7 Conclusion

The disagreement over phenomenal thought stems from a number of global and complex differences rooted in varying historical legacies, different implicit "logical geographies," and the extensive obscurities and subtleties involved. So it is not realistic to expect a ready consensus. Certainly it would be naïve to suppose it could be achieved by simple appeal to introspection. But that needn't make us despair of critical examination, nor need it make us abandon first-person reflection altogether.

To justify a view of these matters, we need to take care in how we configure our framework for thinking about them. This is a complicated business. Ways of thinking can become entrenched because one is interpreting key terms and questions in ways that restrict what one can take seriously. And the need to sustain a certain conception of mind, or a relation to the history of relevant discussions, may set one's thoughts inexorably along certain paths. Part of what is necessary then is to make explicit an interpretation of crucial notions that will help us to formulate our questions carefully and without prejudice, and to raise awareness of intellectual history in a way that frees up thought, and makes us alive to alternatives. And all this needs to be done while minimizing controversial assumptions.

I claim to have made here just such a reasoned case for the view that if we shut conceptual thought out of consciousness, we diminish our experience significantly. While I have made use of first-person reflection, I have done so in the context of arguments. I have not resorted to mere assertion, nor made brute appeal to “intuitions of cognitive phenomenology.” The arguments are not to be blithely dismissed, and can be fairly assessed only with patient reflection on the sorts of cases discussed, as we encounter them in our lives, in the light of relevant distinctions.²¹

References

- Block, N. (2002). “Concepts of consciousness,” in David J. Chalmers (ed.), *Philosophy of Mind: Classical and Contemporary Readings*. Oxford: Oxford University Press.
- Dennett, D. (1991). *Consciousness Explained*. Boston: Little, Brown and Company.
- Georgalis, N. (2003). “The fiction of phenomenal intentionality,” *Consciousness and Emotion*, 4: 243–6.
- Goldman, A. (1993). “The psychology of folk psychology,” *Behavioural and Brain Sciences*, 16: 15–28.
- Horgan, T. and Tienson, J. (2002). “The intentionality of phenomenology and the phenomenology of intentionality,” in David J. Chalmers (ed.), *Philosophy of Mind: Classical and Contemporary Readings*. Oxford: Oxford University Press.
- Husserl, E. ([1900/01] 2001). *Logical Investigations*. J. Findlay (trans.). London: Routledge.
- Jackendoff, R. (1987). *Consciousness and the Computational Mind*. Cambridge, MA: MIT Press.
- James, W. (1890). *The Principles of Psychology*. Cambridge, MA: Harvard University Press.
- McDowell, J. (1994). *Mind and World*. Cambridge, MA: Harvard University Press.
- Moore, G.E. (1962). “Propositions,” in *Some Main Problems of Philosophy*. New York: Collier Books.
- Nichols, S. and Stich, S. (2003). *Mindreading: An Integrated Account of Pretence, Self-Awareness, and Understanding of Other Minds*. Oxford: Oxford University Press.
- Pitt, D. (2004). “The phenomenology of cognition or what is it like to think that p?,” *Philosophy and Phenomenological Research*, 69: 1–36.
- Place, U.T. (1956). “Is consciousness a brain process?” *British Journal of Psychology*, 47: 44–50.
- Prinz, J. (2002). *Furnishing the Mind*. Cambridge, MA: MIT Press.
- Putnam, H. (1981). “Brains in vats,” in *Reason, Truth and History*. Cambridge: Cambridge University Press.
- Robinson, W.S. (2005). “Thoughts without distinctive non-imagistic phenomenology,” *Philosophy and Phenomenological Research*, LXX: 534–61.
- Rorty, R. (1979). *Philosophy and the Mirror of Nature*. Princeton: Princeton University Press.

²¹ I would like to express my gratitude for feedback from audiences for presentations of this material at: The California Phenomenology Circle, The University of Arizona, The University of Fribourg, Rice University, and The 2008 Consciousness and Thought Conference in Dubrovnik. I would particularly like to thank for helpful feedback: Baruch Brody, David Chalmers, Frank Hoffman, Terry Horgan, Uriah Kriegel, Joseph Levine, Jesse Prinz, David Pitt, David W. Smith, Ron MacIntyre, Farid Masrou, Kevin Mulligan, Matjaž Portč, Martine Nida-Rümelin, Susanna Siegel, and two anonymous reviewers for Oxford University Press.

- Ryle, G. (1949). *The Concept of Mind*. Chicago: University of Chicago Press.
- Schwitzgebel, E. (2008). “The unreliability of naïve introspection,” *Philosophical Review*, 117/2: 245–73.
- Sellars, W. (1956). “Empiricism and the philosophy of mind,” *Minnesota Studies in the Philosophy of Science*, 1: 253–329.
- Siewert, C. (1998). *The Significance of Consciousness*. Princeton: Princeton University Press.
- Smart, J.C.C. (1959). “Sensations and brain processes,” *Philosophical Review*, 68: 141–56.
- Strawson, G. (1994). *Mental Reality*. Cambridge, MA: MIT Press.
- Tye, M. (1995). *Ten Problems of Consciousness*. Cambridge, MA: MIT Press.