

Precis of *The Significance of Consciousness*

Charles Siewert

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

csiewert@miami.edu

Copyright (c) Charles Siewert 2000

PSYCHE, 6(12), October 2000

<http://psyche.cs.monash.edu.au/v6/psyche-6-12-siewert.html>

KEYWORDS: Consciousness, First-Person Knowledge, Phenomenal Character, Blindsight, Intentionality, Conscious Thought, Higher Order Representation.

PRECIS OF: Siewert, C. P. (1999). *The Significance of Consciousness*. Princeton University Press. ISBN: 0691027242. 392 pp. \$42.50 hbk.

ABSTRACT: The aims of this book are: to explain the notion of phenomenal consciousness in a non-metaphorical way that minimizes controversial assumptions; to characterize the relationship between the phenomenal character and intentionality of visual experience, visual imagery and non-imagistic thought; and to clarify the way in which conscious experience is intrinsically valuable to us. It argues for the legitimacy of a first-person approach to these issues--one which relies on a distinctively first-person warrant for judgments about one's own experience. Thought experiments are employed in which one is asked to conceive of having various forms of blindsight, so as to make consciousness intellectually conspicuous by its absence in such hypothetical scenarios. It is argued that theories of mind that would commit us to denying either the conceptual or the metaphysical possibility of these scenarios neglect the occurrence of consciousness in this phenomenal sense.

1. Introduction

My book is about the significance of consciousness in three senses. First: I try to clarify a significance (or sense) of the word 'consciousness'--one we can give to it, if not find in it: what is sometimes expressed by speaking of "phenomenal consciousness". Second, my book concerns the relationship between consciousness and the fact that experience "signifies" something, that it is "of" or "about" things, or "directed at" them -- that it has, in the philosopher's sense, "intentionality." Third, I discuss the *importance* consciousness has for us--what value it has for us, and why it merits serious theoretical attention.

In confining myself to these topics, I do not pretend to determine the proper way to characterize the relation between experience and its neurophysiological basis, and I say little about whether the character of the experience we have affects what else happens. And these issues--roughly, the struggles over physicalism and mental causation--have often set the agenda for the philosophy of mind. But I am reluctant to make these the central questions, not because I doubt that they are worth pursuing, rather: I believe that responsible views about these issues and about the focus and goals of "consciousness studies" generally require we to attend to these questions first: Just what do we mean by 'phenomenal consciousness'? How is this related to the capacity mind to be of or about the world? Why does consciousness matter to us? In any case, I do not think that the interest of my discussion rests entirely on its helping to answer (or showing unanswerable) the mind-body problem and its recalcitrant near relations in our textbook list of philosophical problems. For we just may find that the prospect of attaining some clarity regarding the trio of topics on which I do focus exerts its own fascination.

When I say that we first need to clarify what we mean by 'conscious' I can imagine the impatient groans this may provoke, and the weary complaint that philosophers' obsession with defining terms has for millennia only led down blind alleys and stymied the genuine intellectual progress that can be won only by a more direct "empirical" approach. But I would say that careful attention to the way we draw distinctions can have important consequences for how we conduct and assess empirical research, and that where doctrinal conflicts seem particularly entrenched (as they do here), we could do worse than to step back and try to articulate carefully what we mean by some of our key terms, so as to identify better the source of our disagreements.

I certainly do not mean to suggest that the study of consciousness must be prefaced by an account of conditions necessary and sufficient for a state to be phenomenally conscious, phrased entirely in general and undisputed terms, and without illustrative lists. My own clarification makes essential appeal to lists of examples, and to a contrast between cases in which certain types of conscious experience occur and those in which they do not. I concede that, as Ned Block (1995) suggests, where we explain what we mean by 'consciousness' we find ourselves (perhaps inevitably) reduced to an elaborate sort of "pointing" at the phenomenon. And of course it is important to point "in the right way." Whether or not I do this will have to be judged by looking at the specifics of what I have to say.

2. Defending a First-Person Approach

I propose to take a "first-person approach" to explaining what I mean by 'conscious experience.' That is: first, when I illustrate what I mean by 'conscious' with examples of episodes of consciousness, I ask you to consider such types of examples that you know occur in a distinctively first-person way. And second, when I contrast these examples with hypothetical cases in which certain kinds of conscious experience are missing, while other features remain in place, I ask you to conceive of being a subject so described--in that sense, I ask you to conceive of these cases from the first-person point of view.

When I say that you know you have certain kinds of experiences in a distinctively first-person way, I mean that the type of warrant you have for beliefs that you have them--beliefs expressible with the first-person singular pronoun (e.g., 'I feel an itch,' 'It looks to me as if this line is longer than that one')--is a type of warrant had (at least ordinarily) only for first-person beliefs attributing these kinds of experiences, and not for beliefs that other people have them. It is possible to recognize one's commitment to this without endorsing or appealing to any views about the respect in which these two types of warrant differ. For instance, one need not adopt an inner perception theory about the source of first-person knowledge of mind, and one need not suppose that judgments about one's own mind involve some peculiar invulnerability to doubt, error, or challenge. It is enough if you find yourself disposed to assert that you know that you have this or that type of experience, in circumstances where you either have not exhibited for observation, or simply have not yourself observed, the sort of thing that another person would ordinarily need to observe to know that same thing about you. Assuming that if one believes and knows that *p*, then one has warrant for believing that *p*, it follows that in these circumstances you are committed to recognizing you have a type of warrant for your belief about your own experience different in kind from that ordinarily had for the corresponding third-person belief (Sections 1.4-5, 1.7).

One might try objecting that in such cases one enjoys a sort of knowledge and belief without warrant. But consider the consequences. First, it would follow that the only warrant one has for beliefs about one's own experience came from access to the sort of evidence available to others. And then, if one knowingly deceived others about one's current experience, it seems one would have more warrant for believing the lie one told others than for believing what one knew to be true. Also, one would not have warrant for believing that one was deceiving others. These results are unacceptable. Similar considerations should make us resist Wittgensteinian notions that it is out of place to speak of self-knowledge here at all (Section 1.6).

But why bother with a first-person approach, and take on the preliminary burden of clarifying and defending it, instead of getting right down to business? And, if I do bother, why not give a positive account of it instead of coyly leaving so much open? In answer to the first question: I insist on a first-person approach because of a suspicion that one source of the consciousness neglect I decry in later chapters lies in a tendency to theorize about these matters from a third-person perspective. In answer to the second: I do not offer a positive account of this warrant because that is a very complex undertaking in its

own right, and part of what makes it complex is that, in my view, a satisfactory account of first-person knowledge rests on achieving a satisfactory prior understanding of phenomenal consciousness.

However, if I cannot say what first-person warrant consists in until I have relied on it to develop my view about consciousness, prior to this I can and should do more than draw attention (as I do in Chapter One) to our ordinary pre-theoretical commitment to the reality of a distinctive kind of warrant and knowledge for beliefs about one's own experience. I should also defend this commitment against sceptical attacks. This I do in Chapter Two, by considering empirical research of the sort referred to in the much-cited work of Nisbett and Wilson (1977) that is supposed to count against "introspective access to mental processes," as well as that presented by Alison Gopnik (1993), purporting to show that developmental evidence refutes commonsense belief in a special first-person knowledge of intentionality. I argue (Section 2.2) that neither of these in fact present any threat to the assumption that we enjoy a distinctive sort of warrant for our first-person beliefs about what we experience. Further, I argue (Section 2.4) that a reliance on this sort of self-knowledge is not discredited by eliminativist attacks on the mental generally (as in Churchland, P.M. 1981), since these implicitly rely on epistemological presuppositions we have no reason to endorse.

3. What is Meant by 'Phenomenal Consciousness'?

With these epistemological preliminaries out of the way, I begin in Chapter Three to implement my first-person approach to explaining what I mean by consciousness. I first offer several categories of examples of its occurrence, where its recognition is likely to raise the least trouble. I start with the phenomenon of silent speech, but include other forms of imagery, both aural and visual, together with these and other modalities of sense experience--all of which we can know with first-person warrant to share the feature of being episodes of phenomenally conscious experience. Now, while the first-person knowable aural imagery of silent speech is a remarkable and illustrative species of phenomenal consciousness, still, it is only one species, alongside other imagery and sense-experience. And by recognizing this, we keep ourselves from being misled by accounts that begin by taking consciousness to consist in some kind of stream of narrative or commentary, and the associated inclination to think of consciousness as concerned with (at least "virtual") "serial processing."

Now I rely not only on positive examples of the occurrence of consciousness, but negative examples of its absence. Recognition of states of belief that may obtain at such time as one has no experiences conscious in the sense illustrated can help us to be clear that the common feature--phenomenal consciousness--exemplified, is not just that of having representational or intentional content. But still one can construe the relevant commonality in ways that amount, I maintain, to a neglect of genuine phenomenal consciousness.

To see this, we need to consider four (at least conceivable) forms of "blindsight." Blindsight I take to be some form of discriminative response (as for example in verbal judgment) to visual stimuli, of which the subject nonetheless has no visual experience conscious in the sense earlier illustrated--the stimuli do not *look* any way to him or her. The reader's task is to try to conceive of being a subject with each of these forms of blindsight--to conceive of having certain deficits in visual consciousness, while the exercise of certain other capacities remains intact. If one can employ the sense of 'conscious' first explained by means of positive examples to make sense of these hypothetical cases, one can understand what I mean by phenomenal consciousness. The first type of case is one (I believe) that has been found actual in the celebrated case study of Lawrence Weiskrantz (1986), for which he coined the catchy oxymoron, 'blindsight.' This I label (Section 3.5) *prompted* blindsight, as the subject's discriminatory judgments appear only when prompted by forced choice methods (the subject is asked, "Was an X or an O just flashed?").

The second (as far as I know, non-actual) type of blindsight would involve discriminating the same types of figures as prompted blindsight--relatively simple colors and shapes, spatial orientations--but *spontaneously*, unprompted. Here thoughts about what types of stimuli are present just occur, without forced choices, generated by the stimuli, but still without conscious vision of them (Section 3.6).

In the third stage (Section 3.7), I ask the reader to conceive of spontaneous blindsight that matches, in the discriminatory prowess it affords, the conscious visual experience of someone with extremely poor, low acuity "legally blind" vision. (At its extreme, this would amount to nothing more than a visual experience of a blurry expanse of achromatic light--greyish blobs and the like.) Would a blindsight equivalent of this--what I call "*amblyopic* blindsight"--properly be described as seeing? We can, I think, understand a sense of 'visual' in which it is enough for a blindsighter's judgment to count as visual that it be caused by internal states that are typically the effect of light stimuli, of the sort that would make the judgment true. Whether we call this "seeing" or not seems merely a terminological matter.

But, using this sense of 'visual,' it seems we are free to conceive of a blindsighter who is herself inclined to classify her judgment as visual--even spontaneously--though still the stimuli that triggered her judgment do not look any way to her. This "*reflective* blindsight" (Section 3.9) marks the fourth stage of the thought experiment.

If you can perform the little conceptual exercise thus described, then you are ready for my explanation of the notion of phenomenal character (Section 3.10). This I take to be that first-person knowable respect in which experiences that share the feature of being phenomenally conscious yet differ in ways that only phenomenally conscious experiences can. Familiar examples would include: the difference between the way it feels to be nauseous and the way it feels to be in pain, and the manner in which these both differ from the way it feels to have an itch. Generally we might speak of these differences in phenomenal character as differences in the way it seems to one to have various types of experience. And we may speak of its seeming as it does to someone to have an

experience--e.g., its seeming to you as it does for it to look as if there's a green circle on your left--as a *phenomenal feature*.

4. Consciousness Neglect

Now part of the point of this way of explaining the sense of 'phenomenal consciousness' is that it provides us with a way of determining that a given theory of mind or consciousness does not recognize the reality of consciousness in this sense. This is of interest, I believe, because it is one mark of the depth of disagreement among philosophers and psychologists who profess to concern themselves with consciousness that the theories proposed by some leave others suspecting that real consciousness--a.k.a., the subjective character of experience, what it is like to have it--has been left out or excluded, and that something else, to which the term 'consciousness' is affixed, has been substituted. I confess that I myself have harboured such suspicions, and so have tried to find some way of articulating what I understand by 'consciousness' that would help clarify these concerns, while minimizing controversial metaphysical and epistemological assumptions, and avoiding overly vague, misleading or metaphorical formulations.

I do this by summing up my explanation of what I mean by 'phenomenal consciousness' as follows. By 'phenomenal consciousness' I mean "that feature we know with first-person warrant to be shared by episodes of silent speech, other imagery, and sense-experience--a feature whose occurrence one would deny, if one held any of the forms of blindsight earlier described to be strictly impossible." (p. 100) Thus if your theory of consciousness apparently leaves no reasonable alternative to holding that the hypothetical forms of blindsight in my thought experiment are really inconceivable, or in some distinct (but more-than-nomological sense) metaphysically impossible, then your theory fosters what I term "consciousness neglect." So, if your theory is: (1) fairly construed as committed to views about what we do or should mean by speaking of conscious experience, or about what constitutes the feature or property of being a conscious state; and (2) these views do not permit us to conceive of, or rule out the mere possibility of blindsighters such as I have described--then your theory does not recognize, but neglects, phenomenal consciousness.

Of course, having stated this standard, it remains to be seen whether anyone may be fairly convicted of consciousness neglect by its means. That question is not so quickly resolved. I devote Chapter Four to a long and fairly complex discussion of it. My conclusion is that consciousness-neglectful theoreticians of mind include such eminent figures as Daniel Dennett, David Rosenthal, and Michael Tye, and various purveyors of functionalism.

For ease of discussion I frame the issue of consciousness neglect in terms of a hypothetical blindsight subject, Belinda, defined by contrast with another subject, Connie, who in her left field has only rather poor and limited conscious visual experience of some stimuli (the visual experience of an extremely legally blind achromatope, we may suppose). Belinda then, we conceive of as enjoying a spontaneous, amblyopic,

reflective form of blindsight that matches in discriminatory fineness (or crudity) Connie's desperately poor left field vision. If phenomenal consciousness is real, I maintain, then we know with first-person warrant that we have experience conscious in a sense that allows us to preserve in thought the contrast between Belinda and Connie. If a philosophy of mind obliterates this contrast, so as to render Belinda's blindsight either inconceivable or metaphysically impossible, it leaves no room for consciousness.

In this summary I cannot reproduce all the dialectical turns that lead me to find consciousness neglect; I can only give the gist of the argument. Dennett's denial that there is a phenomenon of visual seeming "over and above that of judging in one way or another that something is the case" (1991, p. 364) seems on the face of it to rule out the possibility of spontaneous blindsight. If a spot of gray light on Belinda's left triggers, by visual pathways, a spontaneous judgment that there's one there, then on his view it seems, there cannot be anything else--the *look* of the light to her--that she could be missing. However, it appears that Dennett would allow for some forms of spontaneous blindsight, by maintaining that what constitutes the difference between such blindsight and conscious sight is just that the latter involves information about the stimulus that is relatively richer, less sparse. But this seems to rule out the possibility of Belinda's (amblyopic) blindsight. For since she would have as good (or weak) a spontaneous discriminatory capacity with regard to left field stimuli as her cousin Connie, the information content they have about their visual stimuli would be equally rich (or impoverished). So a difference in Belinda and Connie's visual experience becomes impossible, perhaps inconceivable (Section 4.2).

Similar difficulties attend Michael Tye's (1995) theory of consciousness. On his view, roughly, conscious visual experience consists in the possession of nonconceptual visual information content "poised" to affect one's beliefs, without an exercise of will. But it would seem that someone of Belinda's talents would be just as rightly regarded as a possessor of such information about her left field as Connie would be; so on Tye's view of visual consciousness Belinda could not possibly lack any phenomenal visual experience that Connie had. That is to say, Belinda's blindsight is metaphysically impossible (Section 4.6).

Now it may seem unfair for me to criticize David Rosenthal's (1986, 1993) "higher order thought" theory of what makes conscious states conscious on the grounds that it neglects phenomenal consciousness. For it may seem my "phenomenal consciousness" is not to be accounted for by "higher order thought" at all, but is treated under the rubric of "sensory quality"--about which Rosenthal has other things to say (1991). The problem is that what he has to say about this (for instance, that a sensory quality of pain is such as can be had even while one has no feeling of pain) leads one to think that Rosenthal's sensory quality is not my phenomenal character, and that nothing would prevent his visual sensory qualities from being had by a blindsighter.

So, unless Rosenthal takes care of phenomenal consciousness in the "higher order thought" part of his theory, he does not take account of it at all. And he does not seem to me to do so. For what, on that view, gives Connie a conscious visual experience of a

patch of light on her left where Belinda has none? Should we say: Connie has a higher order thought about her visual representation to the effect that she has it, while Belinda lacks this? But Belinda, being a "reflective" blindsighter, would have the higher order thought that she visually judges that there's a patch of light on her left. So what is the crucial consciousness-making thought that she would lack?

It seems we should not want to say that what makes Connie consciously sighted on her left, while Belinda is not, is just that Connie has the higher order judgment "I consciously see on my left." For one can not constitute a difference (e.g., between conscious vision and blindsight) purely out of the thought that there is one (Section 4.5). So, will we say that what gives Connie the visual experience Belinda lacks is the former's having a higher order thought that she possesses a (Rosenthalian) visual quality (i.e., a sort of visual quality that even a blindsighter might enjoy)? A problem here is that ordinary consciously sighted people (like myself) typically do not, consciously or unconsciously, think that they have such visual qualities, when things look some way to them. So it can not be reasonable to suppose this is what leaves open the possibility of Connie and Belinda (Section 4.4).

I conclude that neither Rosenthal, nor Tye, nor Dennett have the resources in their theories of "consciousness" to leave open this possibility. If one allows no consciousness but what one finds in their theories, one renders the contrast between Belinda and Connie inconceivable or otherwise impossible, and thereby leaves out phenomenal consciousness.

Functionalists have been accused of failing to take into account the "feel" of experience, or of "leaving out qualia," and my own discussion backs this reproach, at least regarding certain forms of this doctrine. Though my reason for thinking functionalists have been guilty of consciousness neglect bears some resemblance to the familiar "absent qualia" argument, the point is not quite the same. I do not claim that Belinda's blindsight involves an internal state that duplicates the functional role of Connie's blurry left field conscious vision. It may well be that there are things that necessarily only conscious visual experience, and no ersatz blindsight stand-in, could do. For instance, it seems plausible to me (Section 4.3) that learning to judge spontaneously, by visual stimulation, how to apply certain terms (such as color or shape terms) contributes to one's learning the meaning of those terms (acquiring the relevant concepts) only if that judgment is mediated by conscious visual experience. Only if there is some way red looks to you, does your learning to say 'red' in response to red stimuli contribute to your understanding what 'red' means. (I do not insist on the truth of this--I say only that it seems plausible to me.)

But even if this is an inevitable functional difference between conscious vision and blindsight, we cannot rightly say that all there is to the difference between Connie's blursight and Belinda's blindsight is that the former affords the capacity to acquire concepts through learning visual discriminations, while the latter does not (Section 4.8). For consider a Belinda-like subject, Linda, who grows up learning blindsighted verbal discriminatory responses to visual stimuli. Shall we say that all we mean by saying she differs from Glenda, who grows up learning her discriminatory responses through

extremely weak Connie-like conscious vision, is that Linda lacks, while Glenda exercises, the capacity to acquire concepts visually? Or shall we say that this alone constitutes the difference between Glenda's conscious vision and Linda's blindsight?

Either way, it seems we would be attributing a difference in capacities for understanding that cannot in any way be displayed--either to the subjects themselves or to others who observe them. For I assume that, if two subjects differ in their ability to understand a given term, evidence of this difference, or any distinguishing mark from which this difference between them followed, would have to be a manifest, not a hidden one. That is, it would need to be a difference such as either the subjects themselves or others could have warrant for attributing to them without the literal observation of anything transpiring hidden beneath the skin. (Hidden differences--e.g., brain lesions--bear on differences in understanding only insofar as they are relevantly related to manifest differences.) Now on the views under consideration the only relevant difference between Glenda's visual experience and Linda's lack of it would be nothing but that between the possession and the lack of an ability to acquire concepts visually. But, if that is all there is to conscious visual experience, then: there will be nothing *manifest* whereby they can display this difference in understanding. However, I assume we do not want to say there could be differences in two speakers' capacity to understand that can be manifest to no one, including themselves. So even if it is inevitable that the consciously sighted and the blindsighted differ in their concept acquiring capacities, it would not do to say that the experiential difference between them is nothing but a difference in how they can come to understand their terms.

Thus for me the crucial question to pose a functionalist theory of mind, regarding consciousness, is not the "absent qualia" question, viz.: Could something other than a conscious experience conceivably play precisely the same functional role? My--slightly more complicated-- question is this. Once we hold constant all the functional likenesses that we cannot reasonably deny could obtain between Belinda and Connie (or subjects very like them), would there be any essential functional differences remaining between Connie's conscious vision and Belinda's blindsight, such that we can justifiably say: *that is all we mean by saying*, or: *that is all there is to* this difference between Connie's having visual experience and Belinda's lacking it? My claim is that, upon consideration of plausible candidates, there is no *manifest* difference in functional role that will return a positive answer to this question. And so any theory according to which differences in mental concepts or mental properties are exhausted by differences in manifest functional roles of subjects' internal states, leaves us no reasonable alternative to denying the conceptual or metaphysical possibility of an experiential difference between Connie and Belinda--which is to say, that theory leads to a neglect of phenomenal consciousness.

The question then naturally arises: can one reasonably identify the difference between the occurrence and absence of conscious experience with that between the occurrence and absence of some hidden feature--whether this be thought of as a functional role property, or otherwise (as perhaps a physical or physiological feature whose identity is not exhausted by the role it plays vis-a-vis other things)? Some possible "hidden" feature functionalisms I think can be seen as no better off than the manifest kind: those that

would locate the crucial hidden functional difference near the periphery (at say the level of receptor or afferent neurons). However, if we locate the hidden difference more internally still, we are no longer--by my "Belinda" test--doomed to theorize ourselves into a neglect of phenomenal consciousness. But we do face several other problems to which I see no satisfactory solution. One is: by identifying consciousness as something buried internally, whose presence is knowable only through probing deep beneath the surface, it would seem we will lose any rationale for maintaining a specifically functionalist kind of physicalism about consciousness. For such deeply hidden "physical-functional" properties would seem to be hardly more amenable to realization in a physically diverse range of entities--scarcely more "multiply realizable"--than physical features that are not "functionalistically" construed. Second, without any analysis of our concept of phenomenal consciousness in manifest functional terms, it becomes very unclear just what is to guide our decision about the physical sine qua non of consciousness, in a way that applies not just to normal actual human subjects, but ranges over all "metaphysically" possible worlds. Third, and relatedly, it becomes difficult to see on what grounds we are to maintain a physicalism that rules out the possibility of David Chalmers' (1996) "zombie world"--a world with the same type distribution of fundamental physical particles in fields of force as our own, but in which, unlike our own, phenomenal consciousness does not occur. It will be harder to dismiss or argue away the logical possibility of such a world, once we have allowed the conceivability of Belinda-like blindsight.

However, my main concern here is not with these issues, but with the prevention of consciousness neglect. And even if no clear and robust form of physicalism is true or warrantable, or some version of property dualism is true, still: someone who would claim that consciousness is some hidden physical feature of the brain, or that consciousness logically supervenes on such features, does not count as a neglecter of consciousness in my book. Let all who will maintain that consciousness just is the very feature specified by some as-yet-undetermined predicate of neuroscience, to be discovered through stabbing and scanning the intricate recesses of our brains; let them claim that consciousness must occur in any possible world in which that hidden feature does (or that it must not occur in any world in which that feature does not). I want to say only: I understand these claims only if consciousness is also taken to be that *manifest* feature to which we can attend by recourse first, to a distinctive first-person knowledge of experience, and second, to our capacity to conceive of certain merely hypothetical situations--as in my Chapter Three thought experiments.

5. Can "Consciousness Neglect" be Defended?

In Chapter Five I try to anticipate and answer several lines of defense that might be offered against the criticisms in Chapter Four--arguments to the effect that the "neglect" I deplore is really no offence at all, since we have reason to think that there is nothing in reality answering to my phrase 'phenomenal consciousness,' to which we might pay attention. The arguments advanced on behalf of eliminative materialism supply one

possible source. Already in Chapter Two I argue that general eliminativist attacks on the use of ordinary ("folk") mentalistic notions depend on a conception of our warrant for applying such notions that is baseless. On this conception, their use is warranted only if it can be justified as part of a general theory of bodily movement from whose description they have been thoroughly expunged. However, once that epistemology is renounced, we no longer have reason to demand talk of experience and attitudes be "vindicated" via some intertheoretic reduction to neuroscience.

But it might seem that an eliminativist argument which is not subject to this criticism could be directed specifically against phenomenal consciousness, as I understand this, an argument made by Patricia Churchland (1988). Roughly: just as scientific advance showed that the range of things (the sun, fireflies, logs) which once upon a time would have been said to be "on fire" were in fact not properly classified in this way, so neuroscience will probably one day show that the paradigms by which I explain my use of 'conscious' (various forms of imagery and sense perception) possess no common feature that we know with first-person warrant that they have, which leaves open the possibility of Belinda-like blindsight.

On examination, the argument seems to rest on this idea: The lack of a unified explanation of apparently similar effects (light, heat) of the sun, fireflies and burning wood required us to give up applying the word 'fire' univocally to all; similarly we also will need to refuse to apply the term 'conscious experience' to the heterogeneous class of phenomena to which I apply it, once we see that no unified neuroscientific explanation is to be had of what these so-called "conscious experiences" typically do. But I reply: to the extent we can identify common effects conscious experiences typically have, we evidently lack reason to think that they will probably find no unified explanation in an ideally completed neuroscience. Furthermore, we have no reason to demand that they do so, once we stop saying that ordinary modes of mentalistic explanation must be reducible to neuroscientific ones (Section 5.2).

A second challenge to my notion of phenomenal consciousness is likely to come from another quarter. Here one focuses, not on the heterogeneity of conscious experience, but on the (metaphysical) possibility of such a one as my Belinda. Again, the argument rests on an analogy to some case of successful theorizing in natural science. It will be argued that, much as science shows us that it may seem,--but is not really--possible that there could be water that was not H₂O, so too science will show (or has shown) that it may seem but is not really possible that there could be a subject of Belinda's discriminatory talents, but who lacked Connie's blurry left-field conscious visual experience. This at any rate is what would seem to be behind appeals to "a posteriori necessities" likely to be invoked in this context (Section 5.3).

Merely stating such an analogy surely does not by itself give us reason to believe that Belinda is metaphysically impossible. Nor would it even be right to say that this shifts some burden of proof onto me to demonstrate that she is possible. For affirming the reality of phenomenal consciousness does not commit me to asserting the metaphysical possibility of Belinda. It only commits me to asserting the occurrence of a feature whose

instantiation one would deny who maintained the metaphysical impossibility of Belinda. And I can assert this while officially abstaining on questions of metaphysical possibility.

The situation then seems to be this. Apparently we can interpret the term 'conscious' so as to conceive of someone in Belinda's shoes. Now we could stipulate that the term not be interpreted this way, but rather in such a way (on some analogy with 'water') that welded the exercise of Belinda's discriminatory prowess to conscious visual experience with the same iron bonds of necessity with which H₂O is water. But this is not to say 'conscious' cannot also be interpreted in a way that leaves open the possibility of Belinda-style blindsight. For that matter, we could easily interpret the word 'water' so that it applied to the same this-world stuff as we take it to apply to, even while leaving open the possibility that, in some possible world, stuff that was not H₂O still could count as water, merely because it behaved on a macro-level, to ordinary appearance and for ordinary uses as water in our world does.

Admittedly, we cannot correctly say that, just as we might interpret 'water' to apply to whatever sufficiently appears like water, so we might interpret 'conscious' to apply to whatever sufficiently appears like conscious experience. For there is no way in which experience appears to us--as distinct from the way in which we judge, think, or believe it to be. But this only makes it harder to uphold the analogy with metaphysical necessities obtaining in connection with natural kinds like water or gold. For it is admitted, I presume, on this picture, that Belinda is in some sense conceivable. Once this is allowed, one who maintains that she is, nonetheless, strictly impossible needs to explain how we can nonetheless be oblivious to this when we perform the thought experiment. How is it that we do not see the impossibility of Belinda, and perhaps even are inclined to assert she is in a broad sense possible (though probably not within the constraints of human neurophysiology)?

We cannot explain this by saying that in conceiving of Belinda we mistake the contingent superficial appearance of consciousness for the thing itself--consciousness has no appearance in the relevant sense. One could try proposing that the problem is that in conceiving of Belinda, we mistakenly suppose her lack of a higher order thought about her vision entails her lack of something essential to phenomenal visual experience. But this would not work. For if that thought attributes a visual state of the sort that a blindsighter could have, I for one certainly do not suppose that thought is essential to conscious vision, since I do not think it's one that even I--consciously sighted though I am--typically have. And though I imagine Belinda not to think she consciously sees things on her left, I do not for a moment suppose (in fact I explicitly deny (Section 6.5)) that the capacity to engage in this kind of higher order reflection is essential to phenomenal visual experience. So it will not explain our obliviousness to the alleged impossibility of Belinda's blindsight to assert that we confuse a lack of higher order judgment about her visual consciousness with a lack of conscious vision itself. Thus it seems to me that the prospects are very dim for arguing for the impossibility of Belinda by appeal to some water/H₂O-like metaphysical necessities.

A third strategy to defend against my charges of neglect might try to attack the possible absences of experience by which I explain consciousness by questioning the *knowability* of such absences. The thought is: the lack of conscious visual experience in Belinda's blindsight would be in principle unknowable by anyone--including herself. And undetectable absences of experience should be deemed impossible. Such a pattern of thought appears to be at work in Dennett's attack on qualia (1988, 1991).

Applying this style of argument to the case at hand, one might say: "Someone with the talents ascribed to Belinda would not be in a position to know that she was blindsighted on her left--and neither would anyone else. And if no one could know that she lacked conscious visual experience on her left, we should deny this is a scenario we can really make sense of." But this argument faces two problems. First, evidently nothing warrants the assumption that the impossibility of some scenario involving consciousness follows from its unknowability. Second, I argue (Section 5.4) that Belinda would indeed be in a position to know that her denials of left field conscious visual experience were correct.

A more general skepticism about the knowability of others' conscious experience also can help foment a neglect of consciousness (even if it is not used explicitly to justify this). I offer two points in response (Section 5.5). First, a denial that one can know other conscious minds may be used to undermine claims to know one's own, only if it is assumed that first-person knowledge is dependent on knowledge of others. But, depending on how this sort of dependence is understood, it may furnish us with an argument, not for a universal skepticism about consciousness that would extend even to the first-person case, but against skepticism even about other's experience. Second, even if knowledge of one's own experience is independent of knowledge of others', plausible considerations appear to stall a move from this to a sound sceptical argument about other minds. Briefly: the ways in which I form beliefs about others' experience would provide very accurate results if used by others to form beliefs about me. Now do I have reason to think others actually differ from me in some manner that would make the ways others could form reliable beliefs about my experience by contrast thoroughly faulty, if directed by me on others? No. Do I then have some reason to think they *do not* differ from me in such manner? Yes. For: first, the picture formed of their experience is no less coherent, explanatorily rich, and conducive to successful anticipation of behavior than is the picture that would be similarly formed of my own. Also, I have reason to think that others have (literally) in their heads what in mine suffices to make conscious experience occur. Taken together this does not prove that it is impossible I am a lonely little island of consciousness in a vast zombie sea. But it appears to give me more warrant for asserting than for denying what I cannot in any case stop believing without going insane--that others also feel, see, and think consciously.

6. Phenomenal Consciousness Distinguished from the Mind's Self-Representation

Having done my part, by the end of Chapter Five, to defend phenomenal character from neglect, I devote most of Chapters Six through Eight to certain fundamentals regarding the relationship of the phenomenal character of experience to its intentionality. For this I need to furnish a conception of intentionality adequate to my purpose. I choose not to base my discussions on some notion of mental or intentional or informational "content" because I find the general notion of content is too unclear to furnish a good starting point.

So instead of relying on talk of "content" to elucidate intentionality I proceed on the basis of this idea: it is sufficient for a feature to be intentional (in the philosopher's sense) that it be a feature in virtue of which its possessor is assessable for truth or accuracy. (For instance: x is assessed for truth in virtue of an intentional feature if it is said that x believes falsely that tomatoes are poisonous, or that the way it looks to x is accurate, when it looks to x as if a is taller than b.) I distinguish between a feature merely *in respect of which* one is assessable for accuracy, and a feature *in virtue of which* one is assessable for truth or accuracy. I may be assessed for truth in respect of the feature: uttering the sentence, 'Oswald shot Kennedy'--someone may say that what I uttered is true, or is false. But I am so assessable for truth in respect of my utterance, only if it is supplied with an interpretation: from the fact (if it is one) that Oswald shot Kennedy, and the fact that I uttered 'Oswald shot Kennedy,' it follows that what I uttered is true, only if some interpreting condition is added, such as: by uttering this, I asserted that Oswald shot Kennedy. Intentional features are never only features in respect of which one is assessable for truth or accuracy (though some of them are that). They are (some of them) features the possession of which entails corresponding assessments for truth or accuracy, when conditions are added, none of which need include interpreting conditions (as in the utterance example)--but only conditions of truth or accuracy.

In Chapter Six I use this conception of what's sufficient to make a feature intentional to frame the issue of how phenomenal consciousness relates to what I call "mentally self-directed" or "self-reflexive" intentionality--the capacity of some minds (ours for instance) to represent themselves. This is one (important) dimension of the issue of how consciousness relates to self-consciousness. Here I continue the discussion begun when I distinguished phenomenal consciousness from higher order thought (in Chapters Three and Four). In Chapter Six, I argue not just for this distinction, but more broadly that no intentional feature whose conditions of truth or accuracy require that its possessor be in this or that state of mind, is either necessary or (more than trivially) sufficient for having a phenomenally conscious experience. So I aim to distinguish phenomenal consciousness clearly not just from thought about one's own states of mind, but from mental self-representation generally--including what has been called "inner scanning," "inner perception," and "inner sense." Phenomenal consciousness is not "the perception of what passes in one's mind."

I contend that we have been sometimes led to think otherwise by shifting too uncritically between talk of one's "conscious experience," and talk of experience one is "conscious of having." We are sometimes willing to characterize conscious mental states as mental states one is conscious (or aware) of being in. And this can appear to support the notion that a state's being conscious involves one's forming a representation of it, because one

assumes that to be conscious of one's experience one must somehow represent it to oneself, perhaps by having a "higher order thought" about it, perhaps by having, via "inner sense," a "higher order perception" of it. But I think we should be wary here. We should cautiously recognize some inclination to say that, generally, we are conscious of having our conscious experience. But then we should ask: under just what interpretation of that claim are we entitled to regard as true?

By saying we are conscious of our conscious experience we may mean no more than that typically we know that we have (phenomenally) conscious experience when we do. But notice that in this sense someone may be conscious of being in a mental state without its being phenomenally conscious. (For one may know that one is in mental states that are not phenomenally conscious.) Moreover, this interpretation leaves open the possibility that creatures other than us incapable of higher order representation can have conscious experiences that they are not conscious of having.

Alternatively, it seems one might interpret one's saying one is conscious of feeling a certain way, or thinking something, in a manner that entails nothing about higher order representation at all. One may interpret 'He is conscious of feeling pain' on analogy with 'This paint is the color of blood,' or 'It has the shape of a mushroom'--that is, as saying what determinate of the determinable 'consciousness' someone possesses. On this interpretation, the 'of' here is not the 'of' of intentionality.

The assumption that one's state of mind is conscious just in case one is conscious of having it feeds a tendency to identify an experience's being conscious with its being the object of some sort of higher order representation. But it's just not clear that this assumption has any truth or plausibility that stands in need of explanation, once it is interpreted as implying that some species of metacognition accompanies all conscious experience (Section 6.3).

Do we then have any warrant for believing that all conscious experiences are targeted by some kind of ongoing higher order representation? It will be admitted for instance, that we are often (maybe usually) in conscious states that we do not then consciously think that we are in. But it is proposed (by Rosenthal) that nonetheless these conscious states are always shadowed by unconscious thoughts about them. However, why should we believe this? We (most adult humans) do usually have the ability to give some report of our conscious experiences, and this implies the ability to think that we have them. But to explain our ability to report our experiences not concurrently consciously thought about, we do not need to suppose that we are actually (unconsciously) thinking about them continuously.

Further, a problem arises even for saying that a disposition to think one is in a given state is necessary to its being conscious. Greater conceptual resources are needed for forming thoughts or beliefs about one's seeing or hearing, than are required merely for seeing or hearing things. But then what about unreflective perceivers--such as sighted children so mentally disabled that they cannot form or wield mentalistic concepts? We would be committed to saying that they would have only a sophisticated form of blindsight--no

color or shape would ever look any way to them. The point is not that such children (and mute sighted animals of many species), who possess no capacity for metacognitive thoughts, clearly do have phenomenal vision. (Though that is certainly plausible). The point is this. If we understand blindsight in such a way as to allow us to conduct the thought experiments of Chapter Three, we can see their incapacity for higher order thoughts simply does not entail that they are blindsighters (Section 6.5).

A natural defense to make to this and similar criticisms is to "dumb down" the prerequisites for having thoughts about one's experience so much that we will be hard pressed to suppose that sighted children and animals might fail to meet them. But the problem then is to specify what this cognitively or conceptually less demanding sort of higher order representation does involve in a way that does not make us say that animals unable to reflect on their own minds must be blindsighters, while also requiring more than mere responsiveness to one's own experience (for surely *that* need not involve thinking anything true or false of it).

Something like this last problem also troubles theories that make some sort of inner perception--as distinct from (conceptual) thought--essential to or constitutive of consciousness. Just what is this less demanding, quasi-sensory, non-conceptual form of mentally self-reflexive intentionality? And why should we suppose that it is trained on all our conscious experience? To meet this challenge, it is not enough to claim--however plausibly--that even without deploying full blown mental concepts, a creature may attend to its own experience. For such attention still is more occasional than conscious experience, and in any case does not evidently require that one represent (in thought or quasi-sensorily) one's own experience--truly or falsely, accurately or inaccurately.

Now we can form a distinction between ("outer") sensory appearance or experience--looking, smelling, tasting, etc.--and the sensory judgments sometimes based upon them. But, I argue (Section 6.6), we find no similar sort of warrant for the notion of an inner sense as distinct from higher order judgment. Consider: I can make sense of visible shapes or figures looking as if they are other than I judge them to be. However, I cannot analogously make sense of my visual experience of them appearing to me to be other than I judge it to be. (Though I can make sense of my visual experience *being* other than I judge it to be). And, though the blindsight thought experiment gives us a way of conceiving of the absence of visual appearance, while a form of visual judgment remains, no analogous thought experiment helps us to make inner sense conspicuous by its absence. On the contrary, this effort only feeds my doubts. For when I try to make this conceptual manoeuvre, I find the absence of inner sense indistinguishable from its alleged presence! Also: we can indeed contrast changes in the sensory appearance of "outer" objects (through changes in perspective, for instance) while we judge the shape, size, location, etc. of things seen, heard, etc. to remain constant. And we can also conceive of changes in how we judge them to be while their appearance is constant. But I cannot find any analogous contrast between changes in the "appearance" of my experience, while I judge its phenomenal character to hold constant. Nor can I see how to judge the phenomenal character to change, while the experience presents to me an

unvarying appearance. I conclude I have no warrant for attributing inner sense to myself. And I have not found a satisfying understanding of what it is supposed to be.

7. Visual Consciousness and Intentionality

So I argue that we need to distinguish phenomenal consciousness clearly from higher order intentionality, partly because consciousness should not be considered the exclusive province of those with the intellectual wherewithal to reflect on their own minds. But I argue that it would also be a mistake to see phenomenal character as irrelevant to differences in level and style of cognitive sophistication or intelligence. More specifically, I argue, in Chapter Seven, that the phenomenal character of visual experience is not a matter of mere sensation or "raw feel," requiring (like pictures or symbols) something on the order of an interpretation for it to be "of" or "about" anything, for it to have intentionality: intentionality (thus intelligence) is already built into it.

I make this case not by arguing that differences in phenomenal character determine differences in intentional content. There may be ways of distinguishing content "externally" such that experiences with different content have the same phenomenal character. Should we say that differences in phenomenal character at least determine differences in "internal" or "narrow" content? I think that, to argue for the inseparability of intentionality from phenomenal character, it is not necessary or desirable to try to justify this sort of distinction between different types of content (internal/external, broad/narrow)--a problematic effort that is anyway vexed by the unclarity of the general notion of content. So I prefer to focus the discussion on the issue of whether, once the phenomenal character of one's visual experience is fixed, anything other than conditions of accuracy need obtain, if assessments of one for accuracy with respect to those phenomenal features are to follow. Does the phenomenal character of visual experience need to have added to it something it might lack, before that experience can be assessable as accurate or inaccurate of anything--does it need interpreting conditions to have intentionality?

I argue (Sections 7.2-6) that we have more reason to say phenomenal experience does not, than that it does, in this way, "stand in need of interpretation." If visual phenomenal consciousness needs to have intentionality bestowed on it somehow, then there are conditions that: (a) might fail to obtain though one's experience had that character, which also (b) must obtain if assessments of one for accuracy with respect to one's having such experience are to follow, but which (c) are distinct from the mere conditions of accuracy of that experience. I find no good reasons for thinking that something about the nature of phenomenal consciousness or intentionality generally shows that some conditions need to be added to visual consciousness to make visual experience intentional. (For sense perception does not involve the appearance of sense-data that need interpreting, and the phenomenal character of sense-experience is not too atomistic to accommodate whatever holism is essential to intentionality.) Further, a broad search for plausible candidate conditions turns up none which we have more warrant for believing do satisfy all of (a)

(b) and (c) than not. (For we lack warrant for thinking that special links of experience to judgment, or intention, or to behavior or the environment are both necessary to sensory intentionality and essentially separable from the phenomenal character of sensory experience.) If this is what we find on examination, we are more warranted in believing such "interpreting" conditions are not needed than that they are. So phenomenal consciousness should not be pushed to the margins of one's picture of the mind, by dissociating it from intentionality and hence from intelligence.

We should also resist another way of belittling consciousness: the claim that it extends much less widely than one might have thought--that consciousness is a sparse, spotty, gappy, or occasional affair. Some remarks from Julian Jaynes (1976) seem to me to illustrate an extreme form of this. Jaynes says, for instance, that in reading we are typically conscious, not of the words on the page, but only of their meaning. We can recognize how far off track he has gone, when we realize that he would apparently have us believe that typically we read only by blindsight. I conjecture that Jaynes is drawn to such absurd extremes by a failure to recognize that objects of visual consciousness--what one sees that one can rightly be said to be conscious of--are very often only vaguely identifiable; there is no saying with much precision exactly what part of an array of visual stimuli one was conscious of when. This is true not only in normal skillful reading, but in a broad range of activities, which one can, in a sense, say are often done "unconsciously"--i.e., without conscious deliberation about how to do them--such as driving and typing, and the use of tools or instruments generally.

I believe it is true that, as we are engaged in much ordinary competent goal-directed activity, there are not many items we can rightly say we were visually conscious of, and we can identify but vaguely the spatio-temporal boundaries of those we can list. But we ought not confuse this admission with the claim that often during such times we are visually conscious of nothing at all, or of scarcely anything, or very little. For we still see consciously, even if we cannot identify with precision which objects we see. And, there is a way in which we can consciously see a lot--a way in which visual experience can be relatively rich--that is not proportional to the number of items that can be named in an inventory of which things one was conscious of. We often consciously see and can visually differentiate many levels of complexity in spatial patterns and configurations (think of the figure in a carpet, the branches in a leafless tree) even if we cannot name commensurately great numbers of distinct entities we were then conscious of, from among the stimulus array. We should not make ourselves oblivious to this kind of wealth in our visual experience of space, on account of our inability to describe verbally the differences in how we see things to be spatially configured and situated, or on account of our acknowledgment of all the countless aspects and differences in visual stimuli that escape our experience of them during a given time (Section 7.8).

Finally, visual recognition (of individuals and types) presents us with a rich realm of differences in phenomenal experience, which go beyond those characterizing our consciously seeing merely where things are--their particular shape, position, movement and contour--and differences in lighting and color (Section 7.9). Recognition of something as a sunflower or a glove, rather than merely as something filling a particular

space, involves differences in the phenomenal look of things--though this kind of difference is most evident perhaps when one can voluntarily alter the type of pattern one sees--as when one goes from seeing a figure as of one type (an M) to seeing it as of another type (a sigma)--or when one suddenly recognizes a person or place that first seemed unfamiliar. This is not a phenomenally invariant accompanying "feeling of familiarity"--but something integral to the many differences in phenomenal character involved when something looks recognizable to one as of this or that type, as this or that individual. Though this kind of experience is common and diverse, its very ubiquity makes it intellectually elusive; the kind of contrasts that bring phenomenal experiences of recognition to the attention of first-person thought often occur only when subnormal conditions of viewing (odd orientations, occlusions) obtain, or lapses of memory make recognition more gradual or more effortful than usual. Still I maintain that the differences in the way it seems to recognize things visually are found in the exercise of many kinds of intelligence--both those that distinguish our cognitive capacities from those of other species, and those that distinguish people who have special talents (technical, artistic) from others who lack these. Visual phenomenal (or "qualitative") character emphatically does not have only to do with "brute," "dumb," or "raw" sensation--on the contrary, it seems that, if there are ways of seeming so devoid of intelligence as to deserve such appellations, they are rather rare, atypical, or abnormal.

8. Conscious Thought

In Chapter 8 I explore one more area where the role of consciousness has been often underestimated. It has often been observed that there is no qualitative aspect to belief, and that consciousness is involved in linguistic understanding, if at all, only as some dispensable "feeling of understanding," or in some extraneous imagery accompanying the use of words. But I argue that it is a mistake to confine phenomenal consciousness to sensory experience and imagery. And it is not necessary to assume that differences in belief content have distinctive "qualitative feels" to show that there are differences in phenomenal character inseparable from our capacity to engage in non-imagistic, conceptual thought or judgment.

But prior to making this point, it is important to get some clarity about just how phenomenally conscious thought that *is* imagistic--visualization of shapes and figures, for instance--relates to intentionality (Section 8.2). I argue that it does make sense to speak of one's visualizing some shape accurately or inaccurately, to some extent or other, independently of assuming that one *intends* to visualize that shape. And we also do not have to assume that one has the *concept* of that type of shape, or of shape in general (where possession of such concepts involves inferential capacities that may be absent in subjects who nonetheless can perceptually recognize and discriminate shapes of those types). Nor is it necessary that one's imagery be secured, via historical causal links, to particular objects, for it to be assessable as a more or less accurate visualization of objects with their type of shape. Once that is clear, we can see that the phenomenal character of visualization experience--how it seems to us to visualize--can determine that we visualize

something as recognizably of this or that shape. Much as it may look to me as if a certain figure is recognizable as either M or sigma shaped, because of the phenomenal character of my visual experience of it, so too I may visualize a figure as recognizably either sigma or M shaped on account of the way it seems to me to visualize a figure. Gestalt phenomena occur in visual imagery as well as in vision. Thus, whatever similarities may exist between visual imagery and pictorial representation, differences in the phenomenal character of visual imagery cannot rightly be likened to differences in the non-representational features of pictures, and so they do not, like lines on paper or blobs of color, need interpretation in order to be accurate or inaccurate of something. And to the extent that the phenomenal character of visual imagery is detachable from conceptual, inferential abilities, intentions to visualize, and causal links, this does not mean it is detachable from its assessability for accuracy, its intentionality. So I conclude that nothing in the way of interpreting conditions needs adding to the phenomenal character of one's visualization experience, before one can become assessable for accuracy with respect to it.

But if this shows that thinking, in the sense in which thinking of something consists in forming an image of it, already has intentionality built into it, still it leaves open the question of what consciousness has to do with the intentionality specific to non-imagistic thinking. I argue, in several stages (Sections 8.3-5), for a closer connection between phenomenology and conceptual thought than is sometimes recognized.

The first step is to notice that there is a difference between the way it seems to you to hear what you or others say, when understanding or following what is being said, and the way it seems to you to hear this uncomprehendingly--and this difference need not be accompanied by any difference in the phenomenal character of one's imagery or sensory experience. Thus there are differences in the phenomenal character of experience--the experience of understanding, hence of consciously thinking--that transcend differences in the way it seems to visualize or sense. The same point can be made by drawing attention to wordless imageless thought--sudden realizations, or turns of thought, that occur unverbally.

The next step is to recognize that these sorts of differences are not phenomenally all the same. These more-than-sensory differences in phenomenal character are many and continuously various, as the sensory ones are. This is illustrated both by what I call the phenomenon of "interpretive switch" as well as by changes in the direction of wordless imageless thought. But if non-imagistic (or "non-iconic") thought is phenomenally conscious in this way, it remains to be clarified just how intimately these differences in the way it seems to us to think are related to the intentionality of thought--to the fact that one is assessable for truth in virtue of what one thinks.

I characterize the relationship as follows. "If during some period of time t , it seems to x as it seems to me on some occasion to think non-iconically that p , then there is some q , such that during t , x thinks that q ." (p. 284) To put this another way: each of us non-iconic thinkers, to whom it seems some way to think, are such that, if it seems some way to someone else the way it thus seems to us (if someone is our counterpart in respect of

"noetic phenomenal features"), then that entity is also a non-iconic thinker. Moreover, for any such noetic phenomenal counterpart *x*, there are conditions such that when they obtain, it follows that what *x* thinks is true--and since these conditions do not constitute an interpretation of *x* or *x*'s features, *x*'s phenomenal noetic features are not merely ones *in respect of which*, but ones *in virtue of which* *x* is assessable for truth--therefore they are intentional features.

I describe things in this manner, instead of more straightforwardly (say, by claiming that if it seems to *x* as it does to me to think that *p*, *x* thinks that *p*) for a couple of reasons. First, I want to make it clear that when it seems to me as it does to think that *p*, there can be others ways of characterizing what I (or my counterpart) think, other than as a thought that *p*, that have just as good title to be reports of what I (or he or she) then thinks. Moreover, at least if we allow that thoughts can be "externally individuated" there will be some values of *p*, such that, when it seems to my counterpart the way it seems to me to think that *p*, it will not be the case that he or she thinks that *p* (as for instance, when I think, of the pen I am holding, that this is out of ink).

I argue for this conception partly by arguing that the sorts of reasons that may lead one to think that noetic phenomenality could not guarantee non-iconic thought do not withstand examination. To those for whom the Wittgensteinian dictum that understanding is not an inner process has the ring of truth, we may grant that there is no essentially isolable experience the speaker or hearer has that guarantees some understanding of what is said. For there is no reason to suppose that its seeming to us the way it does for us to think is an essentially isolable experience. We have no reason to think phenomenal experience is too atomistic to exhibit whatever holism is essential to intentionality. Furthermore, we must be careful to distinguish the idea that it may seem to one as it does to think and understand, while one has only a very flawed or partial understanding of what one is saying, from the idea that it may seem this way to one, while one understands or thinks nothing whatsoever. Also, we must beware of the misleading claim that no "feeling of understanding" can guarantee understanding--for what is at issue here is not a "feeling" at all, in the sense where this suggests either sensation or emotion, and it is not some phenomenal character allegedly common to all cases of understanding.

Now there are "extraphenomenal" conditions (causal links to the natural substances in one's environment, membership in a certain linguistic community) that certain influential arguments (associated with Saul Kripke, Hillary Putnam, and Tyler Burge) tell us are involved in the identity of thought and belief. Is this an indication, either that one's phenomenal twins might be literally devoid of thought, or else that something on the order of interpretation is necessary before what they think is determinate enough to make them assessable for truth? I argue for a negative answer to this question. Let us allow that one has no genuine water thoughts where a causal link to an H₂O infested environment is lacking, and one is surrounded merely by "watery" XYZ. And let us allow that one thinks no bona fide arthritis thoughts when one is no part of a language community whose arthritis experts assert that arthritis can afflict only the joints: one thinks at most "tharthritis" thoughts. Finally, let us allow that this link and this membership cannot be guaranteed by one's phenomenal features. Still, it does not follow that my phenomenal

twin, who lacks the relevant causal links, might think nothing whatsoever when he says "Do not waste water," and "Someone's put water in the whisky bottle," etc. And it does not follow that my phenomenal homologue may, failing membership in an appropriate linguistic community, just express no thought at all, when he says "My dog's got terrible arthritis," "These arthritis pills do not work worth a darn," etc. Nor should we assume that, if extraphenomenal conditions make me and my twin thinkers of distinct thoughts, these conditions have the status of *interpretations* added to phenomenal experience, without which one would not be assessable for truth in respect of one's noetic phenomenal features. We can explain what makes me truth-assessable relative to my environment, while my twin is instead truth-assessable relative to his environment, without assuming that these environments furnish an interpretation to noetic phenomenal character, whose absence would leave what one thinks totally indeterminate.

I conclude that no separable interpreting conditions need adding to our phenomenal lives to make us the thinkers of non-iconic thoughts--so our phenomenality is as inseparable from the intentionality of our conceptual, non-imagistic thinking, as it is from the intentionality of our imagery and sense perception.

9. The Value of Consciousness

Even if in Chapters Seven and Eight I rescue phenomenal consciousness from a belittlement that rests on dissociating it from intelligence and intentionality, or on underestimating its range and diversity, a question about its importance still needs explicit treatment. Although I do make some remarks about the efficacy of consciousness and epiphenomenalism (Section 9.1), my main interest is in what value consciousness has for us, quite apart from what it is good for--quite apart that is, from what other good things it brings us, for the sake of which we value it. What I investigate primarily in Chapter Nine is the sense in which, and the extent to which, we value consciousness for its own sake.

That we often value the look, taste, smell, and feel of things to us, not just for some separable payoff they provide, but for themselves, may seem too clear to merit much examination. But I think it is not entirely obvious just what it means to value experience for its own sake, what shows that we do, and how great this sort of value is. It might seem that to value *x* for its own sake, means to value it, even when nothing else for the sake of which one values it is forthcoming. But it is far from clear that I would still think it better to have certain forms of conscious experience than not, even if everything else I normally assume comes with this, for the sake of which I want experience, is denied to me--say because I am locked forever in some kind of experience inducing virtual reality machine.

I take the question of the intrinsic value of consciousness to be that of whether we value having phenomenal features, but not only for the sake of what other non-phenomenal ones we think will come with them (Section 9.2). Instead of testing this by asking whether there is conscious experience you think it better to have, come what may, you

might, it seems, simply ask yourself, regarding some kind of phenomenal experience you value having--say, the visual experience of form and color, taste experience, sexual experience--do you value having it *merely* for the sake of the non-phenomenal benefits you believe they offer? I take it that if you answer this question sincerely in the negative, then, if you have no reason to think you're mistaken, you have reason to think you're correct, and thus you do value phenomenal experience for its own sake.

What might give you reason to think you're mistaken? Presumably if, when offered the chance to get the nonphenomenal goods either with or without consciousness, you find yourself either indifferent between the two options, or inclined to accept the no-consciousness plan, then it turns out you did not really value the conscious experience for its own sake after all. But what if, on the other hand, you find yourself still going for the "with consciousness" option, even if you thought you could get all the nonphenomenal goods you wanted without it? Then, it seems, not only have you not found evidence against your initial assertion that you value (certain forms of) consciousness intrinsically, you have found more reason to think that you do thus value it. Of course, we do not actually know any way to arrange for this option for anyone (how for example to give them blindsight or blind taste that would get them all the extraphenomenal benefits of conscious vision or gustation). So we can run this test only as an experiment in thought. Even so, it seems you can legitimately use it to bolster the conviction that you do value having certain kinds of conscious experience for its own sake. And while probably few will be unable to verify in this way the intrinsic value they accord their experiences of food, drink and sex, the preceding chapters' elucidation of the ways in which differences in sensory recognition and intelligence as well as in non-imagistic thought are phenomenally manifest, prepare us to see that this sort of value pertains to far more than uncultivated or merely sensory pleasures.

It may seem that in testing our evaluation of conscious experience by asking whether we would miss it even when offered the chance of a partial or total "zombified" existence, replete with same nonphenomenal goods, we must abandon the relative metaphysical austerity to which I earlier limited myself in explaining the notion of phenomenal consciousness and formulating my standard for consciousness neglect. But notice that I ask here merely how you would respond to your options *if you came to think it was possible* for you to maintain as good as normal nonphenomenal capacities, after a total excision of relevant talents for phenomenal experience. To answer this question in a manner that reveals a commitment to valuing consciousness for its own sake, you need not *actually* believe that it is possible, in any sense, for this "zombification" to happen to you, or to anyone else.

This way of framing the issue still leaves unresolved the question of *how much* intrinsic value we accord conscious experience. But with an appropriate variation in the thought experiment, I argue we can see that we intrinsically value consciousness this much: it is on account of our valuing the having of conscious experience for its own sake that we would regard a future in which we lack it altogether as little or no better than no future for ourselves at all (Section 9.3). I do focus most on versions of this value-clarifying thought experiment that are egocentric in character: I ask the reader to ask: "Would I

think it better that I got to keep phenomenal experience, even if I thought I could get all the nonphenomenal goodies for the sake of which I want it, but in the absence of the phenomenal?" However, similar exercises can also reveal the ways in which we each value others' possession of conscious experience for its own sake. Though it seems we can in this way show ourselves to value others' experience intrinsically so much that we would regard their zombification as little or no better than their death, only if we are willing to countenance the idea that others' total and permanent loss of consciousness would entail their death--that is, a complete extinction of the persons they are.

To all of this, one may say: "Fine, that shows we do as a matter of fact value our own and others' possession of phenomenal features for its own sake, but is not the more philosophical question that of whether we *should* so value this?" However, I think where intrinsic value is concerned there is little to do to argue someone into its acceptance from a position of non-acceptance or rejection. But what we can do, is to make clear to ourselves that we do already value consciousness for its own sake, and then to defend this value by considering that we have no reason to abandon it. Such reasons as I can imagine being offered depend only on some misconception about phenomenal consciousness (Section 9.4).

In closing I add that, in addition to whatever valuable work we may think consciousness does for us, and beyond whatever intrinsic value we find in it, phenomenal character offers this further reason to hold our reflective interest: it is a pervasive, real, and multifarious way in which human sensory and conceptual intelligence shows itself. To forfeit interest in it is then to relinquish philosophy's and psychology's longstanding concern to describe and explain human perception, imagery and linguistically expressed thought. If accounts of these do not explain their phenomenal manifestations for us, they are radically incomplete. And if they are incompatible with accurate phenomenological description, they are false (Section 9.5).

While the conception of phenomenal consciousness at which one will have arrived at the close of my book may make one better able to grasp the problems that have often dominated philosophical thinking about consciousness--questions about how consciousness relates to the brain, and how it is woven into the causal order of nature--I do not pretend my view makes these problems any easier to solve. But that does not threaten to justify its rejection. Maybe we *should* see the mind-body problem as peculiarly recalcitrant, or at least, as unanswerable to the sorts of solutions that some, in the name of scientific materialism, have claimed to give. However, if I give aid to the thesis that the goal of fitting consciousness into the explanatory framework of natural science will forever elude us, that is an aid I take no special joy in giving. But I would indeed be glad if my perspective on consciousness made us regard more seriously a range of issues about it that may, I believe, be examined while taking a break from the mind-body problem. Once we are clearer about what we mean by phenomenal consciousness and character, and how they relate to the intentionality of thought and perception, we are better placed for a new and deeper examination of a range of other important and interesting issues. We are, for example, better able to examine: the ways in which emotion and volition are manifested in conscious experience; distinctions between

conceptual and nonconceptual forms of representation, intentionality or intelligence; the nature of attention and its relation to intentionality (both first and higher order); the relative richness or poverty of experience; the unity of consciousness and its relation both to the unity of persons, and to the distinctiveness of first-person knowledge; and we can also better reflect on the relative value of different kinds of experience, and the bearing this has on the value of persons and their productions.

References

Block, N. (1995). On a confusion about a function of consciousness. *Behavioral and Brain Sciences*, 19, 227-247.

Chalmers, D. (1996). *The conscious mind: In search of a fundamental theory* Oxford: Oxford University Press.

Churchland, P.M. (1981). Eliminative materialism and the propositional attitudes. *Journal of Philosophy*, 78, 67-90.

Churchland, P.S. (1988) Reduction and the neurobiological basis of consciousness. In A. Marcel & E. Bisiach (Eds.), *Consciousness in contemporary science*. (pp. 273-304). Oxford: Oxford University Press.

Dennett, D.C. (1988). Quining qualia. In A. Marcel & E. Bisiach (Eds.), *Consciousness in contemporary science*. (pp. 42-77). Oxford: Oxford University Press.

Dennett, D.C. (1991). *Consciousness explained*. Boston: Little, Brown.

Gopnik, A. (1993). How we know our own minds: The illusion of first-person knowledge of intentionality. *Brain and Behavioral Sciences*, 16, 1-14.

Jaynes, J. (1976). *The origin of consciousness in the breakdown of the bicameral mind*. Boston: Houghton Mifflin.

Marcel, A. & Bisiach, E. (Eds.). (1988). *Consciousness in contemporary science*. Oxford: Oxford University Press.

Nisbett, R.E. & Wilson, T.D. (1977). Telling more than we know: Verbal reports on mental processes. *Psychological Review*, 86, 474-97.

Rosenthal, D. (1986/1991). Two concepts of consciousness. In D. Rosenthal (Ed.), *The nature of mind*. (pp. 462-477). Oxford: Oxford University Press.

Rosenthal, D. (1991). The independence of consciousness and sensory quality. In E. Villanueva (Ed.), *Philosophical issues: Vol.1. Consciousness*. (pp. 15-36). Atascadero, CA: Ridgeview.

Rosenthal, D. (1993). Thinking that one thinks. In M. Davies and G. W. Humphreys (Eds.), *Consciousness: Psychological and philosophical essays*. (pp. 197-223). Oxford: Basil Blackwell.

Tye, M. (1996). *Ten problems of consciousness*. Cambridge, MA: MIT Press.

Weiskrantz, L. (1986). *Blindsight: A case study and implications*. Oxford: Oxford University Press.