10 How Many Concepts of Consciousness?

I learned a great deal from the commentators about ways in which what I said was wrong or easily misunderstood. I am grateful for the opportunity to rethink and rephrase in response to the criticisms. A number of superb commentaries also had to be held over to a future issue, and I look forward to grappling with them as well.

R1. Is P-Conscious Content Just Highly Informational Content?

Armstrong accepts P-consciousness but says it is ''a species of representational content of a particularly detailed sort.'' Farah notes that perception in blindsight and other such cases is degraded and concludes (here and in other publications) that P-consciousness depends on quality of information representation. [See also Farah: ''Neuropsychological Inference with an Interactive Brain'' *BBS* 17(1) 1994.] Dennett says that my A/P distinction should really be seen as a continuum, a continuum of richness of content and degree of influence. He says I am ''inflating differences in degree into imaginary differences in kind.'' In Note 16 of the target article I suggested that some functionalists ''will see the distinction between A-consciousness and P-consciousness primarily as a difference in degree rather than in kind.'' But I was alluding to degree of access, not degree of informational content. I think the high degree of information views of P don't get to first base.

To see what is wrong with highly informational representational content as a substitute for P-content, consider the common types of blindness in which the (legally) blind person is able to distinguish a few degrees of light and dark, much as you or I can with our eyes closed. *This is P-conscious content that is relatively informationally poor, not informationally rich*. Furthermore, senses can differ in information richness without differing in phenomenality. Perhaps taste (not including smell via the mouth or mouth feel) is less rich than vision (Kapsalis 1987, p. 66). But is taste any less phenomenal than vision? Or consider orgasm again. Are we supposed to think that orgasm is informationally rich? Tye has made a valiant attempt to characterize (partially) the phenomenal content of orgasm in representational terms: something that is very pleasing (and

changes in intensity) is happening down there. Even if this does capture the phenomenal content of orgasm (which I don't believe for a second), this is not a very informationally rich content. Yet there can be no doubt that orgasm is ''phenomenologically impressive''!

Weiskrantz (1988) notes that his patient DB had better acuity in some areas of the blind field (in some circumstances) than in his sighted field. Suppose a blindsight patient with fairly good acuity in tile blind field were to become near-blind in the sighted field, able to distinguish only a few shades of light and dark. He experiences the light and dark but does not experience the blindsight. The blindsight is informationally richer but (presumably) not phenomenal, whereas the near-blind vision is phenomenal and informationally poorer.[1]

Dennett describes the informational content of blindsight as ''vanishingly'' small. In Dennett (1991), he emphasizes the cases in which the blindsight patient is given a forced-choice, for example, an X or an O. But blindsight patients can exhibit contents that have more informational value than that. In Pöppel et al. (1973), the first human blindsight study, the patients were asked to move their eyes in the direction of the stimuli that they apparently had no experience of seeing. The patients could do so even though they thought they were guessing. In addition, as I mentioned in the target article, blindsight patients can catch a ball thrown in the blind field and can shape their hands appropriately so as to grasp an object presented in the blind field. The information involved in these nonphenomenal activities is surely at least as great as the phenomenal discriminations of the blind people just mentioned or of some sighted people with their eyes closed. In addition, the implicit prosopagnosics mentioned in the target article have the capacity to recognize un-P-consciously the face of, say, John Wayne. I skipped over most of the empirical literature on this topic for lack of space, but let me just mention one phenomenon:

Semantic priming is a phenomenon in which the presentation of one stimulus facilitates the subject's response to a related stimulus. For example, if normal Americans are asked to press a button when a familiar face appears in a series of faces presented rapidly one after another, the subject tends to press the button sooner if a related name has been presented recently. For example, ''Reagan'' facilitates reactions to Bush's face. Likewise, one name primes another and one face primes a ''related'' name. Here is the result: in a few prosopagnosics who have been studied in detail and who exhibit some of the other indications of ''covert knowledge'' of faces, faces prime related names despite the prosopagnosics' insistence that they have no idea whose face it is. For example, Lady Di's face primes Prince Charles's name even though the subject insists that he does not know whose face it is. See also the phenomenon mentioned by Graham. The perceptual content that it is Lady Di's face is moderately informationally rich, but this is not a P-conscious content. So once again, we have moderate informational richness without P, contradicting the point of view of Armstrong, Farah, and Dennett.

Dennett constructs a thought experiment, a superblindsight patient who comes to be able to tell us—effortlessly—very detailed facts about the visual properties in his blind field. He can tell us that there is a bright orange Times Roman italic X on a blue-green background about 2 inches high with a smudge. This superblindsighter says that he knows these sorts of features of stimuli in his blind field, even though he is just guessing, and contrasts what is going on with the real visual experiences of his sighted field. Dennett rightly says that it is hard to swallow that "anybody who could gather that much information from a glance" might have no visual experience. And he adds, imagining another patient, that if all he can tell us about the sighted field is that he saw an X rather than an O, we would be baffled by his claim of P-conscious experience.

Dennett is on to something here, but he has misdiagnosed what it is. There is some plausibility in the idea of high-information representational content as an empirically *sufficient* condition of phenomenality in humans (though not a conceptually sufficient condition). But there is no plausibility at all to the idea that high information content is a *necessary* condition for phenomenality, as is shown by the example of orgasm and the discrimination of a few degrees of light and dark with one's eyes closed. I think that the reason for the plausibility of Dennett's examples is that they illustrate *depictive* or *pictorial* representations. And it is very tempting to believe that pictorial visual representations must be P-conscious. That explains Dennett's first example. In the second example, if we think of the person as having a pictorial representation of an X, it is hard to imagine how the person could see the X without seeing it as having some particular size, typeface, color, and so on (hard only because one naturally but wrongly thinks of images as photograph-like; see Block 1983).

However, even mental images can apparently fail to be P-conscious. Cooper and Shepard (1973) noted that when subjects practiced image rotation to the point of automaticity, they reported that they had no image experiences, yet the rotation data argued for the same images as before (I am indebted here to Baars 1994). Kosslyn (1980; 1994) asked subjects to zoom in on an imagined map to the point where they could only see a part of it. It would be interesting to see whether these subjects could make use of the information in the "invisible" parts of the map better than subjects who had not just had those parts in the mind's eye.

So far I have been criticizing the view that P-content = highly information content, the view of Armstrong and Farah, but not quite Dennett's view. Dennett also mentions access. But how is access supposed to figure? Is P-content content that is *both* highly informational *and* highly accessible? If so, the legally blind contents I mentioned are a counterexample since though they are highly accessible, they are low in information. Or perhaps the mode of combination is disjunction rather than conjunction: P content is content that is high in information *or* high in accessibility. But now the P-unconscious images are a counterexample in the other direction. Furthermore,

everyone has relatively high information but quiescent beliefs and desires that are not P-conscious.

R2. Does P Exist?

Rey says it is strange that I did not explicitly consider any of the accounts that have been offered of P-consciousness in computational terms. This was an explicit strategic decision. Not everything can be discussed in every paper. This one was about some distinctions and how missing them causes trouble. In other papers I have discussed some computational (functional), intentional, and cognitive theories of P-consciousness. Rey says that the assumption of P-consciousness as noncomputational and noncognitive impedes research. That depends on whether P-consciousness *is* noncomputational and noncognitive. If it is, then Rey's assumption that it isn't impedes research. Rey assumes that the right approach to consciousness is that of computational and cognitive psychology (computational in this context = functional; see Block 1994). But why does Rey ignore the neurophysiological approach? The sort of research program described, for example, in Crick (1994) does not accept Rey's assumption, yet it seems to be going somewhere. Rey admits that phenomenal consciousness deflated so as to be amenable to computational analysis of the sort he favors is ''phenomenologically unsatisfying,'' but he persists because he knows of no non-question-begging evidence for the postulation of P-consciousness as distinct from computational notions. I find Rey's stance baffling. Let's look at a specific example. Above, I discussed the idea that P-conscious content is simply highly informational content. I appealed to the evidence of our own experience: when you close your eyes, you can nonetheless distinguish a few degrees of light and dark via truly phenomenal experience, so phenomenal experience can be low on the informational scale. Does Rey propose that we should ignore the evidence of our own experience in cases like this? To do so would be ''phenomenologically unsatisfying'' in a way that carries some weight, namely, it ignores a source of evidence that we all have from our own experience.

Rey is worried about conundra involving epiphenomenalism, zombies and such, and to avoid them he wants to reduce phenomenal consciousness to the computational. If P can be functionally defined, epiphenomenalism is ruled out and zombies are impossible. But this is the wrong way to avoid the conundra. Russell once hypothesized that the world was created five minutes ago with all the evidence of an age of many billions of years. Some philosophers have wanted to avoid this possibility by defining the past in terms of its effect on the present. To say that there were dinosaurs 65 million years ago is to say there are dinosaur-signs now. But this is a foolish metaphysical over-reaction. Better to face the conundra head on as discussed below in section R13.

R3. Is A-Consciousness Consciousness at All?

Graham, Lloyd, Natsoulas, Revonsuo, and the Editorial Commentary question whether A-consciousness is consciousness at all. As Searle (1992) emphasizes, a zombie that is a functional duplicate of us but lacks any P-consciousness is not conscious at all. (This point is made forcefully by Tyler Burge 1997.) But it is a mistake to jump from the idea that a zombie is not conscious in any sense to the idea that A-consciousness is not a form of consciousness. A-consciousness can be a kind of consciousness even if it is in some way parasitic (as Burge, Lloyd, and Revonsuo rightly say) on a core notion of P-consciousness. (A parquet floor is a kind of floor even though it requires another floor beneath it.) A-consciousness can come and go in a background of a P-consciousness person (that is, a person who sometimes has P-conscious states). Suppose a drunk becomes ''unconscious.'' He may have P-conscious states both before and during his episode of unconsciousness; for example, while unconscious he may be seeing stars or having mental images of various sorts. I don't want to try to specify exactly the relation between being unconscious in this ordinary sense and the concepts of P and A, but roughly, I think we count the drunk as unconscious to the extent that he has no A-consciousness of the environment via P-conscious perceptions of it. The drunk is A-unconscious in a way the specification of which involves appeal to P.

We tend to deploy the concept of A-consciousness in describing unconscious phenomena, so it is not surprising that many of the most common uses come in at a somewhat theoretical level. Consider, for example, Freudian unconscious states. Suppose a person is tortured horribly in a cinnabar room (a particular shade of orange-red). The cinnabar color symbolizes the pain and is repressed. He remembers the torture vividly but denies remembering the color of the room. Nonetheless, the memory of the color comes out in slips, associations, dreams, and so on. For example, he dreams of horrid cinnabar things. When he is in a cinnabar room he shudders violently and comes up with an excuse to leave, but does not recognize why. There is nothing in Freudian theory or common sense that precludes repressed phenomenal color images of the room. In fact, we can imagine the patient realizing this himself after years of psychoanalysis. ''I had a cinnabar image all the time that I would not let myself acknowledge.'' Whether or not this actually occurs, it makes sense for there to be a blockage that keeps a phenomenal color image from being informationally promiscuous. So the sense in which repressed memories are unconscious is A-unconscious. The Freudian type of unconsciousness does not *require* P-unconsciousness, but it does require A-unconsciousness.

Similar points apply to neurological syndromes such as prosopagnosia, in which the patient is not ''conscious'' of whose face he is seeing, even though he reveals in a variety of experimental circumstances that the information is unconsciously represented.

The unconsciousness is A-unconsciousness. It is not the presence or absence of a feeling of familiarity that defines prosopagnosia but rather the patient lacking A-consciousness of the information about the identity of the person. As Young notes, the lack of a P-conscious feeling of familiarity is (normally) a consequence of the lack of A-consciousness, but is not a defining feature of the syndrome. This point is nicely illustrated by Young's theory of Capgras's delusion, a syndrome in which patients claim that people whom they know (usually friends and relations) have been replaced by aliens who look just like them. Young (1994c) provides evidence that what is going on is that the subject recognizes (say) his mother, but he gets no feeling of familiarity from the perception, so he supposes the person is not really his mother. (It may be that victims of Cotard's syndrome, in which patients think they have died likewise, lack this feeling of familiarity, but blame it on themselves instead of on the people who don't stimulate the feeling of familiarity.) Suppose Young's suggestion is right. Still, this lack of the feeling of familiarity does not make the patient a prosopagnosic. He recognizes his mother's face and the faces of others despite the lack of the feeling of familiarity. So lack of a P-conscious feeling is not at the heart of prosopagnosia. In sum, though I agree that P is the core notion, A is still a kind of consciousness.

R4. P without A; A without P

Many of the commentators agreed with at least some of my cases of P without A. Revonsuo proposes a new one: dreams. The trouble with dreams as cases of P without A is that dreams often involve substantial rationality. For example, Chomsky tells me he plans papers in his dreams, and there is a well-known phenomenon of lucid dreaming in which the dreamer, knowing he is dreaming, changes the course of the dream. Of course, many dreams are much less rational, on the surface at least, but I would be reluctant to suppose that these dreams are unconscious in a different sense of the term than less overtly sensible dreams. I expect that dreams are unconscious in the sense described in the discussion above of the unconscious drunk.

What about A without P. In the target article I said that such eases were conceptually possible, but I knew of no actual ones. If it is so much easier to find P without A than A without P, that is a striking empirical fact. Humphrey was the only commentator to make any remarks in favor of A without P. He argues that Helen was in some respects such a case, correcting wrong impressions about his unpublished work. But, as he notes, even if Helen is a case of no visual phenomenality, she had no shape recognition of any kind. So we need to know more about cases such as Helen to count her as a good case of A without P.

I have long been troubled by cases from the ''imageless thought'' controversy from the early part of the century (G. Humphrey 1963—not the same Humphrey). For example, pick up an object from your desk, put it down, and pick up another. If they dif-

fer substantially in weight, you may ''just know'' this. My experience here is that I have ''images'' of the weights of the two objects, but apparently not of the *relation* between them. The relation between them appears to be something I just know without any phenomenal experience of it. This is a tricky bit of introspection, just the sort of thing that got psychologists in trouble during this controversy. But it would be foolish to ignore it, since it can guide experimentation. I hope some clever experimentalist figures out how to get an experimental handle on it. Burge (1997) has some very interesting arguments for A without P. One of Burge's examples involves the familiar case of a solution to a problem popping into mind. We often know we've got the solution without actually expressing it in any internal phenomenal clothes such as words or pictures.

I should add that from the point of view of empirical model building, it is very important to distinguish between cases of A without P, like superblindsight, if it exists, and Burge's sort of case. Both count against Model 2 because they show that some highly sophisticated thought (that must be accomplished by the Executive System if there is one such system) is not P. But neither Burge's type of case nor the imageless thought case challenges Schacter's (1989) model, because that model allows for executive processes that are not P. However, superblindsight, if it exists, would challenge Schacter's model, because that model tells us that the only way for perceptual information to get to the Executive is to pass through the P-module. Of course, we already have good reason to reject Model 2. As pointed out in the target article, the intermediate steps in problem solving often fail to be P, and that alone counts against Model 2.

Let us now move to objections to the idea of P without A. Church, Kobes, and Revonsuo criticize my example of the unattended noise as a case of P without A. They say that people adjust the loudness of their speech in response to noise even when not attending to it, and that is a rational action that reflects A-consciousness of the noise. But this noise adjustment does not show A-consciousness in my sense of the term. There is no inferential promiscuity here. If the notion of A-consciousness were to be weakened in this direction (also suggested by Graham and in Flanagan's [1992] critique of my notion of A), the consequence would be to let in cases of A without P. If A-consciousness were watered down in this way, then blindsight and the other ''covert knowledge'' syndromes would be cases of A without P. Of course if you like A without P and hate P without A, you could adopt Graham's and Flanagan's suggestion, but your pleasure would be purely verbal.

Kobes criticizes one of my arguments for conceptually possible cases of P without A. In Note 7 of the target article I worried about an A-unconscious state that caused an A-conscious state with the same content. The first state is ex hypothesis not A-conscious, but it is in virtue of one's having *that* state that its content is inferentially promiscuous, so it seems that it does have to be A-conscious. I avoided the problem

by thinking of the ''in virtue of'' in the definition of A-consciousness (a state is A-conscious if in virtue of one's having the state, a representation of its content is inferentially promiscuous, etc.) as *directly* in virtue of. If state X has an inferentially promiscuous content, but only because it causes state Y, which inherits X's content, then X doesn't count as A-conscious. Kobes thinks this answer gets me in trouble with the perceptual states of the superblindsighter (who, you will recall, is conceptually possible but apparently nonexistent). He says that it is only in virtue of the effects of these perceptual states on the superblindsighter's thoughts that the perceptual states are inferentially promiscuous, and so on, so the perceptual states are *neither A-conscious nor P-conscious*, and the supposed example of A without P dissolves. But why suppose, as Kobes does, that the superblindsighter's perceptual states of ''seeing'' an X have their effects only via the causation of the *thought* that there is an X? If such a perceptual state could only have an effect by causing a specific thought, then it would not be informationally promiscuous and it would not be A-conscious. A genuinely A-conscious perceptual content would be freely available for use in thought.

Kobes criticizes my account of the Sperling (1960) experiment, saying that before the icon fades, the subject is both P- and A-conscious of them all jointly, and after the icon fades the subject is neither P- nor A-conscious of them all jointly, so there is no divergence of P and A. Let me try to state my point better. Consider the distinction between *jointly* and *severally*. A pair of swimmers can be poised to win jointly if, say, they are on the same team in a race of teams. But in a race in which only one swimmer can win, a pair of swimmers would be poised to win severally, not jointly, that is each is poised to win, even though they can't both win. In the Sperling experiment there is never a time in which the letters are all poised to win jointly (become inferentially promiscuous, etc.) because, as with the individual swimmers, they cannot all win. But they can all jointly be P-conscious, or at any rate that's my claim.

Zalla and Palma argue that the tip of the tongue phenomenon (and ''Feeling of Knowing'' states generally) is a case of P without A. You have the feeling of knowing someone's name, but you can't access it. But on the surface, at least, the content that the name is known is *both* P and A, and the specific name content (e.g., ''Blanche'') that leads to knowledge of or correct guesses about features of the name (rhymes with ''ranch'') is *neither* P nor A. I agree with Zalla & Palma's statement that Feeling of Knowing states involve consciousness of ''the *existence* of a content but not *of* the content itself.'' But as far as I can see, this characterizes both P and A, so there is no gap here between P and A. Zalla & Palma argue (if I understand them rightly) that when I have the feeling that I know her name, typically the name itself is P, but not A. ''Blanche'' is a P-content but not an A-content. But I don't see why the facts that they mention about frontal lobe patients are supposed to support this idea. They point out that I have not got a wealth of data to support the idea of P without A. They are right, but I am hopeful for the future.

Baars thinks that P = A, but what is his evidence? In his reply he mentions a few cases in which both P and A are present and other cases in which both P and A are absent. I mentioned in the target article that this sort of evidence is suggestive of a rough correlation, but it can hardly show that P = A. He says that ''implicitly ... we all treat'' P and A as ''empirically inseparable.'' True, but implicitly we treat the earth as stationary as well. Baars studies P by studying A, so if P ≠ A, his research program would have to be reevaluated. Even if P and A are perfectly correlated within some conditions, but not identical, his research strategy would be questionable. One would want to know *why* P and A are correlated. Weight and mass are correlated at the surface of the earth, but studying weight is not studying mass. He says, ''GW theory shows that the equivalence ... is very productive indeed.'' I question that; we know lots about A, almost nothing about P. I think we are more likely to make progress by looking very carefully at cases where P and A seem to diverge. If they do diverge, that's where the interesting results are. Baars just ignores the cases that cause problems for his view instead of subjecting them to scrutiny. Church, Kobes, Revonsuo, and Chalmers (1997) agree with Baars that P and A are correlated, but unlike Baars they take the responsibility to confront the putative exceptions. Baars should consider what Humphrey says about A without P. Baars brings in evolution to argue for P = A, saying that we are unlikely to get two nearly identical organs for one job. ''That is not how the natural world works.'' In our state of ignorance of what P is and how P might be related to access, I don't think such an argument has much weight. One can imagine an uninformed person wondering why evolution needed both sperms and eggs, two things for one job. Should the uninformed person conclude that sperm = egg? No, he should try to find out more about the correlation. And we are no less ignorant about consciousness than this uninformed person. Furthermore, even sophisticated biologists do not agree about why sex evolved. So the question of Why two things for one job? is still a live one. Nonetheless, concluding that sperm = egg would be a joke. Baars coauthored ''Does Philosophy Help or Hinder Scientific Work on Consciousness?'' (Baars & McGovern 1993), arguing that philosophers should get out of the way. We should evaluate this criticism in the light of the fact that philosophers are the ones most likely to raise doubts about Baars's research program.

R5. Why Not Make A Easier to Have?

Church, Graham, and Revonsuo think I set the A-hurdle too high. Church suggests accessibility instead of being poised for access, and Graham wants to set the hurdle low enough so that the kind of access that blindsight patients have is good enough for A. Flanagan (1992) reacted in a similar manner to an earlier version of this paper, proposing that we substitute a notion of ''informational sensitivity'' for A, where blindsight patients are informationally sensitive to information in their blind fields.

Of course, one is free to define ''Access-conscious'' as one chooses. What I was after was a notion of access that I find in common-sense reasoning and that is the best shot at coextension with P. Defining A in terms of informational sensitivity will frustrate that aim. As I mentioned, blindsight will count as a case of A without P. Indeed, one reason for choosing ''poised for access'' instead of ''accessible'' is to avoid classifying as A a familiar kind of inactive or dormant belief. For example, we were all taught facts in elementary school, such as that the sun is 93 million miles away from the earth. Perhaps you were taught this and have believed it ever since, even though you haven't thought of this fact in years. It was an inactive belief. But if we make A a matter of accessibility, then such inactive beliefs will be A but not P, and that makes the failure of coextension of A and P a trivial consequence of a definition. In the view of many of the commentators, A = P or at least A and P are coextensive or can be made so with a little tinkering with definitions. I disagree, but I do not want my disagreement to rest on a triviality.

Similar points apply to the definition of ''P.'' Humphrey prefers to restrict P to the result of irritation of the sensory surfaces. This leaves out, for example, images and the phenomenal aspect of thinking, and would therefore generate A without P. Perhaps Humphrey would want to include images that reflect the same states as are produced by sensory surfaces, but then why leave out the phenomenal aspect of thought?

As Chalmers (1997) notes, I should welcome attempts to tinker with the definitions of ''P'' and ''A'' so as to make them coincide better. I don't want my claim that P $\neq$ A to depend on anything that is merely verbal. So I invite further attempts to improve these definitions.

R6. Why Does A Have to Be Rational?

I defined A using the notion of rationality, and this draws complaints from Graham, Lloyd, Revonsuo, and Warren as well as an excellent account of what I should have said from Kobes. Though I think there are deep connections between consciousness (both P and A) and rationality, I didn't intend to imply that principles of logic or good reasoning are necessary for A or that animals cannot have A. I meant to appeal to the use of a representation in reasoning, even if the reasoning is poor. And I intended a relativization to the capacities of the type of animal involved. As Kobes says, ''Access is not diminished merely in virtue of the creature's having less power to reason or act.'' I apologize to my readers for not being clearer about this.

R7. Is the P/A Distinction Useful?

What is good about the P/A distinction? (1) It is an ordinary concept of consciousness, and so it is relevant to how we think about ourselves. (2) It is the information-

processing image of P and thus a good candidate for what P is in information-processing terms. And (3) the relative ease of finding cases of P without A as compared with A without P suggests the distinction is on to something to do with the joints of nature.

Dixon argues that there is no point in distinguishing P from A. He gives a number of cases that he apparently sees as borderline ones, not clearly A or P. But to conclude from the existence of borderline cases that there is no distinction or that the distinction is not useful is a mistake. There are objects that are borderline cases between a table and a chair—a bit table-like, a bit chair-like, but neither table nor chair. That doesn't impugn the utility of the table/chair distinction. Dixon mentions many cases, but I'll just discuss the one he says casts the greatest doubt on the distinction: Hypnosis induces a hallucination in the subject of the experimenter, but when the subject turns to the side he also sees the real experimenter. Dixon seems to think that there is some sort of conundrum here that casts doubt on the A/P distinction. But if there is a genuine hallucination, then when I'm having it I'm having one experience that is both A and P, and when I see the real experimenter I have a similar experience that is also both A and P. What's the problem? Dixon goes on to argue that I should not have constructed a theory of consciousness on the basis of evidence from brain damage, because these patients may have compensating defects that make them different from normal people, and there are not enough of them for good sample size. But if brain-damage cases show P without A, then its existence is proved whether or not it ever occurs in normal people, and if brain damage does not yield cases of A without P, this is an especially interesting fact given the fantastic wealth of variation in brain-damage cases. These points illustrate why general cautions like, ''You can't show anything by appealing to brain damage'' are so weak. Every source of evidence has its pitfalls—the critic of a particular bit of empirical reasoning must show that the pitfalls have been engaged.

Warren rejects the A/P distinction because it is not defined in a completely clear and unambiguous way. But the demand for such definitions is misplaced. Especially at the beginning of a scientific enterprise there is no alternative to going by the seat of one's pants. I once saw a book that discussed the quality of scientific articles. The authors shared Warren's mistaken view of definition. They felt it was important to define ''scientific quality,'' and they did so in terms of the number of references to the article in the subsequent literature. As anyone can see, that is no good—for example, an article can be referred to as a standard example of a terrible mistake. At an early stage in inquiry, noncircular definition is usually not possible. It took a whole thermodynamic theory to ground the thermodynamic definition of temperature, and further work reducing this theory to statistical mechanics to ground the definition of temperature as mean molecular kinetic energy. Definition and theory must progress together. The demand for definition at an early stage encourages misplaced precision.

R8. Is There a Fallacy at All?

Atkinson and Davies quote Shiffrin and Schneider (1977) giving a theory of P in terms of A, and they say reasonably that there is no conflation here. But the fact that Shiffrin and Schneider don't exhibit a conflation does not show that others don't. The sign of lack of conflation in what they quote is that in a single sentence the authors say they are going to give an information-processing theory of the ''phenomenological feeling of consciousness.'' But there is another route to a theory of P in terms of A: first you conflate P and A, and then you give a theory of A, taking it to be a theory of P. The cases I quoted look quite different from Shiffrin and Schneider, more in the direction of the second story. The important point is that the difference here is not in premises and conclusion. I agree that in that regard Shiffrin and Schneider are more or less the same as a number of the cases I mention. The difference lies not in the premises and the conclusion but in the means of getting from one to the other.

Atkinson and Davies go on to suggest a new argument for explaining A in terms of P. P is relatively intrinsic and categorical compared to A, whereas A is relatively relational and dispositional. They are right about this, but the upshot for matters causal is limited. Atkinson and Davies think the relation between P and A is like that between the chemical basis of solubility and the tendency to dissolve. However, a token thought can be accessible at one time but not another, depending on the whole system and the pathways available. We do not know that P content is a force toward mechanisms of reasoning and reporting. This is of course intuitively plausible, but then blindsight is or was intuitively *im*plausible. Suppose, for example, that Schacter's (1989) model is correct. Then we may be able to explain why P-conscious representations tend to be A-conscious without any appeal to the intrinsic properties of P. It is a property of the model that anything that gets to the P module is close to the Executive System and (perhaps) likely to be sent there.

Tye also argues that there is no fallacy, but on a different ground. P is preconceptual, so how could it involve the Executive System? So it must be that information fails to arrive at the Executive System because it fails to be P; so there is no fallacy. On the substance of Tye's argument: How do we know if P is preconceptual? I used the phrase ''representational'' to describe P-content instead of ''intentional'' to allow for that possibility, but I have seen no convincing argument to the effect that P-content is preconceptual. Furthermore, there is reason to doubt the preconceptual claim. The specialized modules appear to have lots of conceptualized information. For example, there appears to be information in the face module about people's *occupations* (see Sergent and Poncet 1990; Young 1994a; 1994b). On Schacter's (1989) model, all the inputs to the P-module are from sources that contain conceptualized contents. But suppose that Tye is right. Still, this is a new argument. When someone finds a good argument from premises to a conclusion, there is a temptation to suppose that this is what others

have had in mind who have argued, apparently fallaciously, from the same premises to the same conclusion. However, I saw no sign of Tye's argument in the works that I criticized.

R9. Is P at All Representational?

Katz argues that I have not presented sufficient reason to conclude that P-content is at all representational. He notes that even though what it is like to hear a sound from the right ≠ what it is like to hear a sound from the left, one cannot conclude that P-content is representational. What it is like to be married ≠ what is it like to be single; marital status is social, but it would be a mistake to conclude that P-content is social. The argument form is certainly not valid, but it has something to it: the premises call for an account of the difference, in the case of marriage, we have an adequate account of why what it is like to be married ≠ what is it like to be single without assuming that phenomenal content is social. And in the direction case, as Katz says, we could explain the different P-contents on the basis of a difference in attention or orienting. But I believe that there are many cases for which the best explanation of the difference between the P-content of seeing X and the P content of seeing Y appeals to representational features of P.

Consider the P content of seeing a square compared to the P-content of seeing a circle. These P-contents allow one to see that the squares are packable together without gaps, whereas the circles are not. Also, the squares have a small number of axes of symmetry, but the circles have a large number. These examples show that P-content is representational, but they also show something stronger and more interesting, something that must be Katz's real target. Katz's position allows that P-content represents as ink represents, namely, extrinsically, that is, it can be used to represent. But what the examples show is that P-contents represent per se. The P-contents are intrinsically packable (for the square-representing contents) or not packable (for the circle-representing contents). The P-contents alone allow one to see such facts.

Katz notes that my examples are spatial, and suggests that to the extent that I am right, it may be because P-consciousness involves ''clicking'' on a region of a spatial buffer. But nonspatial properties, for example causality, are represented by P-contents. Roll a ball at another so that the first makes the second move. I don't think that very subjective type could be experienced as the second ball acting on the first (see Michotte 1946). Katz points out that for P properties that are subject to spectrum inversion or an analog of it, the P-contents can represent the properties involved in inversion only extrinsically. True, but the possibility of spectrum inversion applies notably to ''secondary'' qualities like colors and not to P-contents, of, for example, shapes or causal properties. Katz also objects that there is some tension between P-content being representational and there being a P-module. But the representational features of P-contents

could depend on processes that occur prior to the P-module (say in the specialized modules) or after it. The idea would be that non-P representations sent to the P-module become representational P-contents within it. Like Farah, Katz sees the P-module as concerned only with intrinsic properties of the representations in it. Katz does not comment on whether, for example, thoughts are P-states, but his views fit best with Humphrey's restriction of P to sensory properties.

R10. Is There an Important Distinction Left Out?

Harman and Lycan have similar criticisms. First, both think that I have left out a distinction. Lycan distinguishes between a quale, for example, a red area of the visual field, on the one hand, and self-monitoring of that quale on the other. The self-monitoring apparently consists in mobilizing internal attention toward the quale or alternatively, in the views of many students of this issue, having a higher order thought about the quale. Lycan says that my term ''P'' comprehends both of these, and he seems to suggest that ''P'' (as with ''what it is like'') is ambiguous between the two. I find it hard to accept the criticism that I left out the P/monitoring distinction, since I explicitly mentioned three forms of internal monitoring consciousness and I explicitly distinguished them from P (sect. 4.2.2, para. 1). Lycan also disagrees with me about whether qualia are entirely representational, that is, whether there is more to P-content than representational content. I say yes, Lycan says no. But Lycan promotes this entirely legitimate disagreement into another sense of P-consciousness. He calls my qualia, the ones that aren't entirely exhausted by their representational properties, ''Q-ualia.'' I don't see why a disagreement about P-content should be blown up into a new sense of P.

Harman says I miss the distinction between ''raw feel'' and ''what is it like.'' Raw feel is Harman's word for Lycan's Q-ualia, the P-contents that I accept, and he rejects that are supposed to outrun their representational content. Harman's what it is like is Lycan's monitoring consciousness. So raw feels are P-contents that are at least in part nonrepresentational and can exist without monitoring, and what it is like, which Harman feels is the proper referent of ''consciousness,'' is entirely representational and has a constitutive connection to the self. As with Lycan, I find it hard to take seriously the criticism that I have conflated these things. I was very explicit about all the ''pieces.'' The issue between me and Harman is one of what the most revealing way of assembling these pieces is.

What separates Harman and Lycan from me is mainly two issues. First, I say the phenomenal content of an experience goes beyond its representational content. They disagree (more on this later). A second source of disagreement has to do with the relation between consciousness and the self and monitoring. My P-content is a kind of phenomenal content that need not be monitored, and I give little emphasis to the con-

nection with the self. (I describe P-content at a few points as having a ''me-ish'' phenomenal quality.) So my major category does not emphasize monitoring or connection with the self, and in fact I mention monitoring and self-consciousness as separate categories. By contrast, Lycan emphasizes monitoring. For example, he says that there is an explanatory gap for monitoring consciousness, but not (or not obviously) for the phenomenal content that is itself monitored. And Harman emphasizes that A-conscious experience is always an experience of a self.

To sum up, (1) I have a substantive disagreement with both Lycan and Harman about whether there is any phenomenal but not entirely representational content. And (2) there is a substantive disagreement with Lycan about the explanatory gap. But (3) there is also a much more diffuse issue between me and them about what is important in the study of phenomenal consciousness. Lycan emphasizes monitoring, Harman emphasizes the self (at least by contrast with me), and I emphasize the phenomenal quality of experience. Because there are some differences between Lycan and Harman, let me discuss them separately.

Lycan in effect criticizes (see especially his Note 4) my claim in the target article that the ''explanatory gap'' applies to P-consciousness. He says there is an explanatory gap for monitoring consciousness (P attended), but not, or not obviously, to P itself. I would like to see Lycan back this up. Attention is as likely to yield to the information-processing theories of cognitive psychology and cognitive neuroscience as is, say, memory or any other cognitive process. It is an active area of research with many competing theories—see, for example, the seven articles in *Attention and Performance XV* (Umiltà and Moscovitch 1994), or the seven articles in *Attention and Performance XIV* (Meyer & Kornblum 1993). By contrast, there are really no theories (nothing that deserves to be called a theory) of P. No one really has any idea about what P is. As mentioned earlier, the typical research program is to study A, hoping A = P (see Baars). Monitoring consciousness is attended P-consciousness, so what is likely to be understood within the confines of current research paradigms is just the part that Lycan thinks adds the mystery.

Harman says that A-conscious experience is always an experience of a self and necessarily involves access to that self, so, trivially, Consciousness is ''access consciousness.'' Is access to the self supposed to involve engagement with mechanisms of reasoning and reporting bringing with them inferential promiscuity, and so on. If so, then Harman owes us some comment on the putative cases of P without A: If not, then I don't think there is a large disagreement here, for Harman's view does not then preclude P without A. (Levine makes essentially this point.) But there is at least a difference in emphasis. I am a Humean about the self (like Dennett and Church), seeing the self-regarding aspect of P-consciousness as being a matter of connection of the P-state to other states. I said in the target article that P-content often represents the state as a state of mine. Part of the self-regarding aspect of P in my view is a further P-attribute

that involves some apprehension of the connection to other states. But I am also willing to countenance P-states in my body that are not fully mine. (I mentioned hypnotic analgesia as a possible example.)

There is one issue that I have not yet mentioned on which Lycan agrees with me rather than Harman. Lycan allows a category of qualia (e.g., a red area of the visual field) that are phenomenal but not necessarily monitored. I would guess that these fit into the category of what Harman calls ''sense data,'' which he takes me (wrongly) as endorsing. I am grateful to Lycan for explicitly not supposing (as he did in Lycan, 1987, and as Harman does here) that the advocate of qualia is committed to sense data or ''phenomenal individuals.'' If any of us is committed to sense data, it is Lycan, Armstrong, Church, Kitcher (and perhaps Harman) and other advocates of monitoring. The rest of us can agree with Harman (1990) that we look *through* our experiences, and that the experiences do not need to be *observed* in order to be phenomenally conscious.

Lycan and Harman think that P-content is entirely representational. They note that I think P-content outruns representational content, and they both appear to conclude that I am therefore committed to some new strange kind of phenomenal content that is entirely nonrepresentational, Lycan's Q-ualia and Harman's raw feels. I did say that the P-content of orgasm represented nothing at all, but this is not a strongly held view. I am happy to say that very little of the phenomenal content of orgasm is representational. Certainly very little of what matters about orgasm is representational. What puzzles me about Lycan and Harman is that they appear to think that the idea that there is more to phenomenal content than what it represents entails some ''weird'' or ''exotic'' realm of sense data that are entirely nonrepresentational and of which one is ''directly aware'' in perception. As reflection on the example of the phenomenal content of orgasm should make clear, the idea that there is more to phenomenal experience than its representational content is just common sense from which it should take argument to dislodge us. Furthermore, why should believing in phenomenal contents that are *partly* nonrepresentational commit one to *wholly* nonrepresentational phenomenal contents (of the sort Katz advocates)? Perhaps Harman and Lycan think that if a P-content is partly nonrepresentational, one can simply separate off the nonrepresentational part and think of it as a separate realm. But once the argument is made explicit it looks dubious. Consider the examples I used in my reply to Katz, say, the example of packability in the case of experiences as of squares contrasted with circles. Is it obvious that there is any separable phenomenal content of that experience that is phenomenal but not representational? I don't think so.

R11. More on Monitoring

Kitcher objects to my contrast between P-consciousness, which applies primarily to states, and monitoring or reflective consciousness, which applies primarily to persons.

A pain is monitoring conscious if (roughly speaking) the person has another state that is about the pain. She notes that monitoring consciousness is a matter of some states being about others, and wonders why I make this distinction. The answer is that if a state of mine is about a pain of yours, your pain is not thereby monitoring conscious. So the notion of a person is crucial. (Someone could argue, as Kitcher does not, that the same is true of A-consciousness.)

Kitcher also says that she cannot see what ''what it is like'' could evoke if not monitoring consciousness and that the explanatory gap applies most obviously to monitored states. She also finds it implausible that there could even be phenomenal consciousness without monitoring. These points should sound familiar, since I just discussed versions of them in the comment by Lycan, and to a slightly lesser extent in the comment by Harman. Church also favors the view. See also Rosenthal (1986) and Burge (1997). I find this difference of opinion far more troubling than any other that comes up about consciousness. I really don't know how to explain the vast divergence we see here. The magnitude of the gulf is apparent from the fact that two of the commentators, Armstrong and Natsoulas, assumed that *I* mean monitoring consciousness to be involved in A or A and P together. Armstrong complains about the term. ''A'' would be better, he says, if it stood for action; and even better: change it in for ''I'' for introspection. My A-consciousness, however, requires no introspection. Natsoulas says—and he says that I agree—that if we have an A-conscious P-state, then we must have another representation of that state. He calls this representation of the phenomenal state ''the required representations,'' since it is supposed to be necessary for A-consciousness. I am not sure that I follow the rest of the argument, but he seems to go on to argue that the required representation itself has to be the object of yet another state.[2]

What can be said in favor of the idea that monitoring is necessary for phenomenal states, or at least for ''sensory experience'' (Lycan). Kitcher mentions that listening to a piece of music requires integration over time. But what reason is there to think that sensory integration requires states that are about other states? It requires memory, of course, but memory images can be linked in the appropriate way without any ''aboutness.'' Lycan appeals to the familiar long-distance truck driver who drives competently but in a daze. He stops at red lights and so must have had a real quale, but for experience, says Lycan, he has to notice the quale, that is be aware of it. Nissan is funding some work at MIT that apparently includes an investigation of this phenomenon, and I have been told some simple preliminary results. If you probe ''unconscious'' drivers, what you find is that they can always recall (accurately) the road, decisions, perception, and so on, for the prior 30–45 seconds, but farther back than that it's all a blank. No one should be surprised by this result. What else would one expect? If you were a subject who was just asked about the last 30 seconds, do you think you would say that you had not experienced the last 30 seconds? If you say yes, you are in the grip of a theory.

This seems a clear case of experience as genuine as any but quickly forgotten, a moving window of memory. The driver is paying some attention—to the road. Otherwise the car would crash. He is not paying attention to his own states, but one rarely is. Of course, more attention to the road or to the experiences themselves would yield different experiences. But the inattentive driver is still experiencing the bends in the road, the red lights, the other cars maneuvering around him. Why should anyone suppose that there is nothing it is like to be that driver or that to the extent that there is an explanatory gap it doesn't apply here?

One way to see what is wrong with the idea that monitoring consciousness is crucial for P-consciousness is to note that even if I were to come to know about states of my liver noninferentially and nonobservationally (as some people know what time it is), that wouldn't make those states P-conscious. Furthermore, even if I were to come to know of states of my mind that way—say, the operation of my language-understanding mechanisms, or Freudian unconscious states—that wouldn't make those states P-conscious. Of course, all this observation shows is that monitoring isn't sufficient for P, but if monitoring is necessary for P, what else is required to get a sufficient condition? Advocates of this view have not provided an answer to this question.

A second point is that monitoring seems too intellectual a requirement for phenomenal consciousness. Dogs and babies may have phenomenal pains without anything like thoughts to the effect that they have them. If we have two dogs, one of which has a pain whereas the other has a similar pain plus a thought about it, surely the latter dog has an A-conscious state even if the former doesn't! Yes, but it is the converse that is problematic. The first dog could be conscious without being conscious of anything.

Kitcher anticipates the dog objection and replies that I make monitoring a sophisticated activity requiring a sense of self. Not so. What I doubt is that a dog that has a phenomenal state need have any further state that is about the first one. I don't require a sense of self.

As observed in the target article, advocates of the higher-order thought perspective (e.g., Rosenthal) note that if I infer my anger from my behavior, that does not make my anger conscious. They therefore include a requirement that the higher-order thought be arrived at noninferentially and nonobservationally. But as Byrne (forthcoming) notes, why should these details of the causation of the monitoring state matter to whether the state that is monitored is *conscious?* Byrne mentions a number of other conundra for the advocates of monitoring having to do with the level of detail of the monitoring state and the question of whether the description included in the monitoring state could be false of the state monitored.

Levine makes some remarks that may help to explain this puzzling difference of opinion with advocates of monitoring. He notes that phenomenal character itself is a "kind of presentation," a presentation to the self. He also says that this brings with it a kind of access that is distinct from A, phenomenal access as distinct from information-

processing access. And he suggests that the existence of two kinds of access is partly responsible for the difficulty in distinguishing A from P. There is at least something right about this. It is often said that phenomenology is self-revealing, that there is something intrinsically epistemic about phenomenology. Perhaps phenomenal access is itself a phenomenal quality, a quality that has some representational features. These representational features represent the state as a state of me. But it does not follow that any kind of information-processing access (such as A) or monitoring is necessary for P.

R12. Does P Outrun Its Representational Content?

Armstrong, Harman, Lycan, and Tye all take the view that P-content is entirely representational. I like Tye's approach best because he doesn't treat it as obvious that representationalism is right, but rather sees a responsibility to say what the representational contents actually are. In the case of orgasm, he specifies that the representational content is in part that something that is intense, throbbing, changing in intensity, and very pleasing is happening down there. OK, I will concede one thing—that there is *some* representational content to orgasm. But this representational content is one that I could have toward another person. Suppose I have perceptual contents about my partner's orgasm without having one myself. The location of ''down there'' might differ slightly from my own orgasm, but why should that matter? Of course, the subject the orgasm is ascribed to is itself a representational matter. But is that the difference between my having one and perceiving yours—that I ascribe it to me instead of you? What if I mistakenly ascribe yours to me? Furthermore, the phenomenal quality of orgasm varies from time to time. Similarly, there are very different phenomenal experiences that fit descriptions like ''in the toe,'' ''intense,'' ''burning,'' and the like.

I had a pain yesterday that is quite different phenomenally from the one I am having now, but not in any way describable in words. Of course, we should not demand that a representationalist be able to capture his contents in words, but we should be told something about the representational difference. Suppose the content is specified in terms of recognitional capacities. That runs into the problem that recognitional capacities can work without P-content, as in blindsight. At this point of the dialectic, the representationalist often appeals to functional role to specify the representational contents. So is the debate about whether phenomenal content is entirely representational just the old debate about functionalism and qualia? Representationalists certainly give the impression that their position is stronger than mere functionalism, that they can accommodate the idea that there are phenomenal contents, but that those contents are representational.

The way in which representationalism is stronger than mere functionalism comes out in Tye's criticism of my example of two kinds of experiences as of something overhead. Tye doesn't just say: sure, the representational difference resides in the

functional difference (though that might be the upshot of the last two sentences of his commentary). Instead, he tries to say what the representational differences are. He argues that the difference will reside in other visual and auditory features. I believe that Tye is wrong about vision but right about audition. In peripheral vision, something can be seen only as having a certain location, without any color, shape, or size. (Try waving your hand near your ear while looking straight ahead.) But without a comparable point in audition, my example will not work, and I know of no auditory analog of peripheral vision. However, my side has another arrow, for the loudness of the sound is irrelevant to its representing something as of overhead. The as-of-overheadness of the visual perception seems independent of color, shape, and so on, and likewise for the auditory perception. The difference seems to reside in the phenomenal character of vision as opposed to audition, and that has not been shown to be a representational difference.

R13. What Is the Relation between A and P?

Shepard, as usual, asks hard questions. How do we know if P peters out as we go down the phylogenetic scale as A peters out? It is a measure of our ignorance about P-consciousness that we have no idea how to go about answering such a question. I think all we can do is investigate P in the creatures we know best and hope that the answer we get throws some light on creatures who are very different from us. Shepard says that what agents do is evidence of A, not of P. I disagree. Sure, purposive action is evidence of A, but it is also evidence, albeit indirect evidence, of P. For example, let us accept for the moment Crick's (1994) current theory of P: that P is a matter of neuronal activity in reverberating cortico-thalamic circuits that run between cortical layer five and the thalamus. Such a theory can only be arrived at on the basis of behavior that indicates A. But once we have the theory (and especially when we understand why that neuronal activity underlies P) we can use it to isolate cases of P without A, or cases, if they exist, of A without P. Of course, we have to explain the discrepancy. Thus, if we find the neuronal activity but no A and hence no outward indication of consciousness, we have two choices: conclude that Crick's theory is wrong, or find some reason why in this particular case there is no A.

This line of thought also supplies my answer to Rey's charge that if P is not identical to anything functional, intentional or cognitive, ''what possible reason could we have to posit it in anyone's case, even our own?'' I think it is always a mistake to suppose that no one could ever find evidence of something (with a few exceptions—e.g., the thing is literally defined in terms of there being no possible evidence for it). This is just an argument from lack of imagination. A neat example is provided by the familiar idea that the world was created five minutes ago complete with all the evidence of an earlier age. It is tempting to argue that no one could find any evidence for or against

such a theory, but that would be a mistake. Steady state cosmology plus the second law of thermodynamics (entropy increases in a closed system) dictate that the relatively ordered state we see around us is a result of a random fluctuation from a steady disordered state. The great fluctuation that created our order happened in the past, but when? Answer: the most likely moment for the fluctuation is the least ordered moment, and that is the most recent moment, that is, *now*. So the evidence *against* steady state theory is evidence *for* the existence of a real past.

Furthermore, in thinking about this sort of possibility, we should not ignore the utility of ordinary considerations of scientific simplicity and ad hocness. For example, one can maintain any theory—even that the earth is flat—if one is willing to adopt all sorts of ad hoc auxiliary hypotheses to explain away recalcitrant observations. In so doing, one could arrive at a totally wacko theory that is observationally equivalent to contemporary physics. But the wacko theory can be ruled out just because it is ad hoc. A further point about the ''epiphenomenal'' possibility is that the epiphenomenalism of Figure 3 (target article) is not the ''philosopher's epiphenomenalism'' in which the epiphenomenal entity has no effects at all. Rather, it is the psychologists' epiphenomenalism that rules out effects only *in a system*. The latter allows for effects, but outside the system. The color of the wires in a computer are epiphenomenal in the psychologist's sense but not in the philosopher's sense, since there are effects on observers. Thus the P-module of Figure 3 could be detectable by physiologists even if it had no psychological function.

I agree with what Morton says about the interdependence of A and P, and I gave a number of similar examples myself. (There is a foreground/background example in sect. 4, para. 3 and three more examples in sect. 4.2, para. 6.) I also agree with the idea that we would not have the concepts we have if not for these facts. But I do not agree that the intuitive idea of there being only one consciousness shows that the concept of consciousness is a cluster concept rather than a mongrel. The distinction, as I intended it, was linked to the concept of a conflation. If conflation is possible, then mongrel; if not, cluster. If the myth of uniqueness is enough to make a cluster, then Aristotle's conception of velocity is a cluster concept. Of course, there is no right or wrong here, only utility. If we adopt Morton's terminology, we shall have to make a distinction within the cluster concepts between those that allow conflation and those that do not.

Farah argues that if superblindsight existed, that would be evidence for Schacter's model, complete with P-module, and the nonexistence of superblindsight is evidence against such a P-module. In other words, she thinks that if the presence or absence of such a module made no difference to perception (but only to whether the subject says he is having experiences) that would be evidence for such a module. This seems to me to be precisely backwards. If a module has some information-processing function—and why else would it deserve a box—then whether it is present or absent should make a

difference. It seems to be an essential feature of a P-module on Farah's idea of it, that it doesn't do much of anything except paint representations with P-paint and promote reports of experience. Sure, if it has little in the way of an information-processing function then its presence or absence shouldn't make much or a difference. But why assume that if there is a P-module it doesn't have much of an information-processing function? For example, perhaps the Executive System can do things with P representations that it can't do with non-P representations.

Farah objects to my suggestion of P-consciousness as the implementation of the function specified by the Phenomenal Consciousness box in Schacter's model. As I mentioned in the target article, the function specified by that box (and there may be others not specified by the box) is that of talking to the specialized modules, integrating information from them, and talking to the Executive System about that information. I suggested that perhaps P-consciousness is part of the implementation of that function. I used an analogy in which this function could be implemented in a number of ways, some involving consciousness, others not involving consciousness. Farah interprets the label on the box as specifying the sole function represented. I tried to cancel that reading in the text by mentioning that the function was to be understood partly in terms of the box, arrow, their relations, and the textual remarks on how these are to be interpreted. Since the label is ''phenomenal consciousness,'' she assumes that that is the one and only intended function. So we were at cross purposes.

Young suggests that P is responsible for confidence. After all, people who ''just know'' what time it is don't have the confidence of people who are looking at a clock that they know to be reliable. This is certainly sensible and compelling. But blindsight raises a doubt about such commonsensical ideas: maybe we could know without P? And if we could know without P, why not confidence without P?

Navon suggests that the function of P may be found primarily in motivation rather than cognition (I made the same suggestion in Block 1991). But there is an evolutionary puzzle that this idea raises, one that derives from William Paley (1964) (via a column by Stephen Jay Gould). Paley pointed out that there is no mystery about why birds copulate—pleasure is the answer. But we can't give the same answer to the question of why the bird sits on the egg. (Paley backs this up with a description of the misery of sitting on the egg.) But why does evolution deploy two such different motivators?

Bachmann notes that I say that P depends on what goes on inside the P-module, and he goes on to indicate that this is incompatible with interaction effects involving representational contents (see the replies to Katz and Farah). But these are not incompatible ideas, and I was careful in the target article to describe a number of respects in which P-consciousness is (in the words of Armstrong describing my views) ''thoroughly interpenetrated'' by representational matters. Note that in Schacter's model, the P-module talks to the Executive System and the specialized modules, so interac-

tions are allowed for. Bachmann mentions the possibility that a P-module might have its activity lowered, but he somehow takes me to be denying this possibility and doubting the possibility of measurement of P. Bachmann mentions a number of fascinating phenomena that may cast some light on the relation between A and P, but I have not investigated these phenomenon sufficiently to comment on them.

R14. Is Consciousness a Cultural Construction?

Dennett says that my critique of his view that consciousness is a cultural construction simply begs the question. I assume the A/P distinction, but he rejects it, he says. ''Because I not only decline to draw any such distinction but argue at length against any such distinction, Block's critique is simply question-begging.'' This is a strange response from Dennett, since he does not actually reject the A/P distinction but rather reconstructs it in terms of information and access. Perhaps he thinks that the reconstructed A/P distinction is so different from what I meant that it is tantamount to rejecting the distinction. Well, then, let's suppose Dennett is completely right. To the extent that there is an A/P distinction, it is a matter of degree of access and information. Dennett's theory of access, you will recall, is that it is a matter of brain representations persevering so as to affect memory, control behavior, and so on. So the P/A distinction is a matter of brain representations' degree of informational content and degree of persevering. Then Dennett ought to be able to tell us to what degree or range of degrees of persevering and informational content his theory applies to. I'm not being very demanding. I don't insist on precision, just some idea of what degrees of information and control make for cultural construction. Perhaps he will say it is the highly informational and accessible contents he is talking about, the rich displays of colors and shapes that appear in his examples (e.g., the Times Roman X on a blue-green background). But we have good reason to think that these types of contents are not very influenced by culture. Long ago, Eleanor Rosch (1973) showed that the Dani, a tribe with only two color words, represented colors much as we do. In sum, my point against Dennett does not depend at all on whether the A/P distinction is a matter of degree or of kind. If it is a matter of degree, he must tell us what band of degrees he is talking about.

Dennett used to be an eliminativist (in ''On the Absence of Phenomenology'' [1979], for example). In recent years, especially since Dennett (1991), he has shifted gears, saying he is a realist about consciousness and at the same time saying that his position is not all that different from what it used to be. He appeals to the truth that the difference between eliminativism and reductionist realism is often purely tactical. However, not all reductionisms are close to eliminativism. Indeed, Dennett's new position is very different from his old one, as many readers have recognized (see Rey, e.g.). In giving what he insists is a theory of consciousness, with such highly substantive claims as that

consciousness is a cultural construction, Dennett has left eliminativism far behind. Now he is a real realist, a reductionist or a deflationist, and the theory is supposed to be true of some deflated version of consciousness or something consciousness is reduced to. The trouble is that he has neglected to make up his mind about which deflated version he wants or what it is that he is reducing consciousness to.

My advice to Dennett is to read Church—that's the view of A and P that best captures his intentions. Church says that my analogy with the development of the concepts of heat and temperature is miscast. I said that we have an intuitive preanalytic concept of consciousness that can be resolved into P, A, monitoring consciousness, and self-consciousness. She argues that P should be seen as the preanalytic concept, and given its confused nature, we should abandon it in favor of ideas such as A, monitoring, and self-consciousness. She gives an interesting argument for the confused nature of P. A P-state must be a state of a self, and given that there is no Cartesian self, being a state of a self must involve relations to other states. Then comes the step that mainly bothers me: according to Church, P (if it exists) is intrinsic. Since P is both intrinsic and involves a relation, P is a confusion. My view is that this step conflates concepts with the properties that they are concepts of. The concept of a color does not involve relations, but color (the property) is highly relational. The concept of water has nothing to do with molecules, but water (or the property of being water) is constituted by being a molecular amalgam of hydrogen and oxygen. Similarly, the concept of P (of a state) has nothing to do with other states, but P itself could turn out to be relational. This point is briefly mentioned in Note 10 of the target article.

Oddly enough, Church is well aware of the concept/property distinction, and pins the incoherence on the concept. Yet in laying out her argument, she shifts to properties, saying that the property of being phenomenal is both intrinsic and relational. I also have some disquiet about the prior step in her argument. I agree that a P state must be a state of the self, and I agree about the deflationary view of selves. But I am not convinced that the way in which P involves the self is incompatible with intrinsciness. At a minimum, the mode of self-involvement could be simple, a phenomenal property, the ''me-ishness'' I described (see Levine and my reply to his comment). Alternatively, the state could in some sense be about other states or about the self (it could represent the state as a state of me, as I said in the target article) but not in a way that would satisfy a functionalist. I can want a sloop even if there aren't any sloops. And the relation to other states or to the self could be like that.

Van Brakel takes me to be an advocate of ''one true taxonomy,'' but I took pains to avoid this characterization. I emphasized repeatedly that there are many notions of access-consciousness with utility for different purposes. My purpose had mainly to do with a notion of access as a surrogate for phenomenal consciousness. Furthermore, there are somewhat different notions of phenomenal consciousness that are legitimate for some purposes, for example, the limitation to bodily sensations suggested by Hum-

phrey (see also Katz). I am perfectly happy to allow that culture affects P-consciousness. I emphasized that intentional and phenomenal content interpenetrate, and I don't think anyone should doubt that culture can affect intentional content. But note the difference between the idea that culture affects phenomenal consciousness and the idea that culture creates it. Culture affects feet—the feet of Himalayan tribesmen who walk barefoot in the mountains are different from the bound feet of nineteenth-century Chinese women. But culture does not *create* feet. I have to admit skepticism about much of van Brakel's evidence, however. Whorfians thought culture affected color and form perception until Berlin and Kay (1969) and Rosch (1973) showed the effects were overrated. Van Brakel's evidence is best evaluated when we know enough about consciousness to see whether it really differs in different cultures.

Notes

Reprinted from *Behavioral and Brain Sciences* 18 (2): 227–247, 1995. Reply to 28 critics of ''On a Confusion about a Function of Consciousness'' *Behavioral and Brain Sciences* 18 (2): 227–247, 1995 (reprinted as chapter 9, this volume).

I am grateful for comments from Alex Byrne, David Chalmers, Leonard Katz, and Daniel Stoljar.

1. Opponents of the inverted-spectrum thought experiment should pay attention to cases like the legal blindness/eyesclosed case. It is much easier to come up with an inverted-spectrum type of thought experiment for a sensory modality with reduced informational content. Tye, for example, has objected to the inverted spectrum hypothesis on the basis of asymmetries in color—blue can be blackish but yellow cannot. But such objects do not apply to vision in the legal blindness/eyes-closed mode.

2. I can surmise that what misled Natsoulas was a remark in Note 11. It would take me too far afield to raise the issue here.

References

Akins, K. (1993) A bat without qualities. In: *Consciousness: Psychological and philosophical essays*, ed. M. Davies & G. Humphreys. Blackwell.

Alston, W. (1967) Religion. In: *The encyclopedia of philosophy*. Macmillan/Free Press.

Anderson, J. (1993) To see ourselves as others see us: A response to Mitchell. *New Ideas in Psychology* 11(3):339–34.

Andrade, J. (1993) Consciousness: Current views. In: *Depth of anesthesia*, ed. J. G. Jones. Little Brown.

Armstrong, D. M. (1968) *A materialist theory of mind*. Humanities Press.

——— (1980) What is consciousness? In: *The nature of mind*. Cornell University Press.

Baars, B. J. (1988) *A cognitive theory of consciousness*. Cambridge University Press.

——— (1994) A thoroughly empirical approach to consciousness. *Psyche: An International Journal of Consciousness Research* 1(2).

Baars, B., and McGovern, K. (1993) Does philosophy help or hinder scientific work on consciousness? *Consciousness and Cognition* 2, 18–27.

Bennett, H. (1988) Perception and memory of events during adequate general anesthesia for surgical operations. In: *Hypnosis and memory*, ed. H. Pettinati. Guilford.

Berlin, B., and Kay, P. (1969) *Basic color terms*. University of California Press.

Block, N. (1983) The photographic fallacy in the debate about mental imagery, *Nous* XVII, 4.

——— (1980) What is functionalism? In: *Readings in the philosophy of psychology, vol. 1*, ed. N. Block. Harvard University Press.

——— (1990a) Inverted earth. In: *Philosophical perspectives, vol. 4*, ed. J. Tomberlin. Ridgeview.

——— (1990b) Consciousness and accessibility. *Behavioral and Brain Sciences* 13:596–59.

——— (1991) Evidence against epiphenomenalism. *Behavioral and Brain Sciences* 14:670–67.

——— (1992) Begging the question against phenomenal consciousness. *Behavioral and Brain Sciences* 15:205–07.

——— (1993) Review of Dennett: Consciousness explained. *The Journal of Philosophy* 4:181–19.

——— (1994) ''Functionalism,'' ''Qualia.'' In: *A companion to philosophy of mind*, ed. S. Guttenplan. Blackwell.

——— (in press) What is Dennett's theory about? *Philosophical topics*.

Bornstein, R., and Pittman, T. (1992) *Perception without awareness*. Guilford Press.

Bowers, J., and Schacter, D. (1990) Implicit memory the test awareness. *Journal of Experimental Psychology: Learning, Memory and Cognition* 16(3):404–41.

Burge, T. (1997) Two kinds of consciousness. In *The Nature of Consciousness*, ed. N. Block, O. Flanagan, and G. Güzeldere. MIT Press.

Byrne, A. (1993) *The emergent mind*. Princeton University, PhD dissertation.

Byrne, R. W. (1993) The meaning of 'awareness': A response to Mitchell. *New Ideas in Psychology* 11(3):347–35.

Carruthers, P. (1989) Brute experience. *Journal of Philosophy* 86.

——— (1992) Consciousness and concepts. *Proceedings of the Aristotelian Society (Supplement)* 66:40–5.

Chalmers, D. J. (1993) *Toward a theory of consciousness*. University of Indiana PhD thesis.

Chalmers, D. (1997) Availability: The cognitive basis of experience? *Behavioral and Brain Sciences* 20:148–149.

Churchland, P. S. (1983) Consciousness: The transmutation of a concept. *Pacific Philosophical Quarterly* 64:80–93.

——— (1986) Reduction and the neurobiological basis of consciousness. In: *Consciousness in contemporary society*, ed. A. J. Marcel & E. Bisiach. Oxford University Press.

Cooper, L., and Shepard, R. (1982) Chronometric studies of the rotation of mental images. In: *Visual information processing*, ed. W. Chase. Academic Press. Reprinted in R. Shepard & L. Cooper, *Mental images and their transformation*. MIT Press.

Coslett, H., and Saffran, E. (1994) Mechanisms of implicit reading in alexia. In: *The neuropsychology of high-level vision*, ed. M. Farah & G. Ratcliff. Erlbaum.

Cowie, A., and Stoerig, P. (1992) Reflections on blindsight. In: *The neuropsychology of consciousness*, ed. B. Milner & M. Rugg. Academic Press.

Crick, F. (1994) *The astonishing hypothesis*, Scribner's.

Crick, F., and Koch, C. (1990) Towards a neurobiological theory of consciousness. *Seminars in the Neurosciences* 2:263–75.

Davies, M. (1992) Perceptual content and local supervenience. *Proceedings of the Aristotelian Society* 92:21–45.

——— (forthcoming) Externalism and experience. In: *Categories, consciousness and reasoning*, ed. A. Clark, J. Exquerro, J. Larrazabal. Dordrecht.

Davies, M., and Humphreys, G. (1993a) *Consciousness: Psychological and philosophical essays*. Blackwell.

——— (1993b) Introduction. In: *Consciousness*, ed. M. Davies & G. Humphreys. Blackwell.

de Lannoy, J. Two theories of a mental model of mirror self-recognition: A response to Mitchell. *New Ideas in Psychology* 11(3):337–33.

Dennett, D. C. (1979) On the absence of phenomenology. In: *Body, mind and method: Essays in honor of Virgil Aldrich*, ed. D. Gustafson & B. Tapscott. Dordrecht.

——— (1986) Julian Jaynes' software archeology. *Canadian Psychology* 27(2):149–15.

——— (1993) The message is: There is no medium. In: *Philosophy and Phenomenological Research III*.

Dennett, D., and Kinsbourne, M. (1992a) Time and the observer: The where and when of consciousness in the brain. *Behavioral and Brain Sciences* 15:183–20.

——— (1992b) Escape from the Cartesian theater. *Behavioral and Brain Sciences* 15:234–24.

Dimond, S. (1976) Brain circuits for consciousness. *Brain, Behavior and Evolution* 13:376–95.

Dretske, F. (1993) Conscious experience. *Mind* 102, 406:263–84.

Dupre, J. (1981) Natural kinds and biological taxa. *Philosophical Review* 90:66–9.

Edelman, G. (1989) *The remembered present: A biological theory of consciousness*. Basic Books.

Etcoff, N. L., Freeman, R., and Cave, K. Can we lose memories of faces? Content specificity and awareness in a prosopagnosic. *Journal of Cognitive Neuroscience* 3.

Etcoff, N. L., and Magee, J. J. (1992) Covert recognition of emotional expressions. *Journal of Clinical and Experimental Neuropsychology* 14:95–9.

Farah, M. (1994) Visual perception and visual awareness after brain damage: A tutorial overview. In: *Attention and performance 15*, ed. C. Umilà & M. Moscovitch. MIT Press.

Flanagan, O. (1991) *The science of the mind*, 2d ed. MIT Press.

——— (1992) *Consciousness reconsidered*. MIT Press.

Gallup, G. (1982) Self-awareness and the emergence of mind in primates. *American Journal of Primatology* 2:237–48.

Gallup, G., and Povinelli, D. Mirror, mirror on the wall, which is the most heuristic theory of them all? A response to Mitchell. *New Ideas in Psychology* 11:327–3.

Gazzaniga, M. (1985) *The social train*. Basic Books.

Ghoneim, M., and Block, R. (1993) Learning during anesthesia. In: *Depth of anesthesia*, ed. J. G. Jones. Little Brown.

Ghoneim, M., Hinrichs, J., and Mewaldt, S. (1984) Dose-response analysis of the behavioral effects of diazepam: 1. Learning and memory. *Psychopharmacology* 82:291–95.

Goldman, A. (1993a) The psychology of folk psychology. *Behavioral and Brain Sciences* 16:15–82.

——— (1993b) Consciousness, folk psychology and cognitive science. *Consciousness and Cognition II*.

Goodale, M. A., and Milner, A. D. (1992) Separate visual pathways for perception and action. *Trends in Neurosciences* 15:20–25.

Harman, G. (1990) The intrinsic quality of experience. In: *Philosophical perspectives, vol. 4*, ed. J. Tomberlin. Ridgeview.

Hauser, M., Kralik, J., Botto-Mahan, C., Garrett, M., and Oser, J. (submitted) Self-recognition in primates: Phylogeny and the salience of species-typical features.

Heyes, C. (1993) Reflections on self-recognition in primates. *Animal Behavior*.

Hilgard, E. R. (1986) *Divided consciousness*, 2d ed. Wiley.

Hill, C. (1991) *Sensations: A defense of type materialism*. Cambridge University Press.

Holender, D. (1986) Semantic activation without conscious identification in dichotic listening, parafoveal vision, and visual masking: A survey and appraisal. *Behavioral and Brain Sciences* 9:1–66.

Humphrey, G. (1963) *Thinking*. Wiley.

Humphrey, N. (1992) *A history of the mind*. Simon & Schuster.

Huxley, T. H. (1866) Lessons in elementary psychology. Quoted in Humphrey, 1992.

Jackendoff, R. (1987) *Consciousness and the computational mind*. MIT Press.

Jackson, F. (1977) *Perception*. Cambridge University Press.

——— (1986) What Mary didn't know. *Journal of Philosophy* 83:291–9.

——— (1993a) Appendix A (for philosophers). In: *Philosophy and phenomenological research III*.

——— (1993b) Armchair metaphysics. In: *Philosophy in mind*, ed. J. O'Leary-Hawthorne & M. Michael. Kluwer.

Jacoby, L., Toth, J., Lindsay, D., and Debner, J. (1992) Lectures for a layperson: Methods for revealing unconscious processes. In: *Perception without awareness*, ed. R. Bornstein & T. Pittman. Guilford Press.

James, W. (1890) *The principles of psychology*. Dover, 1950.

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

Johnson-Laird, P. (1988) *The computer and the mind*. Harvard University Press.

Jones, J. G. (1993) *Depth of anesthesia*. Little Brown.

Kapsalist, J. G. (1987) *Objective Methods in Food Quality Assessment*. CRC Press.

Kihlstrom, J. (1987) The cognitive unconscious. *Science* 237:1445–145.

Kihlstrom, J., Barnhardt, T., and Tataryn, D. (1992) Implicit perception. In: *Perception without awareness*, ed. R. Bornstein & T. Pittman. Guilford Press.

Kihlstrom, J., and Couture, L. (1992) Awareness and information processing in general anesthesia. *Journal of Psychopharmacology* 6(3):410–41.

Kihlstrom, J., and Schacter, D. (1990) Anaesthesia, amnesia, and the cognitive unconscious. In: *Memory and awareness in anaesthesia*, ed. B. Bonke. Swets & Zeitlinger.

Kim, J. (1995) "Supervenience." In: *Blackwell's Companion to Metaphysics*, ed. J. Kim and E. Sosa. B. H. Blackwell.

Kirk, R. (1992) Consciousness and concepts. *Proceedings of the Aristotelian Society* (Supplement) 66:23–4.

Kosslyn, S. (1980) *Image and mind*, Harvard University Press.

——— (1994) *Image and brain*. MIT Press.

Kosslyn, S., and Koenig, O. (1992) *Wet mind: The new cognitive neuroscience*. Free Press.

Kuhn, T. (1964) A function for thought experiments. In: *Melanges Alexandre Koyre, vol 1*. Hermann.

Lackner, J., and Garrett, M. (1973) Resolving ambiguity: Effects of biasing context in the unattended ear. *Cognition* 1:359–37.

Landis, T., Regard, M., and Serrat, A. (1980) Iconic reading in a case of alexia without agraphia caused by a brain tumour: A tachistoscopic study. *Brain and Language* 11:45–53.

Levine, J. (1983) Materialism and qualia: The explanatory gap. *Pacific Philosophical Quarterly* 64:354–36.

——— (1993) On leaving out what it is like. In: *Consciousness: Psychological and philosophical essays*, ed. M. Davies & G. Humphreys. Blackwell.

——— (1994) Review of Owen Flanagan: Consciousness reconsidered. *Philosophical Review*.

Loar, B. (1990) Phenomenal properties. In: *Philosophical perspectives: Action theory and philosophy of mind*, ed. J. Tomberlin. Ridgeview.

Lormand, E. (forthcoming) What qualitative consciousness is like.

Luria, A. (1972) *The man with the shattered world*. Harvard University Press (translation: 1987).

Lycan, W. G. (1987) *Consciousness*. MIT Press.

Mandler, G. (1985) *Cognitive psychology*. Erlbaum.

Marcel, A. J. (1983) Conscious and unconscious perception: An approach to relations between phenomenal experience and perceptual processes. *Cognitive Psychology* 15:238–300.

——— (1986) Consciousness and processing: Choosing and testing a null hypothesis. *Behavioral and Brain Sciences* 9:40–44.

——— (1988) Phenomenal experience and functionalism. In: *Consciousness in contemporary society*, ed. A. J. Marcel & E. Bisiach. Oxford University Press.

Marcel, A. J., and Bisiach, E., eds. (1988) *Consciousness in contemporary science*. Oxford University Press.

McCullough, G. (1993) The very idea of the phenomenological. *Proceedings of the Aristotelian Society* 93:39–58.

McGinn, C. (1991) *The problem of consciousness*. Blackwell.

——— (1993) Consciousness and cosmology: Hyperdualism ventilated. In: *Consciousness: Psychological and philosophical essays*, ed. M. Davies & G. Humphreys. Blackwell.

Melzack, R., and Wall, P. (1988) *The challenge of pain*, 2d ed. Penguin.

Menzel, E., Savage-Rumbaugh, E., and Lawson, J. (1985) Chimpanzee ({Pantroglogdytes}) spatial problem solving with the use of mirrors and televised equivalents of mirrors. *Journal of Comparative Psychology* 99:211–17.

Meyer, D. E., and Kornblum, S. (1993) *Attention and Performance XIV. Synergies in Experimental Psychology, Artificial Intelligence, and Cognitive Neuroscience*. MIT Press.

Michotte, A. (1946) *The Perception of causality*, translated by Miles, T. R. and E., 1963. London.

Milner, B., and Rugg, M., eds. (1992) *The neuropsychology of consciousness*. Academic Press.

Mitchell, R. W. (1993a) Mental models of mirror self-recognition: Two theories. *New Ideas in Psychology* 11:295–32.

——— (1993b) Recognizing one's self in a mirror? A reply to Gallup and Povinelli, de Lannoy, Anderson, and Byrne. *New Ideas in Psychology* 11:351–77.

Moscovitch, M., Goshen-Gottstein, Y., and Vriezen, E. (1994) Memory without conscious recollection: A tutorial review from a neuropsychological perspective. In: *Attention and performance 15*, ed. C. Umiltà & M. Moscovitch. MIT Press.

Nagel, T. (1971) Brain bisection and the unity of consciousness. *Synthèse* 22:396–413.

——— (1979) *Mortal questions*. Cambridge University Press.

——— (1986) *The view from nowhere*. Oxford University Press.

Nathan, P. (1985) Pain and nociception in the clinical context. *Philosophical Transactions of the Royal Society London B* 308:219–22.

Natsoulas, T. (1993) What is wrong with the appendage theory of consciousness? *Philosophical Psychology* 6(2):137–15.

Nelkin, N. (1993) The connection between intentionality and consciousness. In: *Consciousness: Psychological and philosophical essays*, ed. M. Davies & G. Humphreys. Blackwell.

Nisbett, R., and Wilson, T. (1977) Telling more than we can know: Verbal reports on mental processes. *Psychological Review* 84:231–59.

Peacocke, C. (1983) *Sense and content*. Oxford University Press.

——— (1992) *A study of concepts*. MIT Press.

Pendlebury, M. (1992) Theories of experience. In: *A companion to epistemology*, ed. J. Dancy & E. Sosa. Blackwell.

Penfield, W. (1975) *The mystery of the mind: A critical study of consciousness and the human brain*. Princeton University Press.

Plourde, G. (1993) Clinical use of the 40-Hz auditory steady state response. In: *Depth of anesthesia*, ed. J. G. Jones. Little Brown.

Povinelli, D. (1994) What chimpanzees know about the mind. In: *Behavioral diversity in chimpanzees*. Harvard University Press.

Putnam, H. (1975) The meaning of 'meaning'. In: *Mind, language and reality*, ed. H. Putnam. Cambridge University Press.

Reingold, E., and Merikle, P. (1993) Theory and measurement in the study of unconscious processes. In: *Consciousness: Psychological and philosophical essays*, ed. M. Davies & G. Humphreys. Blackwell.

Rey, G. (1983) A reason for doubting the existence of consciousness. In: *Consciousness and self-regulation*, vol 3., ed. R. Davidson, G. Schwartz & D. Shapiro. Plenum.

——— (1988) A question about consciousness. In: *Perspectives on mind*, ed. H. Otto & J. Tuedio. Reidel.

Rosch, E. (1973) On the internal structure of perceptual and semantic categories in T. E. Moore, ed. *Cognitive Development and the Acquisition of Language* I:111–144.

Rosenthal, D. (1986) Two concepts of consciousness. *Philosophical Studies* 49:329–35.

——— (1993) Thinking that one thinks. In: *Consciousness: Psychological and philosophical essays*, ed. M. Davies & G. Humphreys. Blackwell.

Schacter, D. L. (1989) On the relation between memory and consciousness: Dissociable interactions and conscious experience. In: *Varieties of memory and consciousness: Essays in honour of Endel Tulving*, ed. H. Roediger & F. Craik. Erlbaum.

Schacter, D., McAndrews, M., and Moscovitch, M. (1988) Access to consciousness: Dissociations between implicit and explicit knowledge in neuropsychological syndromes. In: *Thought without Language*, ed., L. Weiskrantz. Oxford University Press.

Schacter, D. L., and Worling, J. R. (1985) Attribute information and the feeling-of-knowing. *Canadian Journal of Psychology* 39(3):467–75.

Searle, J. R. (1983) *Intentionality*. Cambridge University Press.

——— (1990a) Consciousness, explanatory inversion and cognitive science. *Behavioral and Brain Sciences* 13:4:585–95.

——— (1990b) Who is computing with the brain? *Behavioral and Brain Sciences* 13:4:632–64.

——— (1992) *The rediscovery of the mind*. MIT Press.

Sergent, J., and Poncet, M. (1990) From covert to overt recognition of faces in a prosopagnosic patient. *Brain* 113:989–1004.

Shallice, T. (1988a) *From neuropsychology to mental structure*. Cambridge University Press.

——— (1988b) Information-processing models of consciousness: Possibilities and problems. In: *Consciousness in contemporary society*, ed. A. J. Marcel & E. Bisiach. Oxford University Press.

Shevrin, H. (1992) Subliminal perception, memory and consciousness: Cognitive and dynamic perspectives. In: *Perception without awareness*, ed. R. Bornstein & T. Pittman. Guilford Press.

Shiffrin, R. M., and Schneider, W. (1977) Controlled and automatic human information processing: II. Perceptual learning, automatic attending, and a general theory. *Psychological Review* 84:127–90.

Shoemaker, S. (1975) Functionalism and qualia. *Philosophical Studies* 27:291–315.

——— (1981a) Absent qualia are impossible—a reply to Block. *Philosophical Review* 90(4):581–59.

——— (1981b) The inverted spectrum. *Journal of Philosophy* 74(7):357–38.

——— (1993) Lovely and suspect ideas. *Philosophy and Phenomenological Research* 3(4):905–91.

——— (1994) *Phenomenal character*. Nous.

Sperling, G. (1960) The information available in brief visual presentations. *Psychological Monographs* 74:11.

Stich, S. (1978) Beliefs and sub-doxastic states. *Philosophy of Science* 45:499–58.

Tye, M. (1991) *The imagery debate*. MIT Press.

——— (1993) Reflections on Dennett and consciousness. *Philosophy and Phenomenological Research* 34:893–98.

——— (forthcoming a) Blindsight, the absent qualia hypothesis and the mystery of consciousness.

——— (forthcoming b) Does pain lie within the domain of cognitive psychology? In: *Philosophical perspectives*, ed. J. Tomberlin.

Umiltà, C., and Moscovitch, M. (1994) *Attention and performance 15*. MIT Press.

Van Gulick, R. (1989) What difference does consciousness make? *Philosophical Topics* 17(1):211–23.

——— (1993) Understanding the phenomenal mind: Are we all just armadillos? In: *Consciousness: Psychological and philosophical essays*, ed. M. Davies & G. Humphreys. Blackwell.

Weiskrantz, L. (1986) *Blindsight*. Oxford University Press.

——— (1992) Introduction: Dissociated issues. In: *The neuropsychology of consciousness*, ed. B. Milner & M. Rugg. Academic Press.

White, S. L. (1987) What is it like to be an homunculus. *Pacific Philosophical Quarterly* 68:148–17.

——— (1991) Transcendentalism and its discontents. In: *The unity of the self*, ed. S. L. White. MIT Press.

Wiser, M., and Carey, S. (1983) When heat and temperature were one. In: *Mental models*, ed. D. Gentner & A. Stevens. Erlbaum.

Young, A. W. (1994a) Covert recognition. In: *The neuropsychology of higher vision: Collected tutorial essays*, ed. M. Farah & G. Ratcliff. Erlbaum.

——— (1994b) Neuropsychology of awareness. In: *Consciousness in philosophy and cognitive neuroscience*, ed. M. Kappinen & A. Revonsuo. Erlbaum.

Young, A. (1994c) Conscious and non-conscious recognition of familiar faces. In Umiltà and Moscovitch, ed., *Attention and Performance 15*. MIT Press.

Young, A. W., and De Haan, E. (1993) Impairments of visual awareness. In: *Consciousness: Psychological and philosophical essays*, ed. M. Davies & G. Humphreys. Blackwell.