

Below is the unedited preprint (not a quotable final draft) of:

Block, N. (1995). On a confusion about a function of consciousness. *Behavioral and Brain Sciences* 18 (2): 227-287.

The final published draft of the target article, commentaries and Author's Response are currently available only in paper.

€ For information about subscribing or purchasing offprints of the published version, with commentaries and author's response, write to: journals_subscriptions@cup.org (North America) or journals_marketing@cup.cam.ac.uk (All other countries). €

ON A CONFUSION ABOUT A FUNCTION OF CONSCIOUSNESS

Ned Block

Department of Linguistics and Philosophy

Massachusetts Institute of Technology

Cambridge MA 02139

block@psyche.mit.edu

Keywords

access, attention, awareness, blindsight, consciousness, function, retrieval, subjective experience.

Abstract

Consciousness is a mongrel concept: there are a number of very different "consciousnesses." Phenomenal consciousness is experience; the phenomenally conscious aspect of a state is what it is like to be in that state. The mark of access-consciousness, by contrast, is availability for use in reasoning and rationally guiding speech and action. These concepts are often partly or totally conflated, with bad results. This target article uses as an example a form of reasoning about a function of "consciousness" based on the phenomenon of blindsight. Some information about stimuli in the blind field is represented in the brains of blindsight patients, as shown by their correct "guesses," but they cannot harness this information in the service of action, and this is said to show that a function of phenomenal consciousness is somehow to enable information represented in the brain to guide action. But stimuli in the blind field are BOTH access-unconscious and phenomenally unconscious. The fallacy is: an obvious function of the machinery of access-consciousness is illicitly transferred to phenomenal consciousness.

INTRODUCTION

The concept of consciousness is a hybrid or better, a mongrel concept: the word 'consciousness' connotes a number of different concepts and denotes a number of different phenomena. We reason about "consciousness" using some premises that apply to one of the phenomena that fall under "consciousness", other premises that apply to other "consciousnesses" and we end up with trouble.

There are many parallels in the history of science. Aristotle used 'velocity' sometimes to mean average velocity and sometimes to mean instantaneous velocity; his failure to see the distinction caused confusion (Kuhn, 1964). The Florentine Experimenters of the 17th Century used a single word (roughly translatable as "degree of heat") for temperature and for heat, generating paradoxes. For example, when they measured "degree of heat" by whether various heat sources could melt paraffin, heat source A came out hotter than B, but when they measured "degree of heat" by how much ice a heat source could melt in a given time, B was hotter than A (Wiser and Carey, 1983). These are very different cases, but there is a similarity, one that they share with the case of 'consciousness'. The similarity is: very different concepts are treated as a single concept. I think we all have some tendency to make this mistake in the case of "consciousness".

Though the problem I am concerned with appears in many lines of thought about consciousness, it will be convenient to focus on one of them. My main illustration of the kind of confusion I'm talking about concerns reasoning about the [function] of consciousness. But the issue of the function of consciousness is more of the [platform] of this paper than its topic. Because this paper attempts to expose a confusion, it is primarily concerned with reasoning, not with data. Long stretches of text without data may make some readers uncomfortable, as will my fanciful thought-experiments. But if you are interested in consciousness, if I am right you can't afford to lose patience. A stylistic matter: because this paper will have audiences with different concerns, I have adopted the practice of putting items that will mainly be of technical interest to part of the audience in footnotes. Footnotes can be skipped without losing the thread. I now turn to blindsight and its role in reasoning about a function of consciousness.

Patients with damage in primary visual cortex typically have "blind" areas in their visual fields. If the experimenter flashes a stimulus in one of those blind areas and asks the patient what he saw, the patient says "Nothing". The striking phenomenon is that some (but not all) of these patients are able to "guess" reliably about certain features of the stimulus, features having to do with motion, location, direction (e.g. whether a grid is horizontal or vertical); in "guessing", they are able to discriminate some simple forms, if they are asked to grasp an object in the blind field (which they say they can't see), they can shape their hands in a way appropriate to grasping it, and there are some signs of color discrimination. Interestingly, visual acuity (as measured, e.g. by how fine a grating can be detected) increases further from where the patient is looking in blindsight, the opposite of normal sight. (Blindsight was first noticed by Poppel, et.al., 1973 and there is now a huge literature on this and related phenomena. I suggest looking at Bornstein and Pittman, 1992 and Milner and Rugg, 1992.)

Consciousness in some sense is apparently missing (though see McGinn, 1991 p 112 for an argument to the contrary), and with it, the ability to deploy information in reasoning and rational control of action. For example, Tony Marcel (1986) observed that a thirsty blindsight patient would not reach for a glass of water in his blind field. (One has to grant Marcel some "poetic license" in this influential example, since blindsight patients appear to have insufficient form perception in their blind fields to pick out a glass of water.) It is tempting to argue (Marcel, 1986, 1988; Baars 1988; Flanagan 1991, 1992; van Gulick 1989) that since consciousness is missing in blindsight, consciousness must have a function of somehow enabling information represented in the brain to be used in reasoning, reporting and rationally guiding action. I mean the "rationally" to exclude the "guessing" kind of guidance of action that blindsight patients [are] capable of in the case of stimuli presented to the blind field. The idea is that when a content is not conscious--as in the blindsight patient's blindfield perceptual contents, it can influence behavior in various ways, but only when the content is conscious does it play a [rational] role; and so consciousness must be

involved in promoting this rational role.

A related argument is also tempting: Robert van Gulick (1989) and John Searle (1992) discuss Penfield's observations of epileptics who have a seizure while walking, driving or playing the piano. The epileptics continue their activities in a routinized, mechanical way despite, it is said, a total lack of consciousness. Searle says that since both consciousness and also flexibility and creativity of behavior are missing, we can conclude that a function of consciousness is to somehow promote flexibility and creativity. These two arguments are the springboard for this paper. Though some variants of this sort of reasoning have some merit, they are often given more weight than they deserve because of a persistent fallacy involving a conflation of two very different concepts of consciousness.

The plan of the paper is as follows: in the next section, I will briefly discuss some other syndromes much like blindsight, and I will sketch one model that has been offered for explaining these syndromes. Then, in the longest part of the paper I will distinguish the two concepts of consciousness whose conflation is the root of the fallacious arguments. Once that is done, I will sketch what is wrong with the target reasoning and also what is right about it, and I will conclude with some remarks on how it is possible to investigate empirically what the function of consciousness is without having much of an idea about the scientific nature of consciousness.

OTHER SYNDROMES AND SCHACTER'S MODEL

To introduce a second blindsight-like syndrome, I want to first explain a syndrome that is not like blindsight: prosopagnosia ([prosop] for face, [agnosia] = neurological deficit in recognizing). Prosopagnosics are unable visually to recognize their closest relatives--even pictures of themselves, though usually they have no trouble recognizing their friends via their voices, or, according to anecdotal reports, visually recognizing people by recognizing characteristic motions of their bodies. Although there is wide variation from case to case, prosopagnosia is compatible with a high degree of visual ability, even in tasks involving faces.

One patient who has been studied by my colleagues in the Boston area is LH, a Harvard undergraduate who emerged from a car accident with very localized brain damage that left him unable to recognize even his mother. His girl-friend began to wear a special ribbon so that he would know who she was. Now, years later, he still cannot identify his mother or his wife and children from photographs (Etcoff, et.al. 1991). Still, if shown a photo, and asked to choose another photo of the same person from a set of, say, five photos presented simultaneously with the original, LH can do almost as well as normal people despite differences between the target and matching photos in lighting, angle and expression.

Now we are ready for the analog of blindsight. There are many indicators that, in the case of some (but not all) prosopagnosics, the information about whose face is being presented is "in there somewhere". For example, some prosopagnosics are faster at matching simultaneously presented faces when the faces are "familiar" (i.e. faces that the patient has seen often--Reagan, or John Wayne, or the patient's relatives, for example). Another measure involves "semantic priming" in which the presentation of one stimulus facilitates the subject's response to a related stimulus. For example, if normal people are asked to press a button when a familiar face appears in a series of faces rapidly presented one after another, the button tends to be pressed faster if a related name has been presented very recently. E.g. 'Prince Charles' facilitates reactions to Lady Di's face. Likewise, one name primes another, and one face facilitates reactions to the other's 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 in the same pattern as in normals despite the prosopagnosics' insistence that they have no idea who the faces belong to. The phenomenon appears in many experimental paradigms, but I will mention only this: It has recently been discovered (by Sergent and Poncet, 1990) that some prosopagnosics are very good at "guessing" as between two names in the same occupational category ('Reagan' and 'Bush') of a person whose face they claim is unfamiliar. (See Young and de Haan, 1993 and Young, 1994a,b for a description of these phenomena.) Interestingly, LH, the patient mentioned above does not appear to have "covert knowledge" of the people whose faces he sees, but he does appear to have "covert knowledge" of their facial expressions (Etcoff, et. al. 1992).

Many such phenomena in brain-damaged patients have now been explored using the techniques of cognitive and physiological psychology. Further, there are a variety of such phenomena that occur in normals, you and me. For example, suppose that you are given a string of words and asked to count the vowels. This can be done so that you will have no conscious recollection or even recognition of the words and you will be unable to "guess" which words you have seen at a level above chance. However, if I give you a series of word-stems to complete according to your whim, your likelihood of completing 'rea-' as 'reason' is greater if 'reason' is one of the words that you saw, even if you don't recall or recognize it as one of the words you saw. See Bowers and Schacter, 1990, and Reingold and Merikle, 1990. The phenomenon just mentioned is very similar to phenomena involving "subliminal perception", in which stimuli are degraded or presented very briefly. Holender (1986) harshly criticises a variety of "subliminal perception" experiments, but the experimental paradigm just mentioned and many others, are in my judgement, free from the problems of some other studies. Another such experimental paradigm is the familiar dichotic listening experiments in which subjects wear headphones in which different programs are played to different ears. If they are asked to pay attention to one program, they can report only superficial features of the unattended program, but the unattended program influences interpretation of ambiguous sentences presented in the attended program. See Lackner and Garrett, 1973.

Recall that the target reasoning, the reasoning I will be saying is importantly confused (but also importantly right) is that since when consciousness is missing, subjects cannot report or reason about the non-conscious contents or use them to guide action, a function of consciousness is to facilitate reasoning, reporting and guiding action. This reasoning is [partially] captured in a model suggested by Daniel Schacter (1989--see also Schacter, et.al., 1988) in a paper reviewing phenomena such as the ones described above. Figure 1 is derived from Schacter's model.

[THESE ARE FIGURE CAPTIONS ONLY: FIGURES THEMSELVES ARE ONLY AVAILABLE IN THE PAPER VERSION]

The model is only partial (that is, it models some aspects of the mind, but not others), and so may be a bit hard to grasp for those who are used to seeing inputs and outputs. Think of the hands and feet as connected to the Response System box, and the eyes and ears as connected to the specialized modules. (See Schacter, 1989, for some indication of how these suggestions are oversimple.) The key feature of the model is that it contains a box for something called "phenomenal consciousness"; I'll say more about phenomenal consciousness later, but for now, let me just say that phenomenal consciousness is experience; what makes a state phenomenally conscious is that there is something "it is like" (Nagel, 1974) to be in that state. The model dictates that the phenomenal consciousness module has a function: it is the gateway between the special purpose "knowledge" modules and the central Executive system that is in charge of direct control of reasoning, reporting and guiding action. So a function of consciousness on this model includes integrating the outputs of the specialized modules and transmitting the integrated contents to

mechanisms of reasoning and control of action and reporting.

I will be using this model as a focus of discussion, but I hope that my endorsement of its utility as a focus of discussion will not be taken as an endorsement of the model itself. I have no commitment to a single executive system or even to a phenomenal consciousness module. One can accept the idea of phenomenal consciousness as distinct from any cognitive or functional or intentional notion while frowning on a modular treatment of it. Perhaps, for example, phenomenal consciousness is a feature of the whole brain.

Many thinkers will hate any model that treats phenomenal consciousness as something that could be accomplished by a distinct system. See, for example, Dennett and Kinsbourne's (1992b) scorn in response to my suggestion of Cartesian Modularism. I should add that in Dennett's more recent writings, Cartesian materialism has tended to expand considerably from its original meaning of a literal place in the brain at which "it all comes together" for consciousness. In reply to Shoemaker 1993 and Tye 1993, both of whom echo Dennett's (1991) and Dennett's and Kinsbourne's (1992a) admission that no one really is a proponent of Cartesian materialism, Dennett 1993 says "Indeed, if Tye and Shoemaker want to see a card-carrying Cartesian materialist, each may look in the mirror..." See also Jackson 1993. I call that feature Cartesian Modularism, by analogy to the Cartesian Materialism of Dennett and Kinsbourne (1992a), the view that consciousness occupies a literal place in the brain. Modules are individuated by their function, so the point of the box's place between the specialized modules and the Executive system is to indicate that there is a single system that has the function of talking to the specialized modules and integrating their outputs, and talking to the Executive System, passing on information from the specialized modules. But there is an additional point in [calling] that system the phenomenal consciousness system, namely to say that phenomenal consciousness is somehow involved in performing that function. The idea is that phenomenal consciousness [really does] something, it is involved somehow in powering the wheels and pulleys of access to the Executive system. This is a substantive claim, one that is distinct from the claims that phenomenal consciousness is [correlated] with that information processing function, or that phenomenal consciousness should be [identified] with that information processing function. The idea is that phenomenal consciousness is distinct (at least conceptually) from that information processing function, but is part of the implementation of it.

Martha Farah (1994) criticizes this model on the ground that we don't observe patients whose blindsight-like performance is up to the standard of normal vision. Blindsight and its analogs are always degraded in discriminatory capacity. Her assumption seems to be that if there is a phenomenal consciousness module, it could simply be by-passed without decrement in performance; and the fact that this is not observed is taken as reason to reject the phenomenal consciousness module. She appears to think that if there is a phenomenal consciousness module, then phenomenal consciousness [doesn't do any information processing] (except, I guess, for determining reports of phenomenal consciousness), for otherwise why assume that it could be bypassed without decrement in performance. But why assume that? For example, phenomenal consciousness might be like the water in a hydraulic computer. You don't expect the computer to just work normally without the water. Even if there could be an electrical computer that is isomorphic to the hydraulic computer but works without water, one should not conclude that the water in the hydraulic system does nothing. I will return to this issue later.

One reason that many philosophers would hate Cartesian Modularist models is that such models may be regarded as licensing the possibility of "zombies", creatures which have information processing that is the same as ours but which have no phenomenal consciousness. If the phenomenal consciousness module could be replaced by a device that had the same information

processing effects on the rest of the system, but without phenomenal consciousness, the result would be a zombie. My view is that we now know so little about the scientific nature of phenomenal consciousness and its function that we cannot judge whether the same function could be performed by an ersatz phenomenal consciousness module--that is, whether an ersatz phenomenal consciousness module could inject its representations with ersatz conscious content that would affect information processing the same way as real conscious content. There is much of interest to be said about this idea and its relation to other ideas that have been mentioned in the literature, but I have other fish to fry, so I leave the matter for another time.

The information processing function of phenomenal consciousness in Schacter's model is the ground of the concept of consciousness that I will mainly be contrasting with phenomenal consciousness, what I will call "access-consciousness". A perceptual state is access-conscious roughly speaking if its content--what is represented by the perceptual state--is processed via that information processing function, that is, if its content gets to the Executive system, whereby it can be used to control reasoning and behavior.

Schacter's model is useful for my purposes both because it can be used to illustrate the contrast between phenomenal and access-consciousness and because it allows us to see one possible explanation of the "covert knowledge" syndromes just described. This explanation (and also Schacter's model) are certainly incomplete and no doubt wildly oversimple at best, but it is nonetheless useful to see the rough outlines of how an account might go. In addition, there is an association between Schacter's model and the target reasoning--though as we shall see there is another processing model that perhaps better embodies the target reasoning.

Consider a blindsight patient who has just had a vertical line displayed in his blind field. "What did you see?" "Nothing", says the patient. "Guess as between a vertical and a horizontal line" says the experimenter. "Vertical" says the patient, correctly. Here's a story about what happened. One of the specialized modules is specialized for spatial information; it has some information about the verticality of the stimulus. The pathways between this specialized module and the phenomenal consciousness system have been damaged, creating the "blind field", so the patient has no phenomenally conscious experience of the line, and hence his Executive system has no information about whether the line is vertical or horizontal. But the specialized module has a direct connection to the response system, so when the subject is given a binary choice, the specialized module can somehow directly affect the response. Similarly, there is a specialized module for face information, which can have some information about the face that has been presented to a prosopagnosic. If the prosopagnosia is caused by damage in the link between the face module and the phenomenal consciousness system, then that prevents the face information from being phenomenally conscious, and without phenomenal consciousness, the Executive system does not get the information about the face. When the prosopagnosic guesses as between 'Reagan' and 'Bush', the face module somehow directly controls the response. (It is assumed that the face-module has information about people--e.g. their names--linked to representations of their faces.) It is interesting in this regard that the patients who do best in these experiments are the ones judged to be the most "passive" (Marcel, 1983, p. 204; Weiskrantz, 1988). One can speculate that in a laid-back subject, the Executive does not try out a guessing strategy, and so peripheral systems are more likely to affect the response.

Alexia is a neurological syndrome whose victims can no longer read a word "at a glance", but can only puzzle out what word they have seen at a rate of, e.g. a second per letter. Nonetheless, these subjects often show various kinds of understanding of the meanings of words that have been flashed far too briefly for them to read in their laborious way. The idea, once again, is that one of

the specialized module is specialized for lexical information, and this module has information about words that the subject cannot consciously read. This information somehow affects responses. Landis et al. (1980) report that such a patient actually became worse at "guesses" having to do with the meanings of "unread" words as his explicit reading ability improved (Young and de Haan, 1993). Again, perhaps once the Executive has more information, it "takes over", preventing peripheral systems from controlling responses. Coslett and Saffran (1994) report that alexics did worse at "guessing" words with longer exposures. An exposure of 250 ms was better than an exposure of 2 sec. Again, longer exposures may give the executive system a chance to try to read letter by letter.

Schacter's model and the explanation I have just sketched are highly speculative; my purposes in appealing to them are heuristic.

TWO CONCEPTS OF CONSCIOUSNESS

First, consider phenomenal consciousness, or P-consciousness, as I will call it. Let me acknowledge at the outset that I cannot define P-consciousness in any remotely non-circular way. I don't consider this an embarrassment. The history of reductive definitions in philosophy should lead one not to expect a reductive definition of anything. But the best one can do for P-consciousness is in some respects worse than for many other things because really all one can do is [point] to the phenomenon (cf. Goldman, 1993a). Nonetheless, it is important to point properly. John Searle, acknowledging that consciousness cannot be defined non-circularly, defines it as follows: < By consciousness I simply mean those subjective states of awareness or sentience that begin when one wakes in the morning and continue throughout the period that one is awake until one falls into a dreamless sleep, into a coma, or dies or is otherwise, as they say, unconscious. [This comes from Searle 1990b; there is a much longer attempt along the same lines in his 1992, p. 83ff.]> I will argue that this sort of pointing is flawed because it points to too many things, too many different consciousnesses.

So how should we point to P-consciousness? Well, one way is via rough synonyms. As I said, P-consciousness is experience. P-conscious properties are experiential properties. P-conscious states are experiential states, that is, a state is P-conscious if it has experiential properties. The totality of the experiential properties of a state are "what it is like" to have it. Moving from synonyms to examples, we have P-conscious states when we see, hear, smell, taste and have pains. P-conscious properties include the experiential properties of sensations, feelings and perceptions, but I would also include thoughts, wants and emotions. But what is it about thoughts that makes them P-conscious? One possibility is that it is just a series of mental images or subvocalizations that make thoughts P-conscious. Another possibility is that the contents themselves have a P-conscious aspect independently of their vehicles. See Lormand, forthcoming. A feature of P-consciousness that is often missed is that differences in intentional content often make a P-conscious difference. What it is like to hear a sound as coming from the left differs from what it is like to hear a sound as coming from the right. P-consciousness is often representational. (See Jackendoff, 1987; van Gulick, 1989; McGinn, 1991, Ch 2; Flanagan, 1992, Ch 4; Goldman, 1993b.) So far, I don't take myself to have said anything terribly controversial. The controversial part is that I take P-conscious properties to be distinct from any cognitive, intentional, or functional property. (Cognitive = essentially involving thought; intentional properties = properties in virtue of which a representation or state is about something; functional properties = e.g. properties definable in terms of a computer program. See Searle, 1983 on intentionality; See Block, 1980, 1994a for better characterizations of a functional property.) But I am trying hard to limit the controversiality of my assumptions. Though I will be assuming that functionalism about

P-consciousness is false, I will be pointing out that limited versions of many of the points I will be making can be acceptable to the functionalist. I say both that P-consciousness is not an intentional property and that intentional differences can make a P-conscious difference. My view is that although P-conscious content cannot be reduced to intentional content, P-conscious contents often have an intentional aspect, and also P-conscious contents often represent in a primitive non-intentional way. A perceptual experience can represent space as being filled in certain ways without representing the object perceived [as] falling under any concept. Thus, the experiences of a creature which does not possess the concept of a donut could represent space as being filled in a donut-like way. See Davies (1992, forthcoming), Peacocke (1992), and finally Evans (1982), in which the distinction between conceptualized and non-conceptualized content is first introduced.

It is of course P-consciousness rather than access-consciousness or self-consciousness that has seemed such a scientific mystery. The magazine [Discover] (November, 1992) devoted an issue to the ten great unanswered questions of science, such as "What is Consciousness?", "Does Chaos Rule the Cosmos?" and "How Big is the Universe?" The topic was P-consciousness, not, e.g. self-consciousness.

By way of homing in on P-consciousness, it is useful to appeal to what may be a contingent property of it, namely the famous "explanatory gap". To quote T.H. Huxley (1866), "How it is that anything so remarkable as a state of consciousness comes about as a result of irritating nervous tissue, is just as unaccountable as the appearance of Djin when Aladdin rubbed his lamp." Consider a famous neurophysiological theory of P-consciousness offered by Francis Crick and Christoff Koch: namely, that a synchronized 35-75 hertz neural oscillation in the sensory areas of the cortex is at the heart of phenomenal consciousness. No one has produced the concepts that would allow us to explain why such oscillations might be the physiological basis of phenomenal consciousness.

However, Crick and Koch have offered a sketch of an account of how the 35-75 hertz oscillation might contribute to a solution to the "binding problem". Suppose one simultaneously sees a red square moving to the right and a blue circle moving to the left. Different areas of the visual cortex are differentially sensitive to color, shape, motion, etc. so what binds together redness, squareness and rightward motion? That is, why don't you see redness and blueness without seeing them as belonging with particular shapes and particular motions? And why aren't the colors normally seen as bound to the wrong shapes and motions? Representations of colors, shapes and motions of a single object are supposed to involve oscillations that are in phase with one another but not with representations of other objects. But even if the oscillation hypothesis deals with the informational aspect of the binding problem (and there is some evidence against it), how does it explain [what it is like to see something as red in the first place]--or for that matter, as square or as moving to the right? Why couldn't there be brains functionally or physiologically just like ours, including oscillation patterns, whose owners' experience was different from ours or who had no experience at all? (Note that I don't say that there [could be] such brains. I just want to know [why not.]) And why is it a 35-75 hertz oscillation--as opposed to some other frequency--that underlies experience? If the synchronized neural oscillation idea pans out as a solution to the binding problem, no doubt there will be some answer to the question of why [those] frequencies, as opposed to, say 110 hertz, are involved. But will that answer explain why 110 hertz oscillations don't underly experience? No one has a clue how to answer these questions. Levine (1983) coined the term "explanatory gap", and has elaborated the idea in interesting ways; see also his (1993). Van Gulick (1993) and Flanagan (1992, p. 59) note that the more we know about the connection between (say) hitting middle C on the piano and the resulting experience, the more we have in the way of hooks on

which to hang something that could potentially close the explanatory gap. Some philosophers have adopted what might be called a deflationary attitude towards the explanatory gap. See Levine (1993), Jackson (1993) and Chalmers (1993), Byrne (1993) and Block (1994).

The explanatory gap in the case of P-consciousness contrasts with our relatively good understanding of cognition. We have two serious research programs into the nature of cognition, the classical "language of thought" paradigm, and the connectionist research program. Though no doubt there are many ideas missing in our understanding of cognition, we have no difficulty seeing how pursuing one or both of these research programs could lead to an adequate theoretical perspective on cognition. But it is not easy to see how current approaches to P-consciousness [could] yield an account of it. Indeed, what passes for research programs on consciousness just [is] a combination of cognitive psychology and explorations of neuropsychological syndromes that contain no theoretical perspective on what P-consciousness actually is.

I mentioned the explanatory gap partly by way of pointing at P-consciousness: [that's] the entity to which the mentioned explanatory gap applies. Perhaps this identification is contingent; at some time in the future, when we have the concepts to conceive of much more about the explanation of P-consciousness, this may not be a way of picking it out. (See McGinn (1991) for a more pessimistic view.)

What I've been saying about P-consciousness is of course controversial in a variety of ways, both for some advocates and some opponents of some notion of P-consciousness. I have tried to steer clear of some controversies, e.g. controversies over inverted and absent qualia; over Jackson's (1986) Mary (the woman who is raised in a black and white room, learning all the physiological and functional facts about the brain and color vision, but nonetheless discovers a new fact when she goes outside the room for the first time and learns what it is like to see red); and even Nagel's view that we cannot know what it is like to be a bat. I know some will think that I invoked inverted and absent qualia a few paragraphs above when I described the explanatory gap as involving the question of why a creature with a brain with a physiological and functional nature like ours couldn't have different experience or none at all. But the spirit of the question as I asked it allows for an answer that explains why such creatures cannot exist, and thus there is no presupposition that these are real possibilities. Levine (1983, 1993) stresses that the relevant modality is epistemic possibility. Even if you think that P-consciousness as I have described it is an incoherent notion, you may be able to agree with the main point of this paper, which is that a great deal of confusion arises as a result of confusing P-consciousness with something else. Not even the concept of what time it is now on the sun is so confused that it cannot itself be confused with something else.

ACCESS-CONSCIOUSNESS

I now turn to the non-phenomenal notion of consciousness that is most easily and dangerously conflated with P-consciousness: access-consciousness. I will characterize access-consciousness, give some examples of how it is at least possible to have access-consciousness without phenomenal consciousness and vice versa, and then go on to the main theme of the paper, the damage done by conflating the two.

A state is access-conscious (A-conscious) if, in virtue of one's having the state, a representation of its content is (1) inferentially promiscuous (Stich, 1978), i.e. poised to be used as a premise in reasoning, and (2) poised for [rational] control of action and (3) poised for rational control of speech. (I will speak of both states and their contents as A-conscious.) These three conditions are

together sufficient, but not all necessary. I regard (3) as not necessary (and not independent of the others), since I want to allow non-linguistic animals, e.g. chimps, to have A-conscious (access-conscious) states. I see A-consciousness as a cluster concept, in which (3)--roughly, reportability--is the element of the cluster with the smallest weight, though (3) is often the best practical guide to A-consciousness. What if an A-[un]conscious state causes an A-conscious state with the same content? Then it could be said that the first state must be A-conscious because it is in virtue of having [that] state that the content it shares with the other state satisfies the three conditions. So the state is A-unconscious by hypothesis, but A-conscious by my definition. (I am indebted to Paul Horwich.) I think what this case points to is a refinement needed in the notion of "in virtue of". One does not want to count the inferential promiscuity of a content as being in virtue of having a state if that state can only cause this inferential promiscuity via another state. I won't try to produce an analysis of 'in virtue of' here.

Although I make a firm distinction between A-consciousness and P-consciousness, I also want to insist that they interact. For example, what perceptual information is being accessed can change figure to ground and conversely, and a figure-ground switch can affect one's phenomenal state. For example, attending to the feel of the shirt on your neck, accessing those perceptual contents, switches what was in the background to the foreground, thereby changing one's phenomenal state. (See Hill, 1991, 118-126; Searle, 1992.)

After further explicating A-consciousness, I will argue that A-consciousness plays a deep role in our ordinary 'consciousness' talk and thought. However, I must admit at the outset that this role allows for substantial indeterminacy in the concept itself. In addition, there are some loose ends in the characterization of the concept which cannot be tied up without deciding about certain controversial issues, to be mentioned below. I have been using the P-consciousness/A-consciousness distinction in my lectures for many years, but it only found its way into print in my "Consciousness and Accessibility" (1990b), and my (1991, 1992, 1993). My claims about the distinction have been criticized in Searle (1990b, 1992) and Flanagan (1992); and there is an illuminating discussion in Humphreys and Davies (1993b), a point of which will be taken up in a footnote to follow. See also Levine's (1994) review of Flanagan which discusses Flanagan's critique of the distinction. See also Kirk (1992) for an identification of P-consciousness with something like A-consciousness. My guide in making precise the A-consciousness/P-consciousness distinction is the purpose of the moment, namely to reveal the fallacy in the target reasoning. The target reasoning (in one form) says that the blindsight patient lacks consciousness of stimuli in the blind field, and that is why he does not use information he actually has about these stimuli, so the function of consciousness must be to harness information for use in guiding action. (Maybe the blindsight patient does not lack P-consciousness of these stimuli, but the target reasoning supposes it, and it is independently plausible. For example, Cowie and Stoerig, 1992, point out that the removal of primary visual cortex in these patients disrupts the Crick and Koch 40 hertz oscillations. That is some reason to believe that the blindsight patient lacks P-consciousness of the stimuli.) I will be pointing out that something [else] is also problematic in blindsight that can equally well be blamed for the blindsight patient's failure, namely the machinery of A-consciousness. Of course, the missing P-consciousness may be responsible for the missing A-consciousness; no fallacy is involved in that hypothesis. Rather, the fallacy is [sliding] from an obvious function of A-consciousness to an un-obvious function of P-consciousness.) For that reason, I choose to adopt a notion of access on which the blindsight patient's guesses don't count as access. There is no right or wrong here. Access comes in various degrees and kinds, and my choice here is mainly determined by the needs of the argument. (I also happen to think that the notion I characterize is more or less the one that plays a big role in our thought, but that won't

really be a factor in my argument.)

I will mention three main differences between P-consciousness and A-consciousness. The first point, [put crudely], is that P-conscious content is phenomenal, whereas A-conscious content is representational. It is of the essence of A-conscious content to play a role in reasoning, and only representational content can figure in reasoning. The reason this way of putting the point is crude is that many phenomenal contents are [also] representational. So what I really want to say is that it is in virtue of its phenomenal content or the phenomenal aspect of its content that a state is P-conscious, whereas it is in virtue of its representational content, or the representational aspect of its content that a state is A-conscious. Some may say that only fully conceptualized content can play a role in reasoning, be reportable, and rationally control action. If so, then non-conceptualized content is not A-conscious.

(In the last paragraph, I used the notion of P-conscious [content]. The P-conscious content of a state is the totality of the state's experiential properties, what it is like to be in that state. One can think of the P-conscious content of a state as the state's experiential "value" by analogy to the representational content as the state's representational "value". In my view, the content of an experience can be both P-conscious and A-conscious; the former in virtue of its phenomenal feel and the latter in virtue of its representational properties.)

A closely related point: A-conscious states are necessarily transitive: A-conscious states must always be states of consciousness [of]. P-conscious states, by contrast, sometimes are and sometimes are not transitive. P-consciousness, as such, is not consciousness of. (I'll return to this point in a few paragraphs.)

Second, A-consciousness is a functional notion, and so A-conscious content is system-relative: what makes a state A-conscious is what a representation of its content does in a system. P-consciousness is not a functional notion. However, I acknowledge the empirical possibility that the scientific nature of P-consciousness is something to do with information processing. We can ill afford to close off empirical possibilities given the difficulty of solving the mystery of P-consciousness. Cf. Loar, 1990. In terms of Schacter's model, content gets to be P-conscious because of what happens [inside] the P-consciousness module. But what makes content A-conscious is not anything that could go on [inside] a module, but rather informational relations [among] modules. Content is A-conscious in virtue of (a representation with that content) reaching the Executive system, the system that is in charge of rational control of of action and speech, and to that extent, we could regard the Executive module as the A-consciousness module; but to regard [anything] as an A-consciousness module is misleading, because what makes content A-conscious depends on informational relations between the Executive and other modules.

A third difference is that there is such a thing as a P-conscious [type] or [kind] of state. For example the feel of pain is a P-conscious type--every pain must have that feel. But any particular token thought that is A-conscious at a given time could fail to be accessible at some other time, just as my car is accessible now, but will not be later when my wife has it. A state whose content is informationally promiscuous now may not be so later.

The paradigm P-conscious states are sensations, whereas the paradigm A-conscious states are "propositional attitude" states like thoughts, beliefs and desires, states with representational content expressed by "that" clauses. (E.g. the thought that grass is green.) However, as I said, thoughts often are P-conscious and perceptual experiences often have representational content. For example, a perceptual experience may have the representational content [that there is a red square

in front of me]. Even pain typically has [some] kind of representational content. Pains often represent something (the cause of the pain? the pain itself?) as somewhere (in the leg). A number of philosophers have taken the view that the content of pain is [entirely] representational. (See Dretske, 1994; Shoemaker, 1994; Tye, forthcoming-b.) I don't agree with this view, so I certainly don't want to rely on it here, but I also don't want to make the existence of cases of P-consciousness without A-consciousness any kind of trivial consequence of an idiosyncratic set of definitions. To the extent that representationalism of the sort just mentioned is plausible, one can regard a pain as A-conscious if its representational content is inferentially promiscuous, etc. Alternatively, we could take the A-conscious content of pain to consist in the content that one has a pain or that one has a state with a certain phenomenal content. On my view, there are a number of problems with the first of these suggestions. One of them is that perhaps the representational content of pain is [too primitive] for a role in inference. Arguably, the representational content of pain is non-conceptualized. After all, dogs can have pain and one can reasonably wonder whether dogs have the relevant concepts at all. Davies and Humphreys (1993b) discuss a related issue. Applying a suggestion of theirs about the higher order thought notion of consciousness to A-consciousness, we could characterize A-consciousness of a state with non-conceptualized content as follows: such a state is A-conscious if, in virtue of one's having the state, its content [would be] inferentially promiscuous and available for rational control of action and speech [if] the subject [were to have had] the concepts required for that content to be a conceptualized content. The idea is to bypass the inferential disadvantage of non-conceptualized content by thinking of its accessibility [counterfactually]--in terms of the rational relations it would have if the subject [were] to have the relevant concepts. See Lormand (forthcoming) on the self-representing nature of pain.

There is a familiar distinction, alluded to above, between 'consciousness' in the sense in which we speak of a state as being a conscious state (intransitive consciousness) and consciousness of] something (transitive consciousness). (See, for example, Rosenthal, 1986. Humphrey (1992) mentions that the intransitive usage is much more recent, only 200 years old.) It is easy to fall into an identification of P-consciousness with intransitive consciousness and a corresponding identification of

access-consciousness with transitive consciousness. Such an identification is oversimple. As I mentioned earlier, P-conscious contents can be representational. Consider a perceptual state of seeing a square. This state has a P-conscious content that represents something, a square, and thus it is a state of P-consciousness of] the square. It is a state of P-consciousness of the square even if it doesn't represent the square [as] a square, as would be the case if the perceptual state is a state of an animal that doesn't have the concept of a square. Since there can be P-consciousness [of] something, P-consciousness is not to be identified with intransitive consciousness.

Here is a second reason why the transitive/intransitive distinction cannot be identified with the P-consciousness/A-consciousness distinction: The [of]-ness required for transitivity does not guarantee that a content be utilizable by a consuming] system at the level required for A-consciousness. For example, a perceptual state of a brain-damaged creature might be a state of P-consciousness of, say, motion, even though connections to reasoning and rational control of action are damaged so that the state is not A-conscious. In sum, P-consciousness can be consciousness of, and consciousness of need not be A-consciousness. Later in this paper I introduce the distinction between creature consciousness and state consciousness. In those terms, transitivity has to do primarily with creature consciousness, whereas in the case of P-consciousness and A-consciousness, it is state consciousness which is basic. See the discussion

at the end of this section.

[A-CONSCIOUSNESS WITHOUT P-CONSCIOUSNESS] Since the main point of this paper is that these two concepts of consciousness are easily confused, it will pay us to consider conceptually possible cases of one without the other. Actual cases will be more controversial.

First, I will give some examples of A-consciousness without P-consciousness. If there could be a full-fledged phenomenal zombie, say a robot computationally identical to a person, but whose silicon brain did not support P-consciousness, that would do the trick. I think such cases conceptually possible, but this is very controversial, and I am trying to avoid controversy. (See Shoemaker, 1975, 1981)

But there is a less controversial kind of case, a very limited sort of partial zombie. Consider the blindsight patient who "guesses" that there is an 'X' rather than an 'O' in his blind field. Taking his word for it, I am assuming that he has no P-consciousness of the 'X'. As I mentioned, I am following the target reasoning here, but as I will point out later, my own argument does not depend on this assumption. I am certainly [not] assuming that lack of A-consciousness guarantees lack of P-consciousness--that is, I am not assuming that if you don't say it you haven't got it.

The blindsight patient also has no 'X'-representing A-conscious content, because although the information that there is an 'X' affects his "guess", it is not available as a premise in reasoning (until he has the quite distinct state of hearing and believing his own guess), or for rational control of action or speech. Recall Marcel's point that the thirsty blindsight patient would not reach for a glass of water in the blind field. So the blindsight patient's perceptual or quasi-perceptual state is unconscious in the phenomenal [and] access senses ([and] in the monitoring senses to be mentioned below too).

Now imagine something that may not exist, what we might call [super-blindsight]. A real blindsight patient can only guess when given a choice from a small set of alternatives ('X'/'O'; horizontal/vertical, etc). But suppose--interestingly, apparently contrary to fact--that a blindsight patient could be trained to prompt himself at will, guessing what is in the blind field without being told to guess. The super-blindsighter spontaneously says "Now I know that there is a horizontal line in my blind field even though I don't actually see it." Visual information from his blind field simply pops into his thoughts in the way that solutions to problems we've been worrying about pop into our thoughts, or in the way some people just know the time or which way is North without having any perceptual experience of it. The super-blindsighter himself contrasts what it is like to know visually about an 'X' in his blind field and an 'X' in his sighted field. There is something it is like to experience the latter, but not the former, he says. It is the difference between [just knowing] and knowing via a visual experience. Taking his word for it, here is the point: the content that there is an 'X' in his visual field is A-conscious but not P-conscious. The super-blindsight case is a very limited partial zombie. The, forthcoming-a argues (on the basis of neuropsychological claims) that the visual information processing in blindsight includes no processing by the object recognition system or the spatial attention system, and so is very different from the processing of normal vision. This point does not challenge my claim that the super-blindsight case is a very limited partial zombie. Note that super-blindsight, as I describe it does not require object recognition or spatial attention. Whatever it is that allows the blindsight patient to discriminate an 'X' from an 'O' and a horizontal from a vertical line will do. I will argue later that the fact that such cases do not exist, if it is a fact, is important. Humphrey (1992) suggests that blindsight is mainly a motor phenomenon--the patient is perceptually influenced by his own motor tendencies.

Of course, the super-blindsighter has a [thought] that there is an x in his blind field that is [both] A-conscious and P-conscious. but I am not talking about the thought. Rather, I am talking about the state of his perceptual system that gives rise to the thought. It is this state that is A-conscious without being P-conscious. If you are tempted to deny the existence of these states of the perceptual system, you should think back to the total zombie just mentioned. Putting aside the issue of the possibility of this zombie, note that on a computational notion of cognition, the zombie has [all] the same A-conscious contents that you have (if he is your computational duplicate). A-consciousness is an informational notion. The states of the super-blindsighter's perceptual system are A-conscious for the same reason as the zombie's.

Is there [actually] such a thing as super-blindsight? Humphrey (1992) describes a monkey (Helen) who despite [near] total loss of the visual cortex could nonetheless act in a somewhat visually normal way in certain circumstances, without any "prompting". One reason to doubt that Helen is a case of super-blindsight is that Helen may be a case of [sight]. There was some visual cortex left, and the situations in which she showed unprompted visual discrimination were ones in which there was no control of where the stimuli engaged her retina. Another possibility mentioned by Cowie and Stoerig (1992--attributed to an unpublished paper by Humphrey), is that there were P-conscious sensory events, though perhaps auditory in nature. Helen appeared to confuse brief tones with visual stimuli. Cowie and Stoerig propose a number of ways of getting information out of monkeys that are close to what we get out of blindsighted humans. Weiskrantz (1992) mentions that a patient GY sometimes knows that there is a stimulus (though not what it is) without, he says, seeing anything. But GY also seems to be having some kind of P-conscious sensation. See Cowie and Stoerig (1992.)

The (apparent) non-existence of super-blindsight is a striking fact, one that a number of writers have noticed. Indeed, it is the basis for the target reasoning. After all, what Marcel was in effect pointing out was that the blindsight patients, in not reaching for a glass of water, are not super-blindsighters. As I mentioned, Farah (1994) says that blindsight (and blind perception generally) turns out always to be degraded. In other words, blind perception is never superblind perception. Actually, my notion of A-consciousness seems to fit the data better than the conceptual apparatus she uses. Blindsight isn't always more degraded in any normal sense than sight. 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. It would be better to understand her "degraded" in terms of lack of access.

Notice that the super-blindsighter I have described is just a little bit different (though in a crucial way) from the ordinary blindsight patient. In particular, I am [not relying] on what might be thought of as a full-fledged [quasi-zombie], a super-[duper]-blindsighter whose blindsight is [every bit] as good, functionally speaking, as his sight. In the case of the super-duper blindsighter, the [only] difference between vision in the blind and sighted fields, functionally speaking, is that the quasi-zombie himself regards them differently. Such an example will be regarded by some (though not me) as incoherent--see Dennett, 1991, for example. But we can avoid disagreement about the super-duper-blindsighter by illustrating the idea of A-consciousness without P-consciousness by appealing only to the super-blindsighter. Functionalists may want to know why the super-blindsight case counts as A-conscious without P-consciousness. After all, they may say, if we have [really high quality access] in mind, the super-blindsighter that I have described does not have it, so he lacks [both] P-consciousness and really high quality A-consciousness. The super-duper-blindsighter, on the other hand, [has] both, according to the functionalist, so in neither case, the objection goes, is there A-consciousness without P-consciousness. But the disagreement about the super-duper-blindsighter is irrelevant to the issue about the super-blindsighter, and the

issue about the super-blindsighter is merely verbal. I have chosen a notion of A-consciousness whose standards are lower in part to avoid conflict with the functionalist. I believe in the possibility of a quasi-zombie like the super-duper-blindsighter, but the point I am making here does not depend on it. There is no reason to frame notions so as to muddy the waters with unnecessary conflicts when the point I am making in this paper is one that functionalists can have some agreement with. One could put the point by distinguishing three types of access: (1) really high quality access, (2) medium access and (3) poor access. The [actual] blindsight patient has poor access, the super-blindsight patient has medium access and the super-duper blindsight patient--as well as most of us--has really high quality access. The functionalist identifies P-consciousness with A-consciousness of the really high quality kind. I am [defining] 'A-consciousness'--and of course, it is only one of many possible definitions--in terms of medium access, both to avoid unnecessary conflict with the functionalist, and also so as to reveal the fallacy of the target reasoning. I choose medium instead of really high quality access for the former purpose, and I choose medium instead of poor access for the latter purpose. Though functionalists should agree with me that there can be A-consciousness without P-consciousness, some functionalists will see the significance of such cases very differently from the way I see them. Some functionalists will see the distinction between A-consciousness and P-consciousness as primarily a difference in degree rather than a difference in kind, as is suggested by the contrast between really high quality access and medium access. So all that A-consciousness without P-consciousness illustrates, on this functionalist view, is some access without more access. Other functionalists will stress [kind] of information processing rather than amount of it. The thought behind this approach is that there is no reason to think that the P-consciousness of animals whose capacities for reasoning, reporting and rational guidance of action are more limited than ours thereby have anything less in the way of P-consciousness. The functionalist can concede that this thought is correct, and thereby treat the difference between A-consciousness and P-consciousness as a difference of kind, albeit kind of information processing.

I don't know whether there are any actual cases of A-consciousness without P-consciousness, but I hope that I have illustrated their conceptual possibility.

[P-CONSCIOUSNESS WITHOUT A-CONSCIOUSNESS] Consider an animal that you are happy to think of as having P-consciousness for which brain damage has destroyed centers of reasoning and rational control of action, thus preventing A-consciousness. It certainly seems [conceptually possible] that the neural bases of P-consciousness systems and A-consciousness systems be distinct, and if they are distinct, then it is possible, at least conceptually possible, for one to be damaged while the other is working well. Evidence has been accumulating for twenty-five years that the primate visual system has distinct dorsal and ventral subsystems. Though there is much disagreement about the specializations of the two systems, it does appear that much of the information in the ventral system is much more closely connected to P-consciousness than information in the dorsal system (Goodale and Milner, 1992). So it may actually be possible to damage A-consciousness without P-consciousness and conversely. Thus, there is a conflict between this physiological claim and the Schacter model which dictates that destroying the P-consciousness module will prevent A-consciousness.

Further, one might suppose (Rey, 1983, 1988; White, 1987) that some of our own subsystems--say each of the two hemispheres of the brain--might themselves be separately P-conscious. Some of these subsystems might also be A-consciousness, but other subsystems might not have sufficient machinery for reasoning or reporting or rational control of action to allow their P-conscious states to be A-conscious; so if those states are not accessible to another system that does have adequate

machinery, they will be P-conscious but not A-conscious.

Here is another reason to believe in P-consciousness without A-consciousness: Suppose that you are engaged in intense conversation when suddenly at noon you realize that right outside your window, there is--and has been for some time--a deafening pneumatic drill digging up the street. You were aware of the noise all along, but only at noon are you [consciously aware] of it. That is, you were P-conscious of the noise all along, but at noon you are both P-conscious [and] A-conscious of it. Of course, there is a very similar string of events in which the crucial event at noon is a bit more intellectual. In this alternative scenario, at noon you realize not just that there is and has been a noise, but also that [you are now and have been experiencing] the noise. In this alternative scenario, you get "higher order thought" as well as A-consciousness at noon. So on the first scenario, the belief that is acquired at noon is that there is and has been a noise, and on the second scenario, the beliefs that are acquired at noon are the first one plus the belief that you are and have been experiencing the noise. But it is the first scenario, not the second that interests me, for it is a pure case of P-consciousness without A-consciousness. Note that this case involves a natural use of 'conscious' and 'aware' for A-consciousness and P-consciousness, respectively. 'Conscious' and 'aware' are more or less synonymous, so calling the initial P-consciousness 'awareness' makes it natural to call the later P-consciousness plus A-consciousness 'conscious awareness'. Of course I rely here on introspection, but when it comes to P-consciousness, introspection is an important source of insight. (There is a misleading aspect to this example--namely that to the extent that 'conscious' and 'aware' differ in ordinary talk, the difference goes in the opposite direction.) This case of P-consciousness without A-consciousness exploits what William James (1890) called "secondary consciousness", a category that he meant to include cases of P-consciousness without attention. Of course, even those who don't believe in P-consciousness at all, as distinct from A-consciousness, can accept the distinction between a noise that is A-conscious and a noise that is not A-conscious. There is a more familiar situation which illustrates the same points. Think back to all those times when you have been sitting in the kitchen when suddenly the compressor in the refrigerator goes off. Again, one might naturally say that one was aware of the noise, but only at the moment in which it went off was one consciously aware of it. I didn't use this example because I am not sure that one really has P-consciousness of the noise of the compressor all along; habituation would perhaps prevent it. Perhaps what happens at the moment it goes off is that one is P-conscious of the change only.

I have found that the argument of the last paragraph makes those who are distrustful of introspection uncomfortable. I agree that introspection is not the last word, but it is the first word, when it comes to P-consciousness. The example shows the conceptual distinctness of P-consciousness from A-consciousness and it also puts the burden of proof on anyone who would argue that as a matter of empirical fact they come to the same thing.

The difference between different concepts of consciousness gives rise to different types of [zombie]. We have already encountered the phenomenal zombies that appear in science-fiction and philosophers' examples--the familiar computers and robots that think but don't feel. Their states are A-conscious, but not P-conscious. However, our culture also acknowledges the concept of voodoo zombies and zombies in [Night of the Living Dead]. If we find that voodoo zombies are cognitively or affectively diminished, say without will, rather than phenomenally diminished, we would not decide that they were not zombies after all. And on seeing the next installment in the "Living Dead" series, we would not feel that our concept of a zombie had been toyed with if it turned out that there is something it is like for these zombies to eat their relatives. (They say "Yumm!") No doubt we have no very well formed zombie-concept, but the considerations just

mentioned motivate the view that a zombie is something that is mentally dead in one respect or another, and the different respects give rise to different zombies.

Kathleen Akins (1993) has argued against the distinction between a phenomenal and a representational aspect of experience. She asks the reader to look around his or her office, noting what it is like to have that experience. Then she challenges the reader to imagine that "a bat's consciousness is just like that--the feel of the scene is exactly the same--except, of course, all those visual sensations mean something quite different to the bat. They represent quite different properties. Imagine that!" She goes on to say "The problem is that you cannot imagine that, no matter how hard you try" (267). Of course, she is right that you cannot imagine that. But the explanation of this fact is not that there is no distinction between the P-conscious and representational aspects of experience. The explanation is that, as I said earlier, many representational differences themselves [make] a P-conscious difference. To repeat the example given earlier, what it is like to hear a sound as coming from the left is different from what it is like to hear a sound as coming from the right. Or suppose that you are taken to what appears to be a town from the Old West; then you are told that it is a backdrop for a film and that what appear to be buildings are mere fronts. This representational difference can make a difference in what the buildings look like to you. A visual experience as of a facade differs from a visual experience as of a building, even if the retinal image is the same. Or consider the difference in what it is like to hear sounds in French before and after you have learned the language (McCullough, 1993).

I am now just about finished justifying and explaining the difference between P-consciousness and A-consciousness. However, there is one objection I feel I should comment on. The contrast between P-consciousness and A-consciousness was in part based on the distinction between representational and phenomenal content. Put crudely, I said, the difference was that P-conscious content is phenomenal whereas A-conscious content is representational. I said this was crude because many phenomenal contents are also representational. Some will object that phenomenal content just [is] a kind of representational content. (Dretske, 1994 and Tye, forthcoming-a,b take this line; Shoemaker, 1994 has a more moderate version.) The representational/phenomenal distinction is discussed in Jackson, 1977, Shoemaker, 1981, and Peacocke, 1983.) My reply is first that phenomenal content need not be representational at all (my favorite example is the phenomenal content of orgasm). Second, suppose I have an auditory experience as of something overhead, and simultaneously have a visual experience as of something overhead. I'm imagining a case where one has an impression only of where the thing is without an impression of other features. For example, in the case of the visual experience, one catches a glimpse of something overhead without any impression of a specific shape or color. (So the difference cannot be ascribed to further representational differences.) The phenomenal contents of both experiences represent something as being overhead, but there is no common phenomenal quality of the experiences in virtue of which they have this representational commonality. Note that the point is [not] just that there is a representational overlap without a corresponding phenomenal overlap (as is said, for example, in Pendlebury, 1992). That would be compatible with the following story (offered to me by Michael Tye): phenomenal content is just one kind of representational content, but these experiences overlap in non-phenomenal representational content. The point, rather is that there is a modal difference that isn't at all a matter of representation, but rather is a matter of how those modes of representation feel. The look and the sound are both [as of something overhead], but the two phenomenal contents represent this via different phenomenal qualities. (There is a line of thought about the phenomenal representational distinction that involves versions of the traditional "inverted spectrum" hypothesis. See Shoemaker, 1981b, 1993; Block 1990a.)

In the next section, I will examine some confluences of P-consciousness and A-consciousness, and in the last section of the paper, I will argue that the target reasoning is fallacious because of such a confluence. In the remainder of this section, however, I will briefly discuss two cognitive notions of consciousness, so that they are firmly distinguished from both P-consciousness and A-consciousness.

SELF-CONSCIOUSNESS. By this term, I mean the possession of the concept of the self and the ability to use this concept in thinking about oneself. A number of higher primates show signs of recognizing that they see themselves in mirrors. They display interest in correspondences between their own actions and the movements of their mirror images. By contrast, monkeys treat their mirror images as strangers at first, slowly habituating. And the same for dogs. In one experimental paradigm, experimenters painted colored spots on the foreheads and ears of anesthetized primates, watching what happened. Chimps between ages 7 and 15 usually try to wipe the spot off (Povinelli, 1994; Gallup, 1982). Monkeys do not do this. Human babies don't show similar behavior until the last half of their second year. Perhaps this is a test for self-consciousness. (Or perhaps it is only a test for understanding mirrors; but what is involved in understanding mirrors if not that it is oneself one is seeing?) But even if monkeys and dogs have no self-consciousness, no one should deny that they have P-conscious pains, or that there is something it is like for them to see their reflections in the mirror. P-conscious states often seem to have a "me-ishness" about them, the phenomenal content often represents the state as a state of me. But this fact does not at all suggest that we can reduce P-consciousness to self-consciousness, since such "me-ishness" is the same in states whose P-conscious content is different. For example, the experience as of red is the same as the experience as of green in self-orientation, but the two states are different in phenomenal feel. See White (1987) for an account of why self-consciousness should be firmly distinguished from P-consciousness, and why self-consciousness is more relevant to certain issues of value.

MONITORING-CONSCIOUSNESS. The idea of consciousness as some sort of internal monitoring takes many forms. One notion is that of some sort of inner perception. This could be a form of P-consciousness, namely P-consciousness of one's own states or of the self. Another notion is often put in information-processing terms: internal scanning. And a third, metacognitive notion, is that of higher order thought: a conscious state in this sense is a state accompanied by a thought to the effect that one is in that state. The thought must be arrived at non-observationally and non-inferentially. Otherwise, as Rosenthal points out, the higher order thought definition would get the wrong result for the case in which I come to know about my anger by inferring it from my own behavior. The pioneer of these ideas in the philosophical literature is David Armstrong (1968, 1980). William Lycan (1987) has energetically pursued self-scanning, and David Rosenthal (1986, 1993), Peter Carruthers (1989, 1992) and Norton Nelkin (1993) have championed higher order thought. See also Natsoulas (1993) Lormand (forthcoming) makes some powerful criticisms of Rosenthal. Given my liberal terminological policy, I have no objection to any of these notions as notions of consciousness. Where I balk is at attempts to identify P-consciousness with any of these cognitive notions.

To identify P-consciousness with internal scanning is just to grease the slide to eliminativism about P-consciousness. Indeed, as Georges Rey (1983) has pointed out, ordinary laptop computers are capable of various types of self-scanning, but as he also points out, no one would think of their laptop computer as "conscious" (using the term in the ordinary way, without making any of the distinctions I've introduced). Since, according to Rey, internal scanning is essential to consciousness, he concludes that the concept of consciousness is incoherent. The trouble here is

the failure to make distinctions of the sort I've been making. Even if the laptop has "internal scanning consciousness", it nonetheless lacks P-consciousness. To be fair to Rey, his argument is more like a dilemma: for any supposed feature of consciousness, either a laptop of the sort we have today has it or else you can't be sure you have it yourself. So in the case of P-consciousness, the focus might be on the latter disjunct.

The concepts of consciousness which this paper is mainly about (P-consciousness and A-consciousness) differ in their logics from the consciousnesses just mentioned, self-consciousness and monitoring consciousness. A distinction is often made between the sense of 'conscious' in which a person or other creature is conscious and the sense in which a state of mind is a conscious state. What it is for there to be something it is like to be me, that is for me to be P-conscious, is for me to have one or more states that are P-conscious. If a person is in a dreamless sleep, and then has a P-conscious pain, he is to that extent P-conscious. For P-consciousness, it is states that are primary. Likewise for A-consciousness. If a state has the three properties mentioned earlier (inferential promiscuity, etc.) it is A-conscious, and a person is A-conscious just in case the person has an A-conscious state. In the case of self-consciousness and reflective consciousness, however, creature consciousness is basic. What it is for a pain to be reflectively conscious is, e.g. for the person whose pain it is to have another state that is about that pain. And it is creatures who can think about themselves. It is not even clear what a self-conscious state would [be].

Perhaps you are wondering why I am being so terminologically liberal, counting P-consciousness, A-consciousness, monitoring consciousness and self-consciousness all as types of consciousness. Oddly, I find that many critics wonder why I would count [phenomenal] consciousness as consciousness, whereas many others wonder why I would count [access] or [monitoring] or [self] consciousness as consciousness. In fact two reviewers of this paper complained about my terminological liberalism, but for incompatible reasons. One reviewer said: "While what he uses ["P-consciousness"] to refer to--the "what it is like" aspect of mentality--seems to me interesting and important, I suspect that the discussion of it under the heading "consciousness" is a source of confusion...he is right to distinguish access-consciousness (which is what I think deserves the name "consciousness") from this." Another reviewer said: "I really still can't see why access is called...access-consciousness? Why isn't access just...a purely information processing (functionalist) analysis?" This is not a merely verbal disagreement. In my view, all of us, despite our explicit verbal preferences, have some tendency to use 'conscious' and related words in both ways, and our failure to see this causes a good deal of difficulty in thinking about "consciousness". This point will be illustrated below.

I've been talking about different concepts of "consciousness" and I've also said that [the] concept of consciousness is a mongrel concept. Perhaps, you are thinking, I should make up my mind. My view is that 'consciousness' is actually an ambiguous word, though the ambiguity I have in mind is not one that I've found in any dictionary. I started the paper with an analogy between 'consciousness' and 'velocity', and I think there is an important similarity. One important difference, however, is that in the case of 'velocity', it is easy to get rid of the temptation to conflate the two senses, even though for many purposes the distinction is not very useful. With 'consciousness', there is a tendency towards "now you see it, now you don't." I think the main reason for this is that P-consciousness presents itself to us in a way that makes it hard to imagine how a conscious state could fail to be accessible and self-reflective, so it is easy to fall into habits of thought that do not distinguish these concepts.

The chief alternative to the ambiguity hypothesis is that there is a single concept of consciousness that is a [cluster concept.] For example, a prototypical religion involves belief in supernatural

beings, sacred and profane objects, rituals, a moral code, religious feelings, prayer, a world view, an organization of life based on the world view and a social group bound together by the previous items (Alston, 1967). But for all of these items, there are actual or possible religions that lack them. For example, some forms of Buddhism do not involve belief in a supreme being and Quakers have no sacred objects. It is convenient for us to use a concept of religion that binds together a number of disparate concepts whose referents are often found together.

The distinction between ambiguity and cluster concept can be drawn in a number of equally legitimate ways that classify some cases differently. That is, there is some indeterminacy in the distinction. Some might even say that [velocity] is a cluster concept because for many purposes it is convenient to group average and instantaneous velocity together. I favor drawing the distinction this way: when there are some occasions of literal use of a word to clearly and determinately express one element of "the cluster", we have ambiguity. When one catches a glimpse of a car flashing by so quickly that it is a blur, calling it 'fast' is plausibly calling it high in instantaneous velocity, since there seems no implicit relativization to a trip with a high average velocity. A similarly pure use in the case of 'consciousness' is provided by the science fiction example I've mentioned of the robot that is computationally like us even though it lacks consciousness. What it is supposed to lack is [P-consciousness].

When I called [consciousness] a mongrel concept I was not declaring allegiance to the cluster theory. Rather, what I had in mind was that an ambiguous word often corresponds to an ambiguous mental representation, one that functions in thought as a unitary entity and thereby misleads. These are mongrels. I would also describe [velocity] and [degree of heat] (as used by the Florentine Experimenters of the 17th Century) as mongrel concepts. This is the grain of truth in the cluster-concept theory.

CONFLATIONS

Conflation of P-consciousness and A-consciousness is ubiquitous in the burgeoning literature on consciousness, especially in the literature on syndromes like blindsight. Nearly every article I read on the subject by philosophers and psychologists involves some confusion. For example, Baars (1988) makes it abundantly clear that he is talking about P-consciousness. "What is a theory of consciousness a theory of? In the first instance...it is a theory of the nature of experience. The reader's private experience of [this] word, his or her mental image of yesterday's breakfast, or the feeling of a toothache--these are all contents of consciousness." (14) Yet his theory is a "global workspace" model of A-consciousness. Shallice (1988a,b) says he is giving an account of "phenomenal experience", but actually gives an information processing theory of A-consciousness. (His 1988b is about an "information-processing model of consciousness".) Mandler (1985) describes consciousness in P-conscious terms like "phenomenal" and "experience" but gives a totally cognitive account appropriate to A-consciousness. Edelman's (1989) theory is also intended to explain P-consciousness, but it seems a theory of access-consciousness and self-consciousness; see Chalmers (1993). Kosslyn and Koenig (1992) say "We will address here the everyday sense of the term ["consciousness"]; it refers to the phenomenology of experience, the feeling of red and so forth." (431-433 I am indebted to Michael Tye for calling this quotation to my attention.) But then they give a "parity check" theory that seems more of a theory of monitoring consciousness or A-consciousness.

One result of conflating P-consciousness with other consciousnesses is a tendency to regard ideas as plausible that should be seen as way out on a limb. For example, Johnson-Laird (1988, p. 360-361) talks of consciousness, using terms like "subjective experience". He goes on to

hypothesize that consciousness is a matter of building models of the self and models of the self building models of itself, and so on. This hypothesis has two strikes against it, as should be obvious if one is clear about the distinction between P-consciousness and self-consciousness. Dogs and babies may not build such complex models, but the burden of proof is surely on anyone who doubts that they have P-consciousness.

Another example: In a discussion of phenomena of implicit perception, Kihlstrom, et.al. (1992) make it clear that the phenomena concern P-consciousness: "In the final analysis, consciousness is a phenomenal quality that may accompany perception..." (42). But they claim that self-consciousness is precisely what is lacking in implicit perception: "This connection to the self is just what appears to be lacking in the phenomena of implicit perception...When contact occurs between the representation of the event--what might be called the "fact node" and the representation of oneself--what might be called the "self node," the event comes into consciousness." (p. 42). But again, as we go down the phylogenetic scale we may well encounter creatures that are P-conscious but have no "self-node", and the same may be true of the very young of our own species. What should be announced as a theory that conflicts with common sense, that P-consciousness arises from representing the self, can appear innocuous if one is not careful to make the distinctions among the consciousnesses.

Andrade (1993) makes it clear that the concern is P-consciousness. For example, "Without consciousness, there is no pain. There may be tissue damage, and physiological responses to tissue damage, but there will not be the phenomenological experience of pain" (13). Considering work on control by a central executive system, Andrade (correctly, I think) takes the dominant theories to "identify" consciousness with central executive control. "Current psychological theories identify consciousness with systems that coordinate lower-level information processing". But there are two very different paths to such an identification: (1) conflating P-consciousness with A-consciousness and theorizing about A-consciousness in terms of the systems Andrade mentions, (2) clearly distinguishing P-consciousness from A-consciousness and hypothesizing that the mechanisms that underly the latter give rise to the former. I doubt that any objective reader of this literature will think that the hypothesis of path 2 is often very likely.

In the writings of some psychologists, assimilation of P-consciousness to A-consciousness is a product of the (admirable) desire to be able to [measure] P-consciousness. Jacoby, et. al., 1992 assimilate P-consciousness to A-consciousness for that reason. Their subject matter is perception without "subjective experience", in normal perceivers in conditions of divided attention or degraded presentations. In other words, perception without P-consciousness, what is often known as subliminal perception. They note that it is very difficult to disentangle conscious perception from unconscious perception because no one has conceived of an experimental paradigm that isolates one of these modes. "We avoid this problem," they say, "by inferring awareness ["subjective experience"--N.B] from conscious control and defining unconscious influences as effects that cannot be controlled." (108) The effect of this procedure is to definitionally disallow phenomenal events that have no effect on later mental processes and to to definitionally type phenomenal events by appeal to judgements made on the basis of them. "Subjective experience", they say, "results from an attribution process in which mental events are interpreted in the context of current circumstances." (112) I am reminded of an article in the sociology of science that I once read that defined the quality of a scientific paper as the number of references to it in the literature. Operational definitions do no good if the result is measuring something [else].

Schacter (1989) is explicit about what he means by 'consciousness' (which, he often calls 'conscious awareness'), namely P-consciousness. He mentions that the sense he has in mind is that

of "phenomenal awareness" ...`the running span of subjective experience'" (quoting Dimond, 1976), and consciousness in his sense is repeatedly contrasted with information processing notions. Nonetheless, in an effort to associate the "Conscious Awareness System" (what I call the phenomenal consciousness system in my labeling of his model in Figure 1) with the inferior parietal lobes, he says that lesions in this area

have also been associated with confusional states, which are characterized by disordered thought, severe disorientation, and a breakdown of selective attention--in short, a global disorder of conscious awareness...Several lines of evidence indicate that lesions to certain regions of the parietal lobes can produce disorders of conscious awareness. First, global confusional states have been reported in right parietal patients...Second, the syndrome of anosognosia--unawareness and denial of a neuropsychological deficit--is often associated with parietal damage...Anosognosic patients...may be unaware of motor deficits...perceptual deficits...and complete unawareness can be observed even when the primary deficit is severe...(1988, p. 371) Here, Schacter reverts to a use of `consciousness' and `awareness' in a variety of cognitive senses. Disordered thought, disorientation and a breakdown of selective attention are not primarily disorders of P-consciousness. Further, anosognosia is primarily a defect in A-consciousness, not P-consciousness. Anosognosia is a neurological syndrome that involves an inability to acknowledge or have access to information about another neurological syndrome. A patient might have anosognosia for, say his prosopagnosia while complaining incessantly about another deficit. Young (1994a) describes a woman who was a painter before becoming prosopagnosic. Looking at portraits she had painted, trying to figure out whom they represented, she laboriously figured out whom each painting was of, reasoning out loud about the person's apparent age, sex, and any significant objects in the picture, plus her verbal memories of the portraits that she had painted. When the experimenter commented on her prosopagnosia, she said that she "had recognized them," and did not think that there was anything odd about her laborious reasoning. Interestingly, she was in many respects much worse at many face-perception tasks than LH (the prosopagnosic mentioned earlier)--she couldn't match photographs of faces, for example. I have noticed that people who know little about anosognosia tend to favor various debunking hypotheses. That is, they assume that the experimenters have made one or another silly mistake in describing the syndrome, because, after all, how could anyone fail to notice that they can't recognize faces, or worse, that they are blind. See Young, et. al., 1993, for a good debunking of the debunking hypotheses.

The crucial feature of anosognosia about prosopagnosia is that the patient's access to information about her own inability to recognize faces is in some way blocked. She cannot report this inability or reason about it or use information about it to control her action. There may also be some defect of P-consciousness. Perhaps everyone looks familiar or more likely, perhaps patients with prosopagnosia no longer have the ability to have visual feelings of familiarity for faces that are distinct from feelings of unfamiliarity. But this is not crucial to the syndrome, as is shown by the fact that we confidently ascribe anosognosia on the basis of the patient's cognitive state--the lack of knowledge of the deficit--without knowing what defects of P-consciousness may or may not be involved. Further, the same defects of P-consciousness could be present in a [non-anosognosic] prosopagnosic without discrediting the patient's status as non-anosognosic. One can imagine such a person saying "Gosh, I don't recognize anyone--in fact, I no longer have a visual sense of the difference between familiar and unfamiliar faces." This would be prosopagnosia [without] anosognosia. To take anosognosia as primarily a defect of P-consciousness is a mistake.

I don't think these confluences cause any real problem in Schacter's theorizing, but as a general

rule, if you want to get anywhere in theorizing about X you should have a good pre-theoretical grip on the difference between X and things that are easily confused with it.

Daniel Dennett (1986, 1991) provides another example of conflation of a number of concepts of consciousness. (See my 1993.) I will focus on Dennett's claim that consciousness is a cultural construction. He theorizes that "human consciousness (1) is too recent an innovation to be hard-wired into the innate machinery, (2) is largely the product of cultural evolution that gets imparted to brains in early training" (1991, p. 219). Sometimes he puts the point in terms of memes, which are ideas such as the idea of the wheel or the calendar. Memes are the smallest cultural units that replicate themselves reliably, viz., cultural analogs of genes. In these terms then, Dennett's claim is that "human consciousness is [itself] a huge complex of memes". (1991, p.210) This view is connected with Dennett's idea that you can't have consciousness without having the concept of consciousness. He says consciousness is like love and money in this regard, though in the case of money, what is required for one to have money is that [someone] have the concept of money (1991, p.24; 1986, p. 152).

I think the reason Dennett says "largely" the product of cultural evolution is that he thinks of consciousness as the software that operates on genetically determined hardware that is the product of biological evolution. Though consciousness requires the concept of consciousness, with consciousness as with love, there is a biological basis without which the software could not run.

Now I hope it is obvious that P-consciousness is not a cultural construction. Remember, we are talking about P-consciousness itself, not the concept of P-consciousness. The idea would be that there was a time at which people genetically like us ate, drank and had sex, but there was nothing it was like for them to do these things. Further, each of us would have been like that if not for specific concepts we acquired from our culture in growing up. Ridiculous! Of course, culture [affects] P-consciousness; the wondrous experience of drinking a great wine takes training to develop. But culture affects feet too; people who have spent their lives going barefoot in the Himalayas have feet that differ from those of people who have worn tight shoes 18 hours a day. We mustn't confuse the idea that culture [influences] consciousness with the idea that it (largely) creates it.

What about A-consciousness? Could there have been a time when humans who are biologically the same as us never had the contents of their perceptions and thoughts poised for free use in reasoning or in rational control of action? Is this ability one that culture imparts to us as children? Could it be that until we acquired the concept of [poised for free use in reasoning or in rational control of action], none of our perceptual contents were A-conscious? Again, there is no reason to take such an idea seriously. Very much lower animals are A-conscious, presumably without any such concept.

A-consciousness is as close as we get to the official view of consciousness in [Consciousness Explained] and in later writings, e.g. Dennett (1993). The official theory of Dennett (1991) is the Multiple Drafts Theory, the view that there are distinct parallel tracks of representation that vie for access to reasoning, verbalization and behavior. This seems a theory of A-consciousness. Dennett (1993) says "Consciousness is cerebral celebrity--nothing more and nothing less. Those contents are conscious that persevere, that monopolize resources long enough to achieve certain typical and "symptomatic" effects--on memory, on the control of behavior, and so forth." (929) Could it be anything other than a biological fact about humans that some brain representations persevere enough to affect memory, control behavior, etc? So on the closest thing to Dennett's official kind of consciousness, the thesis (that consciousness is a cultural construction) is no serious proposal.

What about monitoring consciousness? No doubt there was a time when people were less introspective than some of us are now. But is there any evidence that there was a time when people genetically like us had no capacity to think or express the thought that one's leg hurts. To be able to think this thought involves being able to think that one's leg hurts, and that is a higher order thought of the sort that is a plausible candidate for monitoring consciousness (Rosenthal, 1986). Here for the first time we do enter the realm of actual empirical questions, but without some very powerful evidence for such a view, there is no reason to give it any credence. Dennett gives us not the slightest hint of the kind of weird evidence that we would need to begin to take this claim seriously, and so it would be a disservice to so interpret him.

What about self-consciousness? I mentioned Gallup's and Povinelli's "mark test" evidence (the chimp tries to wipe off a mark on its face seen in a mirror) that chimps are self-conscious. An experiment in this vein that Dennett actually mentions (p. 428), and mentions positively, is that a chimp can learn to get bananas via a hole in its cage by watching its arm on a closed circuit TV whose camera is some distance away (Menzel, et.al, 1985). The literature on the topic of animal self-consciousness is full of controversy. See Heyes, 1993, Mitchell, 1993a, 1993b; Gallup and Povinelli 1993; de Lannoy, 1993; Anderson, 1993; Byrne, 1993. I have no space to do justice to the issues, so I will have to make do with just stating my view: I think the weight of evidence in favor of minimal self-consciousness on the part of chimps is overwhelming. By minimal self-consciousness I mean the ability to think about oneself in some way or other--that is, no particular way is required. Many of the criticisms of the mark test actually presuppose that the chimp is self-conscious in this minimal sense. For example, it is often suggested that chimps that pass the mark test think that they are seeing another chimp (e.g. Heyes, 1993), and since the chimp in the mirror has a mark on its forehead, the chimp who is looking wonders whether he or she does too. But in order for me to wonder whether [I] have a mark on my forehead, I have to be able to think about myself. In any case, Dennett does not get into these issues (except, as mentioned, to favor chimp self-consciousness), so it does not appear that he has this interpretation in mind.

So far, on all the consciousness I have mentioned, Dennett's thesis turns out to be false. But there is a trend: Of the concepts I considered, the first two made the thesis silly, even of animals. In the case of monitoring consciousness, there is a real empirical issue in the case of many types of mammals, and so it isn't completely silly to wonder about whether people have it. Only in the last case, self-consciousness, is there a serious issue about whether chimps are conscious, and that suggests that we might get a notion of self-consciousness that requires some cultural elements. In recent years, the idea of the self as a federation of somewhat autonomous agencies has become popular, and for good reason. Nagel (1971) made a good case on the basis of split-brain data, and Gazzaniga and LeDoux (1978) and Gazzaniga (1985) have added additional considerations that have some plausibility. And Dennett has a chapter about the self at the end of the book that gives similar arguments. Maybe what Dennett is saying is that non-federal-self consciousness, the ability to think of oneself as not being such a federation (or more simply, federal-self consciousness) is a cultural construction.

But now we have moved from falsity to banality. I'm not saying that the proposal that we are federations is banal. What is banal is that the having and applying a sophisticated concept such as being a federation (or not being a federation) requires a cultural construction. Compare chairman-self consciousness, the ability to think of oneself as chairman, as the one who guides the Department, the one who has the keys, etc. It is a banality that a cultural construction is required in order for a person to think of himself in that way, and the corresponding point about federal-self consciousness is similarly banal.

The great oddity of Dennett's discussion is that throughout he gives the impression that his theory is [about P-consciousness], though he concedes that what he says about it conflicts with our normal way of thinking about consciousness. This comes out especially strongly in an extended discussion of Julian Jaynes' (1976) book which he credits with a version of the view I am discussing, viz. that consciousness is a cultural construction which requires its own concept. He says:

Perhaps this is an autobiographical confession: I am rather fond of his [Jaynes'] way of using these terms; ['consciousness', 'mind' and other mental terms] I rather like his way of carving up consciousness. It is in fact very similar to the way that I independently decided to carve up consciousness some years ago.

So what then is the project? The project is, in one sense, very simple and very familiar. It is bridging what he calls the "awesome chasm" between mere inert matter and the inwardness, as he puts it, of a conscious being. Consider the awesome chasm between a brick and a bricklayer. There isn't, in Thomas Nagel's (1974) famous phrase, anything that it is like to be a brick. But there is something that it is like to be a bricklayer, and we want to know what the conditions were under which there happened to come to be entities that it was like something to be in this rather special sense. That is the story, the developmental, evolutionary, historical story, that Jaynes sets out to tell. (149)

In sum, Dennett's thesis is trivially false if it is construed to be about P-consciousness, as advertised. It is also false if taken to be about A-consciousness which is Dennett's official view of consciousness. But if taken to be about a highly sophisticated version of self-consciousness, it is banal. That's what can happen if you talk about consciousness without making the sorts of distinctions that I am urging.

THE FALLACY OF THE TARGET REASONING We now come to the denouement of the paper, the application of the P-consciousness/A-consciousness distinction to the fallacy of the target reasoning. Let me begin with the Penfield-van Gulick-Searle reasoning. Searle (1992) adopts Penfield's (1975) claim that during petit mal seizures, patients are "totally unconscious". Quoting Penfield at length, Searle describes three patients who, despite being "totally unconscious", continue walking or driving home or playing the piano, but in a mechanical way. Van Gulick (1989) gives a briefer treatment, also quoting Penfield. He says "The importance of conscious experience for the construction and control of action plans is nicely illustrated by the phenomenon of automatism associated with some petit mal epileptic seizures. In such cases, electrical disorder leads to a loss of function in the higher brain stem...As a result the patient suffers a loss of conscious experience in the phenomenal sense although he can continue to react selectively to environmental stimuli." (p. 220) Because van Gulick's treatment is more equivocal and less detailed, and because Searle also comments on my accusations of conflating A-consciousness with P-consciousness, I'll focus on Searle.

Searle says: ...the epileptic seizure rendered the patient [totally unconscious], yet the patient continued to exhibit what would normally be called goal-directed behavior...In all these cases, we have complex forms of apparently goal-directed behavior without any consciousness. Now why could all behavior not be like that? Notice that in the cases, the patients were performing types of actions that were habitual, routine and memorized...normal, human, conscious behavior has a degree of flexibility and creativity that is absent from the Penfield cases of the unconscious driver and the unconscious pianist. [Consciousness adds powers of discrimination and flexibility] even to memorized routine activities...one of the evolutionary advantages conferred on us by

consciousness is the much greater [flexibility, sensitivity, and creativity] we derive from being conscious. (1992, p.108-109, italics mine) Searle's reasoning is that consciousness is missing, and with it, flexibility, sensitivity and creativity, so this is an indication that a function of consciousness is to add these qualities. Now it is completely clear that the concept of consciousness invoked by both Searle and van Gulick is P-consciousness. Van Gulick speaks of "conscious experience in the phenomenal sense," and Searle criticizes me for supposing that there is a legitimate use of 'conscious' to mean A-conscious: "Some philosophers (e.g. Block, "Two Concepts of Consciousness") claim that there is a sense of this word that implies no sentience whatever, a sense in which a total zombie could be 'conscious'. I know of no such sense, but in any case, that is not the sense in which I am using the word." (1992, p. 84). But neither Searle nor van Gulick nor Penfield give any reason to believe that P-consciousness is missing or even diminished in the epileptics they describe. The piano player, walker and the driver don't cope with new situations very well, but they do show every sign of [normal sensation]. For example, Searle, quoting Penfield, describes the epileptic walker as "thread[ing] his way" through the crowd. Doesn't he [see] the obstacles he avoids? Suppose he gets home by turning right at a red wall. Isn't there something it is like for him to see the red wall--and isn't it different from what it is like for him to see a green wall? Searle give no reason to think the answer is no. Because of the very inflexibility and lack of creativity of the behavior they exhibit, it is the [thought processes] of these patients (including A-consciousness) that are most obviously deficient; no reason at all is given to think that their P-conscious states lack vivacity or intensity. Of course, I don't claim to know what it is really like for these epileptics; my point is rather that for the argument for the function of P-consciousness to have any force, a case would have to be made that P-consciousness is [actually] missing, or at least diminished. Searle argues: P-consciousness is missing; so is creativity; therefore the former lack explains the latter lack. But no support at all is given for the first premise, and as we shall see, it is no stretch to suppose that what's gone wrong is that the ordinary mongrel notion of consciousness is being used; it wraps P-consciousness and A-consciousness together, and so an obvious function of A-consciousness is illicitly transferred to P-consciousness.

This difficulty in the reasoning is highlighted if we assume Schacter's model. In terms of Schacter's model, there is no reason to doubt that the information from their senses reaches the P-conscious module, but there is reason to doubt that the Executive system processes this information in the normal way. So there is reason to blame their inflexibility and lack of creativity on problems in the Executive system or the linkage between the P-consciousness module and the Executive system. There is an additional problem in the reasoning that I won't go into except here. There is a well-known difficulty in reasoning of the form: X is missing; the patient has lost the ability to do blah-blah; therefore a function of X is to facilitate blah-blahing. In a complex system, a loss may reverberate through the system, triggering a variety of malfunctions that are not connected in any serious way with the function of the missing item. An imperfect but memorable example (that I heard from Tom Bever) will illustrate: the Martians want to find out about the function of various Earthly items. They begin with The Pentagon, and focus in on a particular drinking fountain in a hall on the third floor of the North side of the building. "If we can figure out what that is for", they think, "we can move on to something more complex." So they vaporize the drinking fountain, causing noise and spurting pipes. Everyone comes out of their office to see what happened and the Martians conclude that the function of the fountain was to keep people in their offices. The application of this point to the petit mal case is that even if I am right that it is A-consciousness, not P-consciousness, that is diminished or missing, I would not jump to the conclusion that A-consciousness has a function of adding powers of discrimination, flexibility and creativity. Creativity, for example, may have its sources in the unA-conscious, requiring powers of

reasoning and control of action and reporting only for its expression.

Searle and van Gulick base their arguments on Penfield's claim that a petit mal seizure "converts the individual into a mindless automaton" (Penfield, 1975, p. 37). Indeed, Penfield repeatedly refers to these patients as "unconscious", "mindless", and as "automata". But what does Penfield [mean]? Searle and van Gulick assume that Penfield means P-consciousness, since they adopt the idea that that is what the term means (though as we shall see, Searle himself sometimes uses the term to mean A-consciousness). Attending to Penfield's account, we find the very shifting among different concepts of consciousness that I have described here, but the dominant theme by far involves thinking of the patients as cognitively rather than phenomenally deficient during petit mal seizures. Here is Penfield's summary of the description of the patients: In an attack of automatism the patient becomes suddenly unconscious, but, since other mechanisms in the brain continue to function, he changes into an automaton. He may wander about, confused and aimless. Or he may continue to carry out whatever purpose his mind was in the act of handing on to his automatic sensory-motor mechanism when the highest brain-mechanism went out of action. Or he follows a stereotyped, habitual pattern of behavior. In every case, however, the automaton can make few, if any decisions for which there has been no precedent. [He makes no record of a stream of consciousness.] Thus, he will have complete amnesia for the period of epileptic discharge...In general, if new decisions are to be made, the automaton cannot make them. In such a circumstance, he may become completely unreasonable and uncontrollable and even dangerous.

In these passages, and throughout the book, the dominant theme in descriptions of these patients is one of deficits in thinking, planning and decision making. No mention is made of any sensory or phenomenal deficit. Indeed, in the italicized passage above (italics mine) there is an implicit suggestion that perhaps there are P-conscious events of which no record is made. I could only find one place in the book where Penfield says anything that might be taken to contradict this interpretation: "Thus, the automaton can walk through traffic as though he were aware of all that he hears and sees, and so continue on his way home. But he is aware of nothing and so makes no memory record. If a policeman were to accost him he might consider the poor fellow to be walking in his sleep." (60) But to properly understand this, we need to know what he means by "awareness", and what he thinks goes on in sleep. Judging by Penfield's use of synonyms, by "awareness" he means something in the category of the higher order thought analyses or the self-consciousness sense. For example, in discussing his peculiar view that ants are conscious, he seems to use 'conscious' and 'aware' to mean 'self-aware' (62, 105, 106) Further, he makes it clear that although the mind is shut off during sleep, the sensory cortex is quite active.

My interpretation is supported by a consideration of Penfield's theoretical rationale for his claim that petit mal victims are unconscious. He distinguishes two brain mechanisms, "(a) the [mind's mechanism] (or highest brain mechanism); and (b) the [computer] (or automatic sensory-motor mechanism)" (p.40, italics Penfield's). The mind's mechanism is most prominently mentioned in connection with planning and decision making, for example, "...the highest brain mechanism is the mind's executive...". When arguing that there is a soul that is connected to the mind's mechanism, he mentions only cognitive functions: He asks whether such a soul is improbable, and answers "It is not so improbable, to my mind, as is the alternative expectation--that the highest brain mechanism should itself understand, and reason, and direct voluntary action, and decide where attention should be turned and what the computer must learn, and record, and reveal on demand." (82). Penfield's soul is a cognitive soul.

By contrast, the computer is devoted to [sensory] and motor functions. Indeed, he emphasizes that the mind only has contact with sensory and motor areas of the cortex via controlling the computer,

which itself has direct contact with the sensory and motor areas. Since it is the mind's mechanism that is knocked out in petit mal seizures, the sensory areas are intact in the "automaton".

Searle (1990b) attempts (though of course he wouldn't accept this description) to use the idea of degrees of P-consciousness to substitute for A-consciousness. I will quote a chunk of what he says about this. (The details of the context don't matter.)

By consciousness I simply mean those subjective states of awareness or sentience that begin when one wakes in the morning and continue throughout the period that one is awake until one falls into a dreamless sleep, into a coma, or dies or is otherwise, as they say, unconscious.

I quoted this passage earlier as an example of how a characterization of consciousness can go wrong by pointing to too many things. Searle means to be pointing to P-consciousness. But A-consciousness and P-consciousness normally occur together when one is awake, and both are normally absent in a coma and a dreamless sleep--so this characterization doesn't distinguish them.

On my account, dreams are a form of consciousness,...though they are of less intensity than full blown waking alertness. Consciousness is an on/off switch: You are either conscious or not. Though once conscious, the system functions like a rheostat, and there can be an indefinite range of different degrees of consciousness, ranging from the drowsiness just before one falls asleep to the full blown complete alertness of the obsessive.

Degrees of P-consciousness are one thing, obsessive attentiveness is another--indeed the latter is a notion from the category of A-consciousness, not P-consciousness.

There are lots of different degrees of consciousness, but door knobs, bits of chalk, and shingles are not conscious at all...These points, it seems to me, are misunderstood by Block. He refers to what he calls an "access sense of consciousness". On my account there is no such sense. I believe that he...[confuses] what I would call peripheral consciousness or [inattentiveness] with total unconsciousness. It is true, for example, that when I am driving my car "on automatic pilot" I am not paying much attention to the details of the road and the traffic. But it is simply not true that I am totally unconscious of these phenomena. If I were, there would be a car crash. We need therefore to make a distinction between the [center of my attention, the focus of my consciousness] on the one hand, and the [periphery] on the other...There are lots of phenomena right now of which I am peripherally conscious, for example the feel of the shirt on my neck, the touch of the computer keys at my finger tips, and so on. But as I use the notion, none of these is unconscious in the sense in which the secretion of enzymes in my stomach is unconscious. (All quotes from Searle, 1990b, p. 635, italics mine)

The first thing to note is the [contradiction]. Earlier, I quoted Searle saying that a "totally unconscious" epileptic could nonetheless drive home. Here, he says that if a driver was totally unconscious, the car would crash. The sense of 'conscious' in which the car would crash if the driver weren't conscious is A-consciousness, not P-consciousness. P-consciousness [all by itself] wouldn't keep the car from crashing--the P-conscious contents have to be put to use in rationally controlling the car, which is an aspect of A-consciousness. When Searle says the "totally unconscious" epileptic can nonetheless drive home, he is talking about P-consciousness; when he says the car would crash if the driver were totally unconscious, he is talking mainly about A-consciousness. Notice that it will do no good for Searle to say that in the quotation of the last paragraph, he is talking about creature-consciousness rather than state-consciousness. What it is for a person to be P-unconscious is for his states (all or the relevant ones) to lack P-consciousness.

Creature P-consciousness is parasitic on state P-consciousness. Also, it will do him no good to appeal to the conscious/conscious of distinction. (The epileptics were "totally unconscious", but if he were "unconscious of" the details of the road and traffic the car would crash.) The epileptics were "totally unconscious" and therefore, since Searle has no resource of A-consciousness, he must say that the epileptics were totally unconscious [of] anything. So he is committed to saying that the epileptic driver can drive despite being totally unconscious of anything. And that contradicts the claim that I quoted that if Searle were totally unconscious of the details of the road and traffic, then the car would crash. If Searle says that someone who is totally unconscious can nonetheless be conscious of something, that would be a backhanded way of acknowledging the distinction.

The upshot is that Searle finds himself drawn to using 'consciousness' in the sense of A-consciousness, despite his official position that there is no such sense. Despite his official ideology, when he attempts to deploy a notion of degrees of P-consciousness he ends up talking about A-consciousness--or about both A-consciousness and P-consciousness wrapped together in the usual mongrel concept. Inattentiveness just [is] lack of A-consciousness (though it will have effects on P-consciousness). Thus, he may be right about the inattentive driver (note, the inattentive driver, not the *petit mal* case). When the inattentive driver stops at a red light, presumably there is something it is like for him to see the red light--the red light no doubt looks red in the usual way, that is it appears as brightly and vividly to him as red normally does. But since he is thinking about something else, perhaps he is not using this information very much in his reasoning nor is he using this information to control his speech or action in any sophisticated way--that is, perhaps his A-consciousness of what he sees is diminished. (Of course, it can't be totally gone or the car would crash.) Alternatively, A-consciousness might be normal, and the driver's poor memory of the trip may just be due to failure to put contents that are both P-conscious and A-conscious into memory; my point is that to the extent that Searle's story is right about [any] kind of consciousness, it is right about A-consciousness, not P-consciousness.

Searle's talk of the center and the periphery, is in the first instance about kinds of or degrees of access, not "degrees of phenomenality." You may recall that in introducing the A/P distinction, I used Searle's example of attending to the feel of the shirt on the back of one's neck. My point was that A-consciousness and P-consciousness interact: bringing something from the periphery to the center can [affect] one's phenomenal state. The attention makes the experience more fine-grained, more intense (though a pain that is already intense needn't become more intense when one attends to it). There is a phenomenal difference between figure and ground, though the perception of the colors of the ground can be just as intense as those of the figure, or so it seems to me. Access and phenomenality often interact, one bringing along the other--but that shouldn't make us blind to the difference.

Though my complaint is partly verbal, there is more to it. For the end result of deploying a mongrel concept is wrong reasoning about a function of P-consciousness.

Let me turn now to a related form of reasoning used by Owen Flanagan, 1992 (142-145). Flanagan discusses Luria's patient *Zazetsky*, a soldier who lost the memories of his "middle" past--between childhood and brain injury. The information about his past is represented in *Zazetsky's* brain, but it only comes out via "automatic writing". Flanagan says "The saddest irony is that although each piece of *Zazetsky's* autobiography was consciously reappropriated by him each time he hit upon a veridical memory in writing, he himself was never able to fully reappropriate, to keep in clear and continuous view, to live with, the self he reconstructed in the thousand pages he wrote." Flanagan goes on to blame the difficulty on a defect of consciousness, and he means P-consciousness: "*Zazetsky's* conscious capacities are (partly) maimed. His dysfunction is rooted in certain defects

of consciousness." (144-145) But Zazetsky's root problem appears to be a difficulty in A-consciousness, though that has an effect on self-consciousness and P-consciousness. The problem seems to be that the memories of the middle past are not accessible to him in the manner of his memories of childhood and recent past. To the extent that he knows about the middle past, it is as a result of reading his automatic writing, and so he has the sort of access we have to a story about someone else. The root difficulty is segregation of information, and whatever P-conscious feelings of fragmentation he has can be taken to result from the segregation of information. So there is nothing in this case that suggests a function of P-consciousness.

Let us now move to the line of thought mentioned at the outset about how the thirsty blindsight patient doesn't reach for the glass of water in the blind field. A similar line of reasoning appears in Shevrin, 1992; he notes that in subliminal perception, we don't fix the source of a mental content. Subliminal percepts aren't conscious, so consciousness must have the function of fixing the source of mental contents. (This line of thought appears in Marcel, 1986,1988, van Gulick, 1989 (though endorsed equivocally) and Flanagan, 1989.) The reasoning is that (1) consciousness is missing, (2) information that the patient in some sense possesses is not used in reasoning or in guiding action or in reporting, so (3) the function of consciousness must be to somehow allow information from the senses to be so used in guiding action (Marcel, 1986, 1988). Flanagan (1992) agrees with Marcel: "Conscious awareness of a water fountain to my right will lead me to drink from it if I am thirsty. But the thirsty blindsighted person will make no move towards the fountain unless pressed to do so. The inference to the best explanation is that conscious awareness of the environment facilitates semantic comprehension and adaptive motor actions in creatures like us." And: "Blindsighted patients never initiate activity toward the blindfield because they lack subjective awareness of things in that field". (Flanagan, 1992, p 141-142; the same reasoning occurs in his 1991, 349.) Van Gulick, 1989, agrees with Marcel, saying "Subjects never initiate on their own any actions informed by perceptions from the blind field. The moral to be drawn from this is that information must normally be represented in phenomenal consciousness if it is to play any role in guiding voluntary action." (p. 220)

Bernard Baars argues for eighteen different functions of consciousness on the same ground. He says that the argument for these functions is "that loss of consciousness--through habituation, automaticity, distraction, masking, anesthesia, and the like--inhibits or destroys the functions listed here."Baars, 1988, p. 356. Though Baars is talking about the function of "conscious experience", he does have a tendency to combine P-consciousness with A-consciousness under this heading.

Schacter (1989) approvingly quotes Marcel, using this reasoning to some extent in formulating the model of Figure 1 (though as I mentioned, there is a model that perhaps more fully embodies this reasoning--see below). The P-consciousness module has the function of integrating information from the specialized modules, injecting them with P-conscious content, and of sending these contents to the system that is in charge of reasoning and rational control of action and reporting.

This is the fallacy: In the blindsight patient, both P-consciousness and A-consciousness of the glass of water are missing. There is an obvious explanation of why the patient doesn't reach for the glass in terms of the information about it not reaching mechanisms of reasoning and rational control of speech and action, the machinery of A-consciousness. (If we believe in an Executive system, we can explain why the blindsight patient does not reach for the water by appealing to the claim that the information about the water does not reach the Executive system.) More generally, A-consciousness and P-consciousness are almost always present or absent together, or rather this seems plausible. This is, after all, [why] they are folded together in a mongrel concept. A function of the mechanisms underlying A-consciousness is completely obvious. If information from the

senses did not get to mechanisms of control of reasoning and of rational control of action and reporting, we would not be able to use our senses to guide our action and reporting. But it is just a mistake to slide from a function of the machinery of A-consciousness to any function at all of P-consciousness.

Of course, it could be that the lack of P-consciousness is itself responsible for the lack of A-consciousness. If [that] is the argument in any of these cases, I do not say "fallacy". The idea that the lack of P-consciousness is responsible for the lack of A-consciousness is a bold hypothesis, not a fallacy. Recall, however, that there is some reason to ascribe the opposite view to the field as a whole. The discussion earlier of Baars, Shallice, Kosslyn and Koenig, Edelman, Johnson-Laird, Andrade and Kihlstrom, et.al. suggested that to the extent that the different consciousnesses are distinguished from one another, it is often thought that P-consciousness is a product of (or is identical to) cognitive processing. In this climate of opinion, if P-consciousness and A-consciousness were clearly distinguished, and something like the opposite of the usual view of their relation advanced, we would expect some comment on this fact, something that does not appear in any of the works cited.

The fallacy, then, is jumping from the premise that "consciousness" is missing--without being clear about what kind of consciousness is missing--to the conclusion that P-consciousness has a certain function. If the distinction were seen clearly, the relevant possibilities could be reasoned about. Perhaps the lack of P-consciousness causes the lack of A-consciousness. Or perhaps the converse is the case: P-consciousness is somehow a product of A-consciousness. Or both could be the result of something else. If the distinction were clearly made, these alternatives would come to the fore. The fallacy is failing to make the distinction, rendering the alternatives invisible.

Note that the claim that P-consciousness is missing in blindsight is just an assumption. I decided to take the blindsight patient's word for his lack of P-consciousness of stimuli in the blind field. Maybe this assumption is mistaken. But if it is, then the fallacy now under discussion reduces to the fallacy of the Searle-Penfield reasoning: if the assumption is wrong, if the blindsight patient [does] have P-consciousness of stimuli in the blind field, then [only] A-consciousness of the stimuli in the blind field is missing, so [of course] we cannot draw the mentioned conclusion about the function of P-consciousness from blindsight.

I said at the outset that although there was a serious fallacy in the target reasoning, there was also something importantly right about it. What is importantly right is this. In blindsight, both A-consciousness and P-consciousness (I assume) are gone, just as in normal perception, both are present. So blind-sight is yet another case in which P-consciousness and A-consciousness are both present or both absent. Further, as I mentioned earlier, cases of A-consciousness without P-consciousness, such as the super-blindsight patient I described earlier do not appear to exist. Training of blindsight patients has produced a number of phenomena that look a bit like super-blindsight, but each such lead that I have pursued has fizzled. This suggests an intimate relation between A-consciousness and P-consciousness. Perhaps there is something about P-consciousness that greases the wheels of accessibility. Perhaps P-consciousness is like the liquid in a hydraulic computer, the means by which A-consciousness operates. Alternatively, perhaps P-consciousness is the gateway to mechanisms of access as in Schacter's model, in which case P-consciousness would have the function Marcel, et. al. mention. Or perhaps P-consciousness and A-consciousness even amount to much the same thing empirically even though they differ conceptually, in which case P-consciousness would also have the aforementioned function. Perhaps the two are so entwined together that there is no empirical sense to the idea of one without the other.

[THESE ARE FIGURE CAPTIONS ONLY: FIGURES THEMSELVES ARE ONLY AVAILABLE IN THE PAPER VERSION]

Compare the model of Figure 1 (Schacter's model) with those of Figures 2 and 3. The model of Figure 2 is just like Schacter's model except that the Executive system and the P-consciousness system are collapsed together. We might call the hypothesis that is embodied in it the Collapse Hypothesis. The Collapse Hypothesis should not be confused with Marcel's (1988, p. 135-7) Identity Hypothesis, which hypothesizes that the processing of stimuli is identical with consciousness of them. As Marcel points out, blindsight and similar phenomena suggest that we can have processing without consciousness. Figure 3 is a variant on Schacter's model in which the Executive module and the P-consciousness module are reversed. Schacter's model clearly gives P-consciousness a function in controlling action. Model 3 clearly gives it no function. Model 2 can be interpreted in a variety of ways, some of which give P-consciousness a function, others of which do not. If P-consciousness is literally identical to some sort of information processing, then P-consciousness will have whatever function that information processing has. But if P-consciousness is, say, a by-product of and supervenient on certain kinds of information processing (something that could also be represented by model 3), then P-consciousness will in that respect at least have no function. What is right about the Marcel, et. al. reasoning is that some of the explanations for the phenomenon give P-consciousness a role; what is wrong with the reasoning is that one cannot immediately conclude from missing "consciousness" to P-consciousness having that role.

CAN WE DISTINGUISH AMONG THE MODELS?

I'm finished with the point of the paper, but having raised the issue of the three competing models, I can't resist making some suggestions for distinguishing among them. My approach is one that takes introspection seriously. Famously, introspection has its problems (Nisbett and Wilson, 1977; Jacoby, Toth, Lindsay and Debnor, 1992), but it would be foolish to conclude that we can afford to ignore our own experience.

One phenomenon that counts against the Collapse Hypothesis (Model 2) is the familiar phenomenon of the solution to a difficult problem just popping into P-consciousness. If the solution involves high level thought, then it must be done by high level reasoning processes that are not P-conscious. (They aren't A-conscious either, since one can't report or base action on the intermediate stages of such reasoning.) There will always be disputes about famous cases (e.g. Kekule's discovery of the benzene ring in a dream), but we should not be skeptical about the idea that though the results of thought are both P-conscious and A-conscious, much in the way of the intermediate stages are neither. If we assume that all high-level reasoning is done in the Executive system, and that Model 2 is committed to all Executive processes being P-conscious, then Model 2 is incompatible with solutions popping into P-consciousness. Of course, alternative forms of Model 2 that do not make these assumptions may not make any such predictions.

I think there are a number of phenomena that, if investigated further, might lead to evidence for P-consciousness without A-consciousness and thus provide some reason to reject 2 in favor of Schacter's model (Figure 1). (I also think that these phenomena, if investigated further, might yield some reason to reject 3 in favor of 1, but I cannot go into that here.) I repeat: the phenomena I am about to mention don't show anything on their own. I claim only that they are intriguing and deserve further work.

One such phenomenon--or perhaps I should describe it as an idea rather than a phenomenon--is the

hypothesis, already mentioned, that there could be animals whose P-conscious brain processes are intact, but whose A-conscious brain processes are not. Another is the case mentioned earlier of states of P-consciousness that go on for some time without attention and only become A-conscious with the focusing of attention. (See also Hill, 1991.)

Sperling (1960) flashed arrays of letters (e.g. 3 by 3) to subjects for brief periods (e.g. 50 milliseconds). Subjects typically said that they could see all or most of the letters, but they could report only about half of them. Were the subjects right in saying that they could see all the letters? Sperling tried signalling the subjects with a tone. A high tone meant the subject was to report the top row, a medium tone indicated the middle row, etc. If the tone was given immediately after the stimulus, the subjects could usually get all the letters in the row, whichever row was indicated. But once they had named those letters, they usually could name no others. This experiment is taken to indicate some sort of raw visual storage, the "icon". But the crucial issue for my purposes is what it is like to be a subject in this experiment. My own experience is that I see all or almost all the letters, and this is what other subjects describe (Baars, 1988, p. 15). Focusing on one row allows me to report what letters are in that row, (and only that row) and again this is what other subjects report. Here is the description that I [think] is right and that I need for my case: I am P-conscious of all (or almost all--I'll omit this qualification) the letters at once, that is jointly, and not just as blurry or vague letters, but as specific letters (or at least specific shapes), but I don't have access to all of them jointly, all at once. (I would like to know whether others describe what it is like in this way, but the prejudice against introspection in psychology tends to keep answers to such questions from the journals.) One item of uncertainty about this phenomenon is that responses are serial; perhaps if some parallel form of response were available the results would be different. Ignoring that issue, the suggestion is that I am P-conscious, but not A-conscious, of all jointly. I am indebted to Jerry Fodor here.

It may be that some evidence for P-consciousness without A-consciousness can be derived from phenomena involving hypnosis. Consider the phenomenon known as hypnotic analgesia in which hypnosis blocks a patient's access to pain, say from an arm in cold water or from the dentist's drill. Pain must be P-conscious, it might be said, but access is blocked by the hypnosis, so perhaps this is P without A-consciousness? But what reason is there to think that there is any pain at all in cases of hypnotic analgesia? One reason is that there are the normal psychophysiological indications that would be expected for pain of the sort that would be caused by the stimulus, such as an increase in heart rate and blood pressure. (Melzack and Wall, 1988; Kihlstrom, et. al., 1992) Another (flakier) indication is that reports of the pain apparently can be elicited by Hilgard's "hidden observer" technique in which the hypnotist tries to make contact with a "hidden part" of the person who knows about the pain (Hilgard, 1986; Kihlstrom, 1987). The hidden observer often describes the pain as excruciating and also describes the time course of the pain in a way that fits the stimulation. Now there is no point in supposing that the pain is not P-conscious. If we believe the hidden observer, there is a pain that has phenomenal properties, and phenomenal properties could not be P-unconscious.

One way to think about this situation is that we have different persons sharing some part of one body. The pain is both P-conscious and A-conscious to the system that reports as the "hidden observer". This system doesn't control behavior, but since it can report, it may have that capacity under some circumstances. This reasoning is supported by the idea that if there is a P-conscious state in me that I don't have access to, then that state is not [mine] at all. A different way of thinking about what is going on is that there is one system, [the person], who has some sort of dissociation problem. There is P-conscious pain in there somewhere, but the person, himself or

herself, does not have access to that pain, as shown by the failure to report it, and by the failure to use the information to escape the pain. Only on this latter view would we have P without A-consciousness.

Another phenomenon that could lead to evidence of P without A-consciousness has to do with persistent reports over the years of P-conscious events under general anesthesia. Patients wake up and say that the operation hurt. (A number of doctors have told me that this is why doctors make a point of zapping patients with intravenous valium, a known amnestic, to wipe out memory of the pain. If the patients don't remember the pain, they won't sue.) General anesthesia is thought to suppress reasoning power in subanesthetic doses (Kihlstrom, 1987; See also Ghoneim, M.M., et. al., 1984), thus plausibly interfering with Executive function and A-consciousness, but I know of no reports that would suggest diminished P-consciousness. If P-consciousness were diminished much more than A-consciousness, for example, we could perhaps have analogs of super-blindsight, though I'm not sure how it would manifest itself. So if there are P-conscious states under general anesthetic, they may be states of more or less normal P-consciousness with diminished A-consciousness. Further, Crick and Koch (1990) mention that the aforementioned neural oscillations persist under light general anesthetic. Kihlstrom and Schacter, 1990, Kihlstrom and Couture, 1992, and Ghoneim and Block, 1993, conclude that the phenomenon depends in ways that are not understood on details of the procedure and the anesthetic cocktail, but there do appear to be some methods that show some kind of memory for events under anesthesia. Bennett, et. al. (1988) gave some patients under anesthesia suggestions to lift their index fingers at a special signal, whereas other patients were told to pull their ears. Control groups were given similar procedures without the suggestions. The result: the experimental group exhibited the designated actions at a much higher rate than controls. Of course, even if these results hold up, they don't show that the patients [heard] the suggestions under anesthesia. Perhaps what took place was some sort of auditory analog of blindsight.

An item of more use for present purposes comes from a study done on pilots during WW II by a pair of American dentists (Nathan, 1985, Melzack and Wall, 1988). The unpressurized cabins of the time caused pilots to experience sensations that as I understand it amount to some sort of re-creation of the pain of previous dental work. The mechanism appeared to have to do with stimulation of the sinuses caused by the air pressure changes. The dentists coined the term 'aerodontalgia' for this phenomenon. The dentists were interested in the relation of aerodontalgia to general and local anesthetic. So they did dental work on patients using combinations of general and local anesthetics. For example, they would put a patient under general anesthetic, and then locally anesthetize one side of the mouth, and then drill or pull teeth on both sides. The result (with stimulation of the nasal mucosa in place of the sinus stimulation caused by pressure changes): they found recreation of pain of previous dental work only for dental work done under general anesthesia, not for local anesthesia, whether or not the local was used alone or together with general anesthesia. Of course, there may have been no pain at all under general anesthesia, only memories of the sort that would have been laid down if there had been pain. But if you hate pain, and if both general and local anesthesia make medical sense, would [you] take the chance on general anesthesia? At any rate, the tantalizing suggestion is that this is a case of P-consciousness without A-consciousness.

The form of the target reasoning discussed misses the distinction between P-consciousness and A-consciousness and thus jumps from the fact that consciousness in some sense or other is missing simultaneously with missing creativity or voluntary action to the conclusion that P-consciousness functions to promote the missing qualities in normal people. But if we make the right distinctions,

we can investigate non-fallaciously whether any such conclusion can be drawn. Model 2 would identify P-consciousness with A-consciousness, thus embodying an aspect of the target reasoning. But Model 2 is disconfirmed by the apparent fact that much of our reasoning is neither P-conscious nor A-conscious. I have made further suggestions for phenomena that may provide examples of P-consciousness without A-consciousness, further disconfirming Model 2.

My purpose in this paper has been to expose a confusion about consciousness. But in reasoning about it I raised the possibility that it may be possible to find out something about the function of P-consciousness without knowing very much about what it is. Indeed, learning something about the function of P-consciousness may help us in finding out what it is. I would like to thank Tyler Burge, Susan Carey, Martin Davies, Bert Dreyfus, Paul Horwich, Jerry Katz, Leonard Katz, Joe Levine, David Rosenthal, Jerome Schaffer, Sydney Shoemaker, Stephen White and Andrew Young for their very helpful comments on earlier versions of this paper. I am also grateful to many audiences at talks on this material for their criticisms, especially the audience at the conference on my work at the University of Barcelona in June, 1993.

NOTES

1. See Bowers and Schacter, 1990, and Reingold and Merikle, 1990. The phenomenon just mentioned is very similar to phenomena involving "subliminal perception", in which stimuli are degraded or presented very briefly. Holender (1986) harshly criticises a variety of "subliminal perception" experiments, but the experimental paradigm just mentioned and many others, are in my judgement, free from the problems of some other studies. Another such experimental paradigm is the familiar dichotic listening experiments in which subjects wear headphones in which different programs are played to different ears. If they are asked to pay attention to one program, they can report only superficial features of the unattended program, but the unattended program influences interpretation of ambiguous sentences presented in the attended program. See Lackner and Garrett, 1973.

2. See, for example, Dennett and Kinsbourne's (1992b) scorn in response to my suggestion of Cartesian Modularism. I should add that in Dennett's more recent writings, Cartesian materialism has tended to expand considerably from its original meaning of a literal place in the brain at which "it all comes together" for consciousness. In reply to Shoemaker 1993 and Tye 1993, both of whom echo Dennett's (1991) and Dennett's and Kinsbourne's (1992a) admission that no one really is a proponent of Cartesian materialism, Dennett 1993 says "Indeed, if Tye and Shoemaker want to see a card-carrying Cartesian materialist, each may look in the mirror..." See also Jackson 1993.

3. But what is it about thoughts that makes them P-conscious? One possibility is that it is just a series of mental images or subvocalizations that make thoughts P-conscious. Another possibility is that the contents themselves have a P-conscious aspect independently of their vehicles. See Lormand, forthcoming.

4. I say both that P-consciousness is not an intentional property and that intentional differences can make a P-conscious difference. My view is that although P-conscious content cannot be reduced to intentional content, P-conscious contents often have an intentional aspect, and also P-conscious contents often represent in a primitive non-intentional way. A perceptual experience can represent space as being filled in certain ways without representing the object perceived as falling under any concept. Thus, the experiences of a creature which does not possess the concept of a donut could represent space as being filled in a donut-like way. See Davies (1992, forthcoming), Peacocke

(1992), and finally Evans (1982), in which the distinction between conceptualized and non-conceptualized content is first introduced.

5. Levine (1983) coined the term "explanatory gap", and has elaborated the idea in interesting ways; see also his (1993). Van Gulick (1993) and Flanagan (1992, p. 59) note that the more we know about the connection between (say) hitting middle C on the piano and the resulting experience, the more we have in the way of hooks on which to hang something that could potentially close the explanatory gap. Some philosophers have adopted what might be called a deflationary attitude towards the explanatory gap. See Levine (1993), Jackson (1993) and Chalmers (1993), Byrne (1993) and Block (1994).

6. I know some will think that I invoked inverted and absent qualia a few paragraphs above when I described the explanatory gap as involving the question of why a creature with a brain with a physiological and functional nature like ours couldn't have different experience or none at all. But the spirit of the question as I asked it allows for an answer that explains why such creatures cannot exist, and thus there is no presupposition that these are real possibilities. Levine (1983, 1993) stresses that the relevant modality is epistemic possibility.

7. What if an A-unconscious state causes an A-conscious state with the same content? Then it could be said that the first state must be A-conscious because it is in virtue of having that state that the content it shares with the other state satisfies the three conditions. So the state is A-unconscious by hypothesis, but A-conscious by my definition. (I am indebted to Paul Horwich.) I think what this case points to is a refinement needed in the notion of "in virtue of". One does not want to count the inferential promiscuity of a content as being in virtue of having a state if that state can only cause this inferential promiscuity via another state. I won't try to produce an analysis of 'in virtue of' here.

8. I have been using the P-consciousness/A-consciousness distinction in my lectures for many years, but it only found its way into print in my "Consciousness and Accessibility" (1990b), and my (1991, 1992, 1993). My claims about the distinction have been criticized in Searle (1990b, 1992) and Flanagan (1992); and there is an illuminating discussion in Humphreys and Davies (1993b), a point of which will be taken up in a footnote to follow. See also Levine's (1994) review of Flanagan which discusses Flanagan's critique of the distinction. See also Kirk (1992) for an identification of P-consciousness with something like A-consciousness.

9. Some may say that only fully conceptualized content can play a role in reasoning, be reportable, and rationally control action. If so, then non-conceptualized content is not A-conscious.

10. However, I acknowledge the empirical possibility that the scientific nature of P-consciousness is something to do with information processing. We can ill afford to close off empirical possibilities given the difficulty of solving the mystery of P-consciousness. Cf. Loar, 1990.

11. On my view, there are a number of problems with the first of these suggestions. One of them is that perhaps the representational content of pain is too primitive for a role in inference. Arguably, the representational content of pain is non-conceptualized. After all, dogs can have pain and one can reasonably wonder whether dogs have the relevant concepts at all. Davies and Humphreys (1993b) discuss a related issue. Applying a suggestion of theirs about the higher order thought notion of consciousness to A-consciousness, we could characterize A-consciousness of a state with non-conceptualized content as follows: such a state is A-conscious if, in virtue of one's having the state, its content would be inferentially promiscuous and available for rational control

of action and speech if the subject were to have had the concepts required for that content to be a conceptualized content. The idea is to bypass the inferential disadvantage of non-conceptualized content by thinking of its accessibility counterfactually-- in terms of the rational relations it would have if the subject were to have the relevant concepts. See Lormand (forthcoming) on the self-representing nature of pain.

12. Later in this paper I introduce the distinction between creature consciousness and state consciousness. In those terms, transitivity has to do primarily with creature consciousness, whereas in the case of P-consciousness and A-consciousness, it is state consciousness which is basic. See the discussion at the end of this section.

13. The distinction has some similarity to the sensation/perception distinction; I won't take the space to lay out the differences. See Humphrey (1992) for an interesting discussion of the latter distinction.

14. Tye, forthcoming argues (on the basis of a neuropsychological claims) that the visual information processing in blindsight includes no processing by the object recognition system or the spatial attention system, and so is very different from the processing of normal vision. This point does not challenge my claim that the super-blindsight case is a very limited partial zombie. Note that super-blindsight, as I describe it does not require object recognition or spatial attention. Whatever it is that allows the blindsight patient to discriminate an 'X' from an 'O' and a horizontal from a vertical line will do. I will argue later that the fact that such cases do not exist, if it is a fact, is important. Humphrey (1992) suggests that blindsight is mainly a motor phenomenon--the patient is perceptually influenced by his own motor tendencies.

15. If you are tempted to deny the existence of these states of the perceptual system, you should think back to the total zombie just mentioned. Putting aside the issue of the possibility of this zombie, note that on a computational notion of cognition, the zombie has all the same A-conscious contents that you have (if he is your computational duplicate). A-consciousness is an informational notion. The states of the super-blindsighter's perceptual system are A-conscious for the same reason as the zombie's.

16. Actually, my notion of A-consciousness seems to fit the data better than the conceptual apparatus she uses. Blindsight isn't always more degraded in any normal sense than sight. 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. It would be better to understand her "degraded" in terms of lack of access. Notice that the super-blindsighter I have described is just a little bit different (though in a crucial way) from the ordinary blindsight patient. In particular, I am not relying on what might be thought of as a full-fledged quasi-zombie, a super-duper-blindsighter whose blindsight is every bit as good, functionally speaking, as his sight. In the case of the super-duper blindsighter, the only difference between vision in the blind and sighted fields, functionally speaking, is that the quasi-zombie himself regards them differently. Such an example will be regarded by some (though not me) as incoherent--see Dennett, 1991, for example. But we can avoid disagreement about the super-duper-blindsighter by illustrating the idea of A-consciousness without P-consciousness by appealing only to the super-blindsighter. Functionalists may want to know why the super-blindsight case counts as A-conscious without P-consciousness. After all, they may say, if we have really high quality access in mind, the super-blindsighter that I have described does not have it, so he lacks both P-consciousness and really high quality A-consciousness. The super-duper-blindsighter, on the other hand, has both, according to the functionalist, so in neither case, the objection goes, is there A-consciousness without P-

consciousness. But the disagreement about the super-duper- blindsighter is irrelevant to the issue about the super- blindsighter, and the issue about the super-blindsighter is merely verbal. I have chosen a notion of A-consciousness whose standards are lower in part to avoid conflict with the functionalist. I believe in the possibility of a quasi- zombie like the super-duper-blindsighter, but the point I am making here does not depend on it. There is no reason to frame notions so as to muddy the waters with unnecessary conflicts when the point I am making in this paper is one that functionalists can have some agreement with. One could put the point by distinguishing three types of access: (1) really high quality access, (2) medium access and (3) poor access. The actual blindsight patient has poor access, the super-blindsight patient has medium access and the super-duper blindsight patient--as well as most of us--has really high quality access. The functionalist identifies P- consciousness with A-consciousness of the really high quality kind. I am defining 'A-consciousness'--and of course, it is only one of many possible definitions--in terms of medium access, both to avoid unnecessary conflict with the functionalist, and also so as to reveal the fallacy of the target reasoning. I choose medium instead of really high quality access for the former purpose, and I choose medium instead of poor access for the latter purpose. Though functionalists should agree with me that there can be A-consciousness without P-consciousness, some functionalists will see the significance of such cases very differently from the way I see them. Some functionalists will see the distinction between A-consciousness and P-consciousness as primarily a difference in degree rather than a difference in kind, as is suggested by the contrast between really high quality access and medium access. So all that A- consciousness without P-consciousness illustrates, on this functionalist view, is some access without more access. Other functionalists will stress kind of information processing rather than amount of it. The thought behind this approach is that there is no reason to think that the P-consciousness of animals whose capacities for reasoning, reporting and rational guidance of action are more limited than ours thereby have anything less in the way of P- consciousness. The functionalist can concede that this thought is correct, and thereby treat the difference between A-consciousness and P-consciousness as a difference of kind, albeit kind of information processing.

17. Thus, there is a conflict between this physiological claim and the Schacter model which dictates that destroying the P-consciousness module will prevent A- consciousness.

18. There is a misleading aspect to this example-- namely that to the extent that 'conscious' and 'aware' differ in ordinary talk, the difference goes in the opposite direction.

19. Of course, even those who don't belief in P- consciousness at all, as distinct from A-consciousness, can accept the distinction between a noise that is A-conscious and a noise that is not A-conscious. There is a more familiar situation which illustrates the same points. Think back to all those times when you have been sitting in the kitchen when suddenly the compressor in the refrigerator goes off. Again, one might naturally say that one was aware of the noise, but only at the moment in which it went off was one consciously aware of it. I didn't use this example because I am not sure that one really has P-consciousness of the noise of the compressor all along; habituation would perhaps prevent it. Perhaps what happens at the moment it goes off is that one is P-conscious of the change only.

20. See White (1987) for an account of why self- consciousness should be firmly distinguished from P- consciousness, and why self-consciousness is more relevant to certain issues of value.

21. The pioneer of these ideas in the philosophical literature is David Armstrong (1968, 1980). William Lycan (1987) has energetically pursued self-scanning, and David Rosenthal (1986, 1993), Peter Carruthers (1989, 1992) and Norton Nelkin (1993) have championed higher order thought.

See also Natsoulas (1993) Lormand (forthcoming) makes some powerful criticisms of Rosenthal.

22. To be fair to Rey, his argument is more like a dilemma: for any supposed feature of consciousness, either a laptop of the sort we have today has it or else you can't be sure you have it yourself. So in the case of P-consciousness, the focus might be on the latter disjunct.

23. Interestingly, she was in many respects much worse at many face-perception tasks than LH (the prosopagnosic mentioned earlier)--she couldn't match photographs of faces, for example. I have noticed that people who know little about anosognosia tend to favor various debunking hypotheses. That is, they assume that the experimenters have made one or another silly mistake in describing the syndrome, because, after all, how could anyone fail to notice that they can't recognize faces, or worse, that they are blind. See Young, et. al., 1993, for a good debunking of the debunking hypotheses.

24. There is an additional problem in the reasoning that I won't go into except here. There is a well-known difficulty in reasoning of the form: X is missing; the patient has lost the ability to do blah-blah; therefore a function of X is to facilitate blah-blahing. In a complex system, a loss may reverberate through the system, triggering a variety of malfunctions that are not connected in any serious way with the function of the missing item. An imperfect but memorable example (that I heard from Tom Bever) will illustrate: the Martians want to find out about the function of various Earthly items. They begin with The Pentagon, and focus in on a particular drinking fountain in a hall on the third floor of the North side of the building. "If we can figure out what that is for", they think, "we can move on to something more complex." So they vaporize the drinking fountain, causing noise and spurting pipes. Everyone comes out of their office to see what happened and the Martians conclude that the function of the fountain was to keep people in their offices. The application of this point to the petit mal case is that even if I am right that it is A-consciousness, not P-consciousness, that is diminished or missing, I would not jump to the conclusion that A-consciousness has a function of adding powers of discrimination, flexibility and creativity. Creativity, for example, may have its sources in the unA-conscious, requiring powers of reasoning and control of action and reporting only for its expression.

25. Indeed, in the italicized passage above (*italics mine*) there is an implicit suggestion that perhaps there are P-conscious events of which no record is made. I could only find one place in the book where Penfield says anything that might be taken to contradict this interpretation: "Thus, the automaton can walk through traffic as though he were aware of all that he hears and sees, and so continue on his way home. But he is aware of nothing and so makes no memory record. If a policeman were to accost him he might consider the poor fellow to be walking in his sleep." (60) But to properly understand this, we need to know what he means by "awareness", and what he thinks goes on in sleep. Judging by Penfield's use of synonyms, by "awareness" he means something in the category of the higher order thought analyses or the self-consciousness sense. For example, in discussing his peculiar view that ants are conscious, he seems to use 'conscious' and 'aware' to mean self-aware (62, 105, 106) Further, he makes it clear that although the mind is shut off during sleep, the sensory cortex is quite active.

26. A similar line of reasoning appears in Shevrin, 1992; he notes that in subliminal perception, we don't fix the source of a mental content. Subliminal percepts aren't conscious, so consciousness must have the function of fixing the source of mental contents.

27. Baars, 1988, p. 356. Though Baars is talking about the function of "conscious experience", he does have a tendency to combine P-consciousness with A-consciousness under this heading.

28. The Collapse Hypothesis should not be confused with Marcel's (1988, p. 135-7) Identity Hypothesis, which hypothesizes that the processing of stimuli is identical with consciousness of them. As Marcel points out, blindsight and similar phenomena suggest that we can have processing without consciousness.

29. I am indebted to Jerry Fodor here.

30. I would like to thank Tyler Burge, Susan Carey, Martin Davies, Bert Dreyfus, Paul Horwich, Jerry Katz, Leonard Katz, Joe Levine, David Rosenthal, Jerome Schaffer, Sydney Shoemaker, Stephen White and Andrew Young for their very helpful comments on earlier versions of this paper. I am also grateful to many audiences at talks on this material for their criticisms, especially the audience at the conference on my work at the University of Barcelona in June, 1993.

REFERENCES

Akins, K. (1993) A bat without qualities. In Davies and Humphreys (1993a)

Alston, W. (1967) Religion. In [The Encyclopedia of Philosophy]. Macmillan/Free Press, 140-145.

Anderson, J. (1993) To see ourselves as others see us: a response to Mitchell. [New Ideas in Psychol] 11, 3:339-346

Andrade, J. (1993) Consciousness: current views. In Jones, 1993.

Armstrong, D. M. (1968) [A Materialist Theory of Mind]. Humanities Press

_____ What is consciousness? In [The Nature of Mind]. Cornell University Press

Baars, B.J. (1988) [A cognitive Theory of Consciousness] Cambridge University Press

Block, N. (1980) What is functionalism? In N. Block (ed) [Readings in the Philosophy of Psychology] vol 1. Harvard University Press

_____ (1990a) Inverted earth. In [Philosophical Perspectives] 4 ed J. Tomberlin. Ridgeview

_____ (1990b) Consciousness and accessibility. [Behavioral and Brain Sciences] 13: 596-598

_____ (1991) Evidence against epiphenomenalism. [Behavioral and Brain Sciences] 14 (4):670-672

_____ (1992) Begging the question against phenomenal consciousness. [Behavioral and Brain Sciences]

_____ (1993) Review of D. Dennett, [Consciousness Explained] [The Journal of Philosophy] XC,4:181-193

_____ (1994) "Functionalism", "Qualia". In S. Guttenplan (ed) [A Companion to Philosophy of Mind] Blackwell.

Bowers, J. & Schacter, D. (1990) Implicit memory and test awareness. [Journal of Experimental Psychology: Learning, Memory and Cognition] 16:3: 404-416

- Bornstein, R. & Pittman, T. (1992) [Perception without Awareness.] Guilford Press:New York
- Byrne, A. (1993) [The Emergent Mind], Princeton University Ph.D. thesis.
- Byrne, R.W. (1993) The meaning of 'awareness': a response to Mitchell [New Ideas in Psychol] 11, 3:347-350
- Carruthers, P. (1989) Brute experience. [Journal of Philosophy] 86
- _____(1992) Consciousness and concepts. [Proceedings of the Aristotelian Society, Supplementary Volume LXVI]: 40-59
- Chalmers, D.J. (1993) [Toward a Theory of Consciousness]. University of Indiana Ph.D. thesis
- Churchland, P.S. (1983) Consciousness: the transmutation of a concept [Pacific Philosophical Quarterly] 64: 80-93
- Churchland, P.S. (1986) Reduction and the neurobiological basis of consciousness. In Marcel and Bisiach (1988)
- Coslett, H. and Saffran, E. (1994) Mechanisms of implicit reading in Alexia. In [The Neuropsychology of High-Level Vision], M. Farah and G. Ratcliff, eds. Erlbaum.
- Cowie, A. & Stoerig, P. (1992) Reflections on blindsight. In Milner and Rugg (1992)
- Crick, F. and Koch, C. (1990) Towards a neurobiological theory of consciousness [Seminars in the Neurosciences] 2:263-275
- Davies, M. & Humphreys, G. (1993a) [Consciousness] Blackwell
- Davies, M. & Humphreys, G. (1993b) Introduction. In Davies and Humphreys (1993a)
- Davies, M. (1992) Perceptual content and local supervenience. [Proceedings of the Aristotelian Society] 92:21-45
- Davies, M. (forthcoming) Externalism and experience. In A. Clark, J. Exquerro, J. Larrazabal (eds) [Categories, Consciousness and Reasoning]. Dordrecht.
- de Lannoy, J. Two theories of a mental model of mirror self-recognition: a response to Mitchell. [New Ideas in Psychol] 11, 3:337-338
- Dennett, D. (1986) Julian Jaynes' software archeology. [Canadian Psychology] 27, 2: 149-154
- _____(1991) [Consciousness Explained]. Little Brown
- _____(1993) The message is: there is no medium. In [Philosophy and Phenomenological Research] III, 4
- Dennett, D. & Kinsbourne, M. (1992a) Time and the observer: the where and when of consciousness in the brain [Behavioral and Brain Sciences] 15: 183-200
- Dennett, D. & Kinsbourne, M. (1992b) Escape from the Cartesian theater [Behavioral and Brain

Sciences] 15: 234-248

Dimond, S. (1976) Brain circuits for consciousness. [Brain, Behavior and Evolution] 13: 376-395

Dretske, F. (1993) Conscious experience. [Mind] 102, 406:263-284

Dupre, J. (1981) Natural kinds and biological taxa. [Philosophical Review] 90:66-90

Edelman, G. (1989) [The Remembered Present: A Biological Theory of Consciousness]. Basic Books

Etcoff, N. L. & Freeman, R & Cave, K. Can we lose memories of faces? Content specificity and awareness in a prosopagnosic. [Journal of Cognitive Neuroscience] 3, 1

Etcoff, N.L. & Magee, J.J. (1992) Covert recognition of emotional expressions. [Journal of Clinical and Experimental Neuropsychology] 14:95-96

Evans, G. (1982) [The Varieties of Reference]. Oxford University Press Flanagan, O. (1991) [The Science of the Mind] 2nd ed MIT Press

Farah, M. (1994) Visual perception and visual awareness after brain damage: a tutorial overview. In Umiltà and Moscovitch, 1994

Flanagan, O. (1992) [Consciousness Reconsidered] MIT Press

Gallup, G. (1982) Self-awareness and the emergence of mind in primates. [American Journal of Primatology] 2: 237-248

Gallup, G. & Povinelli, D. Mirror, mirror on the wall, which is the most heuristic theory of them all? A response to Mitchell. [New Ideas in Psychol] 11, 3:327-335

Ghoneim, M., Hinrichs, J., Mewaldt, S. (1984) Dose-response analysis of the behavioral effects of diazepam: 1. Learning and memory. [Psychopharmacology] 82:291-295

Ghoneim, M., Block, R. (1993) Learning during anesthesia. In Jones, 1993

Goldman, A. (1993a) The psychology of folk psychology [The Behavioral and Brain Sciences] 16:1:15-28

Goldman, A. (1993b) Consciousness, folk psychology and cognitive science. [Consciousness and Cognition] II, 3

Goodale, M. and Milner, D. (1992) Separate visual pathways for perception and action

Harman, G. (1990) The intrinsic quality of experience. In [Philosophical Perspectives] 4 ed J. Tomberlin. Ridgeview.

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

Hilgard, E. (1986) [Divided Consciousness] 2nd edition. John Wiley 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

- Hill, C. (1991) [Sensations; A Defense of Type Materialism]. Cambridge.
- Humphrey, N. (1992) [A History of the Mind]. Simon & Schuster
- Huxley, T.H. (1866) [Lessons in Elementary Psychology] 8, p. 210. Quoted in Humphrey, 1992.
- Jackendoff, R. (1987) [Consciousness and the Computational Mind]. MIT Press
- Jackson, F. (1977) [Perception]. Cambridge University Press
- Jackson, F. (1986) What Mary didn't know. [Journal of Philosophy] 83: 291-95
- _____(1993) Appendix A (for philosophers). In [Philosophy and Phenomenological Research] III, 4
- _____(1993) 'Armchair metaphysics'. In J. O'Leary-Hawthorne and M. Michael (eds) [Philosophy in Mind]. Kluwer
- Jacoby, L., Toth, J., Lindsay, D., Debner, J. (1992) Lectures for a layperson: methods for revealing unconscious processes. In Bornstein and Pittman, 1992.
- James, W. (1890) [The Principles of Psychology], 2 vols. Dover, 1950
- Jaynes, J. (1976) [The Origin of Consciousness in the Breakdown of the Bicameral Mind] Houghton-Mifflin
- Jones, J. G. (1993) [Depth of Anesthesia] Little Brown:Boston
- Kihlstrom, J. (1987) The cognitive unconscious. [Science] 237:1445-1452
- Kihlstrom, J. & Schacter, D. (1990) Anaesthesia, amnesia, and the cognitive unconscious. In B. Bonke (ed) [Memory and Awareness in Anaesthesia]. Swets & Zeitlinger
- Kihlstrom, J & Barnhardt, T. & Tataryn, D. (1992) Implicit perception. In Bornstein and Pittman, 1992.
- Kihlstrom, J. & Couture, L. (1992) Awareness and information processing in general anesthesia. [Journal of Psychopharmacology] 6(3) 410-417
- Kirk, R. (1992) Consciousness and concepts. [Proceedings of the Aristotelian Society, Supplementary Volume LXVI]: 23-40
- Kuhn, T. (1964) A function for thought experiments. In [Melanges Alexandre Koyre] Vol 1. Hermann:307-334
- Lackner, J. & Garrett, M. (1973) Resolving ambiguity: effects of biasing context in the unattended ear. [Cognition] 1: 359-372
- 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]

_____ (1993) On leaving out what it is like. In Davies and Humphreys (1993a)

_____ (1994) Review of Owen Flanagan's [Consciousness Reconsidered] In [The Philosophical Review]

Loar, B. (1990) Phenomenal properties. In J. Tomberlin (ed) [Philosophical Perspectives: Action Theory and Philosophy of Mind]. Ridgeview.

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

Lycan, W. (1987) [Consciousness] MIT Press

Mandler, G. (1985) [Cognitive Psychology]. Erlbaum, Chapter 3.

McGinn, C. (1991) [The Problem of Consciousness]. Blackwell

_____ (1993) Consciousness and cosmology: hyperdualism ventilated. In Davies and Humphreys (1993a)

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. [The Behavioral and Brain Sciences] 9: 40-41

_____ (1988) Phenomenal experience and functionalism. In Marcel and Bisiach (1988)

Marcel, A. J. & 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] XCIII:39-58

Melzack, R. & Wall, P. (1988). [The Challenge of Pain], 2nd edition. Penguin

Menzel, E., Savage-Rumbaugh, E., Lawson, J. (1985) Chimpanzee ([Pan troglodytes]) spatial problem solving with the use of mirrors and televised equivalents of mirrors. [Journal of Comparative Psychology] 99, 211-217.

Milner, B & Rugg, M. (1992) (eds) [The Neuropsychology of Consciousness]. Academic Press

Mitchell, R. W. (1993a) Mental models of mirror self-recognition: two theories. In [New Ideas in Psychology] 11, 295-325

Mitchell, R. W. (1993b) Recognizing one's self in a mirror? A reply to Gallup and Povinelli, de Lannoy, Anderson, and Byrne. In [New Ideas in Psychology] 11: 351-377

Moscovitch, M., Goshen-Gottstein, Y., Vriezen, E. (1994) "Memory without conscious recollection: a tutorial review from a neuropsychological perspective." In Umiltà and Moscovitch, 1994

Nagel, T. (1974) What is it like to be a bat? [Philosophical Review] _____(1979) [Mortal Questions] Cambridge University Press

_____(1986) [The View from Nowhere] Oxford University Press

Nathan, P. (1985) Pain and nociception in the clinical context. [Phil. Trans. R. Soc. Lond. B] 308: 219-226

Natsoulas, T. (1993) What is wrong with the appendage theory of consciousness? [Philosophical Psychology] VI,2: 137-154

Nelkin, N. The connection between intentionality and consciousness. In Davies and Humphreys (1993a)

Nisbett, R. and Wilson, T. (1977) Telling more than we can know: verbal reports on mental processes, [Psychological Review] 84

Peacocke, C. (1983) [Sense and Content] Oxford University Press

_____(1992) [A Study of Concepts.] MIT Press

Pendlebury, M. (1992) Experience, theories of. In J. Dancy & E. Sosa, [A Companion to Epistemology]. 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 Jones, 1993

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 Putnam's [Mind, Language and Reality] Cambridge University Press

Reingold, E. & Merikle, P.(1993) Theory and measurement in the study of unconscious processes. In Davies and Humphreys (1993a)

Rey, G. (1983) A reason for doubting the existence of consciousness. In [Consciousness and Self-Regulation], vol 3. R. Davidson, G. Schwartz, D. Shapiro (eds). Plenum

_____(1988) A question about consciousness. In [Perspectives on Mind], H> Otto & J. Tuedio (eds). Reidel

Rosenthal, David (1986) Two concepts of consciousness. [Philosophical Studies] 49: 329-359

_____(1993) Thinking that one thinks. In Davies and Humphreys (1993a)

Schacter, D. (1989) On the relation between memory and consciousness: dissociable interactions and conscious experience. In: H. Roediger & F. Craik (eds), [Varieties of Memory and Consciousness: Essays in Honour of Endel Tulving] Erlbaum

Searle, J. (1987) [Intentionality] Cambridge

_____(1990a) Consciousness, explanatory inversion and cognitive science. [Behavioral and Brain Sciences] 13:4:585-595

_____(1990b) Who is computing with the brain? [Behavioral and Brain Sciences] 13:4: 632-642

_____(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.

Shallice, T. (1988b) Information-processing models of consciousness: possibilities and problems. In Marcel and Bisiach, 1988.

Shevrin, H. (1992) Subliminal perception, memory and consciousness: cognitive and dynamic perspectives. In Bornstein and Pittman

Shoemaker, S. (1975) Functionalism and qualia. [Philosophical Studies] 27: 291-315. _____

(1981a) Absent qualia are impossible--a reply to Block. [The Philosophical Review] 90,4: 581-599

_____(1981b) The inverted spectrum. [The Journal of Philosophy] 74, 7:357-381

_____(1993) Lovely and suspect ideas. In [Philosophy and Phenomenological Research] III, 4: 905-910

_____(1994) Phenomenal Character, [Nous]

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

Stich, S. (1978) Autonomous psychology and the belief-desire thesis. [The Monist] 61

Tye, M. (1991) [The Imagery Debate] MIT Press

_____(1993) Reflections on Dennett and consciousness. In [Philosophy and Phenomenological Research] III, 4

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

_____(forthcoming-b) Does pain lie within the domain of cognitive psychology? In J. Tomberlin (ed), [Philosophical Perspectives] 8

Umiltà, C. & Moscovitch, M. (1994) [Attention and Performance XV] MIT Press

Van Gulick (1989) What difference does consciousness make? [Philosophical Topics] 17,1: 211-230

Van Gulick (1993) Understanding the phenomenal mind: are we all just armadillos? In Davies and Humphreys (1993a)

Weiskrantz, L. (1986) [Blindsight] Oxford University Press

_____ (1988) Some contributions of neuropsychology of vision and memory to the problem of consciousness. In Marcel and Bisiach (1988).

_____ (1992) Introduction: Dissociated issues. In B. Milner & M. Rugg (1992)

White, S. L. (1987) What is it like to be an homunculus. [Pacific Philosophical Quarterly] 68: 148-174

_____ (1991) Transcendentalism and its discontents. In White, [The Unity of the Self], MIT Press.

Wiser, M. & Carey, S. (1983) When heat and temperature were one. In D. Gentner and A. Stevens (eds) [Mental Models] Lawrence Erlbaum

Young, A. W. (1994a) Covert recognition. In M. Farah & G. Ratcliff (eds) [The Neuropsychology of Higher Vision: Collected Tutorial Essays]. Erlbaum

Young, A. W. (1994b) Neuropsychology of awareness In M. Kappinen & A. Revonsuo (eds), [Consciousness in philosophy and cognitive neuroscience] Erlbaum

Young, A. W. & De Haan, E. (1993) Impairments of Visual Awareness. In Davies and Humphreys (1993a)