

## The Naturalists versus the Skeptics: The Debate Over a Scientific Understanding of Consciousness

Valerie Gray Hardcastle

*Virginia Polytechnic Institute and State University*

There are three basic skeptical arguments against developing a scientific theory of consciousness: (1) theory cannot capture a first person perspective; (2) consciousness is causally inert with respect to explaining cognition; and (3) the notion "consciousness" is too vague to be a natural kind term. Although I am sympathetic to naturalists' counter-arguments, I also believe that most of the accounts given so far of how explaining consciousness would fit into science are incorrect. In this essay, I indicate errors my colleagues on both sides of the fence make in thinking about this issue, as well as outline data relevant to distinguishing conscious states from unconscious ones empirically.

Skeptical reactions to the task of developing a scientific theory of consciousness fall into three broad categories. The first holds that scientific investigation of consciousness is impossible because there is no way for any theoretical description to capture essentially qualitative feel (Kripke, 1980; McGinn, 1989; Nagel, 1974; Searle, 1980). A second argues that, at best, psychology should be indifferent to consciousness. Science is interested in explaining causal processes, so insofar as consciousness plays no interesting causal role in any model of the mind, it behooves the cognitive sciences to ignore the phenomena (Hannay, 1987; Wilkes, 1984; see also Block and Fodor, 1972). A third points out that those engaged in these debates have a confused picture of the supposed referents of "consciousness," and it is most likely that "consciousness" is not a natural kind term at all. Our notions about consciousness are fuzzy and probably contradictory, and since no one

---

I would like to thank Patricia Churchland for her helpful discussions and Owen Flanagan and an anonymous referee for their helpful comments on an earlier draft of this paper. Requests for reprints should be sent to Valerie Gray Hardcastle, Ph.D., Department of Philosophy, Virginia Polytechnic Institute and State University, Blacksburg, Virginia 24061-0126.

has devised a suitable replacement for the folk-psychological term, talk about scientific theories of consciousness is premature (P.S. Churchland, 1983; Dennett, 1982; Rey, 1983, 1988; Wilkes, 1988).

However, I believe that those who scoff at ever getting a scientific account of what it is like to be a person do not understand the abstract relation between a theory and the phenomena it explains. We can think of a scientific theory as an abstract picture of some relations among variables, and a particular model within a theory as a picture with all of its parameters fixed. To accept a scientific theory as explanatory then is to believe that what its model describes accurately reflects the data we believe exist in the world around us (see also van Fraassen, 1980, 1989). Since we are conscious, it falls well within the scope of the scientific enterprise to account for our observations concerning consciousness in a theoretical model. Likewise, those who dismiss the scientific study of consciousness because they think that raw feels play no causal role in our mental economy misunderstand how scientific models deal with causal histories. After all, we do posit phenomenological phenomena to explain perceptual reports regardless of the role they may or may not play in information processing. Uncovering the causal history which leads to these perceptual phenomena surely falls within the domain of science (cf., Salmon, 1984). If we view theories as empirically adequate models of the world, then no empirical phenomenon is beyond the pale of scientific theorizing. As Rorty (1982) writes, the real problem is not how to study consciousness empirically, but to convince others that, from a methodological point of view, studying it is just like studying anything else.

However, even if one does subscribe to this pragmatic understanding of science, one still has to face the charge that the concept consciousness is too simple or too vague to help frame the generalizations in a mature science of mind, either as explanans or as explanandum. One common tactic in developing arguments of this sort is using our intuitions to identify consciousness with some psychological process or other and then demonstrating that identifications of this sort are conceptually confused (cf., P.S. Churchland, 1983; Wilkes, 1988). Consciousness then appears to be simply the wrong sort of term we would use in a theory. The virtue of this strategy is that we can discuss consciousness cogently because we eliminate consciousness from our psychological ontology except as a place-holder for some other information process. We can thereby avoid mentioning qualia, the big conundrum of consciousness, except obliquely. Of course, the obvious problem with strategies of this sort is exactly Nagel's point: consciousness in fact does have some sort of phenomenological feel, and we generally take that feel to be *prima facie* both a necessary and sufficient condition for consciousness. All the arguments show is that qualia cannot be identified with a psychological process. Indeed, I shall argue that positing conscious experience is needed to

account for some of our behavior, and so ultimately, a scientific theory about consciousness will not only be possible; it will be necessary.

### The Study of Consciousness as a Scientific Enterprise

On the one hand, we know that brains somehow cause consciousness. On the other hand, we so thoroughly fail to understand this causal connection that *any* link between the two seems utterly mysterious, even eerie. Having a solution to the mind–body problem would remove the spookiness of the link between the brain and consciousness. More specifically, having a solution means explaining how the subjective aspects of experience depend upon the brain (and possibly its environment).<sup>1</sup>

There is a fairly vague but deep sentiment that even if we could point to some natural attribute as the causal factor in question, something important would still be left out, viz., the first-personness of consciousness. Some argue that this general feeling of uneasiness about incorporating consciousness into science results from misunderstanding what science does. Although I agree that there is something profoundly wrong about letting our untutored intuitions govern a priori the possibility of knowledge, I also believe that most of the accounts given so far of how explaining consciousness would fit into science are incorrect as well. In this section, I indicate some of the errors my colleagues on both sides of the fence make in thinking about this issue.

Skeptics about the feasibility of the entire project worry that any naturalistic or scientific explanation of consciousness must fail to capture the true and fundamental character of personal subjective experience, in virtue of being a third-person “objective” account. Facts about experience are accessible only from particular first-person stances and any “scientific” theory will ignore single points of view (Jackson, 1982, 1986; McGinn, 1989; Nagel, 1974; Searle, 1980, 1990). If any naturalistic story cannot even adequately *describe* half of the psychophysical equation under question, then it stands to reason that whatever property of the brain we point to as the relevant causal factor would fail to capture consciousness adequately as well (except perhaps accidentally).

#### *Reduction of First Person Points of View*

We can see this worry as having two components: (1) any objective science will fail to capture what consciousness is like for any particular creature, and (2) any completed scientific account will fail to analyze consciousness exhaus-

---

<sup>1</sup>I do not view this position as being inherently hostile to something like Sperry's emergent interactionism (e.g., Sperry, 1969, 1987, 1991). Sperry holds that consciousness is a property of the brain, albeit an emergent one (1965, 1991) and so insofar as we can discover the macroproperties of the brain correlated with consciousness's causal powers, then our projects are in accord.

tively (cf., Flanagan, 1985). Supporters of the project of reducing or replacing first-person accounts of conscious phenomena with objective third-person accounts respond that both components of the worry indicate a fundamental misapprehension of the abstractness of scientific theory. They hold that while the first claim may be true, it is irrelevant, and that the second is simply false (Flanagan, 1985; Rorty, 1982; see also Dennett, 1982).

One particular line of response to the skeptic goes as follows: insofar as science is interested in uncovering general laws which govern the behavior of types, not tokens, then individual instances of phenomenological experience are not the explananda of theories. Instead, theories want to explain what is common to all experiencing creatures. Individual subjective experiences and first-person reports of those experiences serve as the evidence for the more encompassing scientific theories. The skeptics simply misunderstand where first-person points of view fit into the game of science.

However, this reply begs the question against the skeptics.<sup>2</sup> If first-person points of view are the *only way* to describe phenomenological experiences adequately, then whether any resultant theory is non-individualistic is inconsequential because theory-building using third-person perspectives could never get off the ground in the first place. The skeptics' point, I take it, is stronger than merely that any completed theory will abandon individual instances. They believe that any attempt to describe the *data* objectively must fail since the gap between the mental and the brain "cannot be bridged just by applying concepts drawn from one side to items that belong on the other side" (McGinn, 1986, p. 356n). And they conclude that any "objective, physical theory will [inevitably] abandon that [single] point of view" (Nagel, 1974, p. 437).

Nevertheless, the skeptics are wrong, for a reductive move is precisely how the gap is bridged. The skeptics fail to understand how first-person accounts can be used in theory-building and the connection between our everyday "folk" descriptions of how the world seems to us and our scientific theories about the workings of that world. Science regularly and nonproblematically redescribes the way the world seems to us from our first-person point of view in third-person objective terms by applying naturalistic concepts to subjectively experienced objects in the world. We know, for example, that objects that appear red to us do so because they reflect a certain wavelength of electromagnetic radiation. We know that surfaces that seem warm to us do so because their mean molecular kinetic energy is above a certain level relative to the MMKE of our skin. And so on.

---

<sup>2</sup>Flanagan suggests that the skeptics' question would not be begged if the naturalists admit that we use token reports of phenomenal experiences to develop an objective type (Flanagan, personal communication). However, I maintain that the skeptics' point is stronger: naturalists who believe that they can build objective types from first person reports are mistaken (see esp. McGinn, 1989).

This form of concept reduction falls out of the broader reduction of one conceptual framework to another. We get a concept reduction when we claim that the reducing predicate in the new conceptual framework can be applied in (at least) all the cases in which the reduced predicate was thought to have applied in the old conceptual framework. And we are licensed to make these identity statements in virtue of a relatively smooth reduction of one conceptual framework to another. The purpose behind any reductions of this sort would be to explain the relative success of our folk conceptual generalizations about the world by reducing those generalizations to a more rigorous and detailed scientific vocabulary. By explicating the physical mechanisms in virtue of which phenomenological events conform to the reduced generalizations, we would thereby exhibit first-person generalizations just as instances of true scientific theories.

In our case, if our folk conceptual framework concerning consciousness reduces to a scientific theory, then all the kind predicates in our first-person descriptions of experience should be more or less coextensive with the kind predicates in the corresponding scientific physical descriptions. Specified bridge principles connect all the primitive terms appearing in the reduced folk conceptual framework with one or more terms in reducing theory. In the strictest case of reduction, the bridge principles taken together exhaust the domain of the reduced framework and we can derive the reduced conceptual framework from the union of the reducing scientific theory and the bridge principles.<sup>3</sup> If our first-person common sense conceptual framework reduces to a scientific theory, then for some version of our every-day phenomenological kind predicates, there will exist co-extensive third-person scientific predicates, and the generalizations which express this co-extension will help explain the accuracy of the folk generalizations.

We can now see how concept reduction would work. A predicate *F* in the old conceptual framework reduces to predicate *G* in some new theory just in case the new theory reduces the old framework, and bridge laws link *F* and *G*. Intuitively, we can say that falling under *F* reduces to falling under *G* if and only if the causal powers of *F*-ness (as outlined in the old conceptual framework) are a subset of the causal powers of *G*-ness (as outlined in the new theory) [cf., P.M. Churchland, 1989, p. 50].

---

<sup>3</sup>Of course, this version of reductionism is really too strong since even the paradigm cases of reduction do not meet these deductive requirements. (For a more complete discussion, see Schaffner, 1967). One suggestion for getting around this difficulty is to derive a theory closely resembling the reduced theory from reducing theory—we “correct” the reduced theory before we derive it from the union of the reducing theory and the bridge principles (see e.g., P.M. Churchland, 1979, 1985; Hooker, 1981; Schaffner, 1967). But even with this weaker notion of reduction we still get some law-like correspondence between the entities of the reducing and the reduced theories, and we can use this correspondence to explain why (some version of) the reduced theory holds.

Skeptics who believe that any scientific account will inevitably fail to capture what consciousness is like believe that F-ness exists and co-occurs with some set of circumstances addressed by the new theory, but F cannot be reduced to any predicate in the new theory. This would occur whenever the new theory does not have the wherewithal to define all the causal powers of the predicate F. Does this happen with consciousness and neuroscience? To convince that it does, the skeptic argues that even if we know everything there is to know about the properties of brain states, there still is something we don't know about the properties of sensations, namely, what it is like to have a sensation.

Paul Churchland (1989) points out, however, that that argument is a non sequitur because it equivocates on "knows about" (see also Lewis, 1983; Nemirow, 1980). Knowledge in the first instance is pretty clearly a matter of mastering a set of sentences, while knowledge in the second seems to be a matter of having some prelinguistic or sub-linguistic representation, or being able to make certain sensory discriminations, or something like this. This equivocation means that, for example, a person who knows everything there is to know about visual cortex but is totally color blind and a person who knows nothing about the nervous system but has experienced many colors have different types of knowledge about exactly the same thing. The difference between the two knowledge bases is in the manner of knowing something, not in the something itself known. We can see this difference clearly if we rewrite the skeptics' arguments to remove the ambiguity: even if we have mastered the complete set of true sentences about the properties of brain states, we still may not represent sensations prelinguistically or be able to make sensory discriminations. This argument does not entail that consciousness is beyond the reach of science; rather it means simply that the brain uses more than storing sentences to represent the world to itself.

If first-person accounts and third-person accounts are just different sorts of descriptions of the same events, then objective accounts must include all the causal powers of subjective accounts. The worry that scientific accounts of consciousness can't capture individual points of view adequately would be false; if both conceptual frameworks can describe the same causal interactions, albeit in different terms, then whatever we see from the first-person point of view, we must be able to describe in the third. Flanagan et al. must be wrong when they claim that the worry is true but irrelevant.<sup>4</sup> The bottom line is that first-person perspectives can tell us nothing that a third-person account could not say one way or another.

---

<sup>4</sup>Flanagan suggests that really what this point turns on are different notions of description and that the skeptics' conditions for an adequate description are unfounded and too strong (Flanagan, personal communication). He may be correct.

*Theory and Observation*

The second worry is that any completed scientific account will fail to analyze consciousness exhaustively. Basically, the concern is that an objective third-person point of view, even if it can reduce first-person accounts, will take us “farther away” from the “real nature” of the phenomena (Nagel, 1974, p. 434). And any move away from individual experience will abstract over something about that experience, hence leaving it out of the objective description of consciousness.

One naturalist’s reply goes as follows (see Flanagan, 1985): Nagel is ambiguous about what he means by “real nature.” If he means real nature to refer to the way things seem to some particular person, then of course a third-person description will be “farther away.” But there is no mystery why first-person accounts are “closer” to the phenomena—people, after all, have the conscious experiences. Each is uniquely connected to his or her own phenomenology; therefore, we each have a special relation to how things seem to ourselves. But given the discussion above, this sort of special connection is nonproblematic. As long as the objective viewpoint can still account for the phenomena, then its “distance” from consciousness is irrelevant.

If, on the other hand, Nagel means real nature to refer to what is really happening in the cognitive system as a whole, then his claim is simply false. What role conscious mental events are actually playing and how they are physically realized, regardless of how it seems to the first-person perspective, are facts that our privileged access cannot unveil. This real nature is best understood from a scientific viewpoint.

Once we grant, as we do above, that a naturalist can understand getting an accurate objective description of the subjective point of view as part of the project of understanding human consciousness, then an exhaustive analysis of consciousness should be forthcoming. As Flanagan writes, a theory of this ilk

will provide a rich autophenomenology, a theory of how the autophenomenology connects up with the actual goings-on, a theory about how conscious mental events—taxonomized into many different classes of awareness—figure in the overall economy of mental life, a theory of how mental life evolved and thereby a theory of which features of mind are the result of direct selection and which features are free-riders, and finally it will provide a neuroscientific realization theory—a theory about how all the different kinds of mental events, conscious and unconscious, are realized in the nervous system. It is hard to see how the analysis could be more exhaustive. (Flanagan, 1985, p. 387)

But again, the naturalist is not taking the skeptics’ worries seriously. If we assume that real nature refers to the way things seem to particular people, then regardless of how well we understand and grant the privileged connection, a move to objectivity must be a move away and will therefore abstract over some aspects of the phenomena. So even if we could get all the things Flanagan outlines, we would still be leaving something out.

Nevertheless, it seems to me that the naturalists still shouldn't have to be concerned. Let us consider the nature of the relationship between theories and the world they describe.<sup>5</sup> The skeptics seem to be claiming that there are two types of sentences involved in understanding consciousness. First, there are sentences that describe our phenomenal experiences whose truth is non-problematically confirmed by our own (introspective) observations. Second, there are the theoretical sentences of science that also describe our phenomenal experiences, but whose truth is problematic since they cannot be directly confirmed by observation. Here we have a version of the observational-theoretical distinction so prevalent in philosophy of science, and for the most part, the skeptics' picture is correct. The ways in which we confirm our first-person reports (our observation sentences) of our conscious experiences is different from the way we confirm any third-person story (the theoretical sentences) about consciousness, and discovering the truth about any third-person theoretical account is problematic. However, the skeptic cannot claim that our first-person sentences are nonproblematic just because they are observational. That just begs the question.

Moreover, that observation sentences are non-problematic is false as well (cf., Suppe, 1989). Whenever we explain or predict some phenomena by using a theory, we assume that the premises about the data are nonproblematic relative to the theory we are using, not that they are nonproblematic *tout court*. If our predictions consistently failed, then we might come to regard our statements about the hard facts as suspect along with our theory. Our observation sentences—even our observation sentences about our own conscious experiences—are theory laden.

Quick evidence for this claim comes from the fact that other languages (including ancient Greek, Chinese, and Croatian, and English before the seventeenth century) have no word for “conscious” or “consciousness” which captures our folk sense of the word (Wilkes, 1988). People from these linguistic communities quite probably have perceptual experiences similar to ours and yet, equally obviously, they cannot talk about their phenomenal experiences like we do since they don't regard the mind in the same way we do. Something about *our* epistemic milieu makes us assume the first-person stance that consciousness exerts some processing power. How to individuate conscious experiences must enter our first-person perspective from what we are told about them at mother's knee.

We can separate the observational/non-observational distinction from the theoretical/non-theoretical (see also P.M. Churchland, 1979; Maxwell, 1962,

---

<sup>5</sup>In what follows, I shall adopt the semantic view of theories (see Suppe, 1989, for review) to discuss the connection because I believe that the notion of theories as models best captures the structure of modern computational theories in the cognitive sciences; however, I have no doubt that the same arguments could be translated into other views of theory structure.



for additional arguments). The skeptics must now maintain that our observation sentences about consciousness, though they be theoretical, still somehow describe more about our experiences than any objective account. Can the skeptics do that? I believe that, *contra* Flanagan, they can, but to do so will not help their skeptical challenge.

Science does not intend to deal with all the intricacies of its subject matter. Instead, it uses only a small number of parameters abstracted away from the phenomena in question to explain and predict something (though not everything) about the phenomena. To take a clear example, classical particle mechanics uses only mass, velocity, distance traveled over time, etc., in characterizing falling bodies; the color of the body and date of its falling are ignored. And the abstraction goes even farther. Particle mechanics does not concern itself with actual velocities, but with velocities in a frictionless environment with an extensionless object. It predicts behavior which depends only on the position and moment of extensionless points interacting in a vacuum. Following Suppe (1989) we can call this imaginary domain a "physical system," an abstract replica of the phenomena that characterizes how the phenomena would have behaved had the idealized conditions obtained.

We can see this same picture of how science works in the cognitive and biological sciences as well. (I take the following examples from Suppe, 1989, p. 66.) The genetic theory of natural selection characterizes evolutionary phenomena in terms of changes in the distributions of genotypes across a population over time as a function of the rate of reproduction, the frequency of crossover, and so on. Behaviorist theories in psychology describe the behavior of idealized organisms as a function of stimulus-response patterns and reinforcement schedules. Other examples include Chomskian theories of competence in linguistics, theories of cell firing in neurophysiology, functionalist theories in cognitive psychology, and *computational theories of consciousness*, all of which describe the behavior of abstract mechanisms under ideal conditions which only approximates the behavior of real phenomena under normal conditions.

Science then does not apply its law directly to phenomena, but rather uses laws to explain the behavior of physical systems abstracted from the phenomena in such a way that the behavior can be correlated with the phenomena (see also Bogen and Woodward, 1988).<sup>6</sup> The observation sentences must be altered such that they represent what would have been observed if only the few relevant parameters of the phenomena existed under ideal conditions. These altered statements, perhaps in conjunction with some idealized bound-

---

<sup>6</sup>Though Bogen and Woodward argue for a point similar to mine, they use different terminology. What I am calling "phenomena," they call "observation," and what I shall call data, they refer to it as "phenomena."

any conditions, are used with the theory to make predictions about the physical system. The predictions are then converted into statements about the corresponding phenomena by just reversing the procedure for altering the original observation statements.

We have two moves from the raw observations to predictions of the theory: we translate phenomena into idealized data and then calculate the theory's postulates from the physical system, boundary conditions, and data. The first move is fundamentally counterfactual and always involves pruning away some aspects of the phenomena actually observed. Here is where we lose something from our first-person accounts of phenomenological experience. It is not in translating from a first-person account to an objective perspective; rather, it is in using these observational sentences in a scientific theory.

The skeptics are correct when they say that something is left out of scientific accounts of consciousness. And they believe that what is left out is important. Any account of consciousness which abstracts away from the actual phenomenal experiences in the manner outlined above will lose the intrinsic property of our sensations as discriminated by introspection. Our scientific theories will not be able to discriminate the very means by which we discriminate one kind of perception from another.

We can see this worry in the problem of the inverted spectrum.<sup>7</sup> Imagine that someone or something exists which has its color spectrum inverted relative to yours, so that when you perceive redness, it would see blueness, and when you saw blueness, it would perceive redness. Further, when it sees what you would call "red," it would call it "blue," and vice versa, so that there is no way to tell from outside the head the quale of the other's sensation, so to speak. There is nothing incoherent about supposing this, and that fact, maintain the skeptics, shows that the connection between how things seem to us and any scientific theory is contingent. And if how things seem to us is an essential aspect of conscious experience, then science will omit something very important indeed.

The skeptics' intuitions here stem from a tendency to think that our common sense individuation of these psychological states demarks natural kinds. (See Churchland and Churchland, 1981, for similar arguments.) After all, this individuation does play an active and fairly successful role in our every-

---

<sup>7</sup>As far as I know, this thought experiment appeared in the philosophical literature first with Locke, when he examined the possibility that "*the same Object should produce in several Men's Minds different Ideas at the same time; v.g. the Idea, that a Violet produces in one Man's Mind by his Eyes, were the same that a Marigold produced in another Man's, and vice versa*" (1689/1975, p. 389). This idea reappeared in the early years of this century during the logical positivist's movement as an attack on the verificationist's theory of meaning. Classic formulations appear in Black (1949), Lewis (1929), Reichenbach (1938), Schlick (1959), Smart (1963), Wisdom (1952), and Wittgenstein (1968). For more contemporary discussions, see Block (1980), Dennett (1988), and Shoemaker (1975, 1982).

day interactions and the conceptual framework in which it is embedded does have all the earmarks of a type of empirical theory (see above, P.M. Churchland, 1979; Sellars, 1956). But the consideration is misled. Whether our common sense individuation of our sensations constitutes a natural kind is an empirical matter, however powerful the intuition. And empirically speaking, it is too soon to tell what the natural kind entities of the species consciousness are exactly (see also below). Introspection alone does not give us reason to accept the taxonomy. We've picked out certain aspects of our qualitative experiences as the important aspects for identifying our sensations, and we use those aspects as hooks upon which to hang our folk psychological labels. We've centered on those particular aspects because they allow us to converse successfully with our fellow human beings. Had the taxonomy we chose been a failure, we certainly would have tried a different tack. The point is that our conscious inner world does not exist already carved into kinds any more than the external world does.

A scientific explanation of consciousness does require that sensations and perceptions have some intrinsic property that plays some causal role but it is indifferent to what that attribute may be as long as the property coincides fairly well with our introspective discriminations. Our qualitative experiences are token-identical with whatever physical states realize them, so it should be no problem construing some attribute of our brain as the intrinsic property of a given conscious experience. And the property could very well be the objective referent of our introspective reports. The property which you crudely discriminate as "having a sensation of redness" may be precisely, for example, the spiking frequency of a certain cluster of cells. Just because introspective judgments are automatic and easy does not mean that they are privileged in the ways the skeptics need. Nothing important will be left out. The skeptics' second complaint, though true, is irrelevant.

### Explanation and Causal Histories

In this section, we shall examine a second similar sort of objection. Proponents of this position too hold that first person accounts of consciousness are the wrong sort of thing to capture in a computational model, but they support this conclusion by arguing that conscious states do not fall naturally into the taxonomy of psychological functional states at all. The problem is that consciousness crosscuts functionalism.

Arguments of this sort are fairly interesting since they are accepted and used by people on both sides of the naturalizing fence. Naturalists argue that the domain of conscious states is orthogonal to any functionalist categorization; therefore, consciousness must be irrelevant to computational theories of the mind. For those skeptical about science being able to account for all

important aspects of the mind, the fact that conscious states have no place in functional theories of cognitive science only all too vividly illustrates how truly meager the cognitive sciences are. However, I shall conclude that neither side is correct since both either implicitly or explicitly rely on unfounded views of explanation in cognitive science. I shall suggest that we should focus our energies on developing a causal history of consciousness, not on trying to use consciousness as an explanans for previously outlined psychological explananda.

To see how both sides use this skeptical argument, consider the paradigm example of qualia: pain. Scientific naturalists do not deny that the subjective experience of pain exists. However, they would claim that there may be good reasons for believing that the relevant factor(s) for explaining pain behavior (or pains themselves) belong not only to the class of felt pains, but to a larger class including other psychological phenomena, e.g., pains which are not consciously felt, but which influence behavior nonetheless. Because of the shared functional features, felt pains and unfelt pains would belong to one explanatory kind in the general taxonomy of mental states in psychological explananda. It is entirely possible that felt pains will not be a natural kind which forms a discoverable part of some functional regularity. The fact that some pains are conscious may not have any independent explanatory role in scientific explanations (Churchland and Churchland, 1981; Wilkes, 1984).

In the opinion of those upholding the skeptical challenge, the reductionists are trying to cripple explanations involving consciousness from the outset because they expect consciousness to fit neatly into their already established patterns of explanation for their already established explananda for intelligent behavior. They maintain that some properties inherent in conscious states are of fundamental importance in establishing any intentional explanation at all and the cognitive sciences overlook this crucial point because they have parsed the phenomena to be explained incorrectly (Searle, 1990).

However, the naturalists' stance, as well as the skeptical counter-position, lean too hard on the previously established taxonomy of psychological events worth explaining. They both entirely overlook the fact that what we might want to explain is that some pains are conscious in the first place. Consciousness may not be located in the explanans of psychology, but rather as an explanandum. Regardless of whether the factors that make some mental states conscious enter into a functionalist taxonomy, there is something which entails that some states are conscious and that some are not. Sorting out this causal history of consciousness could be a legitimate scientific enterprise. Science need not ignore the phenomena of consciousness; it just may not require it for giving functionalist explanations of behavior.

Outlining a causal story strikes me as the appropriate approach to take in this instance because the brain is too complicated and the cognitive sciences

are not yet advanced enough for us to expect the quantitative theories needed for explicit derivation of explananda from explanans immediately. In the more mathematical sciences, we can see examples of explanations which are very close to the classic deductive-nomological model: we literally derive the explananda from the laws describing the behavior of the system and particular initial and boundary conditions. However, in the softer mind/brain sciences, the mathematical techniques needed for true deductions are not yet available. But to explain any data about consciousness adequately, it is not enough to say that they are due to some general psychological or brain mechanism, which is then left unspecified. (See also Bogen and Woodward, 1988, pp. 323–325.) One must develop systematic dependency-relations from the data to the relevant factors of the general mechanism invoked in the explanans.

In our case, systematic dependency relations could be explicated in terms of tracing the causal mechanisms of the data to be explained. To take a parallel example in neurobiology, when accounting for the choeric movements of Huntington's disease, we get a fairly detailed picture of the functioning of the cholinergic neurons and the neurons which synthesize gamma-amino-butyric acid (GABA-ergic neurons) in the striatum, followed by a history of their death in Huntington's patients using data based on CAT scans, PET scans, and magnetic resonance imaging. This profound loss is then connected to a disinhibition of the nigrostriatal dopaminergic system and an over-excitation of the remaining striatal neurons, and a resulting abnormal pallidal output to the thalamus. We finally get an outline of the connection of the thalamus to motor output via the basil ganglia, a disturbance of which is tied to abnormal choeric movements (cf. Kandal and Schwartz, 1985, pp. 531–532). Here we do not see equations which indicate systematic connections between the pharmacology of the brain and involuntary movement; rather, the general mechanism of the disorder (an imbalance in the dopamenergic-cholinergic-GABA-ergic loop) is unpacked in a series of more or less detailed causal stories documenting systematic connections between factors of the mechanism with specific characteristics of the disease. These causal accounts play the same role as the mathematical derivations in physics or chemistry—they both illustrate a systematic dependency-relation in detail.

But even if we could successfully develop a detailed causal history of consciousness such that we find factors in our neurophysiological life which are positively correlated with the existence of a conscious state, I fear that neither the naturalists nor the skeptics would be very interested. They both accept that functionalist accounts of human cognition work perfectly well without using a theoretical construction such as consciousness, and even though the skeptics tollens where the naturalists ponens, they both seem locked into accepting that consciousness itself must have no interesting scientific causal effects on our behavior (cf., Jackendoff, 1987; Searle, 1980,

1990; Velmans, 1991). The naturalists believe that the events which explain our behavior crosscut the domain of conscious states and the skeptics hold that a mechanistic account of behavior, though adequate for science, necessarily ignores subjective phenomena.

However, both the naturalists and the skeptics seem to be overlooking the possibility that consciousness could simply be outside the domain which cognitive psychologists are *currently* trying to capture. That psychologists can and, in many cases, do distinguish the information processes relevant to their studies from phenomenal awareness (thus ipso facto irrelevant) does not mean that the models they create of those processes entail that consciousness is not a topic for scientific scrutiny (the naturalists' position) or that they are now somehow fundamentally inadequate (the skeptics' position). The explananda and the explanans of science change as our understanding of the world evolves. The warmth and phlogiston of yesterday became the heat and kinetic energy of today. So too may attention and meaningful behavior of today become the consciousness and memory stores of tomorrow (see P.M. Churchland, 1979; Hardcastle, 1992, for related discussion). And whether an information processing model is incomplete depends on which phenomena it is supposed to be explaining. Since good explanatory models do not try to account for every observational event, that current models omit conscious phenomena is irrelevant to psychology and cognitive science as a whole.

Moreover, since functionalists individuate mental states in terms of algorithmic information transformations and the causal role mental states (or brain states) play with respect to the inputs and outputs of the system, and since how to taxonomize our explanatory state space is up to us, we could consider consciousness to be the *output* of some transformation or other. (Just as color could be considered an output when determining an object's reflective properties, but is irrelevant when determining an object's mass.) If so, then it could be included in a functionalist's ontology. Or, and perhaps more likely, since the naturalists do believe that consciousness must be a physical state in the brain, albeit a brain state which may not influence the development of subsequent functional brain states or processes, these conscious states may be more relevant to understanding cognition than mere noise. Although not part of the causal chain currently under investigation, some states, like consciousness, may index certain processes, e.g., focal attention or certain memory processes, which *are* causally relevant in the circumscribed domain.<sup>8</sup> A causal history account of this index would then broaden our understanding of how brain events correlate with mental events.

---

<sup>8</sup>Of course, for science, regardless of whether consciousness is used as an *explananda* or an *explanans*, it would still be a theoretical construction. We would use it to account for or index some aspect of a causal chain we have posited in order to explain (an idealized and abstracted version of) human cognitive behavior.

### Natural Kinds and Qualia as Explananda

We have now seen that even if consciousness does not fit into the current psychological taxonomy, developing a causal history account of conscious phenomena could be a legitimate scientific enterprise either as its own end, or in a future revised taxonomy. It may also turn out that conscious states index other phenomena which are used in our psychological explanations, and so developing a causal history account of consciousness would be a useful extension to our psychological understanding of the mind.

However, this position does assume that the category of conscious states is well-formed and coherent. We still must consider a final sort of skeptical argument directed against developing a scientific account of consciousness which holds that "conscious" and "consciousness" are too ill-defined to serve as fruitful explananda. Notions about consciousness are too imprecise and too riddled with unsubstantiated intuition to act as a natural kind of science. Nevertheless, as I shall argue in defense of the naturalists' project, we do find at least one common element among all our conscious experiences and this should be enough to start a proper empirical investigation into the mysteries of consciousness.

In general, science studies natural kinds, groups of objects whose members are governed by the same set of laws and whose properties are not relativized to any particular personal interest. For example, "cell" is a natural kind, but "germ" is not; "gold" is, but "gem" is not. It is not always immediately clear whether some terms form a natural kind, even though they are used in science as such. For example, "memory" in psychology encompasses many diverse phenomena and scientists make several different fundamental distinctions within the category. Currently, we have short-term memory, long-term memory, working memory, as well as procedural, declarative, semantic, episodic, iconic, non-cognitive, and somatic memories . . . . So many different types that it seems possible, even likely, that the term "memory" will come to be replaced by a set of systematically more useful descriptions. (I get this example from Wilkes, 1988.) These descriptions would then pick out the relevant natural kinds, while the term memory would become obsolete in psychology (though it would probably still be useful in other scientific domains, such as computer science).

Science quite often co-ops everyday folk terms into its more rigorous conceptual framework as natural kind terms. Consider "force" or "mass." Since everyday language is notoriously vague, flexible, and context dependent, everyday terms must be adapted to fit into the more precise and economical conceptual apparatus of science. No one really doubts that if "consciousness" is ever to be used scientifically, it will most likely be refined to some degree. What is under dispute and in doubt is how much we would have to adapt

“consciousness” to fit any scientific framework: amended enough, it would cease to be the same term as the everyday one. If too much refinement is required, the skeptics would be correct in their assessment that (the folk term) consciousness does not denote anything suitable for scientific inquiry.

I should pause here to point out that what we are looking for is not a list of necessary and sufficient conditions for consciousness. Since any definition in science must be empirically motivated, words will come to have a more precise meaning as they become more embedded in the framework of empirical theory (see P.S. Churchland, 1988, for discussion). The question before us is, given empirical discoveries about awareness, how significant will the meaning change have to be.

I take it for granted that there is now substantial evidence that our common sense beliefs will have to be revised at least somewhat, supporting the distinct possibility that consciousness is not a natural kind. We know, for instance, our common assumption that introspection of our conscious states gives us reliable and transparent information about our motives and experiences is largely false. (I take the following examples from P.S. Churchland, 1983.) Nisbett and Wilson (1977), among others, have demonstrated that unconscious factors influence our judgments about preference and quality. They showed that if subjects are evaluating candidates for a job by examining the candidates' files and the subjects are told that they will get to meet some particular applicants, the subjects will issue more favorable judgments for those candidates than for those they were told they would not get to meet. Hess (1975) agrees that many unconscious influences surround our conscious judgments. He discovered that the size of a person's pupils figures crucially into how one perceives the other person. The larger the pupils, the more friendly the person appears.

Contra the thesis that conscious states are transparent, Supa, Cotzin, and Dallenbach (1944) found that echo-location is the mechanism certain blind people use in maneuvering around objects, even though these people have no idea how they avoided running into things. Moreover, the documentation of blindsight (Weiskrantz, Warrington, and Saunders, 1974) shows that perceptual judgments do not require conscious activation. Some subjects with lesions in the primary visual cortex, which results in blindness in the left visual field, can still make accurate perceptual discriminations when told to guess about the location of objects nonetheless.

Finally, our enormous propensity for confabulation leaves little doubt that our conscious states give us little understanding about what of the external world we actually experience and why we behave the way we do. Patients with Anton's syndrome, a form of blindness denial, do not recognize that they cannot see. They deny their visual deficit and persist in behaving as though they can see perfectly well. When asked why they bump into things,



or why they answer questions about their physical environment incorrectly, they often just invent patently false reasons (Critchley, 1979). Split-brain patients also present good examples of confabulation. When the severed hemisphere which is nondominate for language initiates some physical activity, the hemisphere which does control speech will invent a story to explain what its body is doing: generally speaking, it does not express puzzlement at a behavior which appears uncontrolled from that hemisphere's viewpoint, nor does it admit that the other hemisphere must be forcing the action (Gazzaniga and LeDoux, 1978).

And confabulation is not limited to abnormal brains. Nisbett and Wilson (1977) have shown that unbeknownst to most, we have a preference for objects on the right. They placed identical pairs of hose next to one another and asked passers-by to rate them. The subjects typically chose the pair on the right, even though they explained their choice in terms of superior strength, sheerness, color, and so on (see also Goethals and Rechman, 1973). This sort of confabulation appears to be a normal part of daily life when theorizing about one's own behavior; nothing pathological is involved. In sum, our conscious states are neither reliable nor transparent sources of information about the world or our interactions in it.

However, that some of our common sense beliefs about awareness need to be revised does not mean that consciousness is not a natural kind. We can think of natural kinds as lying in a continuum from a strict natural kind in which virtually all the relevant laws or principles apply to virtually all the members of the group, to "cluster" natural kinds in which some laws are true of all instances and others apply to only constituent sub-classes, to groups which have little more coherence than a mere set, in which the laws which govern the subclasses are structurally analogous and there are no overarching principles at all. (See Wilkes, 1988, for discussion.) That what we generally accept as the overarching principles of consciousness are in fact wrong would at best only suggest that consciousness is not a strict, or unitary, natural kind. Given the diversity of conscious phenomena, that conclusion should surprise no one. But even if we stop looking for The Unifying Principle which governs all and only conscious phenomena, we can still meaningfully discuss whether consciousness is a cluster phenomena, or whether it should be dismissed as merely a set term.

How do we decide whether consciousness marks a natural kind? Part of the answer lies in the usefulness of the principles which govern part or all of the phenomena. A matter of degree, it is based on pragmatic considerations. But after we decide that a set of proposed principles are promising enough to warrant scientific investigation, we still have to decide whether the group of principles cull out a true natural kind. One criterion for separating a cluster natural kind from more amorphous sets is that the principles which relate the

subclasses not have their content completely exhausted by a simple conjunction of the subordinate principles describing each subclass. Cartwright, for example, believes that the superlaws in physics are not profitable for just this reason: “what the [superlaws] . . . dictate should happen, happens *because* of the combined action of laws from separate domains, like the law of gravity and Coulomb’s law” (1983, p. 71, as quoted in Wilkes, 1988). If we want to understand how things in the physical world operate, we should look at the contributory laws—the superlaws add little.

Would the principles concerning conscious phenomena be analogous to the superlaws in physics? To answer this question negatively, we need to discover relationships among the collection of conscious phenomena which are *systematic, nonreducible, and useful for science*. When Wilkes looked at the various types of conscious phenomena, she concluded that scientifically interesting criteria for being a conscious state do not exist.

Consider how varied are our experiences which we use as the referents for conscious. We say we are conscious when we are awake instead of being asleep, an apparently simple distinction made more complicated by dreaming, hypnosis, comas, epileptic automatism, etc. Conscious here refers to a state of being. We also say that we are conscious of various bodily sensations which exhibit a range of intensity: itches, pins and needles, pains, etc. Here conscious is a count noun which refers to a type of physical state. We also say we are conscious of different sensory experiences—sight, hearing, taste, smell, touch, and kinaesthesia. This conscious is more a mass term to which very different adjectives apply than the consciousness of pains. Our itches and the like are locatable; they can be “stabbing” or “throbbing” or “better than before.” On the other hand, our visual experience of a blue cup does not seem to be clearly localized anywhere inside us, and if it is “stabbing,” or “throbbing,” or “better than before,” then it is for very different reasons than those we use to ascribe attributes to our pains. Finally, we also use conscious in conjunction with some propositional attitudes: beliefs, desires, and so on. Here conscious is a delineating attribute—we have some conscious desires and some unconscious ones.

Wilkes believes that once we understand how different all these categories of conscious things are, we will have no choice but to conclude that dropping the term conscious would omit nothing from our investigation of the four different groups of phenomena. Nevertheless, it does seem that, as different as are the four categories, they do all have one thing in common: for each category, it is *like something* to have the relevant phenomena. There is something like being awake, having your foot fall asleep, hearing a bird sing, and wanting an ice-cream sundae. Perhaps somehow this phenomenological “something” points toward a unifying principle.

Still, we must also be aware that our folk psychological expressions are generally tied to appearances and not to true essences (see also above). Our folk term "fire" used to refer to burning wood, the activity of stars, lightning, the Northern lights, and the flicker of fire-flies. As we now understand the world, only some of those processes are truly oxidation, that is, are truly "fire," while the others involve fusion, electrical incandescence, spectral emission, and phosphorescence. And there are also oxidation processes which are very different from our folk understanding of fire, like rusting, tarnishing, and metabolism (P.S. Churchland, 1988, p. 285). Only when we discover the underlying nature of a phenomena can we understand how skewed are our intuitive categories. How are we to make sense then of our intuition that consciousness in fact does have some sort of phenomenological feel, and in our everyday interactions, we generally take that quale to be a *prima facie* necessary and sufficient condition for consciousness?

Recent work on the different memory systems of the brain clearly shows that conscious perceptions affect behavior differently from unconscious perceptions. We know from experiments in cognitive psychology that manipulating the type of encoding an input receives has a large influence on conscious memory, as seen in recall or recognition tasks, while it has little or the opposite influence on tests of unconscious memory as seen in masked priming effects (Graf and Mandler, 1984; Graf, Mandler, and Haden, 1982; Jacoby, 1983; Jacoby and Dallas, 1981; Schacter, 1990; Schacter and Graf, 1986; Winnick and Daniel, 1970); that changing the modality between study and test has no effect on conscious memory, while it impairs unconscious memory results (Graf, Shimamura, and Squire, 1985; Jacoby and Dallas, 1981; Kirshner, Milech, and Standon, 1983; Kirshner and Smith, 1974; Roediger and Blaxton, 1987; Schacter and Graf, 1986); that unconscious memory primes last substantially longer (on the order of weeks) than do primes for conscious recall or recognition (Jacoby and Dallas, 1981; Komatsu and Ohta, 1984; Tulving, Schacter, and Stark, 1982); that proactive and retroactive interference affect unconscious and conscious memory tasks differently (Graf and Schacter, 1987; Sloman, Hayman, Ohta, and Tulving, 1988); and that success on an unconscious memory task is stochastically independent of success on a conscious task (Eich, 1984; Graf and Mandler, 1984; Jacoby and Witherspoon, 1982; Schacter, Cooper, and Delaney, 1990; Tulving et al., 1982; see also Musen and Treisman, 1990). Something about a state *being conscious* (as evidenced by verbal report) alters our range of behavioral responses, so insofar as our reports and our behavior are independent of one another, we can separate our judgments from whatever it is that forces the change.

So, we need to posit something besides our ability to issue reports to account for our behavior. And whatever this something is seems to be directly apprehended in consciousness. At least, we can separate it from our judg-

ments concerning conscious experience, which are (fallibly) mediated by memory. We have a directly accessible "black box" that appears to be tied to our intuitive notion of qualia, which can be indexed by our reports about what it is like to have a particular experience, and what is necessary for a complete psychological explanation. It makes sense for the moment to let our "folk" expression "qualia" to stand as the name for this psychological place-holder, since the two are (at least for now) intimately related. Though much of our folk-understanding of consciousness will probably turn out to be false, one aspect does seem to index a phenomena needed to explain some of our behavior.

### Conclusion

To recap, our first person subjective accounts of what it is like to be conscious should be able to be translated more or less faithfully into third person objective language. This translation would preserve all the (important) causal connections among relevant phenomena, but the redescription would abstract over some particular aspects of consciousness. However, though what is left out is important to our folk connection of consciousness, it may not be important to a scientific theory. This difference only underscores that what is needed for communicative success in personal discourse is not the same as what is needed for predictive success in science.

That consciousness does not now fit into our current psychological taxonomy is really irrelevant to the possibility of someday developing a scientific theory of consciousness. Far from entailing that consciousness is epiphenomenal, it only means that consciousness can't be used today as a psychological explanans. And positing conscious experience does appear to be needed to account for all of our behavior (regardless of what is currently popular in the cognitive sciences). Though this theoretical posit may not rely on the particular (folk psychological) qualitative attributes of individual experience, it does seem to be intimately tied to qualitative experience as a whole. That is, we need to posit some entity (or set of entities) which our verbal reports about our conscious experience index fairly well, though the underlying structure of the entity has yet to be uncovered.

Indeed, as we search for the pertinent structure which underlies conscious experience, we must be careful to guard against our folk intuitions that mental states are phenomenological billiard balls without any interesting or relevant physical structure. In the end, there may be no real way to abstract a consciousness from the brain (or relevantly similar brain-like structure) containing it. If we remove the functionalist leanings from our interpretations of mind which force us to disregard the mind's underlying structure, it may become easier to explain how to understand a scientific conception of con-

sciousness. Consciousness could refer to the particular neuronal dynamical firing patterns which are directly correlated with our qualia. And having a certain quale would be the brain exhibiting some particular dynamical structure of other.

### References

- Black, M. (1949). *Language and philosophy*. Ithaca, New York: Cornell.
- Block, N. (1980). Are absent qualia impossible? *The Philosophical Review*, LXXXIX, 257–274.
- Block, N. (1985). Advertisements for a semantics for psychology. In P.A. French, T.E. Uehling, and H.K. Wettstein (Eds.), *Midwest studies in philosophy*, Vol. IX. Minneapolis: University of Minnesota Press.
- Block, N., and Fodor, J. (1972). What psychological states are not. *The Philosophical Review*, LXXXI, 159–181.
- Bogen, J., and Woodward, J. (1988). Saving the phenomena. *The Philosophical Review*, XCVII, 303–352.
- Cartwright, N. (1983). *How the laws of physics lie*. Oxford: Clarendon Press.
- Churchland, P.M. (1979). *Scientific realism and the plasticity of mind*. New York: Cambridge University Press.
- Churchland, P.M. (1981). Eliminative materialism and the propositional attitudes. *Journal of Philosophy*, 78, 67–90.
- Churchland, P.M. (1985). Reduction, qualia, and the direct introspection of brain states. *Journal of Philosophy*, 82, 8–28.
- Churchland, P.M. (1989). *A neurocomputational perspective: The nature of mind and the structure of science*. Cambridge, Massachusetts: The MIT Press.
- Churchland, P.M. (1984). *Matter and consciousness*. Cambridge, Massachusetts: The MIT Press.
- Churchland, P.M., and Churchland, P.S. (1981). Functionalism, qualia, and intentionality. *Philosophical Topics*, 12, 121–145.
- Churchland, P.S. (1983). Consciousness: The transmutation of a concept. *Pacific Philosophical Quarterly*, 64, 80–95.
- Churchland, P.S. (1988). Reduction and the neurobiological basis of consciousness. In A. Marcel and E. Bisiach (Eds.), *Consciousness in contemporary science* (pp. 273–304). Oxford: Clarendon Press.
- Critchley, M. (1979). *The divine banquet of the brain*. New York: Raven.
- Dennett, D. (1982). How to study consciousness empirically or nothing comes to mind. *Synthese*, 53, 158–180.
- Dennett, D. (1988). Quining qualia. In A. Marcel and E. Bisiach (Eds.), *Consciousness in contemporary science* (pp. 42–77). Oxford: Clarendon Press.
- Derrucher, P., and Baveux, D. (1989). Correlation analyses of explicit and implicit memory performance. *Memory and Cognition*, 17, 77–86.
- Eich, J.E. (1984). Memory for unattended events: Remembering with and without awareness. *Memory and Cognition*, 12, 105–111.
- Flanagan, O. (1985). Consciousness, naturalism, and Nagel. *The Journal of Mind and Behavior*, 6, 373–390.
- Gazzaniga, M., and LeDoux, J. (1978). *The integrated mind*. New York: Plenum.
- Goethals, G., and Rechman, R. (1973). The perception of consistency in attitudes. *Journal of Experimental Social Psychology*, 9, 491–501.
- Gordon, B. (1988). Preserved learning of novel information in amnesia: Evidence for multiple memory systems. *Brain and Cognition*, 7, 257–282.
- Graf, P., and Mandler, G. (1984). Implicit and explicit memory for new associations in normal and amnesic subjects. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 11, 501–518.
- Graf, P., Mandler, G., and Haden, P. (1982). Simulating amnesic symptoms in normal subjects. *Science*, 218, 1243–1244.

- Graf, P., and Schacter, D. (1987). Selective effects of interference on implicit and explicit memory for new associations. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 12, 45-53.
- Graf, P., Shimamura, A.P., and Squire, L.R. (1985). Priming across modalities and priming across category levels: Extending the domain of preserved function in amnesia. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 11, 385-395.
- Hannay, A. (1987). The claims of consciousness: A critical survey. *Inquiry*, 30, 395-434.
- Hardcastle, V.G. (1992). Explaining epiphenomenalism. *Behavioral and Brain Sciences*, 14, 680.
- Hess, R. (1975). The role of pupil size in communication. *Scientific American*, 233, 110-118.
- Hooker, C. (1981). Towards a general theory of reduction. *Dialogue*, 20, 38-59; 201-236; 496-529.
- Jackendoff, R. (1987). *Consciousness and the computational mind*. Cambridge, Massachusetts: The MIT Press.
- Jackson, F. (1982). Epiphenomenal qualia. *Philosophical Quarterly*, 32, 127-132.
- Jackson, F. (1986). What Mary didn't know. *Journal of Philosophy*, 83, 291-295.
- Jacoby, L.L. (1983). Remembering the data: Analyzing interactive processes in reading. *Journal of Verbal Learning and Verbal Behavior*, 22, 485-506.
- Jacoby, L.L., and Dallas, M. (1981). On the relationship between autobiographical memory and perceptual learning. *Journal of Experimental Psychology: General*, 13, 456-463.
- Jacoby, L.L., and Witherspoon, D. (1982). Remembering without awareness. *Canadian Journal of Psychology*, 36, 300-324.
- Kandel, E., and Schwartz, J. (1985). *Principles of neural sciences* (second edition). New York: Elsevier Science Publishing Company.
- Kirshner, K., and Smith, M.C. (1974). Modality effects in word identification. *Memory and Cognition*, 2, 637-640.
- Kirshner, K., Milech, D., and Standon, P. (1983). Common and modality-specific processes in the mental lexicon. *Memory and Cognition*, 11, 621-630.
- Komatsu, S.I., and Ohta, N. (1984). Priming effects in word-fragment completion for short- and long-term retention intervals. *Japanese Psychological Research*, 26, 194-200.
- Kripke, S. (1980). *Naming and necessity*. Cambridge, Massachusetts: Harvard University Press.
- Levine, J. (1987). *The Nature of Psychological Explanation* by Robert Cummins, a critical notice. *The Philosophical Review*, XCVI, 249-274.
- Lewis, C.I. (1929). *Mind and the world order*. New York: Scribner.
- Lewis, D. (1983). Postscript to "Mad Pain and Martian Pain." *Philosophical papers, Volume 1*. New York: Oxford.
- Locke, J. (1975). *An essay concerning human understanding*. P. Nidditch (Ed.). New York: Oxford. (Originally published in 1689)
- Mandler, J.M. (1984). Representation and recall in infancy. In M. Moscovitch (Ed.), *Infant memory, its relation to normal and pathological memory in humans and other animals* (pp. 75-101). New York: Plenum Press.
- Maxwell, G. (1962). On the ontological status of theoretical entities. In H. Fiegal and G. Maxwell (Eds.), *Minnesota studies in the philosophy of science, Volume III*. Minneapolis: University of Minnesota Press.
- McGinn, C. (1989). Can we solve the mind-body problem? *Mind*, 48, 349-366.
- Musen, G., and Treisman, A. (1990). Implicit and explicit memory for visual patterns. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 16, 127-137.
- Nagel, T. (1974). What is it like to be a bat? *The Philosophical Review*, LXXXIII, 435-451.
- Nemirow, L. (1980). Review of Thomas Nagel's *Mortal Questions*. *The Philosophical Review*, LXXXIX, 473-477.
- Nisbett, R., and Wilson, T. (1977). Telling more than we know: Verbal reports on mental processes. *Psychological Review*, 84, 321-359.
- Reichenbach, H. (1938). *Experience and prediction*. Chicago: University Press.
- Rey, G. (1983). A reason for doubting the existence of consciousness. In R.J. Davidson, G.E. Schwartz, and D. Shapiro (Eds.), *Consciousness and self-regulation, Volume 3* (pp. 1-38). New York: Plenum Press.

- Rey, G. (1988). A question about consciousness. In H.R. Otto and J.A. Tuedio (Eds.), *Perspective of mind* (pp. 5–24). Dordrecht: Reidel.
- Roediger, H.L. III, and Blaxton, T.A. (1987). Retrieval modes produce dissociations in memory for surface information. *Memory and cognitive processes; The Ebbinghaus Centennial Conference*. Hillsdale, New Jersey: Erlbaum.
- Rorty, R. (1982). Comments on Dennett. *Synthese*, 53, 181–187.
- Salmon, W. (1984). *Scientific explanation and the causal structure of the world*. Princeton, New Jersey: Princeton University Press.
- Schacter, D. (1985). Multiple forms of memory in humans and animals. In N.M. Weinberger, J.L. McGaugh, and G. Lynch (Eds.), *Memory systems of the brain* (pp. 351–379). New York: The Guilford Press.
- Schacter, D. (1990). Perceptual representations systems and implicit memory: Toward a resolution of the multiple memory systems debate. Development and neural bases of higher cognition. *Annals of the New York Academy of Sciences*, 608, 543–571.
- Schacter, D., and Graf, P. (1986). Effects of elaborate processing on implicit and explicit memory for new associations. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 12, 432–444.
- Schacter, D., and Moscovitch, M. (1984). Infants, amnesics, and dissociable memory systems. In M. Moscovitch (Ed.), *Infant memory, its relation to normal and pathological memory in humans and other animals* (pp. 173–216). New York: Plenum Press.
- Schacter, D., Cooper, L.A., and Delaney, S.M. (1990). Implicit memory for unfamiliar objects depends on access to structural descriptions. *Journal of Experimental Psychology: General*, 119, 5–24.
- Schaffner, K. (1967). Approaches to reduction. *Philosophy of Science*, 34, 137–147.
- Schlick, M. (1959). Positivism and realism. In A.J. Ayer (Ed.), *Logical positivism* (pp. 92–95). New York: Free Press.
- Searle, J. (1980). Minds, brains, and programs. *Behavioral and Brain Sciences*, 3, 417–458.
- Searle, J. (1990). Consciousness, explanatory inversion, and cognitive science. *Behavioral and Brain Sciences*, 13, 417–424.
- Sellars, W. (1956). Empiricism and the philosophy of mind. In H. Feigl and M. Scriven (Eds.), *The foundations of science and the concepts of psychology and psychoanalysis. Minnesota studies in the philosophy of science, Vol. 1* (pp. 127–196). Minneapolis: University of Minnesota Press.
- Shoemaker, S. (1975). Functionalism and qualia. *Philosophical Studies*, 27, 291–315.
- Shoemaker, S. (1982). The inverted spectrum. *Journal of Philosophy*, 79, 357–381.
- Slooman, S.A., Hayman, C.A.G., Ohta, N., and Tulving, E. (1988). Forgetting and interference in fragment completion. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 14, 223–239.
- Smart, J.J.C. (1963). *Philosophy and scientific realism*. New York: Humanities.
- Sperry, R.W. (1965). *Brain bisection and mechanisms of consciousness*. In J.C. Eccles (Ed.), *Brain and conscious experience* (pp. 298–313). New York: Springer Verlag.
- Sperry, R.W. (1969). A modified concept of consciousness. *Psychological Review*, 76, 532–536.
- Sperry, R.W. (1987). Structure and significance of the consciousness revolution. *The Journal of Mind and Behavior*, 8, 37–65.
- Sperry, R.W. (1991). In defense of mentalism and emergent interactionism. *The Journal of Mind and Behavior*, 12, 221–245.
- Squire, L.R. (1987). *Memory and brain*. New York: Oxford University Press.
- Supa, M., Cotzin, M., and Dallenbach, K.M. (1944). "Facial vision": The perception of obstacles by the blind. *American Journal of Psychology*, 57, 133–183.
- Suppe, F. (1989). *The semantic conception of theories and scientific realism*. Chicago: University of Illinois Press.
- Tulving, E., Schacter, D., and Stark, H.A. (1982). Priming effects in word-fragment completion are independent of recognition memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 8, 336–342.
- van Fraassen, B. (1980). *The scientific image*. Cambridge, Massachusetts: The MIT Press.
- van Fraassen, B. (1989). *Laws and symmetries*. Cambridge, Massachusetts: Oxford University Press.

- Velmans, M. (1991). Is human information processing conscious? *Behavioral and Brain Sciences*, 14, 651–726.
- Weiskrantz, L., Warrington, E., and Saunders, M. (1974). Visual capacity in the hemianopic field following a restricted occipital ablation. *Brain*, 97, 709–728.
- Wilkes, K. (1984). Is consciousness important? *British Journal for Philosophy of Science*, 35, 223–243.
- Wilkes, K. (1988). \_\_\_\_, Yishi, Duh, Um, and Consciousness. In A. Marcel and E. Bisiach (Eds.), *Consciousness in contemporary science* (pp. 16–41). Oxford: Clarendon Press.
- Winnick, W.A., and Daniel, S.A. (1970). Two kinds of response priming in tachistoscopic recognition. *Journal of Experimental Psychology*, 84, 74–81.
- Wisdom, J. (1952). *Other minds*. New York: Philosophical Library.
- Wittgenstein, L. (1968). Notes for lectures on “private experience” and “sense data.” R. Rhees (Ed.). *The Philosophical Review*, LXXVII, 275–320.