

The resistance against a natural-science approach to conscious experience reminds one also of the great debate a century ago between Darwinians and creationists. Darwin's opponents clearly believed they were attacking not a mere scientific hypothesis, but a conception of human nature that would tear away the last remaining shreds of human dignity. The contemporary resistance by Velmans and others to a straightforward natural-science approach to conscious experience may be driven by a similar anxiety. But Darwin did not deprive us of human dignity; treating conscious experience as a normal topic of psychology and neuroscience will not do so either.

Indeed, one can make the opposite argument: that denial of first-person conscious experience in other people may lead to a profound kind of dehumanization. It comes down to saying that other people are not capable of joy or suffering, that as far as the outside observer is concerned, we are not to see others as they see themselves. The consequence of this prohibition against the first-person perspective is a kind of mechanization of other people. Psychology under the thumb of behaviorism did indeed display this kind of dehumanizing, mechanistic thinking. It is only when we acknowledge the reality of conscious experience in the minds of others that we can recognize their full humanity.

## Evidence against epiphenomenalism

Ned Block

Department of Linguistics and Philosophy, Massachusetts Institute of Technology, Cambridge, MA 02139  
 Electronic mail: block@psyche.mit.edu

Velmans argues for epiphenomenalism, the view that consciousness has no function. I claim, however, that the survey of data he presents actually provides interesting evidence *against* epiphenomenalism.

**Senses of "consciousness."** As Velmans sees, the word "conscious" is used to mean many different things. (Unfortunately, "aware" – which he uses to explicate "conscious" – is ambiguous in much the same way.) The most important distinction from the point of view of the target article is the distinction between *phenomenal consciousness* and various kinds of *cognitive consciousness*. If we are computational devices, I can nonetheless wonder whether *what it is like to be* a computationally equivalent silicon-based device is the same as what it is like to be me (Nagel 1974). What it is like = phenomenal consciousness, which I will write here as consciousness<sub>p</sub>. As Velmans notes, some writers use "conscious processing" to mean focal-attentive processing. This is one of many *cognitive* sense of "consciousness." Other senses, for example, involve second-order states, or self-consciousness, or internal soliloquies, or monitoring. Velmans usually has consciousness<sub>p</sub> in mind, and his conclusion, the one I wish to argue he actually provides evidence against, is that consciousness<sub>p</sub> plays no causal role in information processing. (I will leave it to others to criticize Velmans's views on the complementarity between the first- and third-person perspectives.)

**Does consciousness<sub>p</sub> exist?** As Velmans notes, many writers on this topic (and I'm sure many of the commentators in this issue) don't believe in consciousness<sub>p</sub>, or, more sympathetically, they don't believe in consciousness<sub>p</sub> as distinct from one or another of the cognitive notions they hope can satisfy the demands that lead some of us to acknowledge a distinct consciousness<sub>p</sub> (e.g., Dennett 1978; 1979; 1988; Harman 1990; Nelkin 1989; Rey 1988; Rosenthal 1990). This is not the place to debate the issue, but for clarity's sake, I wish to emphasize that I'm on Velmans's side here (see Block 1978; 1990).

**The argument for epiphenomenalism.** Velmans's argument for epiphenomenalism is simple. He considers a large variety of types of information processing, asking of each one whether

consciousness<sub>p</sub> is required for it. He considers semantic analysis of novel word combinations, on-line analysis of speech, learning, memory, identification, planning and control of complex novel action, and others, arguing in each case that these things *can* be done without consciousness<sub>p</sub>. (This survey is the bulk of his paper.) He concludes that there is no type of information processing for which consciousness<sub>p</sub> is required, and therefore that consciousness<sub>p</sub> plays no causal role in and does not "enter into" information processing, and is in that sense epiphenomenal.<sup>1</sup>

The fallacy of this argument is equally simple. Even if Velmans is right that consciousness<sub>p</sub> is not *required* for any particular sort of information processing, it does not follow that consciousness<sub>p</sub> does not causally enter into information processing. Consciousness<sub>p</sub> might *actually* enter into and causally influence information processing even if it is not *essential*, even if something else could substitute for it. An analogy: A business school degree is not required for getting rich, but it does often serve that function. All processing Velmans describes is processing in specialized modules outside the central executive system. Executive system processing itself is not required for processing in the specialized modules, so if Velmans's reasoning were right, it would show that executive system processing is epiphenomenal. (Here, as in what follows, I assume that there is a central executive system that is in charge of focal-attentive processing and its role in control of action, reasoning, and reporting.)

Velmans is concerned with an "essential function of consciousness," that is, whether consciousness "plays some essential role." His concern is whether consciousness<sub>p</sub> does something that cannot be done without it. Hence his one-by-one examination of types of information processing activity (identification, semantic analysis, etc.) to show that these things can be done without consciousness<sub>p</sub>. However, the epiphenomenalism issue as Velmans describes it (i.e., whether consciousness "enters into or causally influences" information processing) is not an issue of *essential* function in this sense, but rather *actual* function. The heart does not have the essential function (in Velmans's sense) of pumping blood, if blood vessels could be constructed that squeeze the blood along on their own. But the heart does have the *actual* function of pumping blood. The issue of epiphenomenalism is not whether consciousness is essential in this sense to information processing, but rather whether it plays an actual role – even if something else could substitute for it.

(Digression: Why do we *have* consciousness<sub>p</sub> if it is not essential? The answer is that evolution often picks one of several options, any of one of which would do the job, depending on how handy these options are. My favorite speculation about the ultimate evolutionary origin of consciousness<sub>p</sub> is its role in motivation.)

As I just mentioned, all the cases of information processing without consciousness<sub>p</sub> that Velmans discusses are plausibly construed as information processing in specialized modules rather than the central executive system. Indeed, Velmans himself would seem to agree, since he emphasizes that focal attentive processes (which are executive processes in my terminology) are not involved in the nonconscious<sub>p</sub> activities he surveys. So Velmans can be right that the processes he surveys do not require consciousness<sub>p</sub> even if consciousness<sub>p</sub> does have an effect on information reaching the central executive.

My point can be seen more clearly if we contrast the two models illustrated below. Model 1 is adapted from Schacter (1989). It gives consciousness<sub>p</sub> an information processing role in integrating the outputs of specialized modules and transmitting the resulting signals to the executive system. Model 2 is an inverted version of Schacter's model intended to capture Velmans's idea that consciousness<sub>p</sub> does not enter into information processing.

If all the sophisticated information processing that Velmans surveys is done by the specialized modules, then perhaps he is

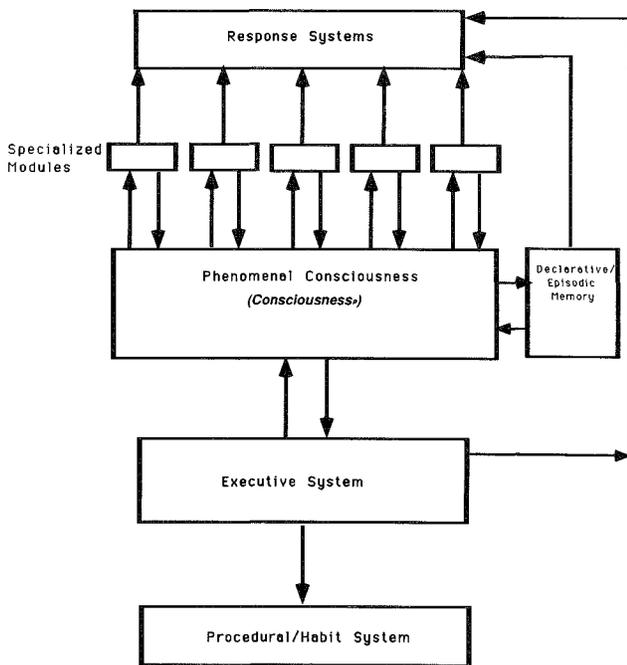


Figure 1 (N. Block). Model 1: Consciousness<sub>p</sub> plays a causal role.

right that consciousness<sub>p</sub> is not essential for sophisticated processing. But even if correct, this claim does not distinguish between Model 1 and Model 2. Consciousness<sub>p</sub> might function as an integrator of information from the specialized modules and the gateway to the executive (as in Model 1), even if consciousness<sub>p</sub> is not essential for any of the wonderful information processing in the specialized modules.

It should be noted that the two models I am contrasting both subscribe to what might be called Cartesian modularism, the view that consciousness<sub>p</sub> has a home in a distinct system. Dennett and Kinsbourne (in press) and Dennett (in press)

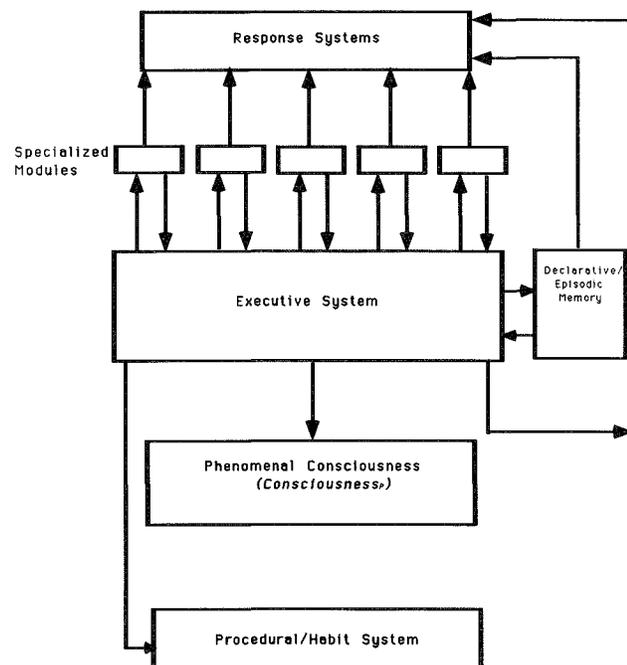


Figure 2 (N. Block). Model 2: Consciousness<sub>p</sub> is epiphenomenal.

inveigh against Cartesian materialism, the view that there is a place in the brain where consciousness<sub>p</sub> happens. Cartesian modularism is a far more plausible view for friends of consciousness<sub>p</sub>, but I don't know whether Velmans would buy into it.

**The evidence against epiphenomenalism.** Note the following difference between Models 1 and 2: if Model 1 is correct, we should never get information reaching the executive system (which you will recall is in charge of focal-attentive processing, and its role in control of action, reasoning, and reporting) without first passing through the consciousness<sub>p</sub> module. Nothing gets to the executive that isn't conscious<sub>p</sub>. In the case of Model 2, by contrast, the executive system could for one reason or another process information without passing it to the consciousness<sub>p</sub> module.

Velmans has combed the literature looking for evidence of sophisticated information processing without consciousness<sub>p</sub>. The most interesting thing about the evidence that he has collected is that nothing he mentions makes a plausible case for information reaching the executive without being in consciousness<sub>p</sub>. According to every model I have ever heard of that postulates a central executive system, the executive system is the home of the most sophisticated processing (even if some sophisticated processing is accomplished elsewhere, as Velmans maintains), so if a search for sophisticated processing without consciousness<sub>p</sub> reveals no executive processing without consciousness<sub>p</sub>, then the search confirms Model 1, arguing against epiphenomenalism, not for it. Note that my claim is not that all information in the executive is sent to consciousness<sub>p</sub>. That is not a commitment of either model. The point is that the structure of Model 1 precludes any information from any of the specialized modules reaching the executive system without passing through consciousness<sub>p</sub>.

I can't make this case in detail, since it would require going over all the data in the target article, but I will briefly mention some of the cases that might be taken to go against my claim.

First, the one case that Velmans claims involves focal attention without consciousness<sub>p</sub> involves Nissen and Bullemer's (1987) experiments on amnesics. Amnesic patients learn a sequence that requires focal attention, but they don't report noticing the pattern. Velmans takes this to be a case of focal-attention without consciousness<sub>p</sub>. Now I am all for taking testimony seriously as evidence for consciousness<sub>p</sub>, but for goodness sake, these are amnesics we are dealing with; they can't put short term memory information into long-term memory. Perhaps they are briefly conscious<sub>p</sub> of the pattern and then forget it. For this and other reasons that I don't have the space to go into, this case does not provide a reason for thinking information can reach the executive system without reaching consciousness<sub>p</sub>.

A second case is emergency action in auto accidents. Velmans says sophisticated evasive action happens too fast for consciousness<sub>p</sub>. In addition, after the event, the driver may show no knowledge of exactly what he did. But once again we must consider the possibility of quick forgetting.<sup>2</sup>

Velmans mentions a number of famous cases in which people say that a sophisticated idea has just "popped into their heads," and we all know this sort of thing from our own cases. Because the ideas that pop into consciousness<sub>p</sub> are often quite sophisticated, it is plausible that they are the product of executive activity, so there can be executive system processing that is unconscious<sub>p</sub> except for the conscious<sub>p</sub> result of the processing, that is, the sophisticated idea itself. So we have executive activity without consciousness<sub>p</sub>, and an apparent problem for Model 1. Apparent, but not real. Model 1 precludes information from the specialized modules reaching the executive system without passing through consciousness<sub>p</sub>, but it does not dictate that everything that happens in the executive system is passed to consciousness<sub>p</sub>. Indeed, this "popping into consciousness<sub>p</sub>" phenomenon provides another confirmation of Model 1 over Model 2. Model 1 postulates two-way talk between the consciousness<sub>p</sub> box and the executive box, whereas Model 2 allows

only one-way flow, as it must, being a form of epiphenomenalism. If consciousness<sub>p</sub> did something, it wouldn't be epiphenomenalism. Because we can report on ideas popping into consciousness<sub>p</sub>, and because the executive system controls reporting, Model 1 yields an explanation, but Model 2 does not.

Further relevant data come from the "disconnection" syndromes surveyed by Schacter et al. (1988) and Young and de Haan (1990). First, consider blindsight. [See Campion & Latto: "Blindsight" *BBS* 6(3) 1983.] Suppose blindsight patients could be trained to spontaneously use information in their blind fields. "I somehow know that there is a cup on the left," they would say, "even though I can't see it." It is interesting that this has never been reported. Blindsight patients do find out what is in their visual fields from noticing what they are inclined to guess, but this process (noticing what they are inclined to guess and reasoning that it must be right) is conscious<sub>p</sub>. Some may be overimpressed by the fact that at this crude stage of cognitive neuroscience our evidential base for both executive and consciousness<sub>p</sub> activity is the same, namely testimony. It may look as if we could never have evidence of one without the other. But this conclusion is mistaken, since as just illustrated there is a kind of testimony that blindsight patients could – but don't – give that would be evidence of executive activity without corresponding conscious<sub>p</sub> activity.

Prosopagnosia is a neurologically caused inability to recognize faces. Starting in 1984, it was discovered that many prosopagnosics showed "covert" knowledge of the faces that they said they couldn't recognize. For example, Lady Di's face primed recognition of Prince Charles's name in these patients, as with normals. (See Young & de Haan, 1990, for a description of this syndrome.) Sergent and Poncet (forthcoming) report showing a patient whom they classify as a covert prosopagnosic eight pictures of people who were all famous actors well known to the patient. She was not able to identify any of them. But when told that they all had the same occupation, she was able to get the occupation and achieve something like normal recognition, apparently with the normal feeling of familiarity. These results again suggest that information that reaches the executive also reaches consciousness<sub>p</sub>, which supports Model 1 rather than Model 2.

Finally, on Model 1, we would expect the possibility of consciousness<sub>p</sub> with diminished executive activity to the extent that the connection between consciousness<sub>p</sub> and the executive is weakened or the executive system is not functioning. Plausibly, this occurs when people "lose consciousness," the lost consciousness being cognitive consciousness – e.g., access to the executive system – rather than loss of consciousness<sub>p</sub>. When we say a drunk has "lost consciousness" we are committed to loss of reasoning and control of action; what we say is compatible with "seeing stars," imagery, and other phenomenal events. (We would not expect to know about completely missing executive activity from first hand reports.)

**Conclusion.** My conclusion is that the evidence supports Model 1 better than Model 2, and to that extent disconfirms epiphenomenalism.

One qualification: Although as I said the evidence supports Model 1 better than Model 2, it would be foolish to suppose that the evidence gives much support to any theory. For one thing, it is easy to think of other possible models (for example, a model in which consciousness<sub>p</sub> is "emergent" rather than modular). A second and more important point is that there are no doubt many important ideas about consciousness<sub>p</sub> that we have not yet thought of, and hence theories of which we cannot now conceive. Nagel (1974) observes that our ability to conceive of consciousness being a physical property is like a cave-man's ability to conceive of matter being a form of energy. I believe that this humble attitude should be present in any reasoning about consciousness<sub>p</sub>. No doubt, all thinking about models of consciousness<sub>p</sub> at our current state of knowledge will be seen to

be just thrashing around when people finally get a handle on consciousness<sub>p</sub>, but one hopes that this thrashing around is part of the process of getting somewhere.<sup>3</sup>

#### ACKNOWLEDGMENT

I would like to thank Michael Anthony and Martin Davies for helpful discussion of these issues.

#### NOTES

1. This sense of "epiphenomenal" should be distinguished from a slightly different sense common in philosophy, according to which an epiphenomenal state has causes but not effects. The color of the wires in a computer would be epiphenomenal in Velmans's and my sense since the colors don't have any causal role in the computer, they don't enter into the computer's computations. These colors are not epiphenomenal in the philosophers' sense, however. See Dennett (forthcoming) for more on this.

2. Dennett and Kinsbourne (in press), and Dennett (1988; in press) argue that in a wide variety of cases, consciousness<sub>p</sub>-plus-forgetting cannot in principle be experimentally distinguished from non-consciousness<sub>p</sub>. Work by Potter (1975; 1976) and her colleagues seems to me to provide the beginning of an approach to teasing these apart.

3. See McGinn (1990) for a more pessimistic view.

## (Un)conscious influences in everyday life and cognitive research

Kenneth S. Bowers

Department of Psychology, University of Waterloo, Waterloo, Ontario N2L 3G1, Canada

Electronic mail: bowers@watdcs.bitnet

This commentary makes two main points: (1) first-person experience of consciousness as causal is often more compelling than valid; (2) the research invoked in Velmans's target article has limited ability to discriminate conscious from unconscious processing.

Even without the benefit of the research reviewed in the target article, it seems clear that unconscious influences on thought and action are far more pervasive and profound than conscious experience suggests (Bowers 1984; 1987a; 1987b; Bowers & Meichenbaum 1984, especially the introduction; Bowers, in press). Humans develop over an extended period of time, and they can consciously remember (with questionable accuracy) only a tiny portion of all the episodes of that developmental course. Surely, however, the accumulated unremembered events of one's personal history are not for that reason inconsequential, nor is it reasonable to assume that childhood experience has consequences limited to childhood. Rather, the accumulated events of one's personal history skew how life events are consciously experienced (Wachtel 1977).

Consider, for example, two ordinary citizens who hear the same political speech. One of them views the speech as a threat to civil liberty, the other as a call to patriotic duty. Actions resulting from such divergent perceptions are apt to differ considerably. Nevertheless, both people invoke the speech as justification for their behavior, while remaining unaware of why (or even that) they responded idiosyncratically to it. Obviously, what each individual brings to the speech in terms of personal history contributes importantly to hearing the speech as menacing or patriotic.

In summary, although a first-person account may regard ongoing conscious experience as causal, such experience tends to be blind and deaf to the events of a lifetime that help shape it. If it were possible to eliminate the impact of past history on the correlation between conscious experience and action that seems to flow from it, the relationship would often be small indeed.

Although Velmans overestimates the validity of first-person claims for the causal efficacy of consciousness, he completely

dismisses its causal efficacy on the basis of research findings. This dismissal is uninformed by the limits of the research he cites, most of which is constrained in two ways. First, it relies heavily on experimental methodology. Effects that emerge from experimental investigations are more unambiguously attributable to the manipulated antecedent than to whether the antecedent is perceived (un)consciously. The conditions for disambiguating the latter question have been notoriously difficult to achieve (e.g., Reingold & Merikle 1988) – in part because (un)conscious perception is not, per se, a manipulable antecedent. Consequently, whether a manipulated variable was (un)consciously processed is very much a matter of interpretation. Such interpretations are underdetermined by data, and subject to interpretative biases of the investigator (cf. Holton 1973). For example, Velmans concludes on the basis of empirical research that consciousness is not causal, whereas Holender (1986) invokes some of the same research in support of a contrary claim that there is no convincing evidence for the causal impact of unconsciously perceived information.

The second limitation of the preferred research paradigm for studying (un)conscious processing is that it tends to focus on “basic processes,” which typically means studying reaction-time in milliseconds to simple stimuli, stripped of context and personal significance (cf. Neisser 1982) – evidently in the hope that something like a Wundtian physics of the generalized human mind will result. However, the issue of (un)conscious influences is not reducible to whether information is noticed (Bowers 1984) or crosses a subjective threshold of awareness (Cheesman & Merikle 1986). It also includes a consideration of whether noticed information is appreciated as influential (Bowers 1984; 1987; in press). This question has traditionally been more the bailiwick of social psychology than of cognitive psychology. Cognitive psychologists tend to disregard such work, however, in part because it is less concerned with basic processes than with issues of how people understand and explain their own and other's behavior.

In section 6, Velmans is only partly right when he argues that problem solving is not a result of conscious processing. Consider a problem-solving task in which subjects are required to give a one-word response to each of 15 cumulatively presented clue words (Bowers et al. 1990). Typically, subjects' responses associatively converged toward the solution word (a low associate of each clue). This convergence, however, often occurred without subjects realizing that they were getting associatively closer and closer to the solution (Mermiges 1990). When a hunch finally surfaced, it often did so quite discontinuously, in the form of a sudden insight. Evidently, “the suddenness with which insight sometimes occurs . . . represents an abrupt awareness of a mental product or end state generated by more continuous, sub rosa cognitive processes” (Bowers et al. 1990, p. 95).

The pre-hypothesis stage of inquiry seems to involve automatic spreading activation (Collins & Loftus 1975), a state of affairs consistent with Velmans's suggestion that the achievement of insight is not consciously achieved. Once a hunch emerged, however, the nature of subsequent information processing was quite different: People began testing the hunch against past and subsequently presented clues. The number of clues required to engender a hunch (the context of discovery) was completely uncorrelated with the number of clues required to test it (the context of justification) – suggesting that the cognitive processes underlying these two phases of problem solving were quite different. But did the context of justification involve *conscious* processing of information?

I find it difficult to understand how science and problem solving in general could proceed without the context of justification, and I find it equally improbable that any progress could be made in this context if hypotheses were unconscious. In other words, being conscious of a scientific hypothesis seems a neces-

sary condition for testing it. If I am right, then Velmans's research-based rejection of consciousness as causal is wrong.

Velmans contends that first-person and research-based views of consciousness are complementary and incommensurable; however, his conclusion assumes that these two views are each valid and mutually contradictory. The target article does not clearly establish the warrant of either assumption, so his conclusion does not follow.

### Consciousness and content in learning: Missing or misconceived?

Richard A. Carlson

Department of Psychology, The Pennsylvania State University, University Park, PA 16802

Electronic mail: [cvy@psuvm.bitnet](mailto:cvy@psuvm.bitnet)

There is little doubt that many processes of interest to psychologists occur “outside of consciousness.” At least some of these processes allow (and perhaps require) information-processing accounts. Furthermore, at least some of these accounts appear to be cognitive, in the sense that essential parts of the story refer to the role of knowledge in information processing. Velmans reviews such accounts of a variety of phenomena, concluding that many cognitive processes occur without consciousness. He repeatedly reminds the reader of the “complexity” and “sophistication” of these processes. Generalizing from these examples, Velmans concludes that “no human information processing is conscious,” “in the sense that consciousness enters into or causally influences the process.”

Despite agreeing that much of what appropriately interests cognitive theorists occurs nonconsciously, I could hardly disagree more strongly with Velmans's general conclusion. The most serious problem with his argument is that he makes no attempt to provide an explicit theoretical description of the term *consciousness*. This is unfortunately typical of contemporary discussions of consciousness, despite the availability of substantial theoretical resources in cognitive science for developing such a description. For example, one might start with foundational discussions of concepts such as representation (Lloyd 1989), intentional state (e.g., Searle 1983), or working memory (e.g., Anderson 1983). Velmans suggests replacing the notion of consciousness with the notion of focal-attentive processing, but he rejects the identification of consciousness with focal attention. The “easily understood” concept of focal-attentional processing is also not developed, however, and is hardly less vague than the concepts of consciousness or awareness. It is therefore very difficult to know just what hypothesis about consciousness is being rejected. The basic problem is revealed in the repeated remark that “outputs” or “contents” may be conscious, but processes are not. Imagine a “completely conscious” theory of cognition. Such a theory would specify lawful relations among conscious states, describing their succession and their consequences for observable behavior. We would not expect these lawful relations as such to enter consciousness, however. A theoretical description of a cognitive process must include some terms that do not paraphrase the contents of cognitive states (conscious or not), and to judge all of “human information processing” unconscious because outputs or contents but not processes can be reported is either trivial or absurd. Note also that *some* terms of *most* theoretical descriptions – for example, the dependent variables in many of the studies cited by Velmans – make sense only as paraphrases of conscious contents.