

Caveat Emptor

DANIEL C. DENNETT

Center for Cognitive Studies, Tufts University, Medford, Massachusetts 02155

What I find particularly valuable in the juxtaposition of these three essays on my book is the triangulation made possible by their different versions of much the same story. I present my view as a product of cognitive science, but all three express worries that it may involve some sort of ominous backsliding toward the evils of behaviorism. I agree with Baars and McGovern when they suggest that philosophy has had some baleful influences on psychology during this century. Logical positivism at its best was full of subtle softenings, but behaviorist psychologists bought the tabloid version and sold it to their students in large quantities. George Miller's account of those dreary days is not exaggeration, and the effects still linger in some quarters. (Philosophers are often amused—but they should really be disconcerted—to note that the only living, preaching logical positivists today are to be found in psychology departments.)

The cognitive revolution has triumphed over that barren period, but the three essays do not quite agree on the moral of the story, and they caution psychologists to be wary of the gifts philosophers bear. Since a large portion of the message of my book is to warn cognitive scientists away from the seductive claims of some other philosophers, I heartily agree! So *caveat emptor*, and—that said—let me try once more to sell you my wares.

“The average psychologist is likely to be a fuzzy physicalist with functionalist tendencies,” Baars and McGovern note, suggesting that this fuzziness is pragmatically useful, since our “average cognitive psychologist would surely boggle” at the consequences of pursuing a more hard-edged functionalism of the sort I endorse. This does indeed describe the standard ethos in the field, and one of the main points I was trying to make in my book is that, although that postponement of the mind-boggling implications may have served a useful purpose in the early days, it is now beginning to haunt the field, as we close in on our quarry. Ill-considered Cartesian leanings that once could be gracefully tolerated or ignored are now positively distorting the imaginations of theorists, who will not be able to take the *next* step in creating a theory of consciousness without coming to terms with these residual metaphors and images. Beneath the relatively calm surface of cognitive science, there is much less consensus about these matters than the participants may suppose.

I have been fascinated—overjoyed, in fact—with the confirmation of that claim to be found in the raucous and contrary reactions my book has provoked in the field. For every committed physicalist who applauds, there seems to be a fence-sitter who is rendered uncomfortable. (One says: “I now see that you have made

Reply to commentaries on D. Dennett (1991). *Consciousness explained*. Boston: Little, Brown.

a pact with the devil.”) It is harder to be a good physicalist than many had thought, and confronted with the choice, some find that dualism does not look so bad to them after all. Others, such as Baars and McGovern, make it clear that they are reluctant to commit to such “radical physicalism,” and still others are now self-professed Cartesian materialists, explicitly endorsing the view I claimed was only tacitly presupposed by much theorizing. Then there is Mangan, who confirms so delightfully what I have long suspected: much of the covert appeal of “cognitive science” to some people in this postbehaviorist era is that it seems to let them cling to “an educated common sense view of consciousness.”

Dennett would have us give up intuitions about consciousness that are not only natural and appealing, but have already shown their ‘cash value’ by informing a successful and stable research program.

I wonder how widespread this belief is. Far from “informing” progress in cognitive science, these intuitions have been, I argue, a chief source of artifactual puzzles. The myth that cognitive science *requires* or *confirms* the mentalistic intuitions of Granny has been vigorously fostered by Chomsky and Fodor, and a few other ideologues, and they have had their work cut out for them, over the years, deploring all the progress in cognitive science that erodes our conviction in those intuitions. (See Dennett, 1991, for a historical overview of this campaign.)

What would I do without Mangan! A number of critics have accused me of knocking down strawmen (in particular, my fictional character, Otto), so it is gratifying to be able to present a real, live exponent of the views I have been attacking. I especially appreciate the passion with which he expresses his opinions. I would never have dared let Otto get so het up. (Otto’s words, by the way, are nearly verbatim transcriptions of discussions I had with several of the readers of the penultimate draft.) I plan to share Mangan’s essay widely, since it shows the virulence of the confusions against which my theory is arrayed. Here I will concentrate on the major points he raises, but first I take pleasure in confirming one of his charges.

He accuses me of deliberately concealing my philosophical conclusions until late in the book, of creating “a presumptive mood,” of relying on “rhetorical devices” rather than stating my “anti-realist” positions at the outset and arguing for them. Exactly! That was my strategy, and my conviction that it was necessary has been strengthened immeasurably by the protests of those who, like Mangan, are furious that this has foiled their attempt to head off at the pass all *consideration* of the subversive hypotheses I wanted my readers to take seriously. Had I opened with a frank declaration of my final conclusions, I would simply have provoked a chorus of ill-considered outrage, and *that* brouhaha would have postponed indefinitely any remotely even-handed exploration of the position I want to defend. The fundamental strategic idea of my book was that consciousness was a knotty problem on which no progress was being made largely because people—philosophers, mainly, but not exclusively—too impatiently demanded that one begin with a sort of Declaration of Principles, which then engendered such heated debate that one never got to see the details. So I set out deliberately to postpone philosophical reflection of such ultimate issues until I could get

enough empirical detail and theory-sketch into the arena to *shift the perspective*. Most people have a hard time taking seriously the ideas of functionalism—they just stop short at some simple-minded version of it and dismiss it with a wave of the hand. I had to try to cajole them into pausing long enough to see its strengths in detail. Mangan goes so far as to call this tactic “a kind of intellectual bait-and-switch operation,” but that is an unwarranted slur. As he himself says, I announce at the outset that I will be defending a brand of functionalism, and I never, ever pretend to be engaged in any enterprise other than the one I say at the outset I am engaging in. I also warn—repeatedly, as Mangan notes—that I am setting out to subvert some of the reader’s most fondly held intuitions.

Years ago John Searle felt the need to issue an explicit warning to my readers: “in these discussions, always insist on the first-person point of view. The first step in the operationalist sleight of hand occurs when we try to figure out how we would *know* what it would be like for others” (1980, p. 451). Mangan continues this tradition of the Thought Police—so worried that people will be swept up by the third-person point of view and will no longer have the critical faculties needed to cling to their cherished beliefs. From their perspective I am the sinister dope peddler, offering free samples and saying “Try it, you’ll like it!” From my perspective, I see my target audience as trapped by a tightly constructed wall of stultifying, imagination-crippling intuitions; my only hope of breaking them free is to quietly insert some subtle wedges and tap, tap, tap—until all of a sudden things give way and a new possibility can be seen. Subversive? Yes indeed, but my methods are aboveboard.

Dennett’s . . . heterophenomenology . . . would have us slough off much useful phenomenological information, and this would encourage habits of mind that all but killed off consciousness research earlier in this century.

I go to some lengths in my book to explain that heterophenomenology is nothing other than the scientific method applied to the phenomena of consciousness, and thus the way to *save* the rich phenomenology of consciousness for scientific study. I did not invent the heterophenomenological method; I just codified, more self-consciously and carefully than before, the ground rules already tacitly endorsed by the leading researchers. I challenge Mangan to name a single respectable experiment in cognitive science that does not fit comfortably within the heterophenomenological method. We can study mental imagery, dreaming, disgust, color vision, taste, memory, and every other subjective aspect of consciousness, all without falling into the introspectionist traps that “all but killed off consciousness research earlier in this century.” If Mangan thinks the cognitivist revolution against behaviorism signalled anything like a return to introspectionism, he should look more closely at the methods and presuppositions of contemporary research.

Consider in this regard one of the most productive and interesting controversies in cognitive science, the Pylyshyn–Kosslyn–Shepard (et al.) debates about mental imagery. The principal participants in this debate have sometimes talked past each other, but they have all recognized that here were legitimate, empirically explorable issues *that could not be settled by just taking subjects at their word*.

Pylyshyn knew that it *seemed* to subjects as if they were rotating mental images (for example), but he insisted—and Kosslyn and Shepard agreed—that that was not conclusive. (Reports about conscious experiences made by human subjects are, as Mangan says, the “primary data” in consciousness research, but that means just what it says: the reports *are* the *data*—they are not reports *of* data.) Whether there *really were* mental images was something to be settled by coming up with theories that (1) honored the heterophenomenology of subjects (by taking on the burden of explaining why it seemed to subjects the way it did) while (2) differing precisely on which features or elements of that heterophenomenology were taken seriously in models of internal process. Different models made different predictions that could be tested. Good, objective science made out of subjective raw materials—the essence of cognitive science. The controversies were there before I came along. I did not create the doubts that Mangan finds so subversive; I just described the neutral ground from which to explore them.

Mangan does not want neutrality; he wants a commitment *in advance* to the reality of a certain class of objects. He claims to detect a shift in my book “to the point that we are told that even experimentally controlled, quantitative and reliable reports about our inner life . . . are to be taken as wildly [*sic*] dubious, although not, perhaps, utterly false.” But heterophenomenology truly is neutral, and so it should be. The systematic neutrality of heterophenomenology is not our normal interpersonal mode, as I point out, and it is amusing to note that Mangan is described in anticipation in my book:

Any subject made uneasy by being granted this constitutive authority might protest: “No, really! These things I am describing to you are *perfectly real*, and have exactly the properties I am asserting them to have!” (p. 83).

And, true to form, I reply that I do not for a moment doubt Mangan’s sincerity. I include it as part of what must be explained by any theory of his consciousness. Mangan misses the import of my comparison to anthropology. I do allow “ontological conclusions” to be drawn in the end; consider the ontological conclusions I claim are open for the anthropologists to make about Feenoman: (1) he is real, but just an ordinary mortal, (2) he is *just* a figment of the believers’ imaginations, and even (3) he is a real divinity. One does not want to start with the *presumption* that (3) is going to be confirmed—that is the neutrality of science in action. We *should* refuse to “go along with the standard realist transition” in this case, for this is no more a “standard” case than Feenoman is. Applied to phenomenology, we do not want to start with the presumption that every *apparent* object and feature of our conscious lives is *really there*—a real element of experience. And I do give examples of how we can sometimes “move from heterophenomenology back to real experiences,” summarized in the passage that to Mangan’s overheated imagination appears to be a “begrudging” admission. As I say, “Sometimes the unwitting fictions we subjects create can be shown to be true after all . . .”

Toribio expresses a more cautious (but less clear) skepticism about the limitations of heterophenomenology. She distinguishes externalism from the third-person point-of-view by declaring that the former is silent about “internal struc-

ture,” but I find the distinction problematic when we turn to ask: what sort of internal structure is at issue? Consider the notoriously externalist Turing test. It is routinely challenged, as Toribio notes, by the utterly imaginary—as she says “(in)famous”—possibility of giant look-up tables. Turing knew perfectly well that no such internal structure was *actually* possible. He knew that the only way a finite, compact entity could pass his test would be by containing structures that *composed* their responses on the fly, in real time. And such composition, if it was to be appropriately responsive to its input, would (barring miraculous coincidence) *have* to be accomplished by structures that somehow useably reflected the meanings of the inputs and outputs, that represented features of the world, etc. So the Turing test *indirectly* constrains internal structure: an intelligence can be of any structure whatever just so long as it makes use of an articulated and accessible representation of the world and the meanings of its interactions with that world. Now is this still an externalist view or did Turing come up with an elegantly minimal way of imposing the right internalist constraint—neither too lax nor too “chauvinistic?”

Another way of asking this question—and getting clearer about what Toribio means by “internal” and “subjective”—is this: does a zombie have real subjectivity in virtue of having such an internal structure? I say yes, but I take Toribio to claim, via her Guinness example, that I am unable to support this answer because my model, lacking detail about just how the mechanisms underlying the virtual machine do their work, is not yet able to *explain* the relation between the lower and higher levels. She says: “But we do not know why those physical mechanisms have as a result such-and-such phenomena—consciousness—instead of completely different ones.”

I grant that my sketch lacks details—on purpose. I deliberately backed off from a lot of the details I could have provided, for two reasons:

(1) I would not want to be “hung” for the wrong options. This is a familiar tactic in science, not just a philosopher’s move.

(2) My larger purpose was to create a new vision, and for many people, too many details get in the way of seeing a new vision. You can’t please everybody.

But even so, I think there is enough detail in my model to respond directly to Toribio’s challenge. I offer, in effect, explanations of *why* a zombie organized in accord with my model would “experience” color phi, or metacontrast, under the conditions that we experience these phenomena. Moreover, the ground is laid for explaining why zombies would exhibit the same reactions to Guinness as regular folk do. (Ah, but does what goes on inside them have the same features? Yes, if the features in question are those structural features that must be there to give the zombies their dispositional similarity to us in a nonmiraculous way. Pray what other internal features did you have in mind?) Or consider one of my favorite examples of third-person cognitive science defying the “upwards opacity” (as Wilks calls it) of conscious experience: Ramachandran’s (1987) invention of a new kind of paradoxical motion, never experienced before. Appreciating something about the different roles played by the magnocellular and parvocellular pathways in the visual system, he predicted that under conditions of isoluminance, he could make a stationary blotch of color “move” without changing

location. Here we do not just explain previously known phenomenology; we predict unexpected phenomenology. Just as organic chemists have recently synthesized novel organic molecules, predicting some of their macroproperties in advance, Ramachandran has synthesized novel phenomenology—a telling demonstration that our explanations are not all mere correlations. Ramachandran's demonstration is particularly telling because of its use of physiological detail, and if this is all Toribio is asking for, I readily acknowledge that we should ask for more, and more—much more than I have provided in my sketch. But my point remains that all this can be accomplished within the (familiar, already adopted) constraints of heterophenomenology; the only *internal* matters we need attend to are those objectively discernible happenings in the brain.

It is interesting to compare Baars and McGovern to Mangan on the zombie question. In different ways they wish the issue did not come up. Mangan deplors my insistence on maintaining neutrality regarding the consciousness of talking robots, while Baars and McGovern urge that “we can develop ideas about consciousness much more easily” if we set aside the zombie question while pursuing the method of contrastive analysis restricted to everyday human beings. I agree that contrastive analysis of the sort they describe is an excellent method. I join them in recommending it, and I also recommend that they simply ignore anybody who claims that it is a defect of their method that it would be unable to distinguish conscious beings from zombies if such there be. As Baars and McGovern note, in my role as philosopher I am professionally obliged to respond to these objections, but the rest of you may avert your eyes—or be spectators if you have a taste for these blood sports. But just bear in mind that it is not serious, no matter how loud the contestants yowl.

Do you know what a zagnet is? It is something that behaves exactly like a magnet, is chemically and physically indistinguishable from a magnet, but is not really a magnet! (Magnets have a hidden essence, I guess, that zagnets lack.) Did you know that physical scientists adopt methods that do not permit them to distinguish magnets from zagnets? Are you shocked? Do you know what a zombie is? A zombie is somebody (or better, something) that behaves exactly like a normal conscious human being and is neuroscientifically indistinguishable from a human being, but is not conscious. I do not know anyone who thinks zagnets are even “possible in principle” and I submit that cognitive science *ought* to be conducted on the assumption that zombies are just as empty a fiction. If instead we start out with the contrary assumption that the zombie hypothesis is serious, we may be honoring an intuition that is “natural and appealing,” but its “cash value” to cognitive science to date has been zero, and now it promises to put us in the red unless we abandon it.

Certainly the most useful exposure of covert Cartesianism in Mangan's essay is his discussion of consciousness as a medium. He is right on target about the intuition I want people to reject, and then, bless him, he defends it explicitly. Consciousness, he proposes, is “a distinct information-bearing medium.” He points out that there are many physically different media of information in our bodies: the ear drum, the saline solution in the cochlea, the basilar membrane, each with its own specific properties, some of which bear on its capacity as an

information medium (e.g., the color of the basilar membrane is probably irrelevant, but its elasticity is crucial). “If consciousness is simply one more information-bearing medium among others, we can add it to an already rather long list of media without serious qualms.”

But now consider: all the *other* media he mentions are *fungible* (replaceable in principle without loss of information-bearing capacity, so long as the relevant physical properties are preserved).¹ As long as we are looking at human “peripherals” such as the lens of the eye, or the retina, or the auditory peripherals on Mangan’s list, it is clear that one could well get by with an artificial replacement. So far, this is just shared common sense; I have never encountered a theorist who supposed an artificial lens or even a whole artificial eye was impossible; getting the artificial eye to yield vision *just* like the vision it replaces might be beyond technological feasibility, but only because of the intricacy or subtlety of the information-bearing properties of the biological medium.

And here is Mangan’s hypothesis: when it comes to prosthetic replacements of media, all media are fungible in principle *except one*: the privileged central medium of consciousness itself, the medium that “counts” because representation in *that* medium is conscious experience. What a fine expression of Cartesian materialism! I wish I had thought of it myself. Now neurons are, undoubtedly, the basic building blocks of the medium of consciousness, and the question is: are they, too, fungible? The question of whether there could be a conscious silicon-brained robot is really the same question as whether, if your neurons were replaced by an informationally equivalent medium, *you* would still be conscious. Now we can see why Mangan, Searle, and others are so exercised by the zombie question: they think of consciousness as a “distinct medium,” not a distinct system of content that could be realized in many different media.

Is consciousness—could consciousness be—a distinct medium? Notice that although *light* is a distinct medium, *vision* is not; it is a distinct content system (a distinct information system) that could in principle be realized in different media. What makes something vision is that it carries information about distal objects and properties via the peripheral *transduction* of electromagnetic radiation in the “visible spectrum”—you leave the light at the doorway, in other words. What makes something audition is that its content arrives via sound waves (but is then transduced and processed in media that are in principle fungible). The content systems of vision and audition have different logical spaces (and are different in different species, of course), due to the specific information-bearing properties of the initial media (light and sound waves) and the different physical structures that happen to be employed by different brains.

But someone might want to object that this leaves out something crucial: there

¹ *Fungibility* is a term from the law. Most debts are fungible. If I borrow \$5 from you, I can repay the debt with *any* 5 dollar bills (or 20 quarters or . . .), but some loans are not fungible: if you loan me your Cezanne landscape, I must return *that very one*, not another painting of similar market value, or even another Cezanne of similar aesthetic value. The term is an etymological cousin of *functionalism* (via *fungor*, to perform) and nicely captures the hallmark of functionalism, its commitment, at some level, to multiple realizability. As Block and others have noted, functionalism is a *kind* of behaviorism: handsome is as handsome does.

is not really any vision or audition at all—not any *conscious* vision or audition—until the information that moves through the fungible peripheral media eventually gets put into the “distinct medium” of consciousness. This is the essence of Cartesianism—Cartesian materialism if you think there is something special about a particular part of the brain (so special that it is not fungible, even in principle), and Cartesian dualism if you think the medium is not just one more physical medium. The alternative hypothesis, which looks pretty good, I think, once these implications are brought out, is that, first appearances to the contrary, consciousness *itself* is a content system, *not* a medium. And that, of course, is why the distinction between a zombie and a really conscious person lapses, since a zombie has (by definition) exactly the same content systems as the conscious person.

Finally, Mangan provides a vivid example of a familiar Janus-faced reaction to my book: “Ho hum, what else is new?” accompanied by remarks that manifest either outrage or a bland failure to distinguish the crucial theses—or both! We have already noted the outrage; here is the bland failure:

Multiple Drafts turns out to largely refer to *unconscious* processes.

That’s how it turns out? No, that’s the assumption that was made, tacitly, by all those sources who had uncovered various elements of the Multiple Drafts model, but not seen their implications. As Mangan’s review of Neisser, Mandler, and others makes plain, there is nothing original about the idea of multiple concurrent unconscious processes or its contrast with serial processes. But almost everyone who has adopted or even contributed to that vision has operated under the (usually tacit) presumption that these parallel processes then *send their results* to some more central *conscious* arena of serial activity. That is the assumption that is attacked by the Multiple Drafts model, which *is* a model of conscious events, not their unconscious “predecessors.” Since the Multiple Drafts model is a model of consciousness, it is new and revolutionary not just to me (as Mangan suggests), but to those who shared some of the ideas in it but who thought of them as referring only to unconscious precursors of “later” or “more central” conscious events.

The recognition of this novelty is manifested in the opposition (or even just skepticism) expressed in response to the pivotal implication of the Multiple Drafts model: since there is no finish line, there is no fact of the matter to distinguish Orwellian from Stalinesque content revisions. Almost no one finds this to be old hat! Baars and McGovern see this as a typical philosopher’s “impossibility proof” and hence as a “verificationist” move that should be resisted with the usual “realist” defiance, but that is not the only way to conceive of the proposal I have made, nor is it the right way. Consider the following case:

You go to the racetrack and watch three horses. Able, Baker and Charlie, gallop around the track. At pole 97 Able leads by a neck; at pole 98 Baker, at pole 99 Charlie, but then Able takes the lead again, and then Baker and Charlie run ahead neck and neck for awhile, and then, eventually, all the horses slow down to a walk and are led off to the stable. You recount all this to a friend, who asks “Who won the race?” and you say, “Well, since there was no finish line, there’s no telling. It wasn’t a real race, you see, with a finish line. First one horse led and then another, and eventually they all stopped running.” The event you witnessed was not a real race, but it was a real event—not some mere illusion or figment of your

imagination. Just what kind of an event to call it is perhaps not clear, but whatever it was, it was as real as real can be.

Notice that verificationism has nothing to do with this case. You have simply pointed out to your friend that since there was no finish line, there is no fact of the matter about who “won the race” because there was no race. Your friend has simply attempted to apply an inappropriate concept to the phenomenon in question. That is just a straightforward logical point, and I do not see how anyone could deny it. You certainly do not have to be a verificationist to agree with it. I am making a parallel claim: the events in the brain that contribute to the composition of conscious experiences all have locations and times associated with them, and these can be measured as accurately as technology permits, but *if* there is no finish line in the brain that marks a divide between preconscious preparation and the real thing—if there is no finish line relative to which preexperienced editorial revision can be distinguished from postexperienced editorial revision—the question of whether a particular revision is Orwellian or Stalinesque has no meaning.

It is not that “the time when the stimulus becomes conscious is forever impossible to know” (Baars and McGovern, 1993) but rather that the definite description “the time when the stimulus becomes conscious,” like “the time at which Able won” or “the amount of time by which Able beat Charlie” has no reference—despite first appearances. We can certainly time the arrival of each horse at the track, and at each location on the track, and we can time the onset and duration of each horse’s period of being in the lead. As we scan these lists of times, none of them can be picked out as the time at which any horse “finished.” In parallel fashion, we can time the arrival at the retina of a stimulus pattern, the various transductions, reactions, decisions, revisions, bindings, rebindings to the effects of other stimulus patterns, until the trail fades in one sort of quiescence or another, but unless we can motivate the drawing of a line (marking a “change of media” in Mangan’s terms) across that sequence of events as the line marking the onset of *consciousness* of the content in question, “the time at which the stimulus becomes conscious” is an empty definite description.

Now here is where push comes to shove. Some people are so sure that there *has* to be a fact of the matter about “the time when the stimulus became conscious” that they are willing—indeed eager—to suppose that there is such a line to be drawn in space and time. To them I say: show us where and why to draw this line. I am claiming that there is no scientific ground for it, but only an ancient (Cartesian) prejudice. I think we *already* know enough about the brain (and the various psychological phenomena I discuss under this heading) to disconfirm this hunch, so I view the Multiple Drafts hypothesis as not only testable but tested (in this minimal, but crucial, regard—the many details of nonsketchy versions of the hypothesis await formulation and testing).

Baars and McGovern suggest something different: that results such as Libet’s shed light on the hypothesis. Libet’s interpretations of his experiments, however, have always been question-begging and sometimes even self-contradictory. He has recently (CIBA Foundation workshop, London, June, 1992, forthcoming)

interpreted his cortical stimulation experiments as showing the possibility of "temporal referral" (a cortical event occurring at time t is assigned some earlier time t' as its subjective time location) while concomitantly giving an interpretation of his volition-timing experiments that depends on denying this possibility of temporal referral. He claims to have invented a method for fixing the time of onset in consciousness of an intention, by having subjects note the appearance of a rotating clock hand at the instant of conscious decision, but the only way the *subjective* simultaneity of a visual experience of a clock dial with the experience of intending could give us evidence of the objective time of conscious intending would be if the visual experience (or the experience of intending) was incapable of temporal referral; if either might be temporally referred forward or back, their subjective simultaneity would show nothing about the actual timing of the "conscious intention." You cannot have it both ways; you cannot show that the Orwellian/Stalinesque distinction is testable by experiments that you interpret in a self-contradictory or question-begging way, but you sure can illustrate the sorrows that confront those who cling uncritically to the Cartesian assumptions of their upbringing.

To summarize: I have presented a theory of the mind and a heterophenomenological method that, I claim, do justice in detail to the best work in cognitive science and lay the foundation for the future by dissolving certain pseudo-problems that have infected the imaginations of theorists. If I am wrong, there ought to be a way of showing it that does not simply appeal to the traditional obviousness of the intuitions I claim must go. I have given reasons for dismissing these intuitions in the face of their longstanding popularity; other philosophers have yet to defend them with anything but a pledge of allegiance. I sympathize with psychologists and other cognitive scientists who do not want to be burned by fast-talking philosophers the way their behaviorist elders were, and I submit to them that they should be particularly wary of thinking they can *avoid* philosophy by clinging to the good old-fashioned home truths thundered from the pulpits by such philosophers as Fodor and Searle. Do not take their word for it, and do not take mine. If you want to avoid being taken, you will just have to think it through for yourselves.

REFERENCES

- Dennett, D. C. (1991). Granny's campaign for safe science. In B. Loewer & G. Rey (Eds.), *Meaning in mind: Fodor and his critics* (pp. 87-94). Oxford: Blackwell.
- Libet (forthcoming). CIBA workshop volume.
- Ramachandran, V. S. (1987). Interaction between colour and motion in human vision. *Nature*, **328**, 645-647.
- Searle, J. (1980). Author's reply to commentaries on Minds, Brains, and Programs. *Behavioral and Brain Sciences*, **3**, 417-58.
- Baars, B. J., & McGovern, K. (1993). Does philosophy help or hinder scientific work on consciousness. *Consciousness and Cognition*, **2**, 18-27.