

resolution with a Selection function restricted to Definite clauses) resolution, is (refutation) complete for Horn clauses, standard Prolog implementations are incomplete due to the choice of a depth-first control strategy (Lloyd 1987). Going the other direction, more may be derivable than is implied logically by the database and rules using procedural attachment (as arithmetic is handled in Prolog), so not everything potentially explicit need be implicit either.

Logical consequence need not be determined by standard first-order logic; any variant for which logical consequence is definable will do. We may choose the logic to fit the epistemic subject. Given a choice of logic, the derivational powers ascribed to the knower may be further limited, say, by limiting the length of claims of inference (Konolige 1986, p. 24). In this way, the gap between implicit and potentially explicit reflects design in light of space/time limitations (as in Prolog), but it is not an implementational artifact.

Dennett's exposition of tacit knowledge seems to me ambiguous. There is a sense of "representation" that does not require representations: for example, the sense in which I am represented in Washington. The laws of arithmetic and Boolean algebra are "represented" in arithmetic-logical units (ALU's) in that ALU's are designed so that their input-output behavior conforms to these laws, but an ALU in operation no more contains their representations than a paper airplane in flight contains representations of the laws of aerodynamics.

Representation without representations might be all that Dennett means by tacit representation. If so, an important mode of representation has been omitted: what is in computer science called "procedural," as opposed to "declarative," representations. Winograd calls this distinction "an artificial intelligence incarnation of the old philosophical distinction between 'knowing that' and 'knowing how'" (Winograd 1975, p. 186). In terms of our model, the declarative knowledge of a program is contained in its database (closed under logical consequence/derivation to include implicit and potentially explicit knowledge), and the procedural knowledge is contained in the rules.

Dennett often refers to explicit representations of rules, however, suggesting that the rules of logic programs and production systems are just the sorts of things he would call explicit representations. In our model, however, the only representations of rules are the rules themselves. Just to have a separate term for what these rules are representations of, let's say that rules are procedural representations of *procedures*. They are explicit from the programmer's point of view, but, since the rules are not in the database, nor entailed by or derivable from the database using the rules, the procedures they represent are not explicitly, implicitly, or potentially explicitly known by the program. Were the rules not explicit from the programmer's point of view, that is, if the rules did not appear as more-or-less separable components at some level of implementation, then, although their procedures might be *represented* in the weak sense described above, there would be no *representations* of them.

Unless, of course, we represented the procedures *declaratively* in the database. Just as whatever can be declaratively represented can be represented through procedural attachment, so whatever can be represented as a rule can also be represented in the database as what we shall call a *plan*. Dennett's categories could be extended to appropriate plan representations, but it is more interesting to consider their procedural analogues. For example, in the case where a rule is replaced by a plan and a matching metalevel rule, might we call the original rule implicit or potentially explicit?

The primacy of tacit knowledge, a point Dennett credits to Ryle (p. 217), comes to this: Plans by themselves are inert; only rules can put them into action. Ultimately, despite the interchangeability of procedures and data from the programmer's perspective, from the machine's point of view only procedures run. I should not say that my Sun computer *knows how* to execute compiled C programs except metaphorically, however,

because this ability is built in at a level where, if it is legitimate to speak of representation at all, it is only register locations and memory addresses that are represented.

Dennett describes how "transient tacit representations" might yield to explicit representations to handle increased complexity (pp. 223–24). As an example of the former, Dennett offers an automatic elevator that can switch between workday and night/weekend programs. (The reading of tacit representation as representationless seems dictated here.) Dennett next imagines a bundle of routines driven by perceptual cues, reminiscent of those in Dennett 1973 (cf. Smith 1986). These are like production systems without an internal database. So understood, their problem is obvious: There is no way to connect the various productions. The insertion of a declarative element between rules allows for what Ryle (1949) calls "multitrack dispositions," like belief.

Dennett's evolutionary account fits the received wisdom that a declarative component provides flexibility and economy in complex systems (cf. Winograd 1975, pp. 187–88). But, he claims, such a component would be limited to an "internal semantics" of memory addresses, internal operations, and so forth (p. 224). If this means that a program inherits the semantics of the microcode executing it, then this just seems to be a confusion of levels, or a Rylean category mistake (Ryle 1949, p. 16; cf. Dennett 1984a, pp. 74–75). More likely, Dennett intends a substitutional semantics, where terms in a program refer only to other parts of the program. I don't understand this claim, and wonder whether Dennett would say that the words in this commentary have a merely internal semantics too, until supplied with a "globally defined role" (p. 224).

#### ACKNOWLEDGMENT

This research was supported in part by the National Science Foundation under grant DCR-8516243. The author is indebted to Bruce T. Smith for helpful discussion.

## Connectionism, Realism, and realism

Stephen P. Stich

Department of Philosophy, University of California, San Diego, La Jolla, Calif. 92093

About ten years ago, Dennett began using the term "instrumentalism" in describing his account of intentional descriptions and explanations. He now considers the choice of that term to have been "a tactical error" (p. 72). For, he tells us, he has always been a sort of realist about intentional states, though he is not a full-blooded Fodorian "Realist with a capital R" (p. 71) (for the most recent version of Fodorian Realism, see Fodor 1987). The latter view is that the actual internal states that cause behavior will be functionally individuated in roughly the way that belief/desire psychology carves things up. And Dennett, true to his Rylean roots, has always insisted that this sort of Realism about belief/desire psychology is singularly implausible. It is my view that there is an argument against Dennett's stand to which, until recently, there was no good reply. However, recent achievements in connectionist modeling of cognitive processes have made Dennett's skepticism about the Realist interpretation of common sense psychology significantly more plausible than it was earlier. In the first part of this review I will give a brief sketch of the role I think connectionist research plays in the debate between Dennett and the Realists. In the second part, the focus will shift to Dennett's version of realism – realism with a little r. My theme there will be that the connectionist models that lend support to Dennett's anti-Realism seem to leave no room for his (little r) realism. As I see it, there is no logical space between a full-blooded (big R) Realism and an unequivocal Eliminativism.

**1. Connectionism and anti-Realism.** It is agreed on all sides that folk psychology is a useful, powerful predictive tool.

Though exactly how we do it is far from obvious, we are clearly very good at attributing beliefs and desires to people, and at deriving predictions about their behavior based on these attributions. Although far from flawless, the predictions are often correct, and in our day-to-day lives we could hardly get on without them. But *why* does this folk psychological strategy of prediction work so well? The answer urged by the Fodorian Realist is that it works because the beliefs and desires are Real, functionally individuated states of the brain, and that the principles we tacitly invoke in extracting our predictions are good approximations to the principles that actually govern the causal interactions among perceptions, beliefs, desires, and behavior. One very natural way to elaborate on this suggestion, as Fodor has long argued, is to build models of human cognition that posit more or less sentence-like representations to play the role of beliefs and desires. The models also posit formal processes that manipulate representations along the lines presupposed by folk psychology. Dennett, like Ryle before him, has consistently resisted the suggestion that the internal workings of the brain are functionally organized in a way that mirrors the posits and processes of folk psychology. But to the Realists, this resistance seemed perversely unscientific, since neither Dennett nor anyone else was prepared to offer any serious suggestions about *alternative* mechanisms that might account for the success of folk psychology in predicting behavior. In the absence of serious alternatives, the Realist hypothesis was “the only straw afloat.” And, as Fodor often reminded us, in other sciences, when there is only one plausible hypothesis that can explain a phenomenon, the inference standardly drawn is that the hypothesis describes the way the world Really is.

It is at this point that connectionist models become relevant to the argument. For, as Dennett notes in his Reflection, “The Language of Thought Reconsidered,” these models have had some remarkable success in simulating behavior that, prior to connectionism, no one had any idea how to stimulate without invoking a set of sentence-like representations and a system of formal rules for manipulating them. Moreover, the connectionist models exhibit some properties (like graceful degradation) which natural cognitive systems exhibit, and which rule-and-representation simulations notoriously do not. I agree with Dennett that “it is too early for verdicts” (p. 229) about the ultimate success of connectionism in modeling the mechanisms underlying the behavior that folk psychology predicts. But for opponents of Realism about folk psychology, the importance of connectionist research is that it suggests there are *lots* of powerful computational architectures whose properties and limitations we have not yet even begun to explore, which might be used in modeling human cognition and behavior. There are *lots* of straws afloat. Since models in the folk psychological mold now have lots of serious competition, no methodological injunction entails that we should be Realists about the posits of folk psychology.

**2. Connectionism and (little r) realism.** Let me turn now to the prospects for Dennett’s (little r) realism, which holds that while (big R) Realism is false, claims invoking beliefs and desires are perfectly true and scientifically respectable simply because folk psychology’s “intentional stance” does a decent job at predicting behavior. As I see it, there are three possibilities that we might have to contend with. The first is that connectionist theories and their kin will turn out to be inferior to the older rules-and-representations theories inspired by folk psychology. If this should happen, then (little r) realism will lose out to (big R) Realism, not because theories based on folk psychology are the only game in town, but because they are the *best* game in town. A second alternative is that the best theories for explaining human behavior are connectionist theories, but that theories in the folk psychological paradigm can be more or less smoothly *reduced* to these connectionist theories, much as certain theories in chemistry can be reduced to theories in physics. This is the scenario favored by those who urge that connectionism should be viewed as a theory aimed at explaining how rules-and-

representations models are implemented. On this alternative too, (little r) realism will lose out to (big R) Realism, since reduction renders the states and processes reduced no less Real. The third alternative is that the best theories will be connectionist theories (or some variation on that theme), and that no plausible way can be found to reduce folk theories to connectionist theories. In this case, it seems to me that there is no alternative but to bite the Eliminativist bullet, and concede that the ontology of folk psychology is just mistaken: There are no such things as beliefs and desires.

As I read Dennett’s most recent writings, this third alternative is the one he considers most likely. Yet he is reluctant to endorse the “extreme eliminativism” (p. 233) that it seems to entail. The two reasons he offers for his reluctance strike me as unconvincing. The first is that “until the rest of the world catches up with [the eliminativists] and shares their world view” there will be no alternative but to continue using the language of folk psychology in the law courts and in everyday life. This is right, of course, but seems irrelevant to the eliminativist’s thesis. Even in these secular times, the law continues to speak of certain events being “acts of God.” But surely those of us who are atheists do not compromise our convictions when we choose to talk in this way. If Dennett’s realism about beliefs and desires is like the atheist’s “realism” about God, then it is a view I find difficult to distinguish from Eliminativism. Dennett’s second reason is that folk psychology has “tremendous – if flawed – predictive power” (p. 234). Here too the relevance of the point is unclear, since as Churchland notes in his accompanying *BBS* review, discredited theories with discredited ontologies *typically* have considerable – though flawed – predictive power. If they didn’t, they would probably never have been accepted in the first place. Perhaps the best way to end this review is to invite Dennett to say a bit more about how realism (with a little r) differs from Eliminativism (with a big E).

## What really matters

Charles Taylor

*Department of Political Science, McGill University, Montreal, Quebec, Canada N3A 1G5*

“This book is about how to talk about the mind,” says Dan Dennett on page 1. But among many interesting subjects in this excellent book, this is one he never quite gets around to talking about. The central issues are all elided, glided by. Just how this happens is very important to understand. Dennett writes from within a complex of interlocking philosophical doctrines which have the effect of interpreting away the central phenomena of mind. So what emerges from the laborious dialectic is another demonstration of the quasitautology: If you can’t see something, you probably won’t know when you’ve failed to explain it.

One of the interlocking pieces of the screen is the stance announced on page 7: “I propose to see . . . just what the mind looks like from the third-person, materialistic perspective of contemporary science.” This is the stance which was specifically designed to exclude the crucial phenomena of “mind.”

What are these phenomena? Not primarily those of “consciousness,” but what I want to call “mattering”: the fact that things matter to ourselves and animals in a way they don’t to your calculator or word-processor. The class of matterings is broader than and includes that of intentional states. States in this class include not only my belief that it’s snowing or my hope that the ski-lift will run, but also the pain in my tooth and the butterflies in my stomach.

Now mattering is what the screen of interlocking doctrines interprets away, or what is the same thing, the distinction between beings for whom things matter and the others. Dennett does this by treating mattering states nonrealistically. This echoes the move of earlier positivist writers in relation to the