6 Two contrasts: folk craft versus folk science, and belief versus opinion

Daniel C. Dennett

Let us begin with what all of us agree on: Folk psychology is not immune to revision. It has a certain vulnerability in principle. Any particular part of it might be overthrown and replaced by some other doctrine. Yet we disagree about how likely it is that that vulnerability in principle will turn into the actual demise of large portions – or all of folk psychology. I believe folk psychology will endure, and for some very good reasons, but I am not going to concentrate on that here. What I want to stress is that for all its blemishes, warts, and perplexities, folk psychology is an extraordinarily powerful source of prediction. It is not just prodigiously powerful but also remarkably easy for human beings to use. We are virtuoso exploiters of not so much a theory as a craft. That is, we might better call it a folk craft rather than a folk theory. The *theory* of folk psychology is the ideology about the craft, and there is lots of room, as anthropologists will remind us, for false ideology.

What we learn at mother's knee as we are growing up, and what might be to some degree innate, is a multifarious talent for having expectations about the world. Much of that never gets articulated into anything at all like propositions. (Here I am in partial agreement with the new Paul Churchland. He now wants to say that folk psychology is a theory; but theories do not have to be formulated the way they are formulated in books. I think that is a good reason for not calling it a theory, since it does not consist of any explicit theorems or laws.) But now, what is this thing that is folk psychology, if it is not a theory? What kind of a craft is it? I certainly have had my say about that, in *Brainstorms* (Dennett 1978) and in *The Intentional Stance* (Dennett 1987), and I shall not try to telescope all I say there into a summary here. Instead, I shall expand on the similarities between folk psychology and folk physics – two crafts that repay attention, and that *should* be studied with the methods of anthropology, not just the informal methods of philosophers.

If we look at folk physics, we discover some interesting anomalies. Folk physics is as effortless, as second-nature as folk psychology, and it keeps us one step ahead of harsh reality most of the time. A pioneering analysis of a portion of folk physics is found in Patrick Hayes's (1978, 1979) work on what he calls the "naive physics of liquids." Consider how robust and swift our anticipations are of the behavior of liquids under normal circumstances.

For instance, if you and I were seated at a table, and I happened to overturn a full glass of water, almost certainly you would push your chair back rather quickly, because you would expect that if you didn't, the water would run off the edge of the table onto your clothes. We do these things *almost* without thinking. But in fact, if the circumstances were slightly different – if there were a lip on the table, or a towel where the liquid was pouring -- we would effortlessly have very different expectations, and behave differently. We know about how towels absorb liquids, and about how liquids don't roll up over the edge of lips under normal conditions. These are part of a huge family of expectations we have about liquids, which we would find very difficult to enunciate in a string of propositions – though that is just what Hayes very ingeniously attempted. He tried to do a formal, axiomatic folk physics of liquids. In the folk physics of liquids, he notes, siphons are impossible. So are pipettes – putting your finger over the top of the straw and drawing up the Pepsi. Hayes views this feature as a virtue of his theory because that is what folk physics declares; it is different from academic physics. There is something counterintuitive about both pipettes and siphons. Therefore, if you want to codify à la anthropology what people actually think and do, you want to make folk physics *predict against* such things as siphons and pipettes.

Now when we turn to folk psychology, we should expect the same thing. We should expect that some deeply intuitive ideas of folk psychology will just turn out to be false. Folk physics would say that gyroscopes are impossible, and that sailing upwind is impossible, but we come to learn that they are not, strange as this seems. We were just wrong about these matters, but even after we learn this, the intuitions don't go away; the phenomena still seem counterintuitive. So we might expect that folk psychology, under the pressure of advanced academic psychology and brain science, will similarly come a cropper. Certain deeply intuitive principles of folk psychology, perhaps never before articulated, may have to be given up. (I presume that folk physics never

articulated the principle that siphons were impossible until siphons were well known to be possible; when siphons were observed, people perhaps said, "Hey wait a minute! Things shouldn't happen that way!") So it would be surprising if we had already articulated the principles of folk psychology that academic psychology is going to undermine – if it undermines any. Rather, we shall find ourselves beset with extremely counterintuitive clashes, and something will have to give. And what very likely will give is parts of folk psychology. That is to say, the craft itself will come to be adjusted to acknowledge the existence of perplexities and peculiarities and contrary predictions that the craft had never before made.

I want to distinguish between craft and ideology, between what we learn to do, and what our mothers and others have actually *told* us the craft was all about when they enunciated the lore, for what the anthropologists tell us is that craft and ideology are often quite distinct. If you ask the native potters how they make their pots, they may tell you one thing and do another. It is not a question of lack of sophistication. Jet airplane pilots tell their students, "This is how you fly a jet plane." They even write books about how they fly jet planes, but often that is not how they fly. They often do not know what they are doing. Now, if you want to study that sort of thing, you should bring the particular talents of the anthropologist to the study, but pending that research I shall hazard some informal observations. I suppose that if we look carefully at the ideology of folk psychology, we find it is pretty much Cartesian – dualist through and through. Perhaps there are also other problems and perplexities within the ideology as it has come down to us through the tradition. But notice that nobody in philosophy working on folk psychology wants to take seriously that part of the ideology. We are all materialists today, so the issue about the future of folk psychology is not whether or not some Cartesian view will triumph. We have apparently just decided that dualism (if it really is, as some have argued, presupposed by "common sense") is an expendable feature of the ideology. The question that concerns us now is whether there are other, less expendable features of the ideology.

Consider what can happen: Fodor, for instance, looks at the craft of folk psychology and tries to come up with a theory about why it works. His theory is that it works because it is good natural history. It is actually an account of what is going on in the head, he thinks. All the things that seem to be salient in the practice of the craft actually have their isomorphs or homomorphs in the head. So he comes up with what he calls 'intentional realism.' He notices that people say things like

Tom believes that *p*

and

Sam desires that *q* (actually, nobody ever uses this form, but bend a little!)

From this, we note that there is an attitude slot and a *p*-or-*q* slot, a propositional slot. We have two different sources of variation, two different knobs that turn in the folk craft. The reason it works, Fodor thinks, is that in the head there are things that line up with those knobs in a nice way. If the attitude knob has fourteen settings on it, there have to be fourteen different state types in the brain. If the *p*-or-*q* knob has an infinity of settings, there has to be an infinity of possible different internal states, each of them distinct and discrete. But that is just one theory about why folk psychology works, and several other chapters in this volume demonstrate that there are some good reasons for thinking it is a bad theory.

It is rather curious to say "I'm going to show you that folk psychology is false by showing you that Jerry Fodor is mistaken." Yet that is pretty much the strategy of Ramsey, Stich, and Garon's chapter, "Connectionism, Eliminativism, and the Future of Folk Psychology." It won't achieve its aim if Fodor is wrong about what is the most perspicuous ideology or explanation of the power of the folk craft. If he is wrong about that, then indeed you would expect that the vision he has of what is going on in the brain would not match what the brain people would discover; but that disparity would not mean that folk psychology as a craft was on the way out, or even that no perspicuous account of the ontology or metaphysics of folk psychology as a craft would survive well on a connectionist day of glory. I think John Heil nicely explains in his commentary how one might imagine that happening. One can also see in Ramsey, Stich, and Garon's own account how they are using folk-psychological terms simply in order to motivate their connectionist models, and explain how they are supposed to work. Why do they feel they must do that, if there is no useful relationship between folk psychology and computational neuroscience? Is it just a temporary ladder that we shall learn to discard? If the ladder works so well, why contemplate discarding it?

We should acknowledge, then, that it does not matter if the folk ideology about the craft is wrong – unless the ideology is taken too seriously! In *The Intentional Stance* (1987; p. 114), I comment at one point that Fodor's theory is a little like a curious folk theory of the common cold: A cold is a large collection of sneezes, some of which escape.

Someone might actually think that that is what a cold was, and might wonder how many more sneezes had to escape before it was over. Fodor's theory is similar: It is that there are all these sentences in the head; some of them come out, and some of them go in, but aside from a bit of translation (between "mentalese" and, say, English), basically they never get changed around all that much.

Some others, such as myself, have tried to give a different analysis of what folk psychology as a craft is, in terms of the intentional stance. I have insisted that far from being most perspicuously treated as (1) discrete, (2) semantically interpretable states (3) playing a causal role, the beliefs and desires of the folk-psychological craft are best viewed as *abstracta* -- more like centers of gravity or vectors than individualizable concrete states of a mechanism.

In Chapter 4, Ramsey, Stich, and Garon give some bottom-up reasons for being dubious about this Fodorian triad of views about the nature of psychological states, and I am going to try to give a few slightly different ones – one might call them top-down reasons – by performing a few quasi-experiments drawn from *The Intentional Stance*. (If you have read the book, you are not a naive subject.)

What follows is a joke. See if you get it. ("Newfies" are people from Newfoundland; they are the Poles of Canada – or the Irish of Canada if you are British.)

A man went to visit his friend the Newfie and found him with both ears bandaged. "What happened?" he asked, and the Newfie replied, "I was ironing my shirt, you know, and the telephone rang." "That explains one ear, but what about the other?" "Well, you know, I had to call the doctor!"

The experiment works – with you as a subject – if you get the joke. Most, but not all, people do. If we were to pause, in the fashion of Eugene Charniak, whose story-understanding artificial intelligence (AI) program (Charniak 1974) first explored this phenomenon, and ask what one has to believe in order to get the joke (and here we have a list of propositions or sentences-believed-true, individuated in the standard way), what we get is a long list of different propositions. You must have beliefs about the shape of an iron, the shape of a telephone; the fact that when people are stupid, they often cannot coordinate the left hand with the right hand doing different things; the fact that the hefts of a telephone receiver and an iron are approximately the same; the fact that when telephones ring, people generally answer them; and many more.

What makes my narrative a joke and not just a boring story is that it is radically enthymematic; it leaves out many facts and counts on your

filling them in, but you could fill them in only if you had all those beliefs. Now an absolutely daft theory about how you got the joke – and this is probably not fair to Fodor but a caricature of sorts – is this: Enter some sentences (in the ear – exactly the sentences I spoke), and their arrival provokes a mechanism that seeks all the relevant sentences – all those on our list – and soon brings them into a common workspace, where a resolution theorem-prover takes over, filling in all the gaps by logical inference. That is the sort of sententialist theory of cognitive processing that Fodor has gestured in the direction of, but nobody has produced a plausible or even workable version, so far as I know. The sketch I have just given is an absolutely harebrained theory, which nobody could assert with a straight face. Nobody believes that theory, I trust, but note that even if it and all its near kin (the other sententialist/inference engine theories) are rejected as theories about how you got the joke, our list of beliefs is not for that reason otiose, foolish, or spurious. It actually does describe cognitive conditions (very abstractly considered) that have to be met by anybody who gets the joke. We can imagine running the experiments that would prove this. Strike off one belief on that list and see what happens. That is, find some people who do not have the belief (but have all the others), and tell them the joke. They will not get it. They cannot get it, because each of the beliefs is necessary for comprehension of the story.

II

Ramsey, Stich, and Garon discuss the phenomenon of forgetting one belief out of a list of beliefs. In the connectionist net they display, the way that you forget a belief is different from the way you might forget a belief in a Collins and Quillian network, but the point I would make here is that on either account, we have counterfactual-supporting generalizations of the following form: If you don't believe (have forgotten) that *p*, then you won't get the joke.

I am prepared, in fact, to make an empirical prediction, which relies on the scientific probity of talking about this list of beliefs, even though they don't represent anything salient in the head, but are mere abstracta. The joke that I just told and you just got is on its way out. It is going to be obsolete in a generation or two. Why? Because in this age of wash-and-wear clothing, the kids that are growing up have never seen anybody ironing. Some of them don't know what an iron looks and feels like, and their numbers are growing. For that matter, telephones are changing all the time, too, so the essential belief in the similarity in

shape and heft of the telephone receiver and iron is going to v[illegible]; it is not going to be reliably in the belief pool of normal audiences. So they are not going to get the joke. One would have to explain it to them: "Well, back in the olden days, irons looked and felt like [illegible]" – and then, of course, it would not be a joke anymore. This example could be multiplied many times over, showing that the power of folk psychology or the intentional stance as a calculus of abstracta is not in the least threatened by the prospect that no version of Fodor's intentional realism is sustained by cognitive neuroscience.

Ramsey, Stich, and Garon also cite some examples to show the presumed discreteness of beliefs on the folk-psychological model. They point out that an explanation can cite a single belief, and beliefs can come and go more or less atomistically – according to folk psychology. Heil comments usefully on what is misleading about this interpretation of such phenomena, so I am going to extend his criticism via a different example. Suppose we explain Mary's suddenly running upstairs by citing her belief that she has left her purse behind on the bed, and her desire to take her purse with her when she goes out. According to the realist interpretation of folk psychology, this would be a case in which we were not speaking metaphorically, and if we happen to speak imprecisely (did she really actually believe just that she had left her purse on some horizontal surface in the bedroom?), this is always correctible in principle because there is a definite fact of the matter, say the realists, as to just which beliefs and desires a person has. Now I take it that in a case like this, what creates the illusion of discrete, separate, individuatable beliefs is the fact that we talk about them: the fact that when we go to explain our anticipations, when we move from generating our own private expectations of what people are going to do, to telling others about these expectations we have, and explaining to them why we have them, we do it with language. What comes out, of course, are propositions. Given the demands for efficiency in communication, we have to highlight the presumably central content features of the explanation or prediction that we have in mind. A distilled essence – in fact, as we have seen, typically an enthymematic portion of the distilled essence – gets expressed in words.

In order to make any sense of what we are doing in the classical way that Fodor suggests, we have to carry a lot of further baggage along. Let me quote Davidson (1975, pp. 15–16): "Without speech we cannot make the fine distinctions between thoughts that are essential to the explanations we can sometimes confidently supply. Our manner of attributing attitudes ensures that all the expressive power of language can be used

to make such distinctions." I agree with that. It remains true, however, that in the ordinary run of affairs, large families of beliefs travel together in our mental lives. At one instant, Mary believes that her purse is on the bed *and* believes that her handbag is on some horizontal surface *and* believes that the item containing her comb is supported by the article of furniture one sleeps in, and so forth. Now do all (or many) of these distinct states have to light up and team up to cause Mary to run upstairs? Or is there just one each from the belief family and desire family that are chosen to do the work? If we cling to Fodor's "classical" view of propositional attitudes, these are the only alternatives, and they are exclusive. That is not to say that there could not be overdetermination (e.g., fourteen beliefs and seven desires were ON at that time, but any pair were sufficient to cause the decision), but that there has to be a fact of the matter about exactly which of these discrete beliefs and desires existed at the time, and whether or not it did, or could, contribute to the causation of the decision.

This is related to a point that Heil makes. Folk psychology recognizes, if you like, the holism of belief attribution in everyday life, and in fact boggles at the suggestion that somebody could believe that her handbag was on the bed and not believe any of these other propositions. The idea that you could believe one of these without believing the others -- and not just the obvious logical consequences of it, but all of the pragmatic neighbors of it – is something that folk psychology does not anticipate, because it is not as staunchly realist about beliefs and desires as Fodor is.

So it seems to me that the illusion of realism arises from the fact that we do not just use folk psychology privately to anticipate – each one of us – the behavior of one another. In contrast, if chimpanzees, for instance, use folk psychology (Premack 1986), they do not talk about it. They are *individual* folk psychologists, but we are not. We are *communal* folk psychologists, who are constantly explaining to other people why we think that so and so is going to do such and such. We have to talk; and when we talk, because life is short, we have to give an edited version of what we are actually thinking; thus what comes out is a few sentences. Then, of course, it is only too easy to suppose that those sentences are not mere edited abstractions or distillations from, but are rather something like copies of or translations of the very states in the minds of the beings we are talking about.

The fact that we talk has, I claim, an obvious but interesting further effect: Since we talk, and write, we have all these sentences lying around – our own and other people's. We hear them, we remember them, we write them down, we speak them ourselves, and with regard

to any such sentence in our language that we encounter or create, we have a problem: what to do with it. We can discard it, forget it, or we can decide to put it in the pile labeled "true" or the pile labeled "false." And this, I claim, creates a rather different sort of specialized state, what in *Brainstorms* I called opinions. These are not just beliefs; these are linguistically infected states – only language users have them. Opinions are essentially bets on the truth of sentences in a language that we understand. My empirical hunch is that a proper cognitive psychology is going to have to make a sharp distinction between beliefs and opinions, that the psychology of opinions is really going to be rather different from the psychology of beliefs, and that the sorts of architecture that will do well by, say, nonlinguistic perceptual beliefs (you might say "animal" beliefs) is going to have to be supplemented rather substantially in order to handle opinions. And I think it is confusion on this score – more specifically, the failure to distinguish between believing that a certain sentence of one's natural language is true, and having the sort of belief that that sentence might be used to express – that has given Fodor's intentional realism the run it has had.

It occurs to me that another feature of this line of thought that Churchland and Stich might like is that if I am right about the distinction between beliefs and opinions, then the following dizzying prospect opens up: Scientists (connectionist heroes of the near future) might "come to the *opinion*" that there are no such things as beliefs, without thereby having to *believe* there was no such thing! If connectionists are right, after all, they are just connectionist systems that on occasion make bets on the truth of various sentences of their natural language. All of their science goes on – at least the public communication and confirmation part of it – at the level of opinions. Although they do not have beliefs, they do have opinions, since they are still using sentences and hence committing themselves to the truth of some of them. But they wouldn't have to say they *believe* that a particular sentence is true; they would just *connect* that it was! Because of the settings of their connectionist networks, they would put that sentence in the "true" pile, but putting that sentence in the "true" pile (even in the "true" pile of sentences you keep stored in your head) is distinct from believing – on the folk-psychological model. (Those of us who are not Fodorians about belief can go on talking about what these connectionists *believe*, but the Fodorians and Stichians among them can consistently be of the *opinion* that they never in fact believe anything!)

I want to say more about connectionism, for I want to throw a few more buckets of cold water on the euphoria expressed if not induced by

the connectionist chapters (Churchland, Chapter 2, this volume; Ramsey, Stich, and Garon, Chapter 4, this volume). First, however, I want to make some remarks in favor of connectionism that they do not quite make. Ramsey, Stich, and Garon claim that connectionism is not merely implementation at a lower level of a traditional, hierarchical model, and I want to say something more in favor of that. Here is why I think connectionism is exciting.

Suppose you have what Haugeland (1985) would call a GOFAI (Good, Old-Fashioned Artificial Intelligence) nonconnectionist AI theory: It postulates a certain level at which there are symbolic structures in something like a language of thought, and it has some mechanism for computing over these. Then, indeed, it makes little difference how you implement that. It makes no difference whether you use a Vax or a Cray, a compiled or interpreted language. It makes no difference how you determine the implementation, because all of the transitions are already explicitly stated at the higher level. That is to say, in technical terms, you have a flow graph and not merely a flow chart, which means that *all* the transition regularities are stipulated at that level, leaving nothing further to design, and it is simply a matter of engineering to make sure that the transition regularities are maintained. It makes no sense to look at different implementations, for the same reason that it makes no sense to look at two different copies of the same newspaper. You might get some minor differences of implementation speed or something like that, but that is not apt to be interesting, whereas the relationship between the symbolic or cognitive level and the implementation level in connectionist networks is not that way. It really makes sense to look at different implementations of the cognitive-level sketch because you are counting on features of those implementations to fix details of the transitions that actually are not fixed at the cognitive level. You have not specified an algorithm or flow graph at that level. Another way of looking at this is that in contrast to a classical system, where the last thing you want to have is noise in your implementation (i.e., you want to protect the system from noise), in a connectionist implementation you plan on exploiting noise. You want the noise to be there because it is actually going to be magnified or amplified in ways that are going to effect the actual transitions described at the cognitive level.

This becomes clear if you consider the hidden units in a connectionist network, such as those in the diagrams in Chapter 4. As Ramsey, Stich, and Garon note there, if you subject those hidden units to careful statistical analysis (it is made easier if you view the results in one of Geoffrey Hinton's lovely diagrams showing which nodes are active

under which circumstances), you can discover that a certain node is always ON whenever the subject is (let us say) dogs, and never (or very weakly) ON when the subject is cats, whereas another node is ON for cats and not ON for dogs. Other nodes, however, seem to have no interpretation at all. They have no semantics; they are just there. As far as semantics is concerned, they are just noise; sometimes they are strongly active and at other times weak, but these times don't seem to match up with any category of interest. As many skeptics about connectionism have urged, the former sorts of nodes are plausibly labeled the DOG node and the CAT node and so forth, and so it is tempting to say that we have symbols after all. Connectionism turns out to be just a disguised version of good old-fashioned, symbol-manipulating AI! Plausible as this is (and there must be *some* truth to the idea that certain nodes should be viewed as semantic specialists), there is another fact about such networks that undercuts the skeptics' claim in a most interesting way. The best reason for *not* calling the dog-active node the dog symbol is that you can "kill" or disable that node, and the system will go right on discriminating dogs, remembering about dogs, and so forth, with at most a slight degradation in performance. It turns out, in other words, that all those other "noisy" nodes were carrying some of the load. What is more, if you keep the "symbol" nodes alive and kill the other, merely noisy nodes, the system *doesn't* work.

The point about this that seems to me most important is that at the computational level in a connectionist system, no distinction is made between symbols and nonsymbols. All are treated exactly alike at that level. The computational mechanism doesn't have to know which ones are the symbols. They are all the same. Some of them *we* (at a higher level) can see take on a role rather like symbols, but this is not a feature that makes a difference at the computational level. That is a very nice property. It is a property that is entirely contrary to the spirit of GOFAI, where the distinction between symbol and nonsymbol makes all the computational difference in the world.

Having offered my praise, let me turn to what worries me about connectionism. Both connectionist chapters exhibit connectionist networks that have input nodes, output nodes, and hidden units, but all their discussion is about the hidden units. We should pause to worry about the fact that some of the input units (for instance) look much too Fodorian. It looks, indeed, as if there is a language of thought being used to input *Dogs have fur* across the bottom of the system, for instance. It looks as if the inputs are organized into something altogether too much like Fodorian propositions. Could it be that the only reason we are

not seeing the language of thought is that we are not looking at the much larger cognitive systems of which these bits of memory are just small subsystems?

This worry is analogous to a concern one can have about traditional AI systems. For instance, Hector Levesque (1984) has described a knowledge representation system (in AI) with some lovely properties, but one of its *un*lovely properties is that there is only one way of putting something into the knowledge base, and there is only one thing the knowledge base can do. Everything goes in by an operation called TELL, followed by a statement in the predicate calculus; the only thing the system can do is permit itself to be subjected to an ASK operation. I submit that any model of knowledge that one can update or enrich only by writing a proposition using the TELL function and that one can use only by extracting from it a proposition via the ASK function is a hopelessly Fodorian sententialist model of a robust knowledge system.

But for all that the connectionist chapters show us, that is what we have in their connectionist models too. We have a memory for which a TELL and an ASK are defined. No other way of tweaking it, or utilizing it, or updating it has yet been defined. This is a serious charge, which I should try to defend with a more specific example. Here, finally, is one more little experiment concerning the structure of human memory. The claim I want to substantiate by it is that what the connectionists have offered us is not an architecture for memory but at best an architecture for perhaps a little subcomponent of memory. When we start making the memory more realistic, we are going to have to add some architectural details that will require some quite different principles.

Here are some questions – personal questions about your own memory – which you should attempt to answer as quickly as you can:

Have you ever danced with a movie star?

Have you ever been introduced to a white-haired lady whose first name begins with the letter *V*?

Have you ever driven for more than seven miles behind a blue Chevrolet?

Most people have a swift yes or no answer to the first question and draw a blank on the others. Imagine how different their responses would be to the following:

Have you ever been introduced to a green-haired lady whose first name begins with the letter *V*?

Have you ever driven for more than seven miles behind a pink Rolls Royce?

First of all, according to anybody's theory of memory it is [illegible]se that you have stored as Fodorian sentences "I have never danced with a movie star," "I have never driven more than seven miles behind a pink Rolls Royce," and so forth, because that would lead to combinatorial explosion. Think of all the things you have never done, and know you have never done.

Any remotely sane theory of how you answer these questions has to work the following way: When you hear the question, it stimulates your memory, and either it succeeds in tweaking a recollection of an event meeting the condition or it does not. In the case of the first proposition, if no recollection is made, you draw the metaconclusion that *had you ever done it, you would now be recalling it*, and since you are not now recalling it, the chances are that you never have danced with a movie star. The parallel metaconclusion, however, is simply not plausible in the third case, because there is no reason to suppose that had you ever driven seven miles behind a blue Chevy, you would now be recalling it. In order to make sense of this simple, robust feature of human memory, you have to suppose that human memory is organized in such a fashion that you can unconsciously assess the likelihood that the failure of your memory to produce a recollection for you is a sign – it can be treated on this occasion as a premise or datum – from which you unconsciously "infer" the conclusion "I have never done that." That shows a complexity far beyond ASK and TELL that we can establish quite clearly as a feature of human memory. So a good cognitive psychology will have to model that. How can I build a model of human memory that has that rather nifty, easily demonstrated property? Nobody in *non*connectionist cognitive psychology has a good model of that, so far as I know, but then neither do the connectionists.

And until the connectionists can show that their marvelous new fabrics can be fashioned into larger objects exhibiting some of these molar properties of human psychology, we should temper our enthusiasm.

References

Charniak, E. (1974). "Toward a Model of Children's Story Comprehension." Unpublished doctoral dissertation, MIT, and MIT AI Lab Report 266.

Davidson, D. (1975). "Thought and Talk." In *Mind and Language: Wolfson College Lectures, 1974*. Oxford: Clarendon Press, pp. 7–23.

Dennett, D. (1978). *Brainstorms*. Cambridge, Mass.: MIT Press.

Dennett, D. (1987). *The Intentional Stance*. Cambridge, Mass.: MIT Press.

Haugeland, J. (1985). *Artificial Intelligence: The Very Idea*. Cambridge, Mass.: MIT Press.

Hayes, P. (1978). "Naive Physics I: The Ontology of Liquids." Working Paper 35, Institut pour les Etudes Semantiques et Cognitives, Univ. de Genève.

Hayes, P. (1979). "The Naive Physics Manifesto." In D. Michie, ed., *Expert Systems in the Microelectronic Age*. Edinburgh: Edinburgh University Press.

Levesque, H. (1984). "Foundations of a Functional Approach to Knowledge Representation." *Artificial Intelligence*, 23:155–212.

Premack, D. (1986). *Gavagai! Or the Future History of the Animal Language Controversy*. Cambridge, Mass.: MIT Press.