METHODOLOGICAL REFLECTIONS ON BELIEF

Robert Cummins

In R. Bogdan, ed., Mind and Common Sense, Cambridge University Press, 1991, pp. 53-70.

I. INTENTIONAL REALISM.

Let's suppose there really are such things as propositional attitudes. And let's understand this claim as Putnam (1975) would have us understand claims about natural kinds so that it could turn out that most of what we believe about the propositional attitudes might be mistaken in one way or another. Or, if that is going too far, let's at least understand the claim in such a way that any particular thing we believe about the attitudes might turn out to be false. This, it seems to me, is the way we have to understand the claim that there are such things as propositional attitudes if we to take the claim to be an empirical one on all fours with such claims as that there are elephants, genes, or crystals. Otherwise, we can legislate against surprising empirical discoveries about the attitudes by claiming that the discoverers have merely changed the subject. "Computationalists can't be talking about belief," we might say, "because beliefs are individuated, at least in part, by social factors, and computational states are not. At best, they are talking about 'shmaliefs', not beliefs."

The trouble with this sort of move is that the philosopher who makes it runs the risk that he or she is talking about nothing at all. If philosophers insist, e.g., that beliefs are socially individuated, while good empirical theory insists that computational states are what there are, then it is open to the computationalist to say that the empirical evidence suggests that there are no beliefs as the philosopher chooses to define them.1 Scientists who have no stomach for a fight over words will either cease talking to philosophers, or they will simply invent a new word. The philosophers will quite rightly be perceived as having won a petty squabble over semantic territory at the price of cutting themselves off from serious empirical research.

II. METHODOLOGY: INTENTIONAL REALISM AND THE PHILOSOPHY OF SCIENCE.

Lot's of philosophers, I suspect, resist realism about the attitudes--in practice, if not explicitly--because it threatens to lead to unemployment. How, after all, can a philosopher who toes the realist line hope to discover anything about the propositional attitudes? Doesn't being a realist about the propositional attitudes amount to conceding them to the scientists?

The same thing might be said about space and time. If we are realists about these matters, shouldn't philosophers concede that the

nature of space, time and space-time are empirical matters and stop poaching on the physicist's preserves? Yes, of course. But accepting that answer hasn't put philosophers out of the space and time business. Rather, it has led philosophers to treat the issues as issues in the philosophy of physics. Instead of asking what the nature of space and time are--a question philosophical methodology cannot properly address--the philosopher of physics asks, e.g., how space must be conceived if some well-articulated physical theory or theoretical framework is to deliver the explanatory goods it advertises. How must we conceive of space if General Relativity is to be true and explanatory? Should we be substantivalists? Is space a thing? Or is it a property of things? This kind of inquiry is especially pressing when the target is an explanatory primitive of the theory or theories that harbor it. Hence, a good deal of the philosophy of science is concerned with explicating the "foundations" of this or that theory or framework or research paradigm.

All this applies straight-forwardly to the philosophy of psychology. Philosophers may, without concern for a misfit between their methodology and their conclusions, identify some theory (or theoretical framework or research paradigm) that invokes propositional attitudes and then ask how we should understand them if the invoking theory is to be true and explanatory. Thus, if our target is computationalist theories, and we determine that computational states are not individuated by social factors, then we may conclude that beliefs, at least as invoked by computationalists, had better be a-social too. We can then go on to attempt to construct an appropriate a-social conception of belief.2

III. METHODOLOGY AND BELIEF ATTRIBUTION.

What I've been describing is not what philosophers have been doing. What they have been doing instead is trying to turn semantics into psychology. Here is a familiar example.3

III.1 Burge: belief and linguistic affiliation.

'Brisket' applies to cuts of breast meat generally. Tyler, however, believes that it applies only to breast of beef. Tyler* is just like Tyler except that he lives in a language community in which 'brisket' applies only to breast of beef. Tyler and Tyler*, we may suppose, are molecule for molecule identical, for, though their language communities differ slightly, Tyler has, in point of fact, never encountered the sort of situation that could distinguish him from Tyler*, e.g., reading a dictionary entry for 'brisket'. In spite of their mereological identity, however, it seems we are prepared to attribute beliefs about brisket to Tyler but not to Tyler*. For when Tyler* says, e.g., "Brisket is better than tenderloin," he is making a statement that has a different truth-condition than the statement Tyler makes with the same words. Tyler's statement is about breast of beast; Tyler*, using the same words, makes a statement about breast of beef. But it seems obvious that what goes for the statements must go for the beliefs they express: the belief Tyler expresses when he says, "Brisket is

better than tenderloin," is about breast of beast; the one Tyler* expresses is about breast of beef. It is concluded from all this that beliefs are not entirely "in the head." They are not, in fact, psychological states as psychologists normally conceive them, since they are individuated in part by such extra-psychological factors as the rules of one's language,4 even in cases in which the relevant facts about the language have had no causal impact whatever on the believer.

How exactly is this sort of thought experiment supposed to lead to conclusions about the nature of belief? Can this sort of "intuition pump" (Dennett, 1980) really show that beliefs are not psychological states,4 that psychologists are mistaken to try to use 'believes' to characterize mental processes in a way that abstracts away from such things as the subject's linguistic affiliation?

I'm going to need a name for the methodology implicit in this sort of thought experiment. I'll call it the CRHS methodology for considered response to a hypothetical situation. It is difficult to see how CRHS can deliver the goods about belief. What we've got to go on is just whether we would say that Tyler but not Tyler* has beliefs about brisket. But how could this tell us anything about their beliefs? Isn't this like inferring that objects dropped from an airplane have a straight-line trajectory (relative to the earth) because that's the trajectory people tend to attribute in such cases?5 Or inferring that the rate of free fall is a function of weight from the fact that people's attributions of rate are sensitive to the supposed weight of the falling object? Why assume that people are any better authorities about belief than about the trajectories or rate of free fall of falling objects? Perhaps people can, within limits, tell you what they believe with fair reliability. But even if people were incorrigible about what beliefs they had, it wouldn't follow that they had any special knowledge about the nature of belief.

How could it even seem that investigating the conditions of belief attribution could yield conclusions about belief? The answer seems simple enough: If we know, for example, that

(A) Tyler and Tyler* are psychologically equivalent,
and that

(B) Tyler and Tyler* have different beliefs,
then it follows that,

(C) Individuation of belief is a function of extra-psychological factors.4

The trouble, of course, is B: what entitles us to the claim that Tyler and Tyler* have different beliefs? CRHS serves up B for us: reflecting on the story, we attribute different beliefs to Tyler and Tyler*. But why should we take our belief attributions seriously in a tough case like this, uninformed, as they are, by good empirical theory?

There are two reasons for caution. To see what they are, we need to look more closely at the underlying situation. Here's the plot outline:

(i) It is stipulated that Tyler and Tyler* are psychologically equivalent.

(ii) We are moved, on reflection, to attribute different beliefs to Tyler and Tyler*.

(iii) It follows that our belief attributions are sensitive to extra-psychological factors--linguistic affiliation in this case.

(iv) We conclude that the truth-conditions of belief attributions contain extra-psychological factors, and hence that beliefs are individuated (in part) by extra-psychological factors.

But (iv) doesn't follow from (iii). Extra-psychological considerations may be legitimate evidence for belief attributions, yet not be implicated in their truth-conditions.5 In this case, we are moved by the fact that what Tyler states when he says, "Brisket is better than tenderloin," is different than what Tyler* states in uttering the same words in the same circumstances. People generally believe what they sincerely state to be the case. What people sincerely state therefore provides a good bet concerning what they believe. But what someone states in uttering a sentence is a function of the meanings of the words in the sentence, and that in turn is a function of the language they are speaking. So here we have an uncontroversial case of an extra-psychological factor--word meaning in the language--counting as legitimate evidence for a belief attribution.

But it is only evidence. From the fact that attributions of belief are properly sensitive to extra-psychological factors, it doesn't follow that extra-psychological factors figure in the truth-conditions of belief attributions. All we are entitled to conclude is that extra-psychological factors figure legitimately as evidence for belief attributions. In the particular case under discussion, it is easy to see why we shouldn't move to a conclusion about truth-conditions. Why, after all, should we assume that Tyler means exactly what his words are conventionally suited to express?6 Why, that is, should we assume that what Tyler says is an exact reflection of what he believes? Tyler thinks his words express what he believes, but, of course, he is wrong about this, for it is part of the story that he doesn't know what 'brisket' means in the langauge he is speaking. Corrected about this matter, it seems likely that Tyler would change his story: "Oh, I see. Well, then, it's breast of beef I like, not brisket! I really haven't given any thought to what you call brisket. Quite a mixed bag, after all." Or he might simply pull a Humpty-Dumpty: "That's not what I mean by 'brisket'. What I mean is breast of beef."

The point of the foregoing is not so much to argue that Tyler and

Tyler* don't have different beliefs, but rather to emphasize that we cannot move from premises about our considered belief attributions to conclusions about the truth-conditions of belief attributions. There is, of course, a kind of consistency argument available: If you agree that Tyler and Tyler* have different beliefs, then you must agree that beliefs are individuated by extra-psychological factors. But this won't carry much weight with thoughtful realists, for they will be rightly suspicious of their "intuitions" in cases like this. They will remember that plausible belief attributions may yet be false. They will remember that even justified belief attributions may yet be false.

More radically, they will remember what intelligent, well-informed people who happen to be innocent of mechanics say about things dropped from airplanes. People generally attribute incorrect trajectories to hypothetical objects dropped from hypothetical airplanes. They do this, presumably, because they have false beliefs about the trajectories of objects dropped from airplanes. Subjects with a knowledge of elementary physics don't make this sort of mistake. Similarly, people with false beliefs about belief can be expected to make mistaken attributions of belief in real and hypothetical situations. "Perhaps," the realist will think, "it is best to let well-motivated empirical theory inform our attributions of belief as it ought to inform our attributions of trajectories." Well-motivated empirical theory, of course, is just what CRHS cannot provide. If belief is to be a serious explanatory construct, then we had better put the serious explanations in the driver's seat and let our intuitions go along for the ride. If we are wrong about belief, we will make mistaken attributions. In the end, the only way to know about belief is to study belief. Studying belief attributions will, at best, tell you only what people believe about belief.7

It might seem that the problem with the Burgeian argument is not with the methodology--not with CRHS--but with the mistaken assumption that sincerity is enough to guarantee a match between what someone states and the underlying belief. But this is a problem with the methodology: our considered responses to hypothetical situations will depend on what we consider. It turns out that we should have considered the fact that sincerity isn't enough to match what is stated to what is believed. What else should we consider? This is just the sort of question that science is supposed to answer better than common sense.

The Burgeian argument is instructive in part because we can see that we have been misled. Lacking a good empirical theory of belief, we normally have no really satisfactory way of knowing whether CRHS is yielding questionable attributions. But in this case, we are told that Tyler doesn't know what 'brisket' means in his language. Since we know that what one states is a function of the meanings of the words one uses, we can be pretty sure that Tyler doesn't succeed in stating what he intends to state, so we can be pretty sure that what he does succeed in stating isn't likely to be a perfect reflection of what he believes. Having come

this far, it is tempting to say that, on further reflection, we can see that what Tyler believes is not that brisket is better than tenderloin, but that breast of beef is better than tenderloin. There's no harm in yielding to this temptation, provided that we don't suppose that further reflection has got round the need for empirical theory after all. CRHS isn't going to tell us how to individuate belief, unless the further reflection in question is adequately informed reflection. Only a good empirical theory of belief could seriously justify confidence that our reflections are adequately informed.

III.2 Putnam: belief and environmental affiliation.

It is tempting, as I've just pointed out, to suppose that the difficulty we've uncovered is special to the sort of argument given by Tyler Burge for the relevance of linguistic affiliation to belief. The Burgeian argument simply assumes that what is stated is what is believed:

B-1: If S sincerely states that p, then S believes that p.

But B-1 is false. In the case lately imagined, is seems plausible to suppose that what Tyler states--viz., that breast of beast is better than tenderloin--is not what he intends to state--viz., that breast of beef is better than tenderloin. What he actually states is, of course, a function of the meanings of his words. But he is, by hypothesis, mistaken about the meaning of 'brisket' in his language. Hence he chooses the wrong words to express what he intends to state. The difficulty, one might suppose, is not with the methodology; the problem was just that it was an inadequately informed application of the methodology. Even if we accept the view that adequately informed reflection is reflection informed by good empirical theory, perhaps there is enough good empirical theory in hand to allow CRHS to yield the targeted result, namely true attributions of distinct beliefs to psychologically equivalent subjects. Consider Putnam's Twin Earth case (Putnam, 1975). Twin earth is an exact duplicate of earth except that where we have H2O they have XYZ. When Hilary says, "Water is better than Pepsi," he makes a statement about H2O. When Twhilary (Twin-Hilary) says, "Water is better than Pepsi," he makes a statement about XYZ. Once again, it is concluded that Hilary and Twhilary express different beliefs. But since, by hypothesis, the twins are molecule-by-molecule duplicates, it follows that beliefs are individuated by extra-psychological factors.

In this case, it seems we cannot object that one of the characters in the story fails to state what he intends to state on the grounds that one of them is mistaken about the meanings of the words he uses. It seems that neither character is mistaken about the meaning of 'water' in the language he uses. Both speak sincerely. It seems to follow that what each states is what each intends to state. Since what one sincerely intends to state is what one believes, both state what they believe. Since the statements are

different, so are the beliefs. QED.

Should we concede that the CRHS can deliver the conclusion that beliefs are individuated by extra-psychological factors? How is this possible? Haven't we just reinstated arm-chair science? A realist about belief will surely want to give this argument very close scrutiny.

The argument makes use of a crucial assumption:

(P-1) What the twins state is what they intend to state sincerely.

P-1 is bolstered (though of course not entailed) by P-2:

(P-2) Neither twin is mistaken about the meaning of 'water'.

P-2 looks harmless enough. In the Burge case, it is an explicit assumption that Tyler is mistaken about the meaning of 'brisket', but there is no comparable assumption operative in the Twin earth case. Still, P-2 is not as harmless as it appears.

Let's begin with the fact that the twins are, in a certain sense, ignorant of the reference of their words: Twhilary doesn't know that the referent of 'water' in Twinglish is XYZ. How, one might wonder, can Twhilary sincerely intend to state something about XYZ given that he doesn't know that what he calls water is XYZ? Aren't we simply assuming about intention what we are trying to prove about belief?

It seems that P-2 is wrong if it is understood to imply that the twins are fully informed about the reference of 'water' in their respective languages: While it is surely uncontroversial that, in some sense, Twhilary knows what 'water' refers to in Twinglish, it doesn't follow that he is fully informed about the nature of the stuff he refers to with his uses of 'water'. There is, therefore, a certain sense in which he doesn't know what he is talking about, viz., he doesn't know what proposition he is expressing when he says, "Water is better than Pepsi." It would seem to follow that he cannot sincerely intend to express that proposition in uttering those words. Hence, P-1 cannot be invoked to yield a conclusion about what he believes. Once again, it appears that the argument depends on a dubious psychological assumption linking what one believes to the reference of what one says. Here's how the trick is done:

(1) The meanings of one's words are determined by extra-psychological factors,

hence,

(2) The truth conditions of one's statements are determined by extra-psychological factors.

But,

(3) One's sincere statements express what one believes,

Hence,

(4) The truth-conditions of one's sincere statements are the same as the truth-conditions of one's beliefs.

I have no quarrel with (1) and (2), but (3) is false, as we've seen, and so is (4). Something close to (3) is plausible, viz.,

(3') One believes what one intends to state sincerely.

But (3') merely shifts the burden from belief to intention. It may be true that the truth-conditions of what one intends to state sincerely are the same as the truth-conditions of what one believes. But this won't help you get a conclusion about the truth-conditions of beliefs from premises about the truth-conditions of statements, unless you already have a link between the truth-conditions of intentions and the truth-conditions of statements. Shifting the focus to intention is no help because there is no more reason to say that Twhilary intends to state that XYZ is wet than to say that he believes that XYZ is wet. On the contrary: the shift to intention makes the trick a little easier to spot.

That's not the end of the matter, of course. The Twin Earth case is different from Burgeian cases in that we are still left wondering what Twhilary does believe, whereas we are pretty sure we know what Tyler believes. We will have to come back to this.

IV. SEMANTICS AND PSYCHOLOGY.

IV.1. The psychological reality of semantic values.

There is, of course, an intellectual tradition according to which you can discover the truth-conditions for belief attributions without knowing anything about belief. Formal semantics in the Tarskian tradition espoused by Davidson (1967) is supposed to yield truth-conditions for whatever it is that is truth-valuable in natural languages. But this kind of traffic in truth-conditions isn't supposed to be the sort of thing that could supply fuel for the anti-individualist fire. For that you need to be able to discover such things as that the truth-value of a belief attribution can depend on an extra-psychological factor. If Tarskian semantics is serving up conclusions like that, then it is certainly doing more than telling us about the language, and is therefore overstepping the boundaries set by its original charter. No one in that tradition thought of semantics as the sort of thing that could overlap the psychology of belief.

Yet this crossing of boundaries--doing philosophy of mind, and even psychology, by doing philosophy of language--has become a small industry. What gives the trespassers such confidence?

I think a lot of the trespassing (or what I'm contentiously calling trespassing) began with the hypothesis that semantic values are psychologically real.

(Psychological reality of semantic values--prsv) To understand (or know the meaning of) an expression requires knowing its semantic value.

Examples of prsv are (a) the claim that to understand a statement requires knowing its truth-condition; (b) the claim that to understand a proper name requires knowing its sense; (c) the claim that to understand a proper name requires knowing its reference; (d) the claim that to understand a general term requires knowing what property it expresses; (e) the claim that to understand a statement requires knowing what proposition it expresses. Two clarifying remarks are essential.

First, understanding is to be construed psychologically in the prsv: For example, (b) is to be construed as the idea that one must represent the sense of a name, and represent it as the sense of that name, in order to be a party to successful communicative episodes that employ that name. A corollary of prsv is the assumption that learning a language involves learning the semantic values of its expressions.8

Second, semantic values are, in the first instance, to be understood in the model theoretic sort of way that is central to the Tarskian tradition: they are whatever one needs to assign to things in order to systematically capture entailment relations that hold among things that are truth-valuable. The paradigm case is the assignment of satisfaction conditions to the primitive terms of a first order language in such a way as not only to enable a truth-definition, but to account for the entailment relations among statements. It is the fact that semantics aims for, and is constrained by, a proprietary goal independent of psychology that gives prsv its bite: prsv is the hypothesis that the stuff that tracks entailments is an essential part of the stuff a system needs to know in order to use and understand a language. This is how prsv links semantics with psychology: since the stuff that tracks entailments is the stuff that (in part) drives use and understanding, when we do the semantics we have done a good deal of the psychology as well.

Meaning and semantic values. I suspect that some have been seduced into accepting prsv by the following line of thought: To understand an expression is just to know what it means. But to know what an expression means is to know it's meaning, i.e., to know it's semantic value.

I suppose there is an innocent sense in which people who know English know the meaning of 'elevator' and 'rescue'. Perhaps they can even tell you what these words mean. I've asked a few people about 'elevator' just now. The best answer I got was this: "An elevator is like a closet. You go in, press one of several buttons with numbers on them, and the thing travels up a vertical shaft in the building, stopping at the floor

corresponding to the number on the button you pushed. Faster and easier than stairs." This person certainly knows what 'elevator' means. But does this person know the semantic value of 'elevator'? Does she know which whatnot 'elevator' must be associated with in order to track the entailments of expressions in which 'elevator' occurs as a constituent? Maybe. But the fact that she knows, in this ordinary way, what 'elevator' means does not, on the face of it anyway, show that she does (or must) know, even tacitly, what semanticists want to know about 'elevator'.

Maybe the psychology works like this: when you encounter an 'elevator' you generate a pointer to a frame that allows you to access, more or less reliably, everything you know about elevators. On this rather plausible view, your ability to use and understand 'elevator's is closely related to your ability to use and understand elevators: it rests on what you know about elevators. You have the concept of an elevator, on this view, when you have a frame whose semantic value is (let's say) the property of being an elevator but whose psychological content is just whatever you know about elevators.9 The crucial point is that there is a strong distinction between the psychological content of a concept and its semantic content. The psychological content is the *knowledge10 in (or accessible via) the relevant data structure; the semantic content (what I've been calling the semantic value) is the sort of thing that enters into truth-conditions. I'll call this view of concepts the encyclopedia view of concepts to emphasize the idea that the psychological content of a concept is like an encyclopedia entry rather than a dictionary entry, i.e., an entry that specifies a meaning.

I think there is a good deal of empirical evidence for the encyclopedia theory of concepts. However that may be, the point I want to make is that it is incompatible with prsv. Since the encyclopedia theory of concepts is plainly an empirical theory, so is prsv. It doesn't follow, of course, that philosophers are not allowed to assume prsv, but it does mean that they are required to acknowledge the empirical loan they are taking out when they do assume it. No argument that assumes prsv, for example, can be used to criticize the encyclopedia theory of concepts.

Belief and the prsv. Even if we know something about the semantic values of such things as beliefs or mental representations, nothing will follow about the psychological content of those representations or states unless we assume something like prsv. If you have the semantics straight, you will be able to track the semantic relations among mental states. But it doesn't follow that you will be able to track their psychological relations. We should therefore be cautions of moves from premises about the truth-conditions of belief attributions to substantive conclusions about belief. If the prsv isn't true, then semantic values--e.g., truth-conditions--aren't in the head. The semanticist will be able, in one sense, to draw conclusions about mental states, viz., conclusions about their semantic values. But nothing will follow about the psychological contents of those states. Hence, even if, contrary to fact,

philosophy were in a position to establish that Tyler and Tyler* have beliefs with different semantic contents, it wouldn't follow that they have beliefs with different psychological contents. Hence it wouldn't follow that they have, in any sense of interest to psychology, different beliefs. Empirical science must be left free to individuate its states in whatever way conduces to good theory. Philosophers can insist that states individuated by psychological content rather than semantic content are not beliefs, but this won't be worth taking seriously unless and until they catch psychologists using what they call beliefs in a way that assumes that they are semantically individuated.11

IV.2. The old and the new semantics.

I have already emphasized that if philosophers are to be able to move from semantical theses to psychological theses, there has to be a sense in which semantics is independent from psychology. That is, it had better be the case that semantics has its own proprietary goals and constraints. Only thus will it be possible to establish some semantical thesis and then expose its implications for psychology. Semantics, I have supposed, is the theory of entailment. Semantic values are whatever one needs to assign to statements (or whatever is truth-valuable) to systematically track entailment relations.12 So the idea behind what I have contentiously called philosophical trespassing on psychological turf can be put this way:

(Easement) Get the semantics of belief attribution right--i.e., find out what it takes to track their entailments--and this will put constraints on belief, and hence on psychology.

From this perspective, it is possible to see how the twin earth argument depends on a rather recent conception of semantics. Twenty years ago, it would have been natural to respond to the argument along the following lines.

The old semantics. If the point is to track entailments, then the semanticist should not think of the reference of 'water' in English as H2O, since "This is water," doesn't entail, "This is made up of hydrogen and oxygen." Semantics, like any other science, must choose its theoretical specifications carefully. In this case, specifying the reference of 'water' as H2O is the wrong choice for semantics, though it is the right choice for chemistry. Of course, 'water' does refer to H2O (in English), but that is the wrong way to specify its reference if you are interested in tracking entailments, though it is the right way if you are interested in tracking chemical interactions. Thus, although Hilary and Twhilary are referring to different things when they use the word 'water', this is not a difference that semantics will notice. And what goes for reference goes for truth-conditions: Good semantic theory will not use the term 'H2O' in specifying the truth-conditions for statements containing 'water'. From the fact, then, that two beliefs or statements are about different things, it doesn't follow that good semantics will distinguish them. One shouldn't assume

that a proper scientific specification of what a belief or statement is about is a proper semantic specification. Semantics, in short, is an autonomous discipline.

The new semantics. Now days, most of us would object to the claim that Hilary's statement that water is better than Pepsi, doesn't entail that H2O is better than Pepsi. Of course, you cannot formally derive "H2O is better than Pepsi," from "Water is better than Pepsi," without adding something like "Water is H2O," but entailment is supposed to be a semantic notion, not a formal one. What we need to know, then, is whether "H2O is better than Pepsi," is true in every possible world in which "Water is better than Pepsi," is true. Now, if "Water is H2O," is true in every possible world if it is true at all, as Kripke and Putnam have made plausible, then the entailment is unproblematic.

Here, of course, the sentences in quotation marks are assumed to be bits of English. If they are assumed to be bits of Twinglish, we get quite different results. In Twinglish, "Water is XYZ," expresses something true in every possible world, so in Twinglish, "Water is better than Pepsi," entails "XYZ is better than Pepsi."

I have a lot of sympathy with this line of thought. But notice that if we understand entailment in this way, semantics is no longer an autonomous discipline: our knowledge of what entails what must wait on such things as chemistry. Since we have enough chemistry in hand to know that water is H2O, we know that beliefs about water are beliefs about H2O. Twhilary certainly has no beliefs about H2O, so his beliefs will be semantically distinct from Hilary's--Twhilary believes that XYZ is better than Pepsi--even though he is Hilary's computational (or neurophysiological) twin. From this perspective, the prsv seems preposterous. People don't need to know the chemistry of water to use and understand 'water'. And yet I think there are many who accept some form of the prsv, at least tacitly, while embracing the new semantics. The problem is that we have updated our conception of semantics without noticing that it requires us to abandon prsv, a view that made sense only against the background of a conception of semantics as an autonomous discipline. If you accept the conception of semantics that underlies the Twin Earth argument, you must abandon the prsv. Hence Putnam's claim that meanings (semantic values, not psychological contents) aren't in the head.

The individuation of belief. If we accept the new semantics and reject the prsv, what motivation could we have to individuate beliefs semantically? If we are realists about belief, we will want to individuate beliefs in whatever way conduces best to good psychological theory. Adherents of the new semantics should scoff at the idea that indexing beliefs to propositions will track their psychological relations. They will scoff precisely because it is so obvious that tracking entailments requires factors irrelevant to psychology. "Intentional realism" is therefore a potentially misleading name for contemporary realism about

belief, for it suggests that we should take seriously the idea that psychology should individuate beliefs semantically. Of course, 'intentional individuation' doesn't have to mean semantic individuation, but it is an easy enough identification to make "intentional realism" a dangerous name for realism about belief. For the doctrine we are now considering is the doctrine that propositional attitudes should not be individuated in psychology by propositions. From the fact that p and q are distinct propositions, we cannot conclude that the belief that p is distinct from the belief that q.14

How should we individuate concepts (psychologically construed) and beliefs? You shouldn't ask me: I'm a philosopher. You wouldn't ask a philosopher what water is. Don't ask a philosopher what belief is either.

V. Folk Psychology.

This paper has been mainly a complaint against arm-chair psychology disguised as the philosophy of mind or as the philosophy of language. The recent interest in so-called "folk psychology" is, I think, largely due to the fact that it provides philosophers an excuse to do psychology without having to bother about sound empirical method. If by "folk psychology" you mean the psychology of the folk--a theory on all fours with folk physics (and as likely to be true)--then I have no objection to studying it--that is, investigating it in the way psychologists study folk physics: empirically. There can be no serious talk of vindicating the psychology of the folk until we have some empirically justified account of what it is.

But, of course, this is not what philosophers typically mean by folk psychology. What they mean is some sort of psychological theory--the details are left unspecified because they are irrelevant--that makes central explanatory appeals to belief and desire. Folk psychology in this sense will be "vindicated" if it turns out that empirical psychology has some serious use for belief and desire. Now this looks to be a question for empirical psychology itself, or for the philosopher of science who asks whether any serious extant empirical theory in psychology makes use of something close to belief and desire--GOALS and KNOWLEDGE, say. And that is exactly the sort of question it is.

Yet that is not how philosophers have been treating the question, for what philosophers have been doing is (i) conceptual analyses of belief, and (ii) the semantics of belief attribution. But it is hard to see how (i) and (ii) could bear on psychology at all. For, first, there is no evident reason why a serious empirical psychologist should care what the ordinary concept of belief is any more than a serious physicist should care what the ordinary concept of force is.15 And, second, the good semanticist will want to describe the way scientists use their concepts, not dictate to science how a certain concept must be used. If Putnam is right is his claim that the reference of 'water' waits on chemistry, then the reference of 'belief' must wait on psychology. But it isn't just reference that waits

on science. Physicists created a sense for 'force'; they didn't just discover the reference of an ordinary concept or term. We can expect psychologists to do the same. When philosophers appreciate this point, they will be forced concede that the semantics of ordinary belief attribution has essentially no relevance to psychology at all.

Still, philosophers will want to know what is involved in semantically characterizing mental states. This is fine, provided it isn't construed as a question about the "ordinary concept of belief" or the semantics of belief attribution, but as a question in the philosophy of science analogous to this: what is involved in the non-continuous and non-deterministic characterization of the states of sub-atomic particles? And philosophers will want to know whether belief-desire explanations can be turned into good science. This question, I say, is analogous to this: can alchemy be turned into good science. Chemists solved that question; let the psychologists solve the other one.

FOOTNOTES

1. The situation is comparable to Dretske's (1981) insistence that mental states get their representational content during a learning period when they are perfect indicators of the properties they are said to represent. The theory runs the risk that no mental states have representational content as defined, or anyway no contents worth bothering about, since it might turn out that no mental state is ever a perfect indicator of distal states of affairs.

2. In a similar vein, we might ask what belief must be if folk psychology is to be true and explanatory. To this I have no objection, but two comments are in order.

First, to explicate the concepts of folk psychology one must begin by determining what folk psychology is. In practice, this is usually a matter of gesturing apologetically to some simple examples of alleged folk psychological explanation, examples that are generally conceded to be seriously flawed in some way. The assumption seems to be that we just know what folk psychology is. Psychologists who want to know what folk physics is--e.g., what people would say about the trajectories of falling objects--do some controlled experiments to find out. Philosophers who make claims about folk psychology don't feel the need. Perhaps they are right: perhaps folk psychology is rather more explicit in daily practice than folk physics. But I don't think this is what's really going on. Philosophers don't need to know what folk psychology is because they never appeal to it. It never matters to the philosophical arguments what the principles actually are--only that there are some. This is suspicious in itself: how can you know what the explanatory role of belief is in folk psychology if you don't know or care what folk psychology is?

Second, it is worth knowing what belief must be if folk psychology is to turn out true and explanatory only if there is some reason to think folk psychology might turn out to be true and explanatory. It is amazing how certain philosophers are that folk psycholgoy is true and explanatory given their avowed ignorance of its aleged principles.

3. The following example is taken from Burge (1979).

4. Again, I don't mean to beg any questions by the use of "psychological" here. For a computationalist, the issue could be put this way: If we know that (A) Tyler and Tyler* are computationally equivalent, and that (B) Tyler and Tyler* have different beliefs, then it follows that (C) Individuation of belief is a function of extra-psychological factors.

5. If you thought the CRHS methodology could lead to conclusions about the evidence-conditions for belief attribution, and you assimilate truth-conditions to evidence-conditions, then you would have a route to

conclusions about belief. This route, however, is bound to be unattractive to those who have learned Quine's lesson: If anything can be evidence for anything else (via intervening theory), then reflection on what might justify a belief attribution is not likely to tell you much about the truth conditions for belief attribution.

6. If there is such a thing: Actually, it is just stipulated that Tyler's words mean that breast of beast is better than tenderloin.

7. "At best" means empirically, under controled conditions. CRHS is bound to yield biased conclusions, for your considered responses will depend on what you consider, including your philosophical theories.

8. Since you can't learn an infinite list, the corollary amounts to the claim that learnable languages must be finite based: there must be a finite number of primitive terms, and a finite number of rules, such that learning the semantic value of each of the primitives, and learning the rules, allows one to generate the semantic value of any expression of the language. See Davidson (1965).

9. Or perhaps, whatever knowledge or misinformation in fact drives your use and understanding of `elevator's and elevators. If it is mostly or all misinformation, perhaps we say you don't have the concept of an elevator. Or perhaps we say that everything you believe about elevators is wrong. In that case, we credit you with a frame whose semantic value is appropriate to `elevator'. A consequence will be that whether you have the concept of an elevator will not depend on the psychological content of the frame in question at all (or only minimally--perhaps it must be the frame accessed in response to encounters with `elevator's if not to elevators).

10. By *knowledge I mean something that functions like knowledge, but needn't be true or justified or conscious.

11. This is harder than it sounds. Psychologists seem to acquiesce in the semantic individuation of beliefs, but they don't mean by semantics what philosophers (and some linguists) mean by semantics. When psychologists talk about semantics, they typically have in mind how something is understood, what it means to the subject; they don't have model-theoretic considerations in mind at all.

12. Perhaps there is a better way to characterize the proprietary goal of semantics, though I kind of doubt it. What matters here, in any case, is just that semantics has a proprietary goal that is independent of psychology.

13. This is Fodor's point in Fodor, 1980.

14. Everyone has known for a long time that it doesn't work the other way: from the fact that p and q are the same proposition, it doesn't follow that the belief that p is the same as the belief that q.

15. A serious psychologist might, of course, be interested in how people conceive of belief. But (i) no serious psychologist would employ the method of reflective equilibrium in such an investigation (I hope), and (ii) no one in their right mind would suppose that the way to construct an important explanatory concept in science is to discover what the "ordinary concept" is.

REFERENCES

Burge, Tyler (1979) "Individualism and the mental," in French, Euhling and Wettstein (eds) Studies in the philosophy of mind, Vol. 10, Midwest Studies in philosophy, University of Minnesota Press.

Davidson, Donald (1965) "Theories of meaning and learnable languages," Proceedings of the 1964 international congress for logic, mathematics and philosophy of science, North Holland.

Davidson, Donald (1967) "Truth and meaning," Synthese, 42:3:304-323.

Dennett, Daniel (1980) comment on Searle, "Minds, brains and programs," Behavioral and brain sciences, 3:417-424.

Dretske, Fred (1981) Knowledge and the flow of information. MIT Press: a Bradford Book.

Fodor, J. (1980) "Methodological solipsism considered as a research strategy cognitive science," Behavioral and brain sciences, 3:63-109.

Kaiser, M. K., J. Jonides, and J. Alexander (1986) "Intuitive physics: reasoning on abstract and commonsense problems." Memory and cognition.

McCloskey, M. (1983) "Intuitive physics," Scientific American, 24:122-130.

Putnam, Hilary (1975) "The meaning of 'meaning' , in Putnam, Mind, Language and Reality. Cambridge University Press.