Not The Mind's New Science?

Tim van Gelder
Department of Philosophy
Indiana University
Bloomington IN 47405 USA
tgelder@ucs.indiana.edu

Abstract.

In contemporary philosophy of mind there are at least two competing basic perspectives on the ontology of belief. One is internalist: beliefs are discrete internal representational states (e.g., Fodor 1987). The other is externalist: beliefs are some function of behavior itself (e.g., Ryle 1949; Dennett 1987). One source of evidence in deciding which perspective is correct is cognitive science, and connectionism in particular. If there are no discrete entities corresponding to beliefs in connectionist models, then connectionism implies that internalists were wrong about the ontology of belief -- and not, as Ramsey, Stich and Garon (1991) have argued, that there are no such things as beliefs at all. If connectionism supports an externalist ontology of belief, then it is undermining both (a) the general metaphysical framework within which most cognitive science proceeds, and (b) cognitive science's self-understanding as the science of the mind.

1. The Science of The Mind?

Cognitive science (here broadly construed so as to include cognitive neuroscience) is the study of cognition. Cognition is that internal, knowledge-based processing which is causally responsible for our intelligent behavior. Virtually all cognitive scientists are materialists. They believe that cognition is, ultimately, brain (CNS) activity. That is to say, all cognitive processes are realized or implemented as brain processes, though there is currently significant disagreement over whether it is more profitable to study cognitive processes *as* brain processes or in a more abstract, implementation-independent way. Often, cognitive scientists talk more grandiosely; their subject is not merely cognition but *the mind.* Mental processes are cognitive processes and *vice versa.* And since cognitive processes are brain processes, it is not uncommon for cognitive scientists to claim that their subject is *the mind/brain.*

In identifying mind, cognition and brain processing, cognitive scientists are taking over the metaphysical commitments of some long-standing philosophical traditions. Thus, the identification of mind and brain, which might be called *brain-based materialism,* has been proposed and debated by philosophers for centuries (see, e.g., (Hobbes, 1651/1962). This position re-emerged as the dominant perspective in Anglo-American philosophy of mind in the mid- to late-1950s (Place, 1956; Smart, 1959) -- that is, at approximately the time that contemporary cognitive science is generally thought to have begun. Brain-based materialism has itself standardly come embedded in a wider philosophical tradition which regards mind as that inner realm of representational states and processes which are causally responsible for our sophisticated behaviors - i.e., which identifies mind and cognition. The famous debate over whether mental processes are processes of the brain (brain-based materialism) or of something non-physical (e.g., substance dualism) has, for the most part, taken place *within* the general mind-as-cognition framework.

These big-picture perspectives may seem pretty much established from the perspective of cognitive science. However, there has always been significant disagreement in philosophical circles over whether the brain-based materialist and mind-as-cognition traditions are correct. That disagreement is alive and well today,

notwithstanding the various advances of cognitive science.

2. The Ontology of Belief.

Consider, for example, the ontology of *belief*. What are beliefs? Those working within the mind-as-cognition framework typically regard beliefs as discrete inner representations which are causally implicated in our behavior. More particularly, contemporary brain-based materialists take individual beliefs to be particular states of the brain, perhaps functionally characterized (Armstrong, 1973). We can think of this as an *internalist* perspective on the ontology of belief. Currently there are competing *externalist* perspectives. Externalists think of beliefs, not as internal causes of behavior, but rather as some kind of function of, or abstraction from, behavior itself.

The two most well-known externalists are Ryle and his student Dennett. According to Ryle, beliefs are heterogeneous[1] dispositions to behavior. To have a particular belief *that p* is to tend to act in a wide variety of heavily context-dependent ways which, as Ryle put it, "hang together on a common propositional hook" (Ryle, 1949/1984). Dennett, partly in order to avoid apparent problems with the dispositional account, proposed a revised view. Beliefs are "abstracta" - abstract states attributed to a system as a whole from a special stance ("the intentional stance") for the purpose of better predicting that system's behavior (Dennett, 1987).

Of course, nothing could have beliefs in either of these externalist senses unless they exhibited sufficiently sophisticated behavior, and exhibiting that behavior requires complex internal mechanisms. But we shouldn't confuse (a) the causes of the behavior which warrants belief ascription, with (b) the beliefs themselves.

3. Adjudicating Ontologies.

Internalists and externalists, then, differ radically on the ontology of belief. How can we decide who is right? There are various approaches. One broad, characteristically philosophical strategy is to suppose that, whatever beliefs are, ordinary folk seem to have a pretty good grasp of them. Thus, one should look for clues to the nature of belief in the language, practices or intuitions of ordinary people. Crudely speaking, focusing on language is characteristic of "ordinary language" school of philosophy generally, and Ryle in particular.[2] Contemporary philosophical discussions of "folk" or "common sense" psychology tend to focus on the practices.[3] A heavy reliance on the "intuitions" that one has, simply by virtue of being one of the ordinary folk oneself, is the specialty of the armchair analytic philosopher.

Another broad strategy is to let science decide ontological disputes. In the case of the ontology of belief, this invites the "pop-the-top" approach - open up the skull and take a look inside. If internalists are right, and beliefs really are states of the brain that are causally responsible for our behavior, then we should be able to find them inside the skull. Conversely, if there are no such entities to be found, the internalist must have been wrong.

Of course, just taking a peek isn't going uncover any beliefs. We need a specialized way of looking, something more sophisticated than just an eye or a microscope, a kind of "understandascope" for the brain. This is where cognitive science comes in. Cognitive science is a way of figuring out what states and processes are implemented in our brains and are causally responsible for our behavior. So, one way of adjudicating the debate over the ontology of belief is to ask cognitive science whether it is actually finding any beliefs, as internalists describe them, inside the skull.

---

[1] By "heterogeneous," Ryle meant that there is no simple way in which the behavior characteristic of the disposition is manifested.

[2] E.g., "...'believe,' on the other hand, is a tendency verb..." (Ryle, 1949/1984, p.133).

[3] E.g., "...folk psychology views beliefs as the sorts of things that can be acquired or lost one at a time..." (Ramsey, Stich, & Garon, 1991, p.205). Note that when this general approach is taken, the practices of ordinary folk are usually observed from the armchair. That is to say, the claims that these philosophers make are not usually based on serious sociological or anthropological study.

4. Connectionism Implies Eliminativism?

Connectionism is a variety of cognitive science. It tries to figure out what cognitive processes are like by modeling them with networks of idealized neural units. So we can direct the pop-the-top question - are there any beliefs in there? - to connectionists.

In fact, something along these lines has already been done. In an influential recent article, Ramsey, Stich and Garon (1991) asked whether connectionist models (of a certain familiar stripe) contained anything corresponding to beliefs. The answer was no: there is nothing in standard connectionist models with which beliefs might be identified. Curiously, however, instead of taking this as evidence that internalists are wrong about the *nature* of beliefs, Ramsey *et al.* took this as evidence *that there are no such things as beliefs at all.* What is going on here?

Their argument, which they laid out in an admirably clear way, went like this:

(1) Beliefs (as the folk understand them) are functionally discrete, semantically interpretable internal states that play a causal role in the production of behavior.

(2) A certain important sub-class of Connectionist networks - namely, distributed, subsymbolic cognitive models - do not contain any such states.

(3) Connectionist networks of that kind properly describe the human cognitive architecture.

Therefore, beliefs (as the folk understand them) do not exist.

Not surprisingly, many people are reluctant to accept the startling conclusion. In this case avoiding the conclusion means challenging at least one of the premises. Let us briefly consider them in turn.

An initially tempting line of response is to challenge Premise 2; i.e., to insist that there really are discrete belief-like entities in those distributed, subsymbolic models, you just have to look hard to see them. Thus, a connectionist might insist that his latest high-powered gadgetron analysis techniques reveal aspects of the model, not apparent to the naked eye, with which particular beliefs can be identified. Ramsey *et al.* defend Premise 2 vigorously and apparently successfully.

Another possibility is to deny Premise 3 -- i.e., to reject this kind of connectionism. This approach might be taken by traditional symbolic computationalists who are inclined to reject the whole connectionist enterprise, or even by connectionists who believe that some very different or more sophisticated connectionist approach is more promising. Ramsey *et al.* wisely prefer not to take a stand on the correctness of this kind of connectionism. So to deflect this kind of objection, they weaken their conclusion to a conditional: *if* this kind of connectionism is correct, there are no beliefs.

What about challenging Premise 1? Interestingly, Ramsey *et al.* don't entertain this as a likely *response* to their argument. They do however lay out detailed arguments in favor of Premise 1 before presenting the argument.

5. Reverse Thrust.

One man's *modus ponens* is, as philosophers like to say, another's *modus tollens*. Roughly translated, this means that we can either take the premises as evidence for the conclusion, or the falsity of the conclusion as evidence against one or more of the premises. In this case, we might take it to be manifestly obvious, or otherwise plausible, that beliefs really do exist, and so take the eliminativist conclusion of this argument as evidence against one of the premises. If we follow Ramsey *et al.* in accepting Premise 2 and deferring any stand on Premise 3, then we are left with the following implication: *if* connectionism is right, then beliefs are *not* functionally discrete, semantically interpretable internal states that play a causal role in the production of behavior.

To reverse the thrust of their argument in this way is, of course, just to adopt the pop-a-top strategy to the question of the ontology of belief. It is to enlist cognitive science, and connectionism in particular, in the effort to understand what beliefs actually are. It is to ask science to help settle a philosophical question.

Why do Ramsey *et al.* keep their ontology of belief and accept the eliminativist implication, rather than keeping beliefs and revising their ontology? Part of the explanation may be this. Ramsey *et al.* appear implicitly to hold a quite traditional view of the relation between science and philosophy. According to this view, philosophy and science have fundamentally different domains. Empirical questions may belong to sci-

ence, but metaphysics is still the domain of philosophy. Ontology -- determining the fundamental nature of things -- is a metaphysical enterprise, and hence within philosophy's purview. Thus, Ramsey *et al.* take their purely philosophical evidence and arguments -- analyses of common sense psychology -- as determining the *ontology* of belief, while cognitive science is allowed to establish, at best, the *existence* of the entities so determined.

The idea that philosophy and science have separate domains, and that philosophy has a special privilege to adjudicate ontological questions, has been challenged this century, especially by the philosopher Quine (1953/1961). A better view is that science and philosophy are continuous, and that neither science nor philosophy has any exclusive privilege over ontological questions. Ontology is part of the general philosophical project of understanding, as Sellars put it, "how things in the broadest possible sense of the term hang together in the broadest possible sense of the term." (Sellars, 1962). Doing this properly means drawing on as many kinds of evidence and argument as possible, including any offered by science.

In this case, the fact that connectionist models contain nothing corresponding to beliefs should be taken as evidence that the internalist camp gets the ontology of belief wrong, an idea which has much independent philosophical merit, rather than for the much more disruptive and paradoxical idea that there are no such things as beliefs at all.

6. Paying the Piper.

Whatever the case with beliefs, there's no such thing as a free lunch. If connectionism implies an externalist ontology of belief, then it implies a number of other surprising things as well:

1. Mind is not cognition. The mind-as-cognition tradition is wrong about mind. If beliefs are not internal entities, then not all of mind can be "internal representational states or processes causally implicated in intelligent behavior". Ryle and others have already argued this case in great philosophical detail. We can see connectionism as offering confirmation of this idea from the direction of cognitive science.

Of course, any philosophical tradition as established as the mind-as-cognition approach is unlikely to be *completely* misguided. It seems reasonably clear that at least *some* aspects of the mind *are* internal states and processes causally implicated etc... A more sophisticated view of mind, then, is that mind is ontologically heterogeneous. Mind is made up of a diversity of entities of different ontological categories. Mind is not a simple thing of any kind; rather, it is a structured collection of many different kinds of entities interrelated in complex ways. A useful analogy is an economy. An economy is a totality comprised of a wide range of very different entities consumers, financial instruments, products, factories, rules, exchanges, etc. Likewise, mind is a complex totality; it might be thought of as an ontological super-category, an entity comprised of many entities of diverse ontological kinds.

2. Brain-based materialism is wrong.. If not all mental entities are inner entities causally implicated in our behavior, then not all entities can be identical with brain states or processes, no matter how subtly one construes that identity relation. The simple equation "mind is (the operation of the) brain" is wrong -- and so, therefore, is the philosophical tradition which stretches from Hobbes to Fodor and Churchland.

Whether something appropriately described as *materialism* nevertheless survives is a further question, one that is both interesting and difficult.

3. Cognitive science is not The Mind's New Science. Cognitive science is the science of cognition. Since mind is not identical with cognition, cognitive science is not the science of the mind, any more than the study of factories is the study of an economy. Cognitive science is, at best, the science of those *aspects* of mind which are cognitive processes. In other words, an important implication of connectionism is that, strictly speaking, there needs to be a change in the self-image of cognitive science.

Ultimately, nobody is forced to accept these conclusions. Developing the best overall philosophical picture is a matter of accepting those doctrines which make for the most harmonious total world-view. There is room for disagreement over which doctrines have that effect. One could, I suppose, insist on believing the internalist story about beliefs. Except, perhaps, that if in-

ternalists and connectionists are right, there aren't any beliefs...[4]

References

Armstrong, D. M. (1973) *Belief, Truth and Knowledge*. Cambridge: Cambridge University Press.

Dennett, D. C. (1987) True Believers: The Intentional Strategy and Why it Works. In *The Intentional Stance*. Cambridge MA: Bradford/MIT Press.

Hobbes, T. (1651/1962) *Leviathan*. New York: Collier Books.

Place, U. T. (1956) Is Consciousness a Brain Process? *British Journal of Psychology*, 44-50.

Quine, W. V. O. (1953/1961) On What There Is. In *From a Logical Point of View*. New York: Harper & Row.

Ramsey, W., Stich, S. P., & Garon, J. (1991) Connectionism, Eliminativism, and the Future of Folk Psychology. In W. Ramsey, S. P. Stich, & D. E. Rumelhart ed., *Philosophy and Connectionist Theory*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Ryle, G. (1949/1984) *The Concept of Mind*. Chicago: University of Chicago Press.

Sellars, W. (1962) Philosophy and the Scientific Image of Man. In R. G. Colodny ed., *Frontiers of Science and Philosophy*. Pittsburgh: University of Pittsburgh Press.

Smart, J. J. C. (1959) Sensations and Brain Processes. *Philosophical Review*, 68, 141-156.

van Gelder, T. J. (1993) Connectionism and the mind-body problem: Exposing the distinction between Mind and Cognition. *AI Review*.

[4] For an expanded treatment of many of the issues in this paper, see (van Gelder, 1993).