Philosophical Studies (2006) 127:299–316
DOI 10.1007/s11098-005-4960-z

JOSÉ LUIS BERMÚDEZ

KNOWLEDGE, NATURALISM, AND COGNITIVE ETHOLOGY: KORNBLITH'S *KNOWLEDGE AND ITS PLACE IN NATURE*

ABSTRACT. This paper explores Kornblith's proposal in *Knowledge and its Place in Nature* that knowledge is a natural kind that can be elucidated and understood in scientific terms. Central to Kornblith's development of this proposal is the claim that there is a single category of unreflective knowledge that is studied by cognitive ethologists and is the proper province of epistemology. This claim is challenged on the grounds that even unreflective knowledge in language-using humans reflects forms of logical reasoning that are in principle unavailable to nonlinguistic animals.

As one might expect, Hilary Kornblith's *Knowledge and its Place in Nature* (Kornblith, 2002) is an exercise in naturalized epistemology. Kornblith sets out to show that knowledge is a robust natural phenomenon that can be studied in the same way as any other natural phenomenon. More specifically, knowledge is a natural kind that can be elucidated and understood in scientific terms, so that we can, roughly speaking, do epistemology by doing science. The key claim here, of course, is that knowledge is a natural kind. It is this that gives point and focus to the in itself rather anemic claim that knowledge is a natural phenomenon. Much of the book is devoted to elucidating and defending the thesis that knowledge is a natural kind. Kornblith's strategy is ingenious and challenging. The first step, undertaken in Chapter 2, is to argue that cognitive ethology makes irreducible and ineliminable appeal to *a* type of knowledge in making sense of a wide range of animal behaviors, from navigational abilities to deception behaviors. The second step, undertaken in Chapters 3 and 4, is to argue that the type of knowledge invoked in

cognitive ethology is the type of knowledge that we should be interested in as epistemologists. My primary concern in these comments is with the second step. I will be exploring Kornblith's central idea that there is a single category of knowledge that is both studied by cognitive ethologists and is the proper study of epistemology.

Let me begin by sketching out the broad outlines of the conception of knowledge that Kornblith is presenting. Kornblith is an externalist about knowledge. He devotes a chapter to rejecting internalist accounts of knowledge, taking an internalist account of knowledge to be one that places a reflective requirement upon knowledge. On the view he favors knowledge should be understood in reliabilist terms. But Kornblith's reliabilism is not a standard reliabilism. He does not define knowledge as reliably produced belief. Rather, an organism (human or non-human) has knowledge of a particular domain to the extent that it has cognitive capacities that reliably generate true beliefs about that domain. When we think about knowledge we should not do so in atomistic terms – by asking, for example, what we need to add to a true belief for it to count as knowledge. Rather, we need to think about knowledge in terms of the capacities that give rise to true beliefs. As he puts it, "the standards for knowledge arise from the demands that nature makes on animals if they are to function in their natural environment" (Kornblith, 2002, 164), and what nature demands is that organisms have cognitive capacities that allow them to respond appropriately to the informational demands that the environment imposes upon them. Since Kornblith takes responding appropriately to the informational demands imposed by the environment to require having true beliefs about the environment, it is clear why he thinks that the category of knowledge is essential to our thinking about animal behavior. Once he has made this move his next step is to argue that the very same category of knowledge that we use to think about animal behavior is what we should be interested in as epistemologists. It is a mistake, he argues in Chapter 3, to see knowledge in primarily social terms, as emerging

from the practice of giving and debating reasons. And nor, he continues in Chapter 4, is it right to impose a reflective requirement upon knowledge in the manner of epistemological internalism. Rather we should, as naturalized epistemologists, take as our subject matter those very same cognitive capacities that produce reliable beliefs in non-human animals and in human knowing subjects. Kornblith is explicitly promoting a two-tiered view of human cognition, on which a reflective layer of sophisticated second-order cognition is superimposed upon a more primitive layer of information-gathering systems that we share with non-human animals. Traditional epistemology has concentrated on the reflective tier. His ambitious aim is to reconfigure epistemology to focus on the more primitive capacities, on what he sometimes calls "animal knowledge". This reconfiguration of epistemology will, he thinks, allow epistemologists to draw upon the scientific study of animal knowledge in cognitive ethology.

Let me begin by noting a puzzling feature of Kornblith's argumentative strategy. Cognitive ethologists explicitly characterize their enterprise as trying to apply the conceptual framework of commonsense, belief-desire psychology to make sense of animal behavior. One might wonder, therefore, why Kornblith's argument proceeds via cognitive ethology at all. Why does he not argue directly from the role that commonsense belief-desire psychology plays in cognitive science to the idea that the central categories of belief-desire psychology are natural kinds?

One significant reason is that, for Kornblith, the central explanatory notion of cognitive ethology is knowledge, whereas the central explanatory notion of cognitive science and scientific psychology is belief. This derives, I think, from his idea that cognitive ethology is primarily concerned with explaining the presence of particular behavior patterns and cognitive capacities at the level of the species, rather than at the level of the individual. He thinks that the notion of belief is all we need to explain individual behaviors. It is only when we move beyond individual behaviors to consider the patterns

that they display and the cognitive capacities that make them possible that we find ourselves having to bring the notion of knowledge into play. If we are interested in knowledge, then, we must turn away from cognitive psychology and cognitive science to cognitive ethology.

In the background here must be the idea that, to the extent that our ordinary commonsense psychological notions have an application in cognitive psychology and cognitive science, they do so only in the service of explaining/predicting why people behave in particular ways on particular occasions. But this is at best disputable. Cognitive psychology and cognitive science are not particularly interested in the explanation and prediction of individual behaviors, which is one reason why one finds very few laws in psychology and cognitive science. Cognitive psychology and cognitive science are primarily concerned with identifying and explaining cognitive capacities. This is one reason why philosophers of psychology such as Cummins have identified explanation by functional decomposition as playing such an important role in the cognitive sciences (Cummins, 2003). Even more to the point, commonsense psychological concepts such as knowledge and belief have very little role to play in cognitive psychology. It is a widespread misapprehension in philosophy that they do. If one thinks that the categories corresponding to natural kinds are those that play a robust role in a well-established scientific theory then cognitive psychology and cognitive science are very bad places to look for evidence that belief (and, by extension, knowledge) is a natural kind.

It is helpful in this context to think through the reasons that Kornblith gives for thinking that the category of knowledge is central to cognitive ethology. Here is a representative passage.

> The very idea of animal behavior requires the reception, integration, and retention of information from a wide range of different sources. But this is just to say that any conception of sophisticated animal behavior that makes any sense of it at all will have to see the animal's cognitive equipment as serving the goal of picking up and processing information. And this commits one to the notion of animal knowledge. (Kornblith, 2002, 61)

Kornblith's claim, in brief, is that cognitive ethologists are driven to the category of knowledge because they have no alternative but to explain animal behavior in information-processing terms. But it is just wrong to characterize cognitive ethologists as treating animals as information-processing systems. Here is a clear statement of the aims and guiding assumptions of cognitive ethology from Carolyn Ristau, one of its leading exponents.

> Cognitive ethology has been defined as the study of the mental experiences of animals, particularly in their natural environment, in the course of their daily lives. Data are derived from the observation of naturally occurring behavior as well as from experimental investigations conducted in the laboratory and in the field. By emphasizing naturally occurring behaviors, cognitive ethologists recognize that the problems faced in finding food and mates, rearing young, avoiding predators, creating shelters, and communicating and engaging in social interactions may require considerable cognitive skills, possibly more complex than and different from those usually examined in traditional psychological laboratory studies. The term "mental experiences" acknowledges that the mental capabilities of animals, in addition to unconscious mental processes, may also include conscious states. This affords the animals sensory experiences, pleasure and pain, the use of mental imagery, and involves at least simple intentional states such as wanting and believing. (Ristau, 1999)

In contrast to what one might think of as the essentially *subpersonal* notion of information-processing, cognitive ethologists are committed to *personal-level* notions of belief, desire, experience, and so forth. This is what distinguishes cognitive ethology from comparative psychology and animal learning theory. And it is, arguably, what distinguishes cognitive ethology from cognitive psychology and cognitive science in general, which are generally characterized in the information-processing terms that Kornblith uses to describe the enterprise of cognitive ethology. What is at issue here is really the equation that Kornblith makes between information-processing, on the one hand, and animal knowledge, on the other – where "animal knowledge" is understood in ethological terms. If we are using the term "knowledge" in any sense that is likely to be of interest to epistemology, then this equation may well be fundamentally misconceived.

It seems to me that one's naturalized epistemology will look very different depending on where in the web of science one looks for guidance. It may well be that cognitive ethology makes ineliminable use of notions of belief, desire, and knowledge. But does this show that knowledge is a natural kind, or that cognitive ethology is no more of a science than folk psychology is a science? Kornblith draws the former conclusion, but there is clearly scope for a range of different positions, depending on how one understands the type of information-processing appealed to in cognitive psychology and the cognitive sciences. Unsurprisingly these different possibilities are correlated with taking different areas of this vast and complex area as paradigmatic. One extreme suggestion, a suggestion that starts to look plausible when one takes the cognitive neurosciences as one's paradigm, is that our commonsense psychological notions of knowledge, belief, and so forth, have no explanatory role to play in mature science (with the obvious implication being that cognitive ethology is not, and probably never will be, part of mature science). It is striking that Kornblith does not engage with this type of eliminativism, which is particularly associated with the writings of Paul Churchland (Churchland, 1981 – and see Bermúdez (2005b) for discussion of different forms that arguments for eliminativism might take).

A less extreme suggestion, one that might emerge if one takes computational cognitive science as one's paradigm, is that some ancestor of our commonsense psychological categories might survive in a mature science but in a significantly different form. This in fact is the view adopted by a number of those who have argued in print that we should view commonsense psychological notions as natural kinds. Lycan, for example, has drawn an explicit comparison between our commonsense psychological vocabulary and the natural kind terms studied by Putnam: ''As in Putnam's examples of ''water'', ''tiger'' and so on... the ordinary word ''belief'' (*qua* theoretical term of folk psychology) points dimly towards a natural kind that we have not fully grasped and that only a mature psychology will reveal'' (Lycan, 1988, 32).

It seems clear that Kornblith cannot accept the first of these two suggestions. His version of naturalized epistemology is predicated on the idea that knowledge is a robust natural phenomenon. This is in stark contrast to Quine's naturalized epistemology (Quine, 1969). Quine's project of exploring the genesis of scientific theory does not in any sense commit him to the survival of commonsense psychological notions – and as is well known his overall position contains strong eliminativist pressures. A natural question for Kornblith, then, is what grounds his confidence that eliminativism cannot be the best response to the methodology of naturalized epistemology. Does he have any reasons, independent of the explanatory practices of cognitive ethology, to think that our commonsense psychological concept of knowledge picks out a robust scientific kind? And suppose it were the case that our concept of knowledge completely dropped out of the picture in, say, cognitive neuroscience, how would this stack up against the deployment of the category of knowledge in cognitive ethology? There are deep and important questions here about the relation between the different sciences of the mind – and indeed about what one actually takes to be a science of the mind.

But even when we put the threat of eliminativism to one side, we still need to engage with the possibility of revisionism. Kornblith is trying to persuade epistemologists in general that the phenomenon they have been studying is the very phenomenon that plays an explanatory role in cognitive ethology and, I am sure he thinks, elsewhere in the sciences of the mind. But then he needs to confront the possibility that the scientific category of knowledge will fail to map on to our pretheoretical understanding of knowledge. As the passage from Lycan quoted earlier reminds us, we are familiar with one form of mismatch between scientific category and pretheoretical category. The stereotype and background beliefs that we attach to natural kind terms such as "water" and "gold" are false of many instances of the kinds that those words pick out. These instances are nonetheless correctly picked out by the relevant

kind words because the stereotypes and background beliefs do not serve to fix the reference of those words. Putnam has an account of the mechanics of natural kind terms that allows their reference to be fixed quasi-indexically (Putnam, 1975). What makes it true that a certain object is made of gold is that it has the same microstructure as certain canonical objects that were originally and ostensively identified as paradigms of gold. A second set of questions for Kornblith, then, has to do with whether he thinks that something like this is the case for our commonsense psychological vocabulary. Should we view the workings of our psychological vocabulary in the way that we interpret the workings of the words "water" and "gold"? How much room does the semantics of "knowledge" leave for the possibility that the stereotype we attach to the term is radically mistaken?

It is worth noting that a broadly Putnamian account of the semantics of "knowledge", "belief", and other terms of epistemological interest, allows some sort of reconciliation between the traditional epistemological project of conceptual analysis, on the one hand, and the project of naturalized epistemology, on the other (see Chapter 1 of Bermúdez (2005a) for further discussion). One might say, for example, that traditional epistemological enquiry has focused on clarifying the stereotypes of our concepts of justification and knowledge, while naturalized epistemology offers the prospect of bringing us to an understanding of the "real essences" of the phenomena of knowledge and justification. On this picture, then, naturalized epistemologists and conceptual analysts are not so much in conflict as talking past each other – or, more charitably, exploring different aspects of a complex phenomenon that involves both the category of knowledge and the complex practices within which it is embedded.

It will be no easy matter, however, to work this picture out in the case of knowledge. In the case of terms such as "gold" and "water" it is relatively straightforward to find paradigm exemplars of the category in question that will provide a quasi-indexical way of anchoring the application of the

relevant terms. It is only to the extent that we have canonical instances of gold and water that it is plausible to define the extension of the terms "gold" and "water" in terms of similarity of structure to those canonical instances ("water" just picks out the stuff that has the same structure as *this*). But do we have canonical instances of knowledge that will serve the same indexical reference-fixing role? Disputes in epistemology, as in philosophy more broadly, tend to come about (or at least to be reinforced by) different groups of philosophers taking fundamentally different instances of a given category to be canonical. This is why, for example, the putative counter-examples that seem so persuasive to internalists rarely move externalists – and vice versa. An internalist's canonical example of knowledge is an externalist's outlier. So what are the prospects, then, of identifying canonical instances that will allow us to treat the semantics of "knowledge" in the same way as we treat the semantics of "water" and "gold"?

This brings us, I think, to the key issue in thinking about Kornblith's book and his overall project. Kornblith, although he favors a reliabilist epistemology, is offering us a new set of canonical instances of knowledge – canonical instances that are drawn from cognitive ethology. Kornblith's argument is, essentially, that we need to reconfigure epistemology by focusing on cognitive capacities that we share with non-human animals. These cognitive capacities are amenable to scientific study and hence count as natural kinds. The full weight of Kornblith's argument, then, rests upon the claim that epistemologists and ethologists are essentially studying the same phenomenon. In the remainder of these remarks I will be focusing on this claim.

In order to set the scene we can consider a passage that Kornblith cites from Sosa's well-known essay on "Knowledge and intellectual virtue". Sosa draws the following distinction between "animal knowledge" and "reflective knowledge".

> One has *animal knowledge* about one's environment, one's past, and one's own experiences if one's judgments and beliefs about these are direct responses to their impact – e.g. through perception or memory – with

little or no benefit of reflection or understanding. One has reflective knowledge if one's judgment or belief manifests not only such direct response to the fact known but also understanding of its place in a wider whole that includes one's belief and knowledge of it and how these come about. (Sosa, 1991, 240)

As Kornblith interprets him, Sosa's distinction (despite initial appearances) turns out to be effectively a distinction between all human knowledge, on the one hand, and all non-human animal knowledge, on the other. We can see Kornblith as trying to turn Sosa's position on its head. Kornblith wants to show not simply that Sosa's category of animal knowledge is applicable to humans, but further that animal knowledge provides our canonical instances of knowledge. In opposition to Sosa and those who think like him Kornblith argues that self-conscious reflection on one's beliefs cannot be a requirement upon knowledge. The various arguments that he offers against internalist reflective requirements in Chapter 4 and against internalists requirements of social reason-giving in Chapter 3 are, he thinks, sufficient to show that we cannot distinguish human knowledge from the knowledge of other animals by appeal to reflection and hence, he concludes, "the conception of knowledge that we derived from the cognitive ethology literature, a reliabilist conception of knowledge, gives us the only viable account of what knowledge is" (p. 135).

Let us suppose that Kornblith's arguments against internalists requirements of self-conscious reflection are good ones. It is hard to see how this could fail to show that reflective knowledge, as Sosa understands it, should not be our paradigm of knowledge. But is this sufficient to show that our paradigm should be the conception of knowledge that we derive from the cognitive ethology literature? Well, yes – provided that one accepts two important and related theses. The first is that reflective knowledge derives simply from the addition of reflection to animal knowledge. Kornblith appears to think that what explains the difference between animal knowledge and reflective knowledge is the emergence of some sort of metacognitive or metarepresentational capacity – to the ability to think is added the ability to think about one's own

thoughts. Although Kornblith is not explicit about this, I think that he would be prepared to accept that this type of metarepresentation is unique to language-using humans – but nothing in his argument stands or falls with this claim. What is crucial to his argument is the second thesis, which might be termed the *subtraction assumption*. Roughly, the idea is that what we get when we subtract the metarepresentational component of human knowledge is the very same type of animal knowledge that we find in cognitive ethology. The word "animal" in the expression "animal knowledge" is doing double duty. On the one hand it is to be taken literally, thus yielding the connection between epistemology and cognitive ethology. On the other hand, it is to be taken figuratively, indicating that the knowledge in question is not reflective.

Of course, though, when one finds words doing double duty it is natural to wonder whether there might not be a fallacy of equivocation, or at least some slippage, in the neighborhood. Are there any reasons to think that we need to separate out two different types of animal knowledge, with cognitive ethologists being interested in one and epistemologists in the other? If we were to need to make such a separation then Kornblith's project would be seriously threatened. I will not challenge Kornblith's reasons for thinking that it is a mistake to build reflective or reason-giving requirements into our account of knowledge – not least because I suspect that my co-symposiasts will be pressing him on precisely those issues. But in the time remaining I do want to explore a different way of trying to drive a wedge between two different ways of understanding animal knowledge.

Let me begin by putting the contrast between reflective and unreflective knowledge in a way that I think Kornblith will find congenial. This way of thinking about the contrast emerges when one thinks about the principles governing the evolution of systems of beliefs – and in particular about the principles that govern the ways in which beliefs are revised. The beliefs that are produced by the systems that collectively make up reflective knowledge are no less modifiable than the

beliefs that collectively make up reflective knowledge. And of course, part of what it is to have cognitive capacities that reliably produce true beliefs is that they should function effectively to revise and replace beliefs that are not true. One way of thinking about the difference between reflective and unreflective knowledge is in terms of the norms that govern the different processes of belief revision and belief modification that they involve.

It seems to me that, at the unreflective level, beliefs are subject to what might be termed the *norm of coherence*. That is to say, beliefs are revised in a way that restores consistency and coherence. Conflicts between beliefs, or between beliefs and the evidence of the senses, creates cognitive dissonance and adjustments are made in the system until cognitive dissonance is dispelled. Suppose, for example, that a rat comes to believe as the result of a reinforcement schedule that pressing a lever will produce food, but then that the contingency is changed so that the food is delivered at random in a manner completely unrelated to the lever pressing. The rat's belief will conflict with its perceptual monitoring of the situation. There is cognitive dissonance that must be resolved, most obviously by revising the belief about the dependence of food delivery upon lever pressing.

At the reflective level, in contrast, belief revision is governed by a further and more demanding norm. This is the *norm of truth*. Coherence is a desideratum but not the sole desideratum. What matters above all is that the evidence for a given belief should be such as to make the beliefs that it supports likely to be true. Reflective belief revision concerns itself explicitly with the logical and probabilistic relations between evidence and belief – as well, of course, as the parallel relations holding between individual beliefs and sets of beliefs. It is only when we are dealing with a belief system that is governed by the norm of truth that we are in a realm where internalism is a live option. At the reflective level a belief might be rejected or modified in the absence of countervailing evidence or tension with existing beliefs – it

might, for example, be rejected simply because the believer recognizes that it is not warranted by the evidence. Although the notions of warrant and justification can be applied in an attenuated sense at the level of direct belief revision (most prominently in a broadly externalist sense according to which, roughly speaking, a set of beliefs is warranted to the extent that it is produced by reliable mechanisms and modified according to principles that tend to preserve truth and eliminate error), there is no sense in which the extent to which their beliefs are warranted or justified can be an issue for creatures operating solely at that level. Internalist notions of warrant and justification can get a grip only at the reflective level, because these notions can be applied only to thinkers capable of explicit reflection on the relations between thoughts and perceptions and between thoughts and other thoughts.

We can use this way of thinking about the distinction between reflective and unreflective knowledge to put Kornblith's basic claim as follows. Although (Kornblith would say) we do take ourselves to be governed by the norm of truth and although we are capable of reflective knowledge, it is not here that we should look for our paradigms of human knowledge. Rather, we should look at the domain of unreflective knowledge, governed by the norm of coherence and regulated by mechanisms that work to generate true beliefs and minimize cognitive dissonance. This gives us the basic reconfiguration of the concept of knowledge that Kornblith is composing. Why is the reconfigured category of knowledge best viewed as a natural kind? Because, says Kornblith, we can study the category of unreflective knowledge in cognitive ethology, where it is put to work in the explanation of animal behavior.

But we are now in a position to explore this suggestion with a more critical eye. Suppose we ask why one might think that reflective knowledge is a uniquely human attainment. An obvious proposal would be that reflective knowledge is only available to language-using creatures. This proposal is, I think, correct – for the following reason. It is

only possible to evaluate and reflect upon the extent to which one belief implies another (more precisely: the extent to which acceptance of one belief commits one to acceptance of another), or the extent to which a belief is supported by a particular type of evidence if one is able explicitly to hold those beliefs in mind. And, as I have argued at some length elsewhere (Bermúdez, 2003, Chapters 8 & 9), thoughts can only be the objects of further thoughts if they have consciously accessible linguistic vehicles. A thought can only be "held in mind" in such a way that it can be the object of a further thought if it has a linguistic vehicle. To put it succinctly, intentional ascent requires semantic ascent – we cannot think about thoughts except through thinking about the sentences that express them.

Reflective knowledge, then, requires language because thoughts can only be the objects of thinking if they are clothed in language. But, one might think, this tells us only about reflective thinking. It has no implications for how we think about non-reflective thinking – and none, in particular, for Kornblith's thesis that there is a single type of animal knowledge possessed both by humans and by non-human animals. What is distinctive about human cognition, one might think, is that language opens up the possibility of taking up a reflective stance upon on a common core of non-linguistic thoughts that are available to non-linguistic creatures. I wonder, however, whether this might not be a mistake. Perhaps there are fundamental differences between our two types of animal knowledge.

Suppose both that reflective thinking is the sole province of language-users, and that reflective thinking is directed at a type of thinking that is shared between language-using and non-linguistic creatures. This means that the evidential relations that are the focus of reflective thinking must already be present in the "animal knowledge" to which reflection is focused. Reflective thinking is targeted at a structure of beliefs that stand in logical and probabilistic relations to each other, and to the perceptual evidence in which they, or at

least the majority of them, are grounded. So, one might think, at the unreflective level of animal knowledge beliefs are inferred from each other, either deductively or probabilistically, or formed on the basis of perceptual experience. These logical and probabilistic relations hold between beliefs and experiences at the level of animal knowledge. They are there all along and reflective thinking merely brings them into the open.

This picture is an attractive one. It only makes sense, though, if we make a very substantial assumption about the thinking of non-linguistic creatures – viz. that it is thinking of fundamentally the same kind that we find in language-using creatures, lacking only the capacity for reflection. It is not clear to me, though, that this is the right way to look at what one might term animal reasoning. The same argument that leads to the conclusion that reflective knowledge requires language also yields the result that logic requires language. If this is right then it looks very unlikely that we will be able to identify a single category of animal knowledge applicable both to language-using and non-linguistic creatures – and, by extension, Kornblith's argument from cognitive ethology to the conclusion that knowledge is a natural kind will look rather shaky. Let me sketch out, then, an argument for the dependence of logic upon language.

I suggested earlier that intentional ascent requires semantic ascent. Clearly, therefore, to argue that logic requires language we need establish simply that logical thinking involves intentional ascent. Let me offer some reasons for thinking that the capacity for intentional ascent is required for all thinking that involves compound thoughts with further thoughts embedded in them.

We can start with a basic class of compound thoughts – namely, those involving the basic logical connectives, such as disjunction, conjunction and the material conditional. Consider a disjunctive thought of the sort that might be expressed in the sentence "A or B". What is it to be capable of entertaining such a thought? It is to be capable of understanding

that a certain relation holds between two thoughts – the relation of their not both being false (and, on some understandings of 'or', their not both being true). Of course, the disjunctive thought itself is not a thought about the two thoughts expressed by "A" and "B". It is a thought about the states of affairs that serve as the truth conditions for the thoughts. Nonetheless, the disjunctive thought is not available to be thought by any creature that is not capable of thinking about how the truth-value of one thought might be related to the truth-value of another thought. The same holds of the other truth-functional propositional operators. Understanding a truth-functional operator is understanding how it serves to form compound thoughts whose truth-value is a function of the truth-values of their parts. No creature that was not capable of thinking about thoughts could have any understanding of truth-functional compound thoughts. But if the capacity for intentional ascent is required, then this is a type of reasoning available only to creatures capable of semantic ascent.

The argument is even clearer for types of logical reasoning that exploit the internal structure of thoughts. Consider the inference form of existential generalization. This is the pattern of inference instantiated by the transition from a thought symbolized Fa to one symbolized $\exists xFx$ – that is to say, from an atomic proposition to the effect that a named individual has a given property to the general proposition that at least one individual has that property. The logical operations involved in this transition are clear enough. The first is breaking down the atomic proposition into two components, a predicative component and a nominative component (or, in Fregean terms, a function and an argument). Once the internal structure of the atomic proposition is manifest, the next operation is to replace the nominative component by a variable. The final operation is to bind that variable with an existential quantifier. This sequence of logical operations gives us an important clue as to what is involved in a subject's being able to understand the existential quantifier in a manner that

permits existential generalization. In order to understand how a given proposition can imply an existential generalization a thinker needs to be able to view it as being composed in such a way that the nominative component can be replaced by an arbitrary name (and hence by a variable). The proposition needs to be "broken down" in thought before the existential quantifier can be applied. But this breaking down in thought of an atomic proposition presupposes the capacity for intentional ascent. It involves holding the thought in mind and determining its structure in a way that creates a space for the variable that will be bound by the existential quantifier. The same point holds, *mutatis mutandis*, for thinking involving the universal quantifier.

So, what does all this tell us? The conclusion to draw, I think, is that the reasoning of non-linguistic creatures is fundamentally different from the reasoning of language-using creatures. It does not involve what we would think of as propositional operators or quantificational structures. It cannot exploit either logical connections between thoughts or the internal structure of those thoughts. This is not, of course, to say that there is no such thing as animal reasoning. I have tried elsewhere to identify forms of reasoning at the non-linguistic level and to explain them without assuming that the animal is deploying elementary logical concepts or exploiting the internal structure of a thought (Bermúdez, 2003, Chapter 7). It seems to me, however, that the fact that animal reasoning takes this form makes Kornblith's central hypothesis very problematic. He needs it to be the case that there is a single category of animal knowledge that applies both to non-human animals and to humans. Only if this is so can his argument that knowledge is a natural kind go through. What makes it seem persuasive that there might be such a single category is that his focus is on individual beliefs and the cognitive capacities that produce them, rather than on the relations holding between beliefs that make it the case that those beliefs form a system. It seems to me that when we think about the role of logic in the generation and regulation of

belief (when we think of the web of belief rather than individual beliefs) the assumption that there is a single category of animal knowledge starts to look problematic.

REFERENCES

Bermúdez, J.L. (2003): *Thinking without Words*, New York: OUP.
Bermúdez, J.L. (2005a): *Philosophy of Psychology: A Contemporary Introduction*, London: Routledge.
Bermúdez, J.L. (2005b): 'Arguing for Eliminativism', in B. Keeley (ed.), *Paul M. Churchland*, Cambridge: CUP.
Bermúdez, J.L. (2006): *Philosophy of Psychology: Contemporary Readings*, London: Routledge.
Churchland, P.M. (1981): 'Eliminative Materialism and the Propositional Attitudes', *Journal of Philosophy* 78-2, 67–90.
Cummins, R. (2003): 'How does it Work? vs. What are the Laws? Two Conceptions of Psychological Explanation', in F. Keil and R. Wilson (eds.), *Explanation and Cognition*, Cambridge: MIT Press (reprinted in Bermúdez 2006).
Kornblith, H. (2002): *Knowledge and its Place in Nature*, Oxford: OUP.
Lycan, W. (1988): *Judgment and Justification*, Cambridge: CUP.
Putnam, H. (1975): 'The Meaning of 'Meaning'', in *Mind, Language and Reality*, Cambridge: CUP.
Quine, W.V.O. (1969): 'Epistemology Naturalized', in *Ontological Relativity and Other Essays*, New York: Columbia University Press.
Ristau, C. (1999): 'Cognitive Ethology', in R.A. Wilson and F. Keil (eds.), *The MIT Encyclopedia of Cognitive Science*, Cambridge, MA: MIT Press.
Sosa, E. (1991): *Knowledge in Perspective*, Cambridge: CUP.

Philosophy–Neuroscience–Psychology Program
Washington University
St Louis, MO 63105
USA
E-mail: bermudez@wustl.edu