Philosophical Studies (2006) 127:337–349
DOI 10.1007/s11098-005-5005-3

HILARY KORNBLITH

REPLY TO BERMUDEZ AND BONJOUR

I want to thank José Bermudez and Laurence BonJour for their constructive and challenging remarks. I am pleased to have the opportunity to respond to their comments. Unfortunately, there is far more in these remarks than I can possibly respond to in the space allotted. I will try to focus on the issues I believe to be most central to the issues which divide us.

1. BERMUDEZ

Bermudez begins with some questions about my commitment to the existence of person-level categories such as belief and knowledge. Cognitive ethologists share such a commitment, but, as Bermudez rightly notes, this hardly shows that knowledge and belief are natural kinds or that cognitive ethology is a science. Bermudez thus asks, "what grounds [my] confidence that eliminativism cannot be the best response to the methodology of naturalized epistemology. Do [I] 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?" [305]

My view about this is that our commitments here should be dictated by the practices of successful sciences. Work in cognitive ethology has provided us with illuminating explanations of animal behavior; appeal to talk of belief and knowledge has served to make sense of patterns of behavior in animals for which no other available explanations exist. This is, to my mind, evidence that the kind terms in these explanations refer. Bermudez raises the possibility that such explanations might be undermined by work in the cognitive

neurosciences, and I agree that this is in principle possible. Particular person-level explanations may be undermined by discoveries in cognitive neuroscience, and, far more radically, neuroscience could in principle undermine the entire explanatory picture provided by person-level explanations. But I have seen no reason to believe that these two levels of description must inevitably compete with one another, and I see no more reason to believe that cognitive ethology will be undermined by a successful neuroscience than that biology will be undermined by a successful physics.

Bermudez also raises the possibility – as do I – that viewing knowledge and belief as natural kinds may lead to radical revisions in our commonsense concepts. But while Bermudez acknowledges that there is a story here about how such revision may occur with terms such as 'gold' and 'water,' the case of 'knowledge,' he argues, is more problematic. What I need, in order to make this analogy work, is a rich body of canonical instances of knowledge to ground our use of the term, just as there are canonical instances of water which ground referring uses of that term, even in the face of widespread mistaken beliefs about the stuff itself. But, as Bermudez points out, "Disputes in epistemology, as in philosophy more broadly, tend to come about (or at least 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." [307]

I certainly agree that there are putative examples of knowledge which are so controversial that appeal to these examples can do little to ground a choice between two otherwise well-motivated theories. I have tried to steer clear of such examples. Instead, after arguing that there is a concept of knowledge which emerges from the cognitive ethology literature, I examine the concepts of knowledge on offer in the philosophical literature to see whether these might reasonably be seen as viable alternatives. The idea, after all, that there might be more than one kind of knowledge here – animal

knowledge and human knowledge, for example – has more than a little initial plausibility. It is for this reason that my argument against internalist accounts of justification attempts to show that internalists cannot account for the very instances of justified belief and knowledge which they themselves regard as canonical. If human knowledge is properly regarded as a kind in its own right, internalist accounts fail, I argue, to capture it. Similarly, I argue against those accounts of human knowledge which make social practices, such as giving and asking for reasons, an essential ingredient in knowledge. Here too I argue that instances of knowledge which, on these very accounts are regarded as canonical, are not properly explained by the rivals to my account. (I will have more to say about each of these moves in my response to BonJour.) But in neither case does my argument proceed simply by appealing to cases which rival accounts would categorize differently.

The heart of my disagreement with Bermudez, I believe, is not about this issue. Rather, Bermudez allows that if my argument against internalism is successful, then reflective knowledge, as, for example, Sosa understands it, "should not be our paradigm of knowledge." More than this, this shows that "our paradigm should be the conception of knowledge that we derive from the cognitive ethology literature...provided that one accepts two important theses. The first is that reflective knowledge derives simply from the addition of reflection to animal knowledge... [The second] 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." [308–9] Bermudez is right, I believe, to focus on this issue, and he is right, I believe, that it is far more controversial than it might initially seem.

Human beings have the ability to reflect on their first-order beliefs, an ability which other animals, it seems, do not. As I see it, the knowledge embodied in our first-order beliefs is of the very same kind as the knowledge which non-human animals have. More than that, our ability to reflect on our first-order beliefs, while it gives us a degree of cognitive and

behavioral sophistication which other animals lack, should not be seen as providing us with another kind of knowledge. When two creatures have different sense organs, for example, this allows them to form beliefs in different ways. But this, by itself, should not lead us to think that they have different kinds of knowledge. We may better understand the differences and similarities between such creatures if we regard them as having two different means to achieve one and the same sort of state. Similarly, I believe, with reflection. While the human ability to reflect allows us to form beliefs in ways which other creatures cannot, this should be understood as simply one more means to achieve the very kind of state which we have in common with other animals – namely, knowledge – rather than producing a state which is somehow different in kind.

As Bermudez sympathetically characterizes my view:

> ...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 structures of belief 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...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. [312–3]

Why does Bermudez see this differently? On Bermudez's view, our ability to reflect is dependent upon our linguistic ability, but, in addition, logic itself, he says, is dependent on language.

> Consider a disjunctive thought '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 being both false... ...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... [313–4]

So on Bermudez's view, the possibility of having disjunctive thoughts, and thoughts with logical structure generally, is dependent upon having very sophisticated conceptual

capacities: one must have the concepts of truth and falsity; one must, in particular, be able to think about truth-values, and this requires a language.

Thinking about truth-values requires more than just a language, of course. Young children are able to speak a language, but the concept of a truth-value is a late arrival in their conceptual development. What I don't see is why one should think that the possibility of having disjunctive thoughts is dependent on having the remarkably sophisticated conceptual abilities required for thoughts about truth-values. Indeed, young children learn to use the English word 'or' long before they have the concept of a truth-value, and they make inferences such as disjunctive syllogism long before they show an understanding of what a truth-value is. Sensitivity to simple logical structure does not require having thoughts about logical structure.

The same is true, I believe, in the case of non-human animals. The medievals referred to disjunctive syllogism as "the syllogism of the dog," and they referred to it in this way because, they claimed, if a dog is chasing another animal down a path and it comes to a fork in the path, it will smell one of the two paths to see if the animal has proceeded down that fork; if the result is negative, it will take the other route without bothering to smell it first. Disjunctive syllogism in action! The sophistication of animal cognition and the manner in which non-human animals respond to new information is easily explained if we suppose that the beliefs of these creatures are represented in a form which is sensitive to at least some features of their logical structure. Sensitivity to features of logical structure does not require beliefs about logical structure, nor does it require the concept of a truth-value.

Bermudez's case for viewing human knowledge as different in kind from the knowledge of other animals is dependent on his claim that sensitivity to logical structure requires the possession of an extraordinarily sophisticated conceptual repertiore: we can only have disjunctive thoughts if we can have beliefs about truth-values. But we should not accept this claim.

2. BONJOUR

After defending an externalist account of knowledge which is motivated by considerations from the cognitive ethology literature, I argue that it is knowledge in this very sense that we have all been talking about all along. In support of this claim, I present a number of arguments against internalist accounts of justification, where justified belief is seen as a necessary condition for knowledge. Knowledge does not, I argue, require internalist justification.

One sort of argument I present against internalism appeals to empirical evidence about the likely results of reflection on one's beliefs. The idea I have is quite simple. Internalists say that we are justified in a belief when that belief meets a certain condition: if after reflecting on a belief, we discover that it is supported by appropriate reasons, then the belief is justified. The psychological literature has a good deal of information about what happens when agents reflect on the reasons they have for their beliefs. I thought it would be interesting to see what the upshot of internalist reflection is likely to be in practice. Justified belief is supposed to be a thing worth wanting. What is likely to happen if we reflect in just the way internalists suggest we should? Would we thereby improve our epistemic situation? Would justified belief, in the internalist sense, prove to be something worth having? Is this the sort of thing that should be viewed as a necessary condition for knowledge?

BonJour discusses a number of the studies I present. Let me focus on just one. Take the case of the position effect: individuals shown a number of products which were, in fact, qualitatively identical, were asked which of the products they preferred. Their answers revealed a strong position effect: they had a marked preference for objects on the right. When asked what their reasons were for their preference, they cited just the sort of reasons people tend to have for preferring one product to another: it was better made. Many of them cited specific features of the product. Similar studies have been done by marketing firms on the effects of packaging. When

laundry detergent is put in containers of different colors, people who actually use the product report that the detergent in some of the packages yields far better results than what is, in fact, the identical product, packaged in a differently colored container. The reasons these people sincerely offer, and the reasons which they believe inform their beliefs, have nothing whatever to do with the factors which influence them. Study after study shows that human beings are not very accurate at reporting the source of their beliefs.

So what? I don't know exactly what a full-dress internalist account would look like of the reasons one would need to have to support a judgment that one product is better than another, but if there are ever such reasons, these people have them. Before reflection, these subjects had misguided beliefs. After reflection, of just the sort internalists say is sufficient for justified belief, these subjects have the same misguided beliefs, but they are now, in addition, more confident than they were before. What I argued is that this is not a rare or unusual case. Reflection of the sort internalists endorse is supposed to improve our epistemic situation, but, in a wide range of important cases, all it does is produce greater confidence in the beliefs we already have. If this is what justified belief is, I find it hard to see why it is something we should want.

What is BonJour's response to this case, and the others like it? BonJour suggests that, "It seems plausible to suppose that there would be a temporal gap between the position-influenced opinion and the inventing of supposed reasons (which cannot have been easy to come up with), and that careful reflection would reveal this. And even more obviously, the reasons invented would be extremely unlikely to stand up under careful scrutiny, since there were in fact no significant differences between the items in question." [324] In a different case, BonJour tells us that, if a subject reflects, he "will conclude properly that he has no good reason ...". [324]

The problem with BonJour's response here is that his view about what individuals will discover on reflection commits him to a very controversial empirical claim; indeed, all of the

available evidence supports the view that individuals will not discover these things on reflection. In addition, the reasons subjects offer are not in any way internally defective. While it is true that there is, in fact, no difference between the items inspected, this is a fact external to the agent's conscious reflection. When agents reflect, it seems to them that one of the products is better than the others, and the reasons they offer are of the very same sort that agents offer when there genuinely are differences among a number of products. There is nothing intrinsically wrong with the reasons offered; the problem is that these reasons have no basis in reality. But that is hardly something to which an internalist can appeal.

BonJour's confidence that reflection will lead to the discovery that the reasons were merely trumped up is contrary to the results which have been obtained in the numerous experiments which have been done on this phenomenon. BonJour states that these experiments do not show that subjects are *''incapable''* of discovering their mistakes by way of reflection, only that ''reflection *may* fail to uncover [these mistakes], especially in the short run.'' [326] He concludes, ''There would only be a problem for the internalist if it were shown that sufficiently careful and reflective subjects are unlikely [to spot their errors] even given a longer but still reasonable stretch of time...'' [327].

This issue is worth pursuing. First, many internalists see justification as a property of beliefs at an instant. They have a current time-slice view of justification. It is not surprising that this is the view many internalists defend, since their view derives from a Cartesian perspective which sees epistemological questions as fundamentally first-personal, but, more than this, which sees beliefs about one's *current* mental state as enjoying a certain epistemological privilege over beliefs about anything else. Any belief, if it is to be justified, must be justified from the perspective of the current instant. Descartes certainly held such a view. Many internalist epistemologists still do. BonJour himself seems once to have held such a view.

But BonJour no longer does. So let us consider what our experimental subjects would believe, or what they would discover on reflection, given "a longer but still reasonable stretch of time". At a minimum, the experimental evidence shows that reflection frequently will fail to make the relevant discovery even after an extended period of time. More importantly, subjects in these experimental situations do not have any internal clue that their reflection has somehow failed them. As far as they can tell from the first-person perspective, they have reflected carefully. Given this, what difference could it possibly make to an internalist if, were they to reflect still longer, they would discover their errors? Prior to this additional reflection, surely an internalist must hold that the belief in question is justified since it meets all the internalist requirements. If subjects do reflect for a longer period of time, and then discover that their justification is somehow defective, at that point, their belief will no longer be justified. But prior to that point, the internalist must allow that the belief remains justified. And for most subjects, their beliefs will remain justified for as long as it is held, since the subjects have no reason at all to think that further reflection is required or would be productive.

If this were a rare problem, touching only a few out of the way beliefs, I would not see this as a significant problem for internalists. But the studies I cite in the book show that this problem appears for a very wide range of the processes we employ when we reflect on our beliefs. We have a very striking tendency to see our beliefs as supported by reasons, and when we reflect, we are likely to take ourselves to have good reasons for our beliefs. So the problem remains: Why should we value reflecting on our beliefs if this is what it does for us? And why should we see such reflection as a necessary condition for knowledge? I can't see that we should.

Let me briefly address a second problem which BonJour raises about the kind of position I favor. BonJour raises the concern that the position I favor may be self-undermining. Consider the psychologists who have published the studies

we've just been discussing. BonJour asks, "First, do the investigators in question have any good *reasons*, reasons that they can be reflectively aware of and could in principle cite to others, for believing that the alleged results of the study are in fact probably correct? Second, do those same investigators have any good *reasons* for believing that the methods that were employed in the studies are in fact reliable? Does Kornblith himself have any good *reasons* for believing either of these things?" [329] Now BonJour acknowledges that the investigators, and I myself, may have perfectly good externalist reasons, if we may speak this way. That is, it may be that my beliefs, and those of the various psychological investigators, are in fact reliably produced. But these reasons, what BonJour calls "Kornblithian reasons," are not the kinds of reasons he's talking about. And no wonder. On BonJour's view, Kornblithian reasons are to reasons as rubber ducks are to ducks. What BonJour has in mind, of course, are internalist reasons. And, given my position, I of course want to deny that I *need* to have good internalist reasons in order to have knowledge.

What problem do I face if I fail to have internalist reasons? BonJour puts the objection quite clearly:

> Suppose that the author of one of Kornblith's studies puts forth its conclusion and, when asked for his reasons, offers various other things that he believes. If the inference from those other beliefs to the conclusion in question is reliable, the reason is a good one; if the inference is unreliable, it is not. But if neither the investigator nor the audience have any access to the fact of reliability or unreliability, it is impossible for any of them to tell which of these possibilities is actually realized.... This seems to me to show that while it is *possible* for an externalist to have Kornblithian reasons for his claims and even further Kornblithian reasons for the genuineness of his initial Kornblithian reasons, the whole structure is still from a critical standpoint indiscernible from one in which none of the reasons are genuine. [332]

Consider a belief of mine, say, the belief that p. And let us suppose, just to be charitable, that this belief of mine is reliably produced. Nevertheless, let us also suppose that I do not have internalist reasons for believing that p. I certainly don't

have introspective access to the process by which p was produced, nor do I have introspective or reflective access to the reliability of the process. Now BonJour says that if I don't have access to the reliability of the process – and access here means first-person introspective or reflective access – then it is impossible for me to tell whether the belief is reliably produced or not and thus, I can't really tell whether p is true, so none of my so-called reasons are genuine reasons. But I don't see why "being able to tell whether p is true" requires having this kind of reflective access to the reliability of the process by which the belief is produced. I can tell whether p is true if I am appropriately responsive to situations in which p is true. Nothing more than that is required. BonJour's suggestion that I "can't tell" whether p is true sounds like I'm in a very weak epistemic position. But notice that BonJour will insist that I can't tell whether p is true – in his sense – even in cases where I form the belief that p when and only when p is true (so long as I have no reflective access to the reliability of the process by which my belief is produced). Now this isn't what I mean when I say that someone can't tell whether p. I have in mind, instead, that such a person either is unable to form beliefs about whether p is true at all, or, if such beliefs are formed, they are very frequently wrong about whether p. Being unable to tell, in BonJour's sense, is compatible with having beliefs about p which are perfectly reliable. That doesn't sound so bad to me. Kornblithian reasons require reliable connection with the world. I can understand why I would want that, and what I would be missing if I lacked it. There seems an important distinction here between people who do, and those who do not, have this kind of connection to the world.

What about internalist reasons? BonJour says that Kornblithian reasons are not genuine because they are not internalist reasons. I am willing to countenance the existence of knowledge even without genuine (that is, internalist) reasons, but BonJour is not. Just how much knowledge, and just how much justified belief, is there on BonJour's view? How often

do we have good internalist reasons? In his recent introduction to epistemology, BonJour allows that it is impossible to justify our reliance on memory: that is, we do not have good internalist reasons for relying on memory. Such reliance, he argues, is presupposed by our justificational practice, but it is not something which we are in a position to provide good internalist reasons for.[1] To put the point somewhat differently, we do not have internalist reasons for any claim which requires the use of our memory. That's a pretty broad form of skepticism. But that's not all. In his recent book on the a priori, BonJour argues against radical empiricist views. According to BonJour, unless one countenances the existence of *a priori* justification – and this means the existence of beliefs which have an internalist justification independent of experience – then the result is "tantamount to the repudiation of argument or reasoning generally, thus amounting in effect to intellectual suicide."[2] It is thus particularly interesting to note that, well into that book, BonJour acknowledge that "both demonstrative and intuitive a priori justification turn out after all to have an externalist dimension... This is a stronger concession to exteralism than I have heretofore been willing to make, but one that seems required by the facts of the situation."[3] To put the point somewhat differently, we don't have internalist a priori justification, and the result is intellectual suicide. In short, BonJour offers us an account of what it is to have a reason according to which we can never have one.

The distinction between those beliefs which are supported by reasons and those which are not seems to me to be an important distinction, one we want to have a proper account of. But any account of what this distinction is which has the consequence that no one ever has a reason for his or her belief seems to me, in virtue of that very fact, extremely unlikely to capture the very phenomenon we sought to capture in the first place. My account of what it is to have a reason, and what it is to have knowledge, does not have this consequence. Internalism does. When BonJour rightly points out that knowledge in my sense can occur even without the

presence of genuine internalist reasons, I therefore do not think that this shows a shortcoming in my account of knowledge.

NOTES

[1] *Epistemology: Classic Problems and Contemporary Solutions*, Rowman and Littlefield, 2002, 184–5.
[2] *In Defense of Pure Reason*, Cambridge University Press, 1998.
[3] *Ibid.*, 128.

Department of Philosophy
University of Massachusetts-Amherst
Amherst, MA 01003-0525
USA
E-mail: kornblith@philos.umass.edu