

REASONS, NATURALISM AND TRANSCENDENTAL PHILOSOPHY

Hilary Kornblith

University of Massachusetts, Amherst

There is a view about knowledge, and about the nature of reason, which is widely held among naturalists. Knowledge, on this view, is a natural phenomenon.¹ Human beings, and many other animals, know a great many things, and it is in virtue of this knowledge that they are able to negotiate their environment so successfully. If we want to understand what knowledge is, and how knowledge is possible, then we need to examine this natural phenomenon in the very same way that we examine other natural phenomena: that is, empirically. In order to understand how knowledge is possible, we need to understand certain large-scale features of the mind, as well as certain large-scale features of the environment. When features of the mind dovetail in certain characteristic ways with features of the environment, knowledge becomes possible. Understanding the way in which our minds are structured to pick up information about the environment allows us to appreciate what makes knowledge possible, and to understand how our minds are responsive to reasons. Two brief illustrations will perhaps be useful.

Consider, first, the Chomskian revolution in linguistics. Chomsky argued that language is learnable for humans only in virtue of certain structural features of the mind.² The mind, on this view, has a modular structure. There is a language-learning module, outfitted with certain presuppositions about the structure of humanly learnable first languages. While such languages might, in principle, have had quite a different structure from the one they in fact have, any attempt to discern the structure in spoken language without making quite substantial presuppositions about that very structure would be doomed to failure. This is the poverty of the stimulus argument.

Language learning is possible for us only because the mind does not approach the problem without any presuppositions. Once we see how the mind is specially outfitted for the task of language learning, we see how the linguistic data which we make use of in coming to learn the structure of our first language provides us with all the reasons we need to comprehend the structure of our linguistic environment. The Chomskian picture, of course, is not supported by way of a priori argument. It is, instead, based on experimental evidence about the early linguistic environment of children, about the tempo and mode of language acquisition, about the structure of human languages themselves, and about the structure of the human mind and the human brain.

Consider, second, work on visual information processing.³ Here again, we see the modular structure of the mind at work.⁴ The visual system has certain presuppositions about the environment built into it, and it is only in virtue of approaching the task of visual information processing with these presuppositions about our environment that we are able to form any beliefs at all about the world around us. These presuppositions built into the visual system are not true of every logically possible environment, but they are largely true of the environment we inhabit. We presuppose, for example, that the world is largely inhabited by three dimensional objects having relatively stable boundaries, and it is in virtue of making this presupposition that the visual system is able to detect the geometrical structure of objects around us. It is also in virtue of making this presupposition that the visual system is subject to certain illusions, misperceptions of features of the environment on those rare occasions when the presuppositions are false. Our characteristic pattern of errors may thus be used to help reveal the presuppositions of the visual information processing system, thereby allowing us to understand how knowledge of the visual environment becomes possible for us. In coming to see how the visual system works, we come to understand how the data we take in visually

give us sufficient reason to form the beliefs we do—beliefs which are largely accurate about the environment around us.

While some of the details of these accounts are of interest only to linguists or perceptual psychologists, we may come to appreciate something about knowledge itself, and about how responsiveness to reason is possible, in seeing the general features which are common to these various accounts. On this kind of view, philosophy may ask questions which are more abstract than the special sciences, but the difference between philosophy and science is merely a matter of degree of abstractness, rather than any real difference in kind. An understanding of these abstract issues about the possibility of knowledge and the nature of reasons can only be achieved by way of empirical investigation.

There is, however, another way to view knowledge and the nature of reason, a way of looking at these matters that has its origin in Kant. We see this view in quite a number of contemporary philosophers, many of whom make plain their indebtedness to Kant, and many others on whom the Kantian influence is largely mediated by way of Wilfrid Sellars. Thus, for example, Michael Williams tells us that, “Knowledge is not a natural phenomenon.”⁵ And John Haugeland, endorsing much the same view, tells us that the capacity for knowledge is “not just a biological or ‘natural’ capacity.”⁶ These authors tend to see a very large difference between human and animal cognition, arguing not only that non-human animals are incapable of genuine knowledge, but, more than this, that they are incapable of having genuine beliefs.⁷ Many of these authors stress, in a Kantian spirit, the importance of our ability to reflect on our beliefs and our reasons for them. Our ability to reflect on our own mental states not only sets us apart from other animals, but it reveals the decidedly anti-naturalistic character of this approach. Thus, Richard Moran, in discussing the role that self-

knowledge plays here, remarks that “a non-empirical or transcendental relation to the self is ineliminable.”⁸ This special relation to the self, and its role in understanding knowledge and the nature of reasons, comes to the fore, on this view, when we recognize the importance of epistemic agency. As Christine Korsgaard comments, “It is because of the reflective character of the mind that we must act, as Kant put it, under the idea of freedom.”⁹

There are, of course, many differences among these authors, but there is, as well, a great deal that they have in common. There is a constellation of issues which these authors see as deeply connected, involving reflection, epistemic agency, normativity, and, in the end, an anti-naturalistic account of knowledge and the nature of reasons. In this paper, I attempt to give a sympathetic presentation of this anti-naturalistic picture, but I also want to explain why it is that I reject it.¹⁰

I

Let me begin, then, by sketching an anti-naturalistic picture of some of the distinctive features of human cognition.

Human beings differ from other animals in a number of fundamental ways. First, human beings are language users, while other animals are not. And second, leaving aside, for a moment, the question of whether non-human animals have genuine beliefs, it is quite clear that non-human animals are incapable of reflecting on their mental states and forming beliefs about them. Thus, human beings not only form beliefs about the world around them; they form beliefs about their own beliefs. Non-human animals cannot do this. This ability to reflect on our own mental states marks a crucial difference between humans and other animals.

What is so important about this ability to reflect on one’s own mental states? Let us suppose,

just for the sake of argument, and just for the moment, that non-human animals do have genuine beliefs. Then the difference we are discussing amounts to this: we humans have both first-order beliefs—beliefs about the world around us, for example—and second-order beliefs as well—beliefs about our first-order beliefs; non-human animals have only first-order beliefs. In the case of non-human animals, their first-order beliefs are produced by a variety of different mechanisms, and many of these mechanisms, no doubt, reliably deliver accurate information about the world around them, allowing these animals successfully to negotiate their environments. Of course, these belief-producing mechanisms are not perfectly reliable; they sometimes deliver false beliefs. And more than this, it may well be that some of these mechanisms are not reliable at all; they deliver false beliefs more often, perhaps far more often, than true ones.

Now we humans, of course, have mechanisms that generate first-order beliefs in just this way, and, just as in the case of other animals, many of these mechanisms reliably produce true beliefs, even if they do not do so infallibly. And probably many of these mechanisms are not only less than perfectly reliable; some of them are, in fact, simply unreliable. But while non-human animals are bound, once and for all, to form beliefs by way of these hard-wired mechanisms, human beings are not. Unlike other animals, we are capable of reflecting on our beliefs and the ways in which they are produced, and when we find that our beliefs fail to meet our standards, we are able to intervene in the belief-producing process, making changes in the ways in which we come to form our beliefs. It is thus the ability to reflect on our own mental states which allows for the possibility of changes in the processes by which our beliefs are formed.

This ability to reflect on our own beliefs, and thereby to engage in epistemic self-assessment, thereby makes room for epistemic agency. While other animals are both the unknowing

beneficiaries and the unwitting victims of their native processes of belief acquisition, we are active parties in the business of cognition. Non-human animal cognition is simply passive, but, in the human case, we are able to take an active role in molding our cognitive processes in our own image. It is for this reason that we may sensibly speak of human beings as epistemically responsible agents, and it is in virtue of this that we may hold humans responsible for the beliefs they hold. It is only because we are epistemically responsible that it makes sense to evaluate beliefs as justified or unjustified. While non-human animals may form beliefs which are true or false, the animals themselves play no role in determining how their beliefs are formed, and thus they may not be held responsible for the success or failures of their cognitive processes. Their belief forming processes may work well or badly, but it would simply be a mistake to think of the resulting beliefs as therefore ones which the animals are either justified or unjustified in holding. There can be no justification in animals who are not responsible for their beliefs. Only human beings may hold their beliefs justifiedly or unjustifiedly.

Since knowledge plausibly requires at least justified, true belief, we see that, strictly speaking, only human beings are capable of genuine knowledge. It is true, of course, that we often speak, informally, of non-human animals as knowing various things, but we also speak—obviously metaphorically—of various mechanical objects, which clearly do not even have beliefs, as knowing various things: the automatic door-opener “knows” when someone is approaching the door; the GPS system in your car “knows” when you have made a wrong turn; and so on. Genuine knowledge, however, is the exclusive property of human beings.

Indeed, once we see just how different human cognition is from what goes on in other animals, it starts to become clear that it is not merely talk of animal knowledge which must be seen

as metaphorical. Talk even of animal belief cannot be understood as literally correct. Animals certainly have information-bearing states which often accurately reflect features of the world around them, and it is for this reason that the metaphor of “animal belief” comes so trippingly off the tongue. But we should note that there is a very large difference between having genuine beliefs and merely possessing information-bearing states, even information-bearing states which are implicated in the control of behavior. Thus, consider the phenomenon of phototropism: plants are sensitive to the presence of sunlight. Internal states of plants reliably register the presence of light, and these internal states are instrumental in moving the plant’s leaves in ways so as to expose a larger surface area to the rays of the sun. While these internal states of the plant reliably register information about the presence of sunlight, no one should regard—and almost no one does regard—these internal information-bearing states as beliefs. So the mere fact that non-human animals have internal, information-bearing states which are implicated in the production of behavior cannot, by itself, give us reason to think that non-human animals have genuine beliefs.

In the case of plants, their information-bearing states may bring about certain motions, but these motions do not constitute actions on the part of the plant. The states of the plant which register the presence of sunlight are part of the “space of causes,” as Sellarsians like to put it; they are not part of the “space of reasons.” But now we see that the same is true of the information-bearing states of non-human animals: they too are causes of bodily motions, but they do not provide reasons for animals to act or to believe. In order for a state to provide a reason for action, or for a state to provide a reason for belief, it must be one which the agent is capable of regarding as a reason. Without the ability to reflect on their own mental states, non-human animals do not have this ability. They do not have the concept of a reason, or of a belief, and without these concepts, they therefore

cannot have either reasons or beliefs. Their internal information-bearing states form a network of mere causes. It is only in creatures capable of reflection, capable of standing back from their own first-order processes of belief acquisition, capable of self-assessment, and capable of epistemic agency, that internal information-bearing states may serve as reasons for belief or action, and thus that these states may be regarded as genuine beliefs. Human beings are thus not only the only genuine knowers; we are the only creatures with genuine propositional attitudes of any kind.

On this view, indeed, not all human beings are genuine knowers or even genuine believers. Neonates are not, nor are young children. They, like non-human animals, do not have the concept of belief, or of reason, or of truth, all of which are needed to have the kinds of second-order mental states needed to be epistemically responsible agents. But the relevant concepts are surely had by all normal human adults, and although we do not, of course, self-consciously review every act of belief acquisition to see that it measures up to our standards, nothing like this is necessary in order to be epistemically responsible. We do, on occasion, stop to ask ourselves whether a certain belief is really one which we should endorse, and, in doing so, we thereby exercise our epistemic responsibility. On some occasions, we come to the conclusion that a belief we already hold, or one we are naturally inclined to accept, is indeed worthy of our acceptance. But we also sometimes discover that such beliefs are not worthy. In such cases, we engage in acts of belief revision, or we resist a temptation to believe. More than this, we may undertake self-conscious acts designed to revise the ways in which we form beliefs in the future. We resolve to be more careful, for example, in evaluating evidence; we resolve not to be taken in by a well-spoken or attractive speaker without more carefully examining the reasons offered; we resolve to think things through rather than jump to conclusions; and so on.

As some authors emphasize, we should not think that the exercise of our epistemic responsibility is entirely, or even primarily, a private affair. We not only engage with reason when we reflect on the legitimacy of our beliefs, but in dialogue with others when we give and ask for reasons. We ask, and are asked, why it is that some particular belief is held, and in doing so, we hold each other responsible for our beliefs. Non-human animals, and very young children, do not engage in the activity of giving and asking for reasons, since they do not speak a language. This important dimension of epistemic responsibility is thus not only essentially social, but essentially linguistic.

Between private acts of deliberation and evaluation, and social acts of giving and asking for reasons, adult human beings are able to take charge of their cognitive faculties. Information processing in lower creatures is simply given for them by nature, but genuine cognition involves the kind of self-constitution which can only be found in epistemically responsible agents. Much of the way in which adult human beings come to take on new beliefs is shaped by their epistemic activity and is not simply determined by their biological natures. But even those processes of belief acquisition which survive both private reflection and public review should not be viewed, in adult humans, as merely due to the providence of nature. The very fact that it has survived such review makes us responsible for this part of our cognitive doings, just as much as the part which we actively initiate ourselves.

Cognition in adult human beings is thus far more complicated than the kind of information processing which may be found in non-human animals as well as young children. Human adult cognition crucially involves self-conscious deliberation, self-evaluation, the social practice of giving and asking for reasons, and thus genuine epistemic agency and epistemic responsibility, none of which are found in young children or other animals. Once we recognize this, we see that it is only

in the case of human adults that there is any kind of genuine engagement with reason, rather than the simple workings of some sort of causal mechanism. And it is for this reason that human cognition must be seen, not merely as something more complicated than the information-processing that goes on in lower animals, but as something which cannot be fully captured by a naturalistic world view.

It is this move, of course, which is crucial: the move from the space of causes—which is occupied by the information processing mechanisms present in non-human animals and young children—to the space of reasons—which is occupied only by adult human beings—is a move which shows the limitations of the naturalistic picture. Naturalists regard human cognition as just one more natural phenomenon, no different in kind from the information processing that goes on in other animals, or, for that matter, from the interaction between salt and water when the two are mixed and sodium and chlorine ions go into solution. From a naturalistic perspective, not only are chemical interactions to be regarded as revealing the operation of natural laws, but the same is true of the information processing that goes on in animals, and, most importantly, so too is the operation of human reason. Human reason is seen, on the naturalistic picture, as yet one more natural phenomenon among many, bound, like chemical interactions, by the operation of causal laws. But what the story I have been telling here is meant to reveal is just how misguided the naturalistic picture is. When information processing is reduced to the mechanical operation of the space of causes, we inevitably leave out what is distinctive of human cognition: our ability to take control of our cognitive lives through the exercise of our epistemic agency as revealed in the activity of reflecting on our own cognitive states and processes. It is only because we are epistemically responsible agents that we may be properly understood as engaging with reason at all. To put the point just a bit differently, but perhaps a bit more familiarly, the naturalistic picture of human beings,

in giving a thoroughly descriptive account of the nature of human cognition, inevitably leaves normativity out of the picture. But from the perspective of someone trying to understand what human reason is all about, there is a certain irony to be found here, for the naturalists thereby succeed in providing the kind of account they aspire to only at the cost of eliminating the very phenomenon—the workings of reason—which they seek to accommodate.

II

I hope that the anti-naturalistic picture I have just sketched is neither unfamiliar nor entirely unappealing. Its main elements draw on certain Kantian views about the nature of reasons, and, at the same time, on certain highly commonsensical views about the differences between human cognition and non-human information processing. Here, I would like to highlight what I see as the central claims of the view, and I would also like to make clear (in the footnotes) that I am not attacking a straw man of my own invention. We see important elements of this view in the work of Robert Brandom, Donald Davidson, John Haugeland, Christine Korsgaard, John McDowell, Richard Moran, and Michael Williams.

1. Adult human beings differ from non-human animals and young children in that the former, but not the latter, have the ability to reflect on their mental states and thus the ability to form second-order beliefs.¹¹

2. The mechanisms by which adult human beings form their beliefs change over time with the acquisition of new information, while the mechanisms of information-processing in non-human animals and young children are fixed once and for all.¹²

3. Points 1 and 2 above are not unrelated. Adult human belief-producing mechanisms may

change over time only because we have the ability to reflect on our beliefs and the ways in which we arrived at them. When we find that we have arrived at our beliefs in ways which we do not endorse, we are able to modify the ways in which we subsequently form our beliefs.¹³

4. The mechanisms of belief acquisition may thus change over time only in creatures who have the concept of belief, the concept of reason, and the concept of truth.¹⁴

5. Only adult human beings may thus be properly regarded as epistemically responsible. We alone are epistemic agents, active in the ways in which we arrive at our beliefs, rather than merely passive information-processors.¹⁵

6. Thus, only adult human belief is apt for normative assessment. Our beliefs alone may be justified or unjustified. Our beliefs alone may be responsive to reason.¹⁶

7. Thus, only adult human beings are capable of genuine knowledge.¹⁷

8. Because animal information processing is merely a matter of the workings of certain causal mechanisms operating within the animals, rather than the kind of responsiveness to reason one sees in adult human beings, it is a mistake even to regard non-human animals as having beliefs.¹⁸

9. Thus, while the kind of information processing which takes place in non-human animals may be fully understood by way of the sciences, adult human cognition is not a natural phenomenon.¹⁹

III

Let me now turn to critical discussion of these claims.

Claim 1—that only adult human beings have second-order beliefs—is common ground to all the authors under discussion. Ironically, although the literature in cognitive ethology strongly

supported this view until recently, there is now some reason to be more cautious in endorsing this claim.²⁰ But let us allow this point to stand for the sake of argument. Every last one of the remaining claims, however, should be rejected.

Thus, let us turn to the second claim—that the mechanisms of information processing in non-human animals are fixed once and for all, unlike in the human case. This claim is demonstrably false. There are, of course, some mechanisms like this in non-human animals. Famously, frogs are responsive to small moving objects very close to their faces.²¹ When flies come within close range, the frog's tongue lashes out, grabbing ahold of the fly, and the fly is then swallowed. Similarly, if a BB is rolled close to the frog's face, it will grab ahold of the BB, and swallow it just as well. While there is nothing at all noteworthy about this behavior, since we all make mistakes, it is worth noting that the frog will do this again if a second BB is rolled its way; and it will do it a third time; and a fourth; and so on. The frog simply doesn't learn from its experience with the BBs. Similarly, as Dean Wooldridge notes,²²

When the time comes for egg laying, the wasp *Sphex* builds a burrow for the purpose and seeks out a cricket which she stings in such a way as to paralyze but not kill it. She drags the cricket into the burrow, lays her eggs alongside, closes the burrow, then flies away, never to return. In due course, the eggs hatch and the wasp grubs feed off the paralyzed cricket, which has not decayed, having been kept in the wasp equivalent of deep freeze. To the human mind, such an elaborately organized and seemingly purposeful routine conveys a convincing flavor of logic and thoughtfulness—until more details are examined. For example, the Wasp's routine is to bring the paralyzed cricket to the burrow, leave it on the threshold, go inside to

see that all is well, emerge, and then drag the cricket in. If the cricket is moved a few inches away while the wasp is inside making her preliminary inspection, the wasp, on emerging from the burrow, will bring the cricket back to the threshold, but not inside, and will then repeat the preparatory procedure of entering the burrow to see that everything is all right. If again the cricket is removed a few inches while the wasp is inside, once again she will move the cricket up to the threshold and re-enter the burrow for a final check. The wasp never thinks of pulling the cricket straight in. On one occasion this procedure was repeated forty times, always with the same result.

The second claim amounts to the suggestion that information processing in non-human animals is always like this. This amounts to the suggestion that non-human animals do not learn from their experience.

We need not look to work with primates to see that this is not even close to the truth. Skinner's early work with rats and pigeons²³ shows that they are responsive to changes in their environment, and that they easily learn how to go about attaining a variety of rewards, displaying a subtle sensitivity to the ways in which their environment has changed. The mechanisms of learning in animals are complex and varied, as even the most conservative writers on this topic allow.²⁴ The simple reflexes we see in the case of the frog and the wasp are not the rule in animal information processing; they are not even the rule in frogs²⁵ or wasps.²⁶ Even in the case of fairly stereotyped behaviors—such as the broken-wing display in piping plovers, used to mislead would-be predators—new information may often be integrated with old in ways completely unlike the hard-wired fly-swallowing behavior of frogs and the striking perseveration of *Sphex*.²⁷ The fact of animal

learning has been well-documented for as long as there has been serious work on animal behavior. Animal learning requires the integration of new information with old, and this, in turn, makes itself manifest in the ways in which still later information is processed. All of this is quite prosaic. More interesting, and far more complex, are the phenomena of problem-solving and innovation.²⁸ But we need not examine such subtle phenomena in order to see that the manner in which animals process information may change over time.

All of this puts the lie, of course, to the third claim as well: that the manner in which adult human beings process information may change over time only as a result of our ability to reflect on our own mental states. Human beings, do, of course, sometimes reflect on their beliefs and the manner in which they came about. We do, as a result of such reflection, sometimes change the ways in which we subsequently reason. But just as non-human animals integrate new information with old without the need for second-order beliefs, human beings do the very same thing. To take a single example: when a student makes an appointment to talk with me in my office, I come to form the belief that the student will arrive in my office at roughly the appointed time. But if a particular student makes such an appointment and then fails to keep it, and then does the very same thing again, I no longer form the belief that this particular student will arrive in my office simply because he tells me that he will. When my expectations are defeated, I come to respond quite differently to being told that the student will show up for his appointment. This change in the way I respond—in particular, this change in the inferences I draw—is not typically accomplished by reflecting on my earlier mistake. I don't need to reflect on such things in order to stop drawing the conclusion that the student will be in my office. Instead, my inferential behavior simply changes in response to the information that the student has regularly failed to show up for appointments he has scheduled. To

suggest that a higher-level belief is required here in order for any inferential change to occur is to make the very mistake the Lewis Carroll so picturesquely warned us against.²⁹

It is also worth pointing out here that we surely tend to over-estimate the extent to which reflection serves as a driving force for epistemic change. It is not merely that epistemic change does not require the intervention of reflection on our beliefs and the manner in which they were acquired. Even when we do reflect on our first-order mechanisms of belief acquisition and retention, the manner in which our beliefs are in fact acquired is not transparent to introspection. The beliefs we form about our own mechanisms of belief acquisition are often inaccurate, and our attempts to monitor the ways in which we form our beliefs in order to improve the accuracy with which we reason very frequently results in greater self-confidence, but no greater reliability.³⁰ Introspection provides us with the illusion that we have an extremely active role to play in monitoring and controlling the ways in which we form beliefs, but, in fact, our reflective activity is often epiphenomenal with respect to the ways in which we reason.³¹

This bears, as well, on Claim 4, that the mechanisms of belief acquisition may change over time only in creatures who have the concepts of belief, of reason, and of truth. Notice just how radical this claim is. It is taken for granted by the writers under discussion here that such sophisticated concepts are not possessed by non-human animals, nor are they possessed by young children. This is not unreasonable.³² Indeed, the standard story about cognitive development in children has it that early on they have beliefs about the world around them; only much later are they able to form beliefs about mental states, let alone about reasons qua reasons, or about truth.³³ But if we acknowledge the conceptual limitations of non-human animals and young children while simultaneously insisting that great conceptual sophistication is required in order for the mechanisms

of belief acquisition to change over time, then one is left with an extremely unpalatable dilemma: either insist that, outside of adult humans, the mechanisms of belief acquisition never change—that is, that learning never takes place—which, as we have seen, is manifestly false; or, alternatively, adopt the still more radical view that there simply are no beliefs in any creatures other than adult humans. The writers under discussion adopt this more radical view. We will discuss this issue directly when we turn to Claim 8.

So let us turn to Claim 5, the suggestion that only adult human beings are epistemically responsible agents because only we are active with respect to our belief acquisition, while the information processing which takes place in children and other animals is entirely passive.³⁴ The picture we are supposed to endorse here is that first-order information processing—the sort that goes on in the absence of reflection on one’s own mental states, their origin, and the relationship of their contents to one another—is merely passive, and it is passive precisely because it is nothing more than a causal process that goes on within these limited creatures, rather than something that these creatures actually do. But if the mere fact that first-order information processing is a causal process—or perhaps the fact that it is subsumable under causal laws, as some writers insist³⁵—thereby makes it something passive, something that merely takes place in these creatures rather than something they do, then what exactly are we supposed to believe about the process of reflection in adult human beings? Is the process of reflection somehow acausal? Is it supposed to be something which somehow eludes causal laws? It is very hard to see how this could be so. Should we really believe that first-order information processing is firmly embedded in the causal structure of the physical world, but reflection on one’s mental states somehow takes place outside that causal structure? And where, precisely, is that? This would require that we endorse an extremely radical

metaphysical view. The motivation for such a view is not aided, of course, by the fact that psychologists have found the processes involved in reflection to be just as susceptible to empirical investigation as information processing involving first-order states.³⁶ Those who actually investigate the workings of reflection find that it is no different in kind, no less embedded in the causal structure of the world, than first-order information processing. Richard Moran's suggestion that there is something "non-empirical or transcendental" here, that second-order information processing has some special feature that makes it altogether different from, and less susceptible to naturalization than first-order processing, flies in the face of our best available theories. Reflection is no more or less active than first-order belief acquisition.

It is worth pointing out that the suggestion that the reflective/unreflective distinction tracks the active/passive distinction is not even *prima facie* plausible. While reflection is, at times, actively initiated, it is clearly something which we may simply find ourselves engaged in, something which thus goes on in us passively. And if reflection counts as active on those occasions when it is initiated as a result of voluntary activity, then the fact that first-order processes of visual scanning are often voluntarily initiated--for example, when we choose to turn our heads to look at something--should thereby make certain first-order processes of belief acquisition count as active as well.

This casts doubt, as well, on Claim 6, that only adult human belief is apt for normative assessment; only our beliefs may be responsive to reason. We should not think that responsiveness to reason requires beliefs about reasons. If in order to be responsive to A as a reason for believing B, one must not only believe A, but also believe that A is a reason for believing B, then in order for the belief that A together with the belief that A is a reason for B to be a reason for believing B, one would also have to believe that these two beliefs constitute a reason for believing B. An infinite

regress results, of the very sort alluded to earlier. Although forming beliefs about reasons is one way in which one might prove to be sensitive to reasons, it is not the only way. As the infinite regress argument shows, it could not be the only way. When my dog forms the belief that there is food in his bowl as a result of hearing the food making a characteristic sound as it strikes the metal of the bowl, the belief he forms is responsive to reason: the characteristic sound provides him with good reason to believe that there is food in his bowl, and he is demonstrably sensitive to this very reason. The many animals that engage in problem solving show a remarkably complex sensitivity to reasons. They recognize whether various problem solving strategies have succeeded or failed, and, when they have failed, they are responsive to this failure as a reason to try another strategy. This kind of responsiveness to reason does not require conceptualizing reasons as reasons, even if it is true that such a complex conceptual ability may provide one with greater and more subtle sensitivity to reason.

One might, of course, suggest some sort of debunking interpretation of animals and young children. While my dog responds to the sound of food being poured into his bowl, there are many reasons he is unresponsive to. Even the reasons he does respond to are ones which he responds to imperfectly. For some, this will suggest that talk of reasons here is inapposite: these considerations show that the dog is not ever responding to reasons. Instead, when we think about the proper explanation of animal behavior, we should simply adopt what Dennett calls “the design stance”³⁷: the dog was built, as it were, to respond in certain sorts of ways; he is simply responding in the way he was built to respond. No talk of reasons is really called for.

But if this is one’s reason for adopting the debunking explanation in the case of animals and young children, then one will be forced to offer a debunking account of reason itself. Adult human

beings are not sensitive to all the reasons there are for belief, and even the reasons we are sensitive to are ones we are sensitive to only imperfectly. Gamblers, notoriously, are terribly insensitive to all manner of reasons. And what is true of gamblers is true of the rest of us. When it comes to the debunking account of reasons, what's bad for the goose is bad for the gambler. Surely this argument proves too much.

One further point is worth making here. The suggestion that normative assessment of any kind, such as the assessment of beliefs as justified or unjustified, presupposes some sort of voluntary control is a suggestion which quite a number of authors have made, including many whose motivations are quite different from those of the authors discussed here.³⁸ The authors under discussion wish to tie epistemic assessment to agency, and then argue that only human adults have the requisite kind of agency. But any such move is far too quick. As many authors have now pointed out, there are all manner of cases of normative assessment which presuppose nothing whatever about agency.³⁹ The special case of talk of reasons at issue here is no different.

If information processing in children and non-human animals is thought of as nothing more than an assemblage of the kinds of mechanisms we see in the fly-snapping of frogs and the nest preparation of Spheks, then there is a good case to be made for the suggestion that such creatures should not be thought of in cognitive terms, and thus should not be regarded as responsive to reason. But as we have seen, cognition in animals and children is not at all like this, and the motivation for a debunking account of their apparent responsiveness to reason is undermined. Responsiveness to reason comes in degrees. Just as individual adult humans show a range of responsiveness to reason, different kinds of creatures show a range of such responsiveness. Animals who learn thereby demonstrate reasons responsiveness, and those who demonstrate sophisticated problem-solving skills

manifest quite subtle and skillful patterns of responsiveness to reasons. Conceptualization of reasons qua reasons, and the ability to reflect on one's own cognitive states and processes, manifests a greater cognitive sophistication still, but it is no prerequisite for sensitivity to reason, and thus no prerequisite for the aptness of normative assessment.

Claim 7—that only adult human beings are capable of genuine knowledge—is thus undermined as well. If we reject the suggestion that only adult human beings are fit subjects for normative assessment, as I have argued we must, then the basis for restricting talk of knowledge to such adults is thereby undermined, unless we endorse the still stronger and far more controversial claim—claim 8—that only adult human beings have beliefs.

So what is to be said on behalf of the claim that we should restrict the realm of believers to adult human beings? As John McDowell, who endorses this view, acknowledges, this does present a problem in explaining how it is that every one of us succeeded in making the cognitive transition from childhood to maturity. As McDowell comments,

Now it is not even clearly intelligible to suppose a creature might be born at home in the space of reasons. Human beings are not: they are born mere animals, and they are transformed into thinkers and intentional agents in the course of coming to maturity. This transformation risks looking mysterious.⁴⁰

It is worth expanding on this point. As I noted above, the standard account offered by developmental psychologists has it that children have beliefs, but, prior to approximately age four, they have no beliefs about their own mental states, and certainly no beliefs about reasons qua reasons, or about truth. The concept of enduring objects are ones which they have quite early on, and they form a robust set of beliefs about what is going on in their environment. It is only after coming to

understand a great deal about the physical world around them that they start to have the conceptual sophistication required to form beliefs about their own mental states. Prior to achieving this more advanced conceptual sophistication, however, a great deal of learning goes on. (Remember: we are talking about children prior to the age of four. Late in this period, they have not only learned a great deal about the world around them; they are also talking in extremely sophisticated sentences.) The kind of sophisticated learning and problem solving which goes on during this period is not only most naturally described in terms of the acquisition of beliefs, often by way of fairly complicated inference; no one has ever offered an account of such learning in any other terms.

If we insist, however, that we will not call anyone a believer until they have the more sophisticated conceptual scheme which begins to emerge at age four, then we will need to redescribe the learning that goes on prior to that point without talking about belief acquisition or inference. Now some authors, at this point, start talking about “proto-beliefs” or other sorts of primitive mental representations⁴¹, as if this simple terminological maneuver allows us to avoid the problem. But the strategy of renaming these states is not a solution to the problem, for the fact remains that there is no motivation in the phenomena themselves for regarding the kind of information processing which goes on prior to age four as different in kind than the kind that goes on afterward. More concepts are added to the child’s repertoire, but the kinds of mechanisms and states by way of which information is registered and processed remains the same.⁴²

So what is McDowell’s solution to this problem? If we insist that young children do not have beliefs and are not sensitive to reasons, how exactly do we account for the ever-increasing cognitive sophistication which we see during this period of pre-belief? And how do children make the transition from being creatures utterly lacking in beliefs and insensitive to reason to being creatures

who do have beliefs and are sensitive to reason? Here is what McDowell says:

A mere animal, moved only by the sorts of things that move mere animals and exploiting the sorts of contrivances that are open to mere animals, could not single-handedly emancipate itself into possession of understanding. Human beings mature into being at home in the space of reasons or, what comes to the same thing, living their lives in the world; we can make sense of that by noting that the language into which a human being is first initiated stands over against her as a prior embodiment of mindedness, of the possibility of an orientation to the world.⁴³

So it is language learning, according to McDowell, that allows for this transition. Now there are a number of problems with this. First, as I just noted, children learn a language prior to having the very concepts which McDowell and others see as a prerequisite for having beliefs. So the suggestion that language learning somehow ushers the child into the space of reasons seems to give the child beliefs just a bit too soon. Second, and more importantly, this suggestion seems to ignore, rather than address, the very real problem of accounting for all of the learning which goes on prior to the alleged acquisition of beliefs. And finally, the suggestion seems to be a failure even on its own terms. Even if we suppose, with McDowell, that there are no beliefs prior to language acquisition, and that the child is initiated into the space of reasons in the very acquisition of language, we have merely relocated the problem rather than solved it: what was once a problem for the individual child now becomes a problem for the origin of language itself. If children are only able to acquire intentional states because they are raised in an environment in which they can be initiated into language, exactly how did the first languages come about, since, as McDowell would have it, there were no beliefs prior to the existence of language, but one cannot make the transition from being a

creature without beliefs to being a creature with them without being embedded in a culture where language is spoken? If children can acquire beliefs only because they are enculturated by way of a “prior embodiment of mindedness,” then we seem to be committed to the existence of that prior embodiment all the way back in time. This is a rather high price to pay for insisting that children do not have beliefs until they are conceptually quite sophisticated.

Any suggestion that we are forced somehow to make sense of this by the fact that animal information processing is nothing but the operation of causal mechanisms, and *therefore* cannot involve either reason responsiveness or belief, surely runs afoul of the problem that it proves too much. Adult human belief acquisition is causally mediated; the science of psychology explains it, just as much as the information processing in children and non-human animals, by bringing it under the scope of psychological law. If this is sufficient to undermine the attribution of belief and sensitivity to reasons, then no one has beliefs at all. This may be congenial to eliminativists, but it was not where this argument was supposed to lead.

Finally, let me briefly address Claim 9, the suggestion that human cognition is not a natural phenomenon, that it may not be fully understood by being brought under the purview of the sciences, unlike the sort of information processing which goes on in other animals and young children. I believe the various claims which are meant to support this have already been addressed: that the divide between adult humans and others tracks the distinction between creatures who can learn and those who cannot; that it also tracks the distinction between epistemic agents and those who are not; that it also tracks the distinction between creatures who are responsive to reasons and those who are not, as well as the distinction between creatures whose states are subject to normative assessment and those who are not; and, finally, that it also tracks the distinction between those who have

knowledge and belief and those who do not. With all of these supporting claims undermined, we are left with no good reason at all to believe that human cognition is not a natural phenomenon, capable of being fully explained by the cognitive sciences. Given the great success of the enterprise of cognitive science, this should come as no surprise at all.

IV

I have tried to articulate the common core of a certain Kantian approach to human cognition, an approach which would remove human reason from the natural world and make it somehow immune to the progress of scientific understanding. I have argued that this approach makes a very large number of empirical presuppositions about the differences between the kinds of states and processes which are involved in non-human animal information processing as well as information processing in very young children, on the one hand, and the kinds of states and processes which are involved in adult human cognition, on the other. I have argued that every last one of the crucial empirical presuppositions of this view are false, and that there is thus very good reason to reject it. Human reason is well within the scope of the naturalistic project.

NOTES

1. I have defended my own version of this view in two books: *Inductive Inference and its Natural Ground*, MIT Press, 1993, and *Knowledge and its Place in Nature*, Oxford University Press, 2002.
2. For early presentations of this view, see *Syntactic Structures*, Mouton, 1957, and *Aspects of the Theory of Syntax*, MIT, 1965.
3. See, e.g., David Marr, *Vision*, Freeman, 1982.
4. The importance of modularity in understanding the large-scale structure of the mind was emphasized by Jerry Fodor in *The Modularity of Mind*, MIT, 1983.
5. "Is Knowledge a Natural Phenomenon?," in R. Schantz, ed., *The Externalist Challenge*, de Gruyter, 2004, 194.
6. *Having Thought: Essays in the Metaphysics of Mind*, Harvard University Press, 1998, 2.
7. On this issue, see also, of course, Donald Davidson, especially "Thought and Talk," in his *Truth and Interpretation*, Oxford University Press, 1984, 155-70, and "Rational Animals," in his *Subjective, Intersubjective, Objective*, Oxford University Press, 2001, 95-105.
8. *Authority and Estrangement: An Essay on Self-Knowledge*, Princeton University Press, 2001, 90.
9. *The Sources of Normativity*, Cambridge University Press, 1996, 94.
10. I have broached these issues before in two other papers: "The Metaphysical Status of Knowledge," *Philosophical Issues*, 17(2007), 145-64, and "The Myth of Epistemic Agency," manuscript.
11. This point is taken for granted by all of the authors listed above. It is certainly a commonsense view, and there is, as well, a good bit of support for it in the cognitive ethology literature, although none of the authors discussed here refers to that literature. Relevant work in ethology includes Daniel Povinelli and T. J. Eddy, "What Young Chimpanzees Know about Seeing," *Monographs of*

the Society for Research in Child Development, 61(1996), 1-152; Daniel Povinelli, *Folk Physics for Apes*, Oxford University Press, 2000; Michael Tomasello and Josep Call, *Primate Cognition*, Oxford University Press, 1997.

12. Thus, for example, Michael Williams remarks,

We might say, animals don't need the capacity for epistemic assessment because they don't test hypotheses: they test themselves. But this is why they are not truly sensitive to reasons. They cannot really change their minds, though the information-acquiring and processing capacities of the species can change over time. [*Op. cit.*, 207]

Williams surely implies here that, while "the information-acquiring and processing capacities of the species can change over time," they do not change within the life of an individual animal.

13. See the quote from Williams in note 12 above. But also see Robert Brandom, *Making It Explicit: Reasoning, Representing and Discursive Commitment*, Harvard University Press, 1994, 199-271, and *Articulating Reasons: An Introduction to Inferentialism*, Harvard University Press, 2000, 97-122; Christine Korsgaard, *op. cit.*, 90-130; John McDowell, *Mind and World (With a New Introduction by the Author)*, Harvard University Press, 1996, 108-126.

14. This is, of course, a central theme in Davidson's work (see especially the work cited in note 7 above). See also the work by Brandom, McDowell, and Williams cited in note 13.

15. See the work by Brandom, Korsgaard, McDowell, and Williams cited in note 13. This is also a central theme of Richard Moran's *op. cit.*

16. See, again, the works cited in note 13.

17. See, again, the works cited in note 13.

18. In addition to the works cited in note 13, and the works of Davidson cited in note 7, see

Haugeland, *op. cit.*

19. See the works cited by Brandom, Haugeland, Korsgaard, Moran and Williams. McDowell insists that he is not rejecting naturalism, but only an extreme form of it.

20. See, for example, Michael Tomasello, "Chimpanzees Understand Psychological States: The Question is which Ones and to what Extent," *Trends in Cognitive Science*, 7(2003), 153-6.

21. J. Y. Lettvin, H. R. Maturana, W. S. McCulloch and W. H. Pitts, "What the Frog's Eye Tells the Frog's Brain," in Warren McCulloch, *Embodiments of Mind*, MIT Press, 1965, 230-255.

22. *The Machinery of the Brain*, McGraw Hill, 1963, 82, quoted in Daniel Dennett, *Elbow Room: The Varieties of Free Will Worth Wanting*, MIT Press, 1984.

23. B. F. Skinner and Charles Ferster, *Schedules of Reinforcement*, Appleton Century Crofts, 1957.

24. For one conservative survey, see Sara Shettleworth, *Cognition, Evolution and Behavior*, Oxford University Press, 1998, 95-232.

25. See, for example, Darcy Kelley, "Vocal Communication in Frogs," *Current Opinion in Neurobiology*, 14(2004), 751-7.

26. H. E. Evans, *The Comparative Ethology of the Sand Wasps*, Harvard University Press, 1966.

27. Carolyn Ristau, "Aspects of the Cognitive Ethology of an Injury-Feigning Bird, the Piping Plover," in Carolyn Ristau, ed., *Cognitive Ethology: The Minds of Other Animals*, Lawrence Erlbaum Associates, 1991, 91-126.

28. See, e.g., Simon Reader and Kevin Laland, eds., *Animal Innovation*, Oxford University Press, 2003.

29. "What the Tortoise Said to Achilles," *Mind*, 4(1895), 278-80. I have discussed this problem further in section IV of "The Metaphysical Status of Knowledge," *op. cit.*, and in section II of "What Reflective Endorsement Cannot Do," *Philosophy and Phenomenological Research*, forthcoming.

30. I have argued for these points at length in chapter four of *Knowledge and its Place in Nature*.

31. Again, see *Knowledge and its Place in Nature*, chapter four, for details.

32. But see some reason for caution about mental state concepts in the work cited in note 20. No one that I know of has suggested, however, that non-human animals or young children have concepts of reason or of truth.

33. See, e.g., Alison Gopnik and Andrew Meltzoff, *Words, Thoughts, and Theories*, MIT Press, 1997; Janet Astington, Paul Harris, and David Olson, eds., *Developing Theories of Mind*, Cambridge University Press, 1988.

34. I have discussed this issue in detail in “The Myth of Epistemic Agency.” I summarize the results from that paper here.

35. See, e.g., John McDowell, *op. cit.*, Lecture VI.

36. See, for example, Timothy Wilson, *Strangers to Ourselves: Discovering the Adaptive Unconscious*, Harvard University Press, 2002. For a different approach, but one which makes second-order cognition no less metaphysically tractable, see K. Anders Ericsson and Herbert Simon, *Protocol Analysis: Verbal Reports as Data*, revised edition, MIT Press, 1993.

37. Daniel Dennett, *Brainstorms: Philosophical Essays on Mind and Psychology*, Bradford Books, 1978, esp. chapter 1. Dennett himself, of course, would reject this move.

38. Thus, for example, see the discussion of this issue in William Alston, “The Deontological Conception of Epistemic Justification,” *Philosophical Perspectives*, 2(1988), 257-99; and Alvin Plantinga, *Warrant: The Current Debate*, Oxford University Press, 1993, chapter 2.

39. For discussions of epistemic assessments which make this point, while disagreeing about much else, see Richard Feldman, “Epistemic Obligations,” *Philosophical Perspectives*, 2(1988), 235-56, “Voluntary Belief and Epistemic Evaluation,” in *Knowledge, Truth and Duty: Essays on Epistemic*

Justification, Responsibility and Virtue, Matthias Steup, ed., Oxford University Press, 2001, 77-92, and "The Ethics of Belief," *Philosophy and Phenomenological Research*, 60(2000), 667-95; and Hilary Kornblith, "Epistemic Obligation and the Possibility of Internalism," in *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility*, Abrol Fairweather and Linda Zagzebski, eds., Oxford University Press, 2001, 231-48.

40.*Op. cit.*, 125.

41. See, e.g., Robert Brandom, "Insights and Blindspots of Reliabilism," *Monist*, 81(1998), 391.

42. I have developed this argument in greater detail in "The Metaphysical Status of Knowledge."

43.*Op. cit.*, 125.