MURRAY CLARKE

EPISTEMIC NORMS AND EVOLUTIONARY SUCCESS*

Recent debates concerning the nature of epistemic justification primarily turn on two distinctions: the objective-subjective distinction and the internal-external distinction. Alvin Goldman has defended a reliable process account of epistemic justification that is externalist and objectivist in nature.[1] John Pollock, in contrast, defends an account that is both internalist and subjectivist in nature.[2] There are deep insights into the nature of epistemic justification contained in both views. Lately, Pollock has "scared up" more trouble in Tucson by attempting a novel, naturalized account of epistemic justification. Data from cognitive psychology and biology is, however, radically at odds with Pollock's project, as I shall attempt to show.

1. POLLOCK'S NATURALIZED INTERNALISM

The issue of what it means to be an internalist itself is controversial. Let me begin with what I take to be the usual construal of this distinction. The traditional internalist with respect to epistemic justification, I shall say, is one who holds that typically, a cognizer must give, or be able to give, reasons that warrant belief in a proposition if that person is to be justified in holding that belief. An externalist, on the other hand, is one who claims that a belief need only have the right sort of causal ancestry in order to be justified – for example, that it be the product of a mechanism that generally results in true belief. The cognizer need not provide, or be able to provide, reasons for belief in a proposition in order to be justified.

Standing in direct opposition to externalist accounts of justification are the more traditional accounts of C. I. Lewis and Roderick Chisholm.[3] Such internalists argue, with Plato, that one must be able to give "an account" if true belief is to be properly distinguished from Knowledge.[4] Otherwise, pure guesswork and luck would fit the bill. Obviously, that won't do, since true belief comes and goes, while knowledge, made of sturdier stuff, endures.[5] For some internalists

Synthese **85**: 231–244, 1990.

(though not Pollock) being able to give reasons is a necessary condition for justifiedness. Instead, Pollock argues in his recent paper 'Epistemic Norms' that: ". . . the justifiedness of a belief is a function exclusively of the internal states of the believer" (Pollock 1987, p. 70). By an internal state Pollock means ". . . those states that are directly accessible to the mechanisms in our central nervous system that direct our reasoning" (p. 70). Such states must be "directly accessible", and "the sense in which they are directly accessible is that access to them does not require us to first have beliefs about them" (p. 70). These directly accessible internal states are psychological states, ones that psychologists presumably must identify for epistemologists.[6]

We now have a fairly precise notion of the sense in which justification must be internal: the justification for a belief, i.e., a psychological state, must be directly accessible to a cognizer's production system governing reasoning in order for that belief to be justified. Pollock is arguing, in effect, that one need not be able to give reasons to be justified, but one must, minimally, have direct, non-epistemic access to those internal states that constitute our justification in order to be justified. Such access is non-epistemic in the sense that although one need not appeal to a belief, there must be a suitable psychological state accessed that guides one's reasoning when one arrives at a justified belief.

Results from cognitive psychologists suggest, however, that introspection is a poor method for detecting the etiology of belief acquisition.[7] Another worry concerns the nature of the direct, non-epistemic access that the mechanisms in our central nervous system have to our psychological states. Traditionally, justification had been thought to be internal to us; i.e., the access was epistemic and transparent to the cognizer. Pollock wants to say that it is the mechanisms in our central nervous system that have such non-epistemic, direct access to psychological states, not the self-reflective cognizer.[8] This is surely a difference that makes a difference, for the putative advantages of internalism over externalism have been thought to trade on the direct access to the cognizer of justificatory belief relations.

Whatever benefits that traditional picture entailed are now lost for Pollock's new naturalized type of internalism. For internal belief relations of that sort need not be directly accessible on his account. The crucial question here is: in what sense do psychological states' being directly accessible to internal mechanisms help in answering traditional questions about the nature of epistemic justifiedness? For we need to

be able to identify justified beliefs to provide an account of when we ought to adopt a belief. It is precisely this internal transparency of justifiedness that we lack on Pollock's account. Pollock's view here is that we can tell when we are reasoning properly by means of non-doxastic norms. Because the conformance to internalized norms is introspectible we can recognize instances of correct reasoning. For instance, in a case where we think something is red, the fact we are being "appeared to redly" constitutes a prima facie reason for belief. As he notes:

> This is not just an observation about what actually happens. It is an observation about what we know to do in judging colors, that is, an observation about how our automatic processing system actually guides us in reasoning about colors. It is the introspectibility of conforming to a norm that makes this observation possible (Pollock 1986, p. 172).

Pollock acknowledges that we do not know that we are in such states, since they are not belief states, but argues that this only seems odd if we implicitly assume a particular intellectualist model of the way epistemic norms regulate belief, one where norms must be explicitly articulated and then appealed to. We can tell that we are justified, even if we cannot dredge up reasons for belief, due to our procedural knowledge of epistemic norms. By analogy, Pollock offers the example of learning to ride a bicycle. At first the novice must be told "when the bicycle seems to be leaning to the right, turn the handlebars to the right". Eventually, this norm becomes internalized, the bicyclist now has procedural knowledge of an action-guiding norm. She can tell when she is driving properly and when she is not even though she has no conscious access to the reasoning going on. This automatic process requires three states of affairs:

(1) The Automatic Processing System governing reasoning has direct non-epistemic access to the relevant psychological states, i.e., seeming to be leaning to the right, balance sensations in the inner ear, and so on.

(2) The Automatic Processing System determines, i.e., reasons, that these psychological states fall under a certain norm, e.g., "When the bicycle seems to be leaning to the right, turn the handlebars to the right".

(3) Given (1) and (2), the appropriate unconscious inference is made resulting in a behaviour.

Epistemic norms are analogous to bike riding norms in the sense that both involve unconscious procedural knowledge, we can tell whether we are conforming to these norms even if we cannot explain how we know that we are conforming to them nor even say what these norms are. We display competence by virtue of our ability to discriminate cases of correct reasoning from incorrect reasoning, even if we have the odd bad performance. Pollock writes off the psychological literature that suggests that human belief acquisition procedures are often irrational. He accomplishes this feat, remarkably, in a single footnote (Pollock 1987, pp. 92–93). As Pollock notes,

> The short way with this charge is to note that if we did not know how to reason correctly in these cases, we would be unable to discover that people reason incorrectly. To say that we know how to reason is to invoke a competence/performance distinction (Pollock 1987, p. 92).

All that the drawing of this distinction requires is that ". . . we can, in principle, discover the errors of our ways and correct them" (p. 92). By simply drawing the competence/performance distinction, as mentioned above, all of our worries about the last two decades of research in empirical psychology are to be allayed. But can all of this literature be so quickly dispatched with such a tidy distinction? I doubt it. The issue here is: Can the logical or conceptual point that we must be competent reasoners in order to discern the errors of our experimental subjects prove that we, in fact, possess this cognitive capacity? An alternative would be to suggest that it is only a working assumption that we are competent to make these judgments. This working assumption is itself open to empirical doubt, though not, perhaps, to empirical refutation in any robust sense. I favor the latter analysis. Pollock, I think, elevates a working assumption to the status of something like a tautology. But the claim that we are competent reasoners, can in principle never achieve the status of a tautology anymore than can the ontological argument prove God's existence. As Bertrand Russell once noted about the ontological argument, "postulation has all the advantages of theft over honest toil". The same can be said, I think, about Pollock's defense of the competence/performance distinction. Equally, there is reason to be suspicious of Pollock's claim that we can introspectively determine the justifiedness of our beliefs by appeal to nondoxastic norms. I will address this latter issue in the next section while drawing on some of the psychological literature that Pollock finds questionable.

2. THE 'HALO EFFECT'

Consider the 'halo effect' in interpersonal perception. Nisbett and Wilson (1977) found that people not only failed to recognize important causal sources of their behaviour, but actually inverted the causal direction when asked to explain the origin of their beliefs. Subjects watched a videotaped interview with a Belgian psychology instructor who spoke English with a fairly strong accent. Half of the subjects were presented with a pleasant, agreeable, and enthusiastic individual talking about teaching practice; the rest saw that portion of the tape where he was like an "autocratic martinet", i.e., rigid, distrustful, and intolerant. When asked to judge his likability, those who saw the "warm" interview liked the Professor, those who saw the "cold" interview disliked him. The subjects were also asked to rate his physical appearance, mannerisms, and his accent. It was here that the 'halo effect' emerged. The "warm" instructor was rated attractive in these three categories, while with the "cold" instructor the verdict was irritating, despite the fact that the three features in question were objectively invariant across interviews.

Some subjects were then asked whether their liking the teacher influenced their rating on these three attributes, and the rest were asked the opposite question concerning their dislike of the teacher. The answer was a uniform "No". Like or dislike, they thought, did not influence their assessment of the three attributes. Interestingly, the subjects who saw the "cold" version reported that their dislike of the instructor's attributes caused them to dislike the instructor. As Nisbett remarks: "In short, it would appear these subjects inverted the true causal relationship. Their dislike for the "cold" teacher had made irritating the appearance, mannerisms, and accent that had seemed attractive in the same teacher behaving warmly" (Nisbett and Ross, p. 209). This example illustrates the point that introspection can be an unreliable guide to the causes and justifiedness of beliefs. Examples can be multiplied, the cumulative effect of which points one to the stronger claim that introspection is an unreliable guide with respect to the source and justification for one's beliefs. Rather than introspection being an infallible guide, as Descartes thought, and the nature of the external world being dubious, a Copernican shift has taken place.

Cognitive psychologists believe that they can give us a more objective analysis of our internal states than we can glean from introspection. If

they are right, the internalist dream of an account of justification turning on an appeal to directly accessible introspectible states is beyond our cognitive reach. The causes of, and justification for, the most innocuous beliefs turn out to be epistemically inaccessible by means of introspection. Pollock disassociates himself from Cartesian Internalists in not requiring that the cognizer have directly introspectible epistemic access to psychological states. For instance, one who held that we had to have reflective awareness of psychological states in order to be justified would stumble on a great deal of data from cognitive.psychology. Nisbett and Ross argue, for instance, that direct access to "psychological occurrences" in the central nervous system are as unlikely as would be direct access to the workings of the eyeball or the adrenal gland (Nisbett and Ross, p. 203).

Pollock does well, then, not to claim that we have epistemic access to psychological states. This is surely an advantage of his account over more traditional, Cartesian Internalist accounts of Justification. Pollock does hold, however, that we can very often tell when a belief is justified and when it is not, even if we cannot articulate the epistemic norm(s) that guide our reasoning. But the 'halo' example, and many others, clearly illustrate the point that we can err even in this respect. Pollock's minimal Internalist claim, that we can often tell when a belief is justified, cannot be assumed a priori, and it seems incompatible with the psychological literature. Since this is so, it follows that even if our beliefs are justified, this fact will often elude our cognitive grasp. I take this to be a serious problem for Pollock's account, for now we cannot tell when it is justifiable to hold a belief. And recall that Pollock claims that to explain how we can tell when our beliefs are justified is the most central of epistemological problems.

3. OPEN AND CLOSED INSTINCTS

Justification, for Pollock, consists in holding beliefs that conform to correct epistemic norms. The norms that actually underlie our reasoning are the correct ones and since reasoning is a natural process, Pollock concludes that his analysis is noncircular and naturalistic. Pollock is convinced that we cannot characterize justified beliefs by means of a single general property like reliability that underlies them but claims that it is the reasoning that supports justified beliefs that is important

(Pollock 1987, p. 85). I think that he is mistaken. To see why this is so, just ask yourself: in what sense is reasoning natural to humans? Beyond the bare fact that we engage in it, the answer is that reasoning is an adaptive feature that is crucial to our survival.

The story of our survival as a species is, in part, a story about the transition from the utilization of closed instincts to the utilization of open instincts.[9] Closed instincts are behaviour patterns that are fixed genetically in all respects – for instance, in the bee's honeydance, birdsongs, and the nest-building pattern of weaver birds. This genetic programming does for weaver birds what intelligence does for humans; it helps facilitate survival. Open instincts, on the other hand, are programs with a gap. Parts of such behaviour are innately determined, others are filled in by experience. The innate aspect for many animals includes strong tendencies such as to get home, to seek water, to hide by day and to avoid open spaces. The method used to facilitate these goals, however, is variable or open. Historically, it has been the case that the more complex a creature becomes, the more open instincts it has and the fewer closed instincts it requires. Hunting is an open instinct for cats. Cats, as it turns out, have a tendency to hunt and will do so even if deprived of all example. But the way a cat chooses to hunt varies with the context and creature. It is a skill that improves during its life; it invents new hunting techniques and can actually pick up tips from other cats. Hunting, for cats, is thus both innate and learned, so that the strict nature/nurture dichotomy breaks down. Learning and acquiring novel skills involves thought and reasoning of some primitive sort. In humans, these processes reach their apex, since humans have the fewest closed instincts of any species. The automatic closing of the eyes when a finger moves swiftly towards them is one of a small set of such closed instincts that, in part, links humans with their ancestors.

But what separates humans in degree from other species is the extent to which open instincts guide our behaviour and belief acquisition procedures. It is no accident that human brain size has grown remarkably in conjunction with the transition from closed to open instincts. Closed instincts exemplify precise, highly structured sequences of behaviours. The reliability of such sequences has been life-preserving for many species. In making the transition from closed to open instincts, human reasoning processes were called upon to become as reliable as the part of a closed instinct they replaced. Here one needs to distinguish

two senses of reliability. Closed instincts are survival-reliable, or S-reliable, in the sense that they make reproductive success possible. Open instincts are S-reliable too, but the cognitive processes that are, in part, constitutive of them are also True Belief-reliable, or T-reliable, in that such cognitive processes tend to produce true belief. Our survival as a species, then, depends on the T-reliability of a variety of processes that, in conjunction with open instincts, tend to produce true belief. In saying this, however, I do not mean to suggest that the utilization of open instincts always results in beliefs, for often, only a behaviour will result. The point is that behaviours that result from open instincts will often be, in part, a result of beliefs. Such beliefs, when justified, are justified due to their grounding in a variety of T-reliable processes. Without such accurate beliefs, our survival as a species would be remarkable or even impossible.

So I think that there is reason to hold that there is a single property in virtue of which we acquire justified beliefs and, ultimately, knowledge: namely, T-reliability. Reasoning, for humans, is T-reliable and, once combined with open instincts, has been as S-reliable as closed instincts for birds in producing survival-conducive behaviour; behaviour that depends on having justified beliefs. Pollock notes that, "Of course, unlike most norms our epistemic norms may be innate, in which case there is no process of internalization that is required to make them available for use in guiding our reasoning" (Pollock 1987, p. 63). Now there is no doubt that learning to ride a bicycle is just that, a learned skill. One should not think that most or even many of our norms are internalized in this fashion. But neither is it the case that most of our epistemic norms are innate. The mistake is to set up a false nature/nurture dichotomy. If we can believe ethologists like Lorenz and others, our behaviour is partially innate and partially learned. The procedure is mixed together; we express this insight by talking of open instincts.

Part of the attraction of externalist accounts of justifiedness is a function, if I am correct, of the ability of such accounts to explain a broader range of justified beliefs than internalists. Namely, those beliefs that are not justified by appeal to directly accessible, introspectible or nonintrospectible, reasoning states. The class of internally justified beliefs, at best, constitutes a small subset of such justified beliefs. The lion's share of such beliefs are justified by means of causal processes that are T-reliable; instinctive reasoning tendencies, perception, auditory,

tactile and other reliable sensory mechanisms. All of these mechanisms participate in larger behaviour sequences that we unite under the title 'open instincts'. These T-reliable processes also underlie noninstinctual belief procedures, and are constitutive of such procedures.

4. THREE OBJECTIONS

I now want to consider three objections. Here's the first objection: If Pollock's claims are jeopardized by the empirical data, then surely your reliabilist claims are likewise, jeopardized by other data pointing to the unreliability of human reasoning. I think the form of an answer to this question is now available to us. Reason was selected for in order to facilitate survival in our ancestor's primitive environment. In stark contrast, the experiments that psychologists employ often test our reasoning capacities in inhospitable locales where high culture, modern science, and complex social phenomena are involved. Little wonder, then, that we make reasoning errors in such artificial problem situations.

Unlike Pollock, however, I doubt that one can definitively settle the issue of whether human reasoning is a reliable process or not a priori. One can say this: from the evolutionary perspective the status of human reasoning is radically unlike the status of, for instance, the chin. The chin, it would seem, is an evolutionary by-product that was not selected for directly. Instead, it is an "inevitable architectural consequence" of the fact that other jaw structure features were selected for (Sober 1984, p. 24). In biological terms, the phenotypic characteristic that we call a 'chin' is likely the result of a pleiotropic gene, i.e., one that produces multiple phenotypic effects. It would, however, be an enormous surprise were human reasoning merely an accident of this kind. The same could be said about our perceptual abilities and other sensory modalities. With Darwin, we agree that the ability to do modal logic or calculus was not selected for directly, but more primitive reasoning skills surely did possess adaptive value. Michael Ruse has recently catalogued a wealth of empirical evidence that suggests that proto-reasoning occurs among the higher-primates, our nearest ancestors.[10] For instance, one chimp, Sarah, displayed analogical reasoning ability (Gillan 1981). Studies by Premack (1976) demonstrate the chimpanzee's ability to appreciate the concept of conservation. As Ruse notes, "A chimpanzee which had been trained to understand the concepts of

'same' and 'different' . . . had little difficulty in distinguishing cases where quantities of water or solid matter were altered in shape but not in size, from those cases where there were differences not only in shape but also in size" (Ruse 1986, p. 114). Moreover, despite superficial differences across cultures we find similar senses of logic and mathematics. This cross-cultural uniformity suggests biological roots (Staal 1967, p. 323; Bockenski 1961).

Though there is no direct correlation between cranial capacity and intelligence (or whales would be rocket scientists!), we do know that the earliest hominids, *Australopithecus afarensis*, had a 500 cc capacity while today humans average about 1400 cc (Ruse 1986, pp. 111–13). *A. Afarensis* lived some four million years ago, stood upright and was about two and a half feet tall. Given that the universe is some fifteen or sixteen billion years old, our cognitive growth has been remarkably fast, though not inevitable. Evolutionary theorists have generally not joined Herbert Spencer or Jean Lamarck in their progressionism. That is, there is no inevitability about humans getting better in any sense, whether one is talking about a movement from simplicity to complexity (Spencer) or less intelligent to more intelligent. We evolve due to the local constraints of our environment, by means of the random process of natural selection, mutation, and chromosomal recombination.[11] For this reason, what adaptation requires varies with the particular environment and its demands. As Ruse says: "Darwinian evolution is a string and sealing-wax process . . . . The human upper limb – one of the most wonderful adaptations possible – is no more than a transformed fin" (Ruse, p. 125). Moreover, I think it is plausible to suppose that the transition from the use of closed to open instincts took place during this same four-million-year period in conjunction with the increased manufacturing and use of tools as meat became a central source of protein for the hominids. But it is with the emergence of "modern man" from Neanderthals about 30,000 years ago that we see the real flowering of sophisticated tools, ornaments, and cave paintings. Culture had arrived (Ruse 1986, pp. 118–25).

But what, if anything, follows from all of this concerning the reliability of reason? First, it would seem that reason was selected for. Deductive, inductive, and analogical reasoning have been as beneficial as bipedality for humans. I take it that few would wish to challenge this point as long as one claims that it is the basic reasoning capacity

that one has in mind rather than specific propositional content. A more interesting, though also more difficult, question now presents itself: Was good reasoning selected for, e.g., reasoning that is truth-preserving such as first-order logic? Ruse argues for just this thesis when he says that "There are [secondary epigenetic] rules for approval of modus ponens and consiliences, no less than there is a rule setting up incest barriers" (Ruse 1986, p. 161). But one wonders whether this claim is a reasonable one. Such claims remain, at this stage of empirical research, unsubstantiated.

A third related question is: Why should we expect reasoning to produce truth even in its original domain of application if nature, typically, is a satisficer in other areas? As Quine once noted about induction.

> To trust induction as a way of access to the truths of nature, on the other hand, is to suppose, more nearly, that our quality space matches that of the cosmos. The brute irrationality of our sense of similarity, its irrelevance to anything in logic and mathematics, offers little reason to expect that this sense is somehow in tune with the world – a world which, unlike language, we never made (Quine 1969, pp. 125–26).

I am not sure that there are sufficient grounds for endorsing Quine's pessimism here. It is entirely possible, for instance, that considerations of simplicity favour truth over empirical adequacy at least with respect to non-inferential, perceptual beliefs and, perhaps, concerning inductively-based beliefs. The biological advantage of simplicity is that it makes for expediency in a complex, threatening environment. If humans had to actually think through a variety of possibilities before they crossed the street we would be, as they say in the gangster movies, history.[12] I do not think, however, that there is any open and shut argument for or against the claim that simplicity favoured truth as the output of reason in our ancestors' primitive environment or, for that matter, in our present environment. This issue remains an open question, but it is certainly compatible with evolutionary theory to suggest that basic processes are reliable in the sense of producing a high true-to-false belief ratio. And that is all that is required for the reliabilist agenda. The argument from instinct offered earlier goes farther than this and so must remain speculative however suggestive it may be.

At any rate Pollock, if I am correct, is unable to accurately describe the sense in which most of our beliefs receive their justification. It is,

contrary to what Pollock says, evident that an informed epistemology should not ignore relevant data from the sciences.[13]

NOTES

* I would like to thank Professor John Davis and Professor James J. Leach for providing philosophical inspiration to a somewhat timid undergraduate student at The University of Western Ontario, and many of his classmates, some years ago.

[1] The two main sources of this account are in 'What is Justified Belief?' and, for a more sophisticated treatment, Goldman's magnum opus *Epistemology and Cognition*.

[2] In what follows I shall rely on his most recent article 'Epistemic Norms' and another version of the same paper in his book: *Contemporary Theories of Knowledge*. But Pollock has, of course, defended a more traditional version of internalism in his *Knowledge and Justification*.

[3] See, for instance, Chisholm's *Perceiving: a Philosophical Study*, p. 16, and Lewis's *An Analysis of Knowledge and Valuation*, p. 27.

[4] See the *Meno* 98.

[5] In the language of cognitive psychology, it is true information that is stored in long term memory. See Goldman's 'The Relation Between Epistemology and Psychology', pp. 58–64, for more on this point.

[6] William Alston also identifies this as a reasonable analysis of what internalism ought to be construed as in his 'Concepts of Epistemic Justification', p. 78.

[7] For instance, see Hilary Kornblith's 'Introspection and Misdirection', (forthcoming in the *Australasian Journal of Philosophy*) where he catalogues a variety of ways in which introspection is a poor guide to the etiology of belief acquisition, drawing on data from cognitive science.

[8] At times, Pollock does talk of psychological states as being directly accessible to the *believer* (see page 70, line 13 of 'Epistemic Norms' and pages 172 and 174 of his book), but what he means is that such states are, nondoxastically, directly accessible to the production system governing reasoning, I think he makes this point fairly clear in 'Epistemic Norms', and on page 137 of his book, but see Hilary Kornblith's review (in *Philosophy and Phenomenological Research*) of Pollock's *Contemporary Theories of Knowledge* for discussion of this point. In particular, see note 4 at the end of that review where Kornblith quotes correspondence with Pollock that confirms my interpretation on this matter. Kornblith is surely right to point out that Pollock slides between the epistemic and the nonepistemic reading of direct accessibility; in the book. Either way, I think that Pollock's account of norms is deeply flawed.

[9] See Mary Midgley's *Beast and Man: The Roots of Human Nature*, pp. 52–53.

[10] Recent estimates put the human-ape split at about five million years ago (Ruse 1986, p. 114).

[11] But see Sober's *The Nature of Selection*, Chapter 6 for some qualifications and Holmes Rolston III's *Science and Religion* for a nonorthodox progressivist alternative.

[12] See Ruse, pp. 162–63 and Bach (1984), for more on these points.

[13] Ancestors of the present paper were read at the 1988 Western Canadian Philosophical Association Meetings; my thanks to Bruce Hunter for his comments on that occasion,

and at the 1989 Canadian Philosophical Association Meetings. My thanks go to Sheldon Wein for his comments on that occasion. Peter Miller delivered me from serious error in the early going. But I also wish to thank Bruce Freed, Hilary Kornblith and Tom Vinci for their written comments.

REFERENCES

Alston, William: 1985, 'Concepts of Epistemic Justification', *The Monist* **68**, 56–89.

Alston, William: 1980, 'Level-Confusions in Epistemology', in G. French, T. Uehling, Jr. and H. Wettstein (eds.), *Midwest Studies in Philosophy V*, University of Minnesota Press, Minneapolis, pp. 135–50.

Bach, Kent: 1984, 'Default Reasoning: Jumping to Conclusions and Knowing When to Think Twice', *Pacific Philosophical Quarterly* **65**, 37–58.

Bockenski, I.: 1961, *History of Formal Logic*, trans. I. Thomas, Notre Dame, Indiana.

Chisholm, Roderick: 1966, *Theory of Knowledge*, Prentice-Hall, New Jersey.

Clarke, Murray: 1986, 'Doxastic Voluntarism and Forced Belief', *Philosophical Studies* **50**, 39–51.

Clarke, Murray: 1987, 'Reliability and Two Kinds of Epistemic Justification', in Newton Garver and Peter H. Hare (eds.), *Naturalism and Rationality*, Prometheus Press, Buffalo, pp. 159–70.

Clarke, Murray: 'Knowledge and Reliability', forthcoming in the *Proceedings of the XVIII World Congress of Philosophy*.

Gillan, D. J.: 1981, 'Reasoning in the Chimpanzee: 2 Transitive Inference', *Journal of Experimental Psychology: Animal Behaviour Processes* **7**, 1–17.

Goldman, Alvin: 1979, 'What is Justified Belief?' in G. Pappas (ed.), *Justification and Knowledge*, D. Reidel, Dordrecht, pp. 1–23.

Goldman, Alvin: 1986, *Epistemology and Cognition*, Harvard University Press, Cambridge, Massachusetts.

Kahneman, D., P. Slovic and A. Tversky: 1982, *Judgement Under Uncertainty: Heuristics and Biases*, Cambridge University Press, Cambridge.

Kornblith, Hilary: 1988, 'Introspection and Misdirection', forthcoming in the *Australasian Journal of Philosophy*.

Kornblith, Hilary: 1988, Review of John Pollock's *Contemporary Theories of Knowledge*, *Philosophy and Phenomenological Research* **XLIV**, 167–71.

Lewis, C. I.: 1946, *An Analysis of Knowledge and Valuation*, Open Court, La Salle.

Midgley, Mary: 1978, *Beast and Man: The Roots of Human Nature*, Cornell University Press, Ithaca, New York.

Nisbett, Richard and Lee Ross: 1980, *Human Inference: Strategies and Shortcomings of Social Judgement*, Prentice-Hall, New Jersey.

Nisbett, R. E. and T. D. Wilson: 1977, 'The Halo Effect: Evidence for Unconscious Alteration of Judgement', *Journal of Personality and Social Psychology* **35**, 250–56.

Plato: 1961, 'Meno', in Hamilton and Cairns (eds.), *Plato: The Collected Dialogues*, Princeton University Press, Princeton, New Jersey.

Pollock, John: 1974, *Knowledge and Justification*, Princeton University Press, Princeton, New Jersey.

Pollock, John: 1986, *Contemporary Theories of Knowledge*, Rowman and Littlefield, New Jersey.
Pollock, John: 1987, 'Epistemic Norms', *Synthese* **71**, 61–95.
Premack, D.: 1976, *Intelligence in Ape and Man*, Lawrence Erlbaum, Hillsdale.
Quine, W. V. O.: 1969, 'Natural Kinds', *Ontological Relativity and Other Essays*, Columbia University Press, New York.
Rolston III, Holmes: 1988, *Science and Religion*, Random House, New York.
Ruse, Michael: 1986, *Taking Darwin Seriously*, Basil Blackwell, London.
Sober, Elliott: 1984, *The Nature of Selection*, MIT Press, Cambridge, Massachusetts.
Sober, Elliot: 1981, 'The Evolution of Rationality', *Synthese* **46**, 95–120.
Staal, J. F.: 1967, 'Indian Logic', in A. N. Prior (ed.), *History of Logic*; in Edwards (ed.), *Encyclopedia of Philosophy*, Collier Macmillan, New York, pp. 520–23.

Department of Philosophy
Concordia University
Montreal
Canada H4B 1R6