



Philosophical Review

The Naturalists Return

Author(s): Philip Kitcher

Source: *The Philosophical Review*, Vol. 101, No. 1, Philosophy in Review: Essays on Contemporary Philosophy (Jan., 1992), pp. 53-114

Published by: [Duke University Press](#) on behalf of [Philosophical Review](#)

Stable URL: <http://www.jstor.org/stable/2185044>

Accessed: 11/10/2011 03:47

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



Duke University Press and *Philosophical Review* are collaborating with JSTOR to digitize, preserve and extend access to *The Philosophical Review*.

<http://www.jstor.org>

The Naturalists Return

Philip Kitcher

1.

Ernst Haeckel, intellectual star of late-nineteenth-century Jena, continued a philosophical tradition by drawing on science to address the great questions of epistemology and ethics.¹ Haeckel would have been surprised to learn that one of his relatively obscure colleagues would help to overthrow that tradition. For many Anglo-American philosophers of our century, Jena is hallowed ground because it is the birthplace of *contemporary* philosophy. Frege's investigations are commonly viewed as a decisive turn, one that dethroned epistemology from its central position among the philosophical disciplines and that set the philosophy of language in its place.² In retrospect, we can trace a great lineage from Frege, leading through Russell, Wittgenstein, and Carnap to the professional philosophy practiced in Britain, North America, Australasia and Scandinavia in the postwar years.³ Distinguished by its empha-

¹Glimpses of Haeckel's ambitious philosophical ideas can be found in his *Anthropogenie oder Entwicklungsgeschichte des Menschen: Keimes- und Stammes-Geschichte* (Leipzig: Wilhelm Engelmann, 1874). The full version emerges in *Die Weltratsel: Gemeinverständliche Studien über monistische Philosophie* (Bonn: Emil Strauss, 1899).

²See Michael Dummett, *Frege and the Philosophy of Language* (New York: Harper and Row, 1973), especially chap. 19.

³The delineation of this lineage need not commit us to supposing that Frege's own influence was fundamental to the articulation of analytic philosophy. Undoubtedly, Frege's role in teaching both Wittgenstein and Carnap contributed to the early-twentieth-century reconfiguration of philosophy, but, as one of the editors reminded me, there were other important influences (such as Mach) and independent lines of developing similar positions (both Schlick and Russell arrived at their views with little influence from Frege). I emphasize Frege because he formulated so succinctly some of the cardinal theses of the emerging antinaturalism.

sis on logical analyses, the analytic movement, the “linguistic turn,” differs from earlier philosophical endeavors in its method as well as in its ordering of philosophical problems.⁴ For at least a period, philosophers could be confident of their professional standing, priding themselves on the presence of a method—the method of conceptual analysis—which they, and they alone, were trained to use.

Frege is the emblem of a revolution which overthrew philosophical naturalism, both in the hyperextended forms apparent in Haeckel, and in the more restrained versions of the early modern heroes, Descartes, Locke, Leibniz, Hume, Kant, and Mill.⁵ Pre-Fregean modern philosophy was distinguished not only by its emphasis on problems of knowledge, but also by its willingness to draw on the ideas of the emerging sciences, to cull concepts from ventures in psychology and physics, later still to find inspiration in Darwin.⁶ Frege’s opposition to what he perceived as intrusions from psychology or biology is evident from celebrated passages in

⁴See Richard Rorty’s introduction to *The Linguistic Turn* (Chicago: University of Chicago Press, 1967).

⁵See Hans Sluga, *Gottlob Frege* (London: Routledge and Kegan Paul, 1980); Michael Resnik, *Frege and the Philosophy of Mathematics* (Ithaca, N.Y.: Cornell University Press, 1980); and Elliott Sober, “Psychologism,” *Journal for the Theory of Social Behavior* 8 (1978): 165–91.

⁶These connections have become especially visible in some recent studies. See, for example, Gerd Buchdahl’s wide-ranging *Metaphysics and the Philosophy of Science* (Oxford: Basil Blackwell, 1969), and Patricia Kitcher’s *Kant’s Transcendental Psychology* (New York: Oxford University Press, 1990). A pioneering study of the connections between philosophy and the sciences in the seventeenth-century context is Maurice Mandelbaum’s *Philosophy, Science, and Sense-Perception* (Baltimore: The Johns Hopkins University Press, 1964).

For many early-twentieth-century epistemologists—Carnap, Schlick, and Reichenbach, for example—contemporary physics was an obvious source of inspiration and of problems. Yet, as I interpret these writers, there was an important difference between their use of science and that of Descartes, Locke, Kant, and Mill. Science, especially physics, supplied important instances for illustrating the power of an analytic apparatus (and testing the merits of different ways of applying that apparatus). But the apparatus itself, the machinery of logical analysis, was not informed by empirical researches, and, in particular, there was explicit disavowal of the need to use concepts and results from psychology in posing epistemological problems. Some later versions of logical empiricism and of ordinary-language philosophy went even further, sundering even the link to the sciences as exemplars of human knowledge.

the *Grundlagen*.⁷ The methodological stance he inspired becomes explicit in propositions of the *Tractatus*:

- 4.111 Philosophy is not one of the natural sciences.
- 4.112 Philosophy aims at the logical clarification of thoughts. . . .
- 4.1121 Psychology is no more closely related to philosophy than any other natural science. . . .
- 4.1122 Darwin's theory has no more to do with philosophy than any other hypothesis in natural science.

Twentieth-century historians of philosophy would ultimately reclaim the great early moderns by sanitizing their psychological and other scientific references. Leibniz, Hume, Kant, and the others emerged as analytic philosophers *manqués*—but, it must be conceded, extremely talented analytic philosophers for all of their psychological fumbling.⁸

In recent years, confidence in conceptual analysis and in “first philosophy” has begun to waver. Anglo-American philosophers have explored a wide range of disciplines, using ideas from psychology, biology, political science, economics, and the arts to reformulate traditional questions in epistemology and metaphysics.⁹ Some of their endeavors mark the return of epistemological naturalism, scorned by Frege and labeled as illicit philosophy by Wittgenstein. I shall examine some of the strands in contemporary naturalistic approaches to issues in epistemology and the philosophy of science. As we shall quickly discover, there are several kinds of naturalism. Yet all share an opposition to the Frege-Wittgenstein conception of a pure philosophy above (or below?) the special disciplines. I shall try to map some of the connections among versions of naturalism and, briefly, to motivate those naturalistic theses with which I am most in sympathy.

⁷See *The Foundations of Arithmetic (Die Grundlagen der Arithmetik)*, trans. J. L. Austin (Oxford: Basil Blackwell, 1950), v–vi, viii–x, 3, 38, 105.

⁸For some rightly influential studies that assimilate historical figures to post-Fregean philosophical practice, see Bertrand Russell, *The Philosophy of Leibniz* (London: Allen and Unwin, 1900); Jonathan Bennett, *Kant's Analytic* (Cambridge: Cambridge University Press, 1966); and Peter Strawson, *The Bounds of Sense* (London: Methuen, 1966).

⁹Use of material from the arts is relatively rare, but see Nelson Goodman, *Ways of Worldmaking* (Indianapolis: Hackett, 1978).

Frege's philosophical heirs may well find contemporary versions of naturalism in epistemology as shallow, scientistic, unphilosophical, and wrongheaded as Frege did.¹⁰ By the same token, naturalists might see the movement Frege inaugurated as an odd blip in the history of philosophy, a desertion of philosophy's proper task and proper roots. Be that as it may, it is important for any appraisal to have a clear view of what naturalism is. So that is where I shall start.

2.

Naturalistic epistemologies confront a range of traditional questions: What is knowledge? What kinds of knowledge (if any) are possible? What methods should we use for attaining knowledge, or, at least, for improving the epistemic qualities of our beliefs? Because the sciences appear to be shining exemplars of human knowledge, pursuit of these questions leads easily into the philosophy of science. Naturalistic philosophy of science emerges from the attempt to understand the growth of scientific knowledge.

Epistemological naturalism can be characterized negatively by its rejection of post-Fregean approaches to these investigations. For many Anglo-American philosophers from the 1930s to the present, the epistemological issues I have mentioned reduce to questions of logic, conceptual analysis, or "grammar."¹¹ Knowledge is viewed as a species of true belief, and a primary philosophical task is to specify and analyze the crucial "third condition." This is to be accomplished by identifying which *logical properties of and/or logical relations among propositions* suffice for justification (or for whatever other epistemic property is taken to turn mere true belief into knowledge). Similarly, proper responses to skeptical concerns about the possibility of knowledge devolve upon showing how propositions that formulate skepticism are either innocuous or else

¹⁰See, for example, Hilary Putnam, "Why Reason Can't Be Naturalized," and other essays in volume 3 of his collected papers (Cambridge: Cambridge University Press, 1983).

¹¹As Gila Sher pointed out to me, 'logic' is used in a very broad sense by many writers in the "analytic movement." Those who were influenced by the work of Frege and Russell on mathematics often seemed to think that any analytic work involving standard logical symbols and/or dummy letters should count as a contribution to logic.

overstep the bounds of grammar, how a skeptical idiom could not be employed in meaningful discourse.¹² Finally, with respect to the improvement of belief, analytic philosophers have yearned for a generalization of Frege's strikingly successful analysis of mathematical proof. An ideal logic of science, encompassing elementary forms of induction as well as appeals to simplicity, explanatory power, and other methodological desiderata, would enable us to emulate Frege's achievement across a broader range.¹³ All these enterprises have two important presuppositions: first, following both Frege and the Wittgenstein of the *Tractatus*, they pursue epistemological questions in an apsychologistic way—logic, not psychology, is the proper idiom for epistemological discussion; second, they conceive of the products of philosophical reflection as *a priori*—knowledge is to be given a “logical analysis,” skepticism is to be diagnosed as subtly inconsistent, the improvement of methodology consists in formulating the logic of science.¹⁴

¹²See, for example, A. J. Ayer, *Language, Truth, and Logic* (New York: Dover, 1952); J. L. Austin, *Sense and Sensibilia* (Oxford: Oxford University Press, 1962); and Peter Strawson, *Individuals* (London: Methuen, 1959). For penetrating analyses of the kinds of arguments invoked here, see Barry Stroud, *The Significance of Philosophical Scepticism* (Oxford: Oxford University Press, 1984).

¹³These projects are pursued in the seminal writings of logical positivism and logical empiricism, as well as in the related but distinctive approach of Karl Popper. Representative examples are Rudolf Carnap, *Logical Foundations of Probability* (Chicago: University of Chicago Press, 1950); Ernest Nagel, *The Structure of Science* (London: Routledge and Kegan Paul, 1961); Carl G. Hempel, *Aspects of Scientific Explanation* (New York: The Free Press, 1965); Israel Scheffler, *The Anatomy of Inquiry* (New York: Knopf, 1963); and Karl Popper, *The Logic of Scientific Discovery* (London: Hutchinson, 1958). Central to the idea of a “logic of science” are the contentions that methodological principles can be formulated in ways that emulate Frege's preferred mathematical idiom and that such methodological principles apply independently of subject matter. Attacks on the former contention have been familiar since the 1950s. More recently, some philosophers have begun to express doubts about the latter: see, for example, Elliott Sober “The Principle of the Common Cause,” in *Probability, Causation, and Induction*, ed. J. H. Fetzer (Dordrecht, The Netherlands: Kluwer, 1988), and, especially, Richard Miller, *Fact and Method* (Princeton: Princeton University Press, 1987).

¹⁴The connection between the apriority of philosophical conclusions and the view that these conclusions are the products of “logical analysis” emerges after Frege. Unlike Frege himself, many of his successors rejected the idea of any source of apriority other than language. But it is precisely because the two presuppositions I have identified are stated so clearly by

While contemporary naturalistic epistemologists disagree about which of Frege's two presuppositions is the real epistemological error, naturalistic epistemology, as I shall understand it, is committed to rejecting both.¹⁵ The rejection is part of a broader vision. Drawing on the deliverances of the sciences, naturalists view members of our species as highly fallible cognitive systems, products of a lengthy evolutionary process. How could our psychological and biological capacities and limitations *fail* to be relevant to the study of human knowledge? How could our scientific understanding of ourselves—or our reflections on the history of the sciences—support the notion that answers to skepticism and organons of methodology (or, indeed, anything very much) could be generated *a priori*? More conservative proponents of naturalism are prompted by these questions to reformulate traditional epistemological questions. Their aim is to reflect on the cognitive enterprise (including the ventures of science), on its history and on the capacities of those who participate in it, to achieve *corrigible* formulations of the goals of the enterprise and *corrigible* accounts of promising strategies for achieving those goals. Epistemology and philosophy of science, thus construed, attempt to fulfill traditional normative functions, and conceive themselves as continuous with the methodological reflections of scientific practitioners. Other forms of naturalism are more radical, seeing in the collapse of apriorism the demise of any possibility for normative appraisals (or, at least, the need for relativizing any such appraisals to a specific, local, context).¹⁶ So construed, epistemology and philosophy of science are descriptive dis-

Frege and underlie so much in twentieth-century epistemology that I have chosen Frege as the emblem of the turn away from naturalism.

¹⁵As will become clear below, writers such as Fred Dretske and Alvin Goldman are most concerned to reverse the Fregean ban on psychology within epistemology. Others, such as Larry Laudan, stress the impossibility of doing epistemology *a priori*, specifically without reflection on the ways in which historical figures have actually undertaken the project of inquiry. Some naturalists thus link epistemology more closely to psychology, others connect it with study of the history of science. My own version of naturalism will incorporate both connections, signaling a double departure from Frege.

¹⁶The denial of normative appraisal flows from the relativization of such appraisals, if one also accepts the idea that there are always available changes of context that would reverse any piece of normative advice. See section 8 below.

ciplines, chapters of psychology, neuroscience, sociology, or the history of science.

Before considering the feasibility of preserving the normative enterprise within a naturalistic framework, I shall examine, separately, naturalism's negative claims about the Fregean tradition. Both the reintroduction of psychology into epistemology and the suspicion of the *a priori* are well supported, and there is an important connection between them.

3.

Psychology re-entered epistemology quietly. A central problem in the analysis of knowledge takes for granted a conception of knowledge as justified true belief and seeks to provide an account of justification (foundationalist and coherence theories of justification being the main rivals).¹⁷ In 1963, a short article by Edmund Gettier called this enterprise into question by describing instances in which people have justified true belief but do not seem to have knowledge. Initial responses to Gettier's problem usually followed the apsychologistic orthodoxy, attempting to impose logical conditions on the subject's beliefs that would rule out the problematic examples as cases of knowledge.¹⁸ In the late 1960s, however, a number of authors proposed that a solution to Gettier's puzzling cases must lie in differentiating the causal processes that generate and sustain belief on those occasions where the subject knows.¹⁹

These generic approaches were articulated with the same kind of attention to detail that distinguished apsychologistic attacks on the Gettier problem. Yet, from a naturalistic perspective, their primary

¹⁷For concise presentation, see Roderick Chisholm, *Theory of Knowledge* (Englewood Cliffs, N.J.: Prentice-Hall, 1966; 2d. ed., 1977).

¹⁸Edmund Gettier III, "Is Justified True Belief Knowledge?" *Analysis* 23 (1963): 121–23. Two important early responses are Keith Lehrer, "Knowledge, Truth and Evidence," *Analysis* 25 (1965): 168–75, and Keith Lehrer and Thomas Paxton, Jr., "Knowledge: Undefeated Justified True Belief," *Journal of Philosophy* 66 (1969): 225–37.

¹⁹See Alvin Goldman, "A Causal Theory of Knowing," *Journal of Philosophy* 64 (1967): 357–72; Brian Skyrms, "The Explication of 'X knows that p'," *Journal of Philosophy* 64 (1967): 373–89; and Peter Unger, "An Analysis of Factual Knowledge," *Journal of Philosophy* 65 (1968): 157–70.

significance was their break with the apychologistic tradition.²⁰ Analyses of the concept of knowledge (and, later, that of justification) were no longer confined to specifying the logical relations among propositions believed by the subject but could take into account the processes, including inevitably the psychological sub-processes, that causally generate states of belief. By the mid 1970s a powerful argument for psychologistic epistemology had emerged. Take any set of favored logical relations among propositions that a subject believes. It is nonetheless possible that the subject lacks knowledge and lacks justification because the *psychological* connections among her states of belief have nothing to do with the logical relations. Thus, to take an extreme example, assume that a subject justifiably believes that p , justifiably believes that $p \rightarrow q$, and believes that q . It might seem that the belief that q must be justified because there is an elementary logical inference to q from propositions that are justifiably believed. Nonetheless, it is easy to understand that the causes of the subject's belief may have nothing to do with this elementary inference, that she fails to make the inference and believes that q because of some thoroughly disreputable generative process.²¹ Apsychologistic epistemology can struggle to accommodate such examples by invoking ever more complicated conditions on knowledge and justification, but the accumulation of epicycles serves only to disguise the fundamental point that the epistemic status of a belief state depends on the etiology of the state and, consequently, on psychological facts about the subject.

²⁰Here it is important to be careful about terminology. The epistemological movement within which Gettier's puzzle arose was apychologistic in two senses: it denied not only that the *findings* of psychology are relevant to epistemological questions but also that the *concepts* of psychology are needed to understand what differentiates cases of knowledge or of justified belief. The minimal break with apychologism consisted in suggesting that psychological concepts needed to be invoked in analyzing knowledge and justification. Stronger versions of psychologism, maintaining the relevance of psychological findings, would come later.

²¹Versions of this argument have been given by Gilbert Harman, in *Thought* (Princeton: Princeton University Press, 1973), chap. 2; Alvin Goldman, in "What is Justified Belief?" in *Justification and Knowledge*, ed. G. Pappas (Dordrecht, The Netherlands: Reidel, 1979), 1–23; and Hilary Kornblith, in "Beyond Foundationalism and the Coherence Theory," *Journal of Philosophy* 72 (1980): 597–612.

This argument by no means provided the only motivation for reintroducing psychology into epistemology. Philosophers discussing perceptual knowledge, notably David Armstrong and Fred Dretske, abandoned prevalent analyses that appealed to the “logical character of perceptual statements” or “the logical relations among beliefs” to suggest that one can see that p only if there is some lawful dependency between one’s belief that p and p .²² Perceptual knowledge depends on the right kind of relation between the knower and the facts known. Both Armstrong and Dretske later generalized this approach to full treatments of knowledge that made copious references to the characteristics of the psychological mechanisms of subjects.²³

Yet perhaps none of these internal philosophical developments was as important for the revival of epistemological naturalism as the contemporaneous changes that occurred within psychology itself. Against the background of the (behaviorist) psychology dominant in the 1940s and 1950s, epistemological talk of psychological mechanisms would have appeared not only contrary to Fregean gospel but also quaint. After Noam Chomsky’s trenchant review of B. F. Skinner’s *Verbal Behavior*, and especially after Chomsky’s subsequent development of his ideas about innate knowledge, cognitive psychology began to provide an idiom for discussing epistemological issues.²⁴ The burgeoning work of the 1960s and 1970s

²²See Armstrong, *Perception and the Physical World* (London: Routledge and Kegan Paul, 1961), and Dretske, *Seeing and Knowing* (London: Routledge and Kegan Paul, 1969). As Dretske has reminded me, the roots of the approach go back to H. P. Grice’s “The Causal Theory of Perception,” *Proceedings of the Aristotelian Society* 35 (suppl.) (1961): 121–68.

²³Armstrong, *Belief, Truth, and Knowledge* (Cambridge: Cambridge University Press, 1973); Dretske, “Conclusive Reasons,” *Australasian Journal of Philosophy* 49 (1971): 1–22, and *Knowledge and the Flow of Information* (Cambridge: The MIT Press, Bradford Books, 1981).

²⁴Chomsky, review of *Verbal Behavior*, *Language* 35 (1959): 26–58, *Aspects of the Theory of Syntax* (Cambridge: The MIT Press, 1964), and *Cartesian Linguistics* (New York: Harper and Row, 1966). Chomsky’s provocative claims were much discussed by philosophers. See, for example, the articles collected in *Innate Ideas*, ed. Stephen Stich (Berkeley: University of California Press, 1975). One side effect of the controversy was to make epistemologists far more familiar with some of the suggestions bruited by cognitive psychologists. The connection was further advanced by the writings of Jerry Fodor, particularly *The Language of Thought* (New York: Thomas Crowell, 1975).

on language learning, memory, perception, and problem solving gave epistemologists the opportunity to go beyond vague references to psychological mechanisms.²⁵ By studying the contributions of their colleagues in cognitive psychology they could sometimes offer quite detailed speculations about what these mechanisms might be.

Nonetheless, the post-Gettier arguments and the descriptive work of cognitive psychologists might seem to require only minimal commitment to naturalism. Apparently, it would still be possible to maintain that the substantive work of epistemology and philosophy of science lies in specifying the kinds of processes that *ought* to be instantiated in the psychological lives of subjects. Arguably, Frege, if not Wittgenstein, thought about knowledge and justification in this way.²⁶ Frege's deductive logic, the inductive logics sought by his successors, the proposals of contemporary Bayesians, might all be regarded as descriptions of ways in which people *ought* to think.

I believe that the naturalist point about the reintroduction of psychology into epistemology cannot be so easily accommodated.²⁷

²⁵For a small sample: P. C. Wason and P. N. Johnson-Laird, *Psychology of Reasoning: Structure and Content* (London: Batsford, 1965); A. Newell and H. Simon, *Human Problem Solving* (Englewood Cliffs, N.J.: Prentice-Hall, 1972); Irwin Rock, *The Logic of Perception* (Cambridge: The MIT Press, 1983); David Marr, *Vision* (San Francisco: W. H. Freeman, 1982); Susan Carey, *Conceptual Change in Childhood* (Cambridge: The MIT Press, 1985); Stephen Kosslyn, *Image and Mind* (Cambridge: Harvard University Press, 1981); and Elizabeth Loftus, *Eyewitness Testimony* (Cambridge: Harvard University Press, 1979).

²⁶See the introduction to *The Basic Laws of Arithmetic (Grundgesetze)*, trans. Montgomery Furth (Berkeley: University of California Press, 1965), and many of the papers on the nature of logic in the *Nachgelassene Schriften*, ed. Hans Hermes et al. (Hamburg: Felix Meiner, 1969). I have elaborated an interpretation of Frege along these lines in "Frege's Epistemology," *Philosophical Review* 88 (1979): 235–62.

²⁷Here we move to the stronger version of psychologism, mentioned above in note 20. The difference is clear in Goldman's writings: earlier essays such as "A Causal Theory of Knowing" and "Discrimination and Perceptual Knowledge" (*Journal of Philosophy* 73 [1976]: 771–91) depart from apsychologistic epistemology only by claiming the relevance of psychological notions for the analysis of epistemic concepts; in later writings, "Epistemics: The Regulative Theory of Cognition" (*Journal of Philosophy* 75 [1978]: 509–23) and *Epistemology and Cognition* (Cambridge: Harvard Uni-

What grounds the claim that our favored logical principles are prescriptions for thought? What are the sources of these principles? Do such idealized recommendations really apply to us? A traditional response is to propose that they present conceptual truths about rationality, thereby formulating an ideal at which we aim. For naturalists, however, such prescriptions must be grounded in facts about how systems like us could attain our epistemic goals in a world like ours. Simply asserting that prescriptions unfold our concept of rationality will be beside the crucial point.

The difference in attitude emerges clearly in some twentieth-century discussions of Hume's problem of induction. Some writers have suggested that adopting the inductive practices and principles that we do is constitutive of our concept of rationality.²⁸ But why should we treat our current concept of rationality as privileged? Communities with different practices and principles could mimic our reply to Hume, declaring that their inductive strategies were constitutive of their concept of rationality.²⁹ The real issue is whether employment of our inductive practice and principles is likely to lead us in the direction of our epistemic goals (most obviously truth).

The failure of appeals to conceptual truth, to analyticity, is fully general.³⁰ If an epistemological theory tells us that a particular policy of belief fixation is justified or a particular type of inference is rational, and that these claims are analytic, that they unfold our concepts of justification and rationality, an appropriate challenge is always, "But why should we care about these concepts of justification and rationality?"³¹ The root issue will always be whether the

versity, 1986), Goldman argues for the relevance of psychological results to the formulation of epistemological principles.

²⁸See, for example, Strawson, *Introduction to Logical Theory* (New York: Wiley, 1952), and Carnap, *Logical Foundations of Probability*.

²⁹This point has been made by a number of authors, but with especial vividness by Brian Skyrms, in *Choice and Chance*, 3d. ed. (Belmont, Calif.: Wadsworth, 1986), chap. 2.

³⁰See Stroud, *The Significance of Philosophical Scepticism*, and Harman, "Quine on Meaning and Existence I," *Review of Metaphysics* 31 (1967): 124–51.

³¹Sometimes, of course, the challenge can be turned back, for, when we specify our broader goals, we may be able legitimately to claim that there

methods recommended by the theory are well adapted for the attainment of our epistemic ends, and that cannot be settled by simply appealing to our current concepts.³²

Traditional epistemology has an important meliorative dimension. Bacon and Descartes were moved to epistemological theorizing by their sense of the need to fathom the ways in which human minds can attain their epistemic ends.³³ If analysis of current concepts of rationality and justification, or delineation of accepted inferential practices, is valuable, it is because a clearer view of what we now accept might enable us to do better. Conceptual clarification has a role to play in the advance of inquiry, even when we understand that our current concepts might give way to improved ones. (Modern logic was born in Frege's attempt to expose the structure of inferences underlying mathematical proof, even though what Frege exposed turned out to be inconsistent.) But what exactly do we want? Advice for any reasoning being—for "reason itself"—that would be good no matter what the world is like? Or advice for limited creatures like ourselves that would be effective in the actual world?

As we shall discover, when naturalists claim that epistemological principles offer advice about how to succeed in our world and that this advice is based on what we believe about that world, they invite skeptical objections. Ambitious recipes for universal success would be welcome, but, as we shall see in the next section, there is no alternative to appealing to empirical information. Analytic epistemology either idealizes so far from the human cognitive predic-

are no available rival concepts that would serve our purposes better. See, for example, the discussion of eliminativism with respect to the idioms of "folk psychology" in section 9 below.

³²The essential point goes back over fifty years to W. V. Quine's "Truth by Convention," reprinted in *The Ways of Paradox* (New York: Random House, 1966). It was already prefigured in Kant's responses to Eberhard. See Henry Allison, *The Kant-Eberhard Controversy* (Baltimore: The Johns Hopkins University Press, 1973).

³³See Bacon, *Novum Organon*, ed. Fulton Anderson (Indianapolis: Bobbs-Merrill, 1960), and Descartes, "Regulae," in *The Philosophical Writings of René Descartes*, ed. E. S. Haldane and G. R. T. Ross (Cambridge: Cambridge University Press, 1911–12). The theme is taken up within a naturalistic context by Goldman, "Epistemics: The Regulative Theory of Cognition."

ment that its deliverances are unhelpful, or it tries to disguise substantive principles about how to proceed in a particular kind of world (that which we take ourselves to inhabit) as if they offered universal recommendations.

The most prominent contemporary versions of naturalism formulate the meliorative epistemological project in terms of enhancing the *reliability* of the cognitive processes we employ. The standard invoked is familiar from the most notable psychologistic attempts to resolve Gettier's problem, those of Alvin Goldman and Fred Dretske.³⁴ In Goldman's original version, a process that confers justification is reliable in the sense of belonging to a type that generates true beliefs with high frequency.³⁵

Simple versions of reliabilism run into trouble when they are viewed as potential analyses of knowledge and justification. For someone may come to believe that p by a process of a type that regularly generates true beliefs, but the person's undergoing that process may be caused in quite bizarre ways.³⁶ One may try to block such counterexamples by insisting that, in addition to the reliability of the belief-generating process, the subject must have a second-order disposition to be moved to belief by a reliable process, a disposition that was activated on the present occasion.³⁷ Although imaginative philosophers may be inspired to offer more baroque examples and further refinements, these will be quite irrelevant if our goal is not to analyze our current concepts of justification but to pursue the meliorative project.

Reliabilism offers a plausible standard against which candidate

³⁴See Dretske, *Knowledge and the Flow of Information*, and Goldman, *Epistemology and Cognition*. Both books further elaborate proposals that the authors had made in earlier articles.

³⁵See "Discrimination and Perceptual Knowledge." This version is greatly refined in *Epistemology and Cognition*, and in some of the previously unpublished essays in *Liaisons* (Cambridge: The MIT Press, Bradford Books, 1992).

³⁶In Goldman's original example, a student learns a reliable algorithm from an unreliable teacher. See *Epistemology and Cognition*, 51–52.

³⁷This is Goldman's own line of solution. A different approach is recommended by Christopher Peacocke, *Thoughts: An Essay on Content* (Oxford: Basil Blackwell, 1986), and David Papineau, *Reality and Representation* (Oxford: Basil Blackwell, 1987).

belief-generating processes can be judged. As it stands, it needs qualification and clarification. In the first place, we need a broader conception of cognitive virtue. Attainment of truth can be trivial, and merely settling for processes that yield truth would slight other cognitive ends.³⁸ The problem of providing an account of our epistemic goals comes to the fore within naturalistic epistemology.³⁹ Second, until we have specified the class of contexts within which candidate processes are to operate, the demand that there be a high frequency of delivering the cognitive good makes no sense.

We can easily resolve the vagueness of reliabilism and formulate an exacting standard for appraising cognitive performance. Say that an agent's formation of a belief is *externally ideal* just in case that belief was generated by a process that, among all the processes available to the agent in his context, was of a type whose expected epistemic utility was highest. Here the notion of expected epistemic utility is parasitic on an account of cognitive goals and on an assignment of frequencies of success within a contextually determined class of situations. The meliorative project is to identify processes that are externally ideal. But naturalistic epistemology allows for the possibility that the ideal standard is sometimes (even always) at odds with our own criteria for justification, and that the processes that we undergo are not those that would accord with the external standard.

There is room for other concepts of rationality and justification.⁴⁰ The obverse of the point that someone might fortuitously satisfy the external ideal is that a cognitive agent may do the best

³⁸See Popper, *The Logic of Scientific Discovery*; Isaac Levi, *Gambling With Truth* (Cambridge: The MIT Press, 1973); and Hartry Field, "Realism and Relativism," *Journal of Philosophy* 79 (1982): 553–67.

³⁹It should however be noted that the problem is considered by some epistemologists who are not naturalists, notably Chisholm (*Theory of Knowledge*), Popper (*The Logic of Scientific Discovery*), and especially Levi (*Gambling With Truth*).

⁴⁰Most obviously by taking a "satisficing" rather than a "maximizing" approach, claiming that justification requires expected epistemic utility to be above some threshold value. This kind of approach is prominent in the writings of contemporary reliabilists. See, for example, Goldman, *Epistemology and Cognition*.

she can and still fall short. The external ideal is worth aiming at. But if we reject the idea that epistemological principles specifying how to attain that ideal can be generated *a priori*, then the history of attempts to formulate such principles might reveal a succession of improving views about how to investigate the world, *none* of which satisfies our exacting standard. Nevertheless, we do not want to treat ourselves and our predecessors in an undifferentiated way, simply remarking that all are methodologically imperfect.

Attributions of justification and rationality stem from the idea that the epistemic performances of subjects may be appraised whether or not the beliefs they acquire are true. For unlimited beings such attributions would be pointless: epistemic performance would simply be assessed by the attainment of truth (more exactly, the epistemic good). Cognitively limited beings, however, can do well or badly in trying to overcome their limitations. We cannot think of them as limited only with respect to “matters of fact”; their perspective on how to proceed in forming their beliefs may also be limited. Thus, just as we excuse ourselves and our predecessors for failure to be omniscient, concepts of rationality and justification *used in assessing the performances of others* should also take into account our methodological foibles.⁴¹

The practice of epistemic appraisal is highly ambiguous. Those who fail to take prior probabilities into account when dealing with Bayesian problems fall short of the external ideal, but they may, nevertheless, do the best they can.⁴² A conception of rationality that allows for human limitations may still view their performance as rational. By the same token, the Aristotelians who refused to acknowledge either Galileo’s arguments or the methodological principles underlying them may be appraised as justified, at least in one sense of the term: their own processes of reasoning accepted time-honored methods for forming and evaluating beliefs, and,

⁴¹In his recent writings, Goldman arrives at similar conclusions. See, for example, his “Epistemic Folkways and Scientific Epistemology,” in *Liaisons*, and “Strong and Weak Justification,” in *Philosophical Perspectives*, vol. 2, ed. James Tomberlin (Atascadero, Calif.: Ridgeview, 1988), 51–69.

⁴²This point is akin to one made by Hilary Kornblith, who stresses the importance of epistemic responsibility. See his “Justified Belief and Epistemically Responsible Action,” *Philosophical Review* 92 (1983): 33–48.

with commendable modesty, they did not conceive of themselves as having grasped an epistemological principle that generations of learned predecessors had failed to appreciate.⁴³

An articulation of the external ideal would provide a clear standard, appropriate for the *context of methodological improvement*. How to circumscribe the class of considerations for which we should allow in the *context of epistemic appraisal* is far harder. However, before we venture on this task, it is useful to inquire what the point of epistemic appraisal is. Appraisals of agents' rationality find their home in educational situations, and it is pertinent to ask what purpose they serve when we are focusing on debates in the history of science.⁴⁴ Yet, even when the subjects under study are contemporary students, when epistemic appraisal can be conceived as diagnosis preceding attempts at improvement, the philosophical dichotomies rational/irrational and justified/unjustified may stand in need of replacement rather than analysis. When we note that a student falls short of the external ideal (as we conceive of it), debate about whether the failure to undergo the epistemically optimal process is excusable or not can profitably be sidestepped in favor of a psychologically richer explanation of what occurred. Cognitively inferior performances can be based on laziness, methodological ignorance or misinformation, failure to perceive relevant similarities, lack of imagination, and numerous other kinds of factors.⁴⁵

The enterprise of analyzing justification (or rationality) straddles uncomfortably the projects of articulating the external ideal and of

⁴³Historians of science are constantly perplexed, even irritated, by the philosophical penchant for applying the methodological standards of the present in evaluating the decisions of past scientists. Once one rejects the notion that methodological principles are identifiable *a priori*, then the liberal approach is easily seen as an extension of the policy of allowing for cognitive limitations that prompts employment of the notions of justification and rationality in the first place.

⁴⁴Another place for epistemic appraisals in contemporary situations is in connection with broader assessments of responsibility, as, for example, in ethical, political, and legal contexts. If epistemic appraisals play a role in understanding the history of science, it is because we hope to defend science as a privileged tradition, one that is more worthy of trust and of social support than rival traditions or institutions.

⁴⁵In *Epistemology and Cognition*, Goldman points toward the need to consider a broader class of notions of epistemic evaluation. His suggestions are further developed in some of the essays in *Liaisons*.

exposing the more lenient standards employed in contexts of appraisal.⁴⁶ It seems impossible to pursue both projects in tandem. Reliabilism gives a promising start to formulating a meliorative naturalistic enterprise, even if it is not the panacea for the problem of analyzing justification.

4.

As I have already noted, some contemporary naturalists are less concerned to reintroduce psychology into epistemology than to repudiate the idea of *a priori* epistemological principles. They conceive of the decisive step in the transformation of epistemology and philosophy of science as *metaepistemological*: philosophical problems about knowledge can be satisfactorily addressed only by considering the ways in which historical and contemporary figures actually undertake their projects of inquiry.⁴⁷ Although there are often important links to the ideas of the classical American pragmatists, I view the main motivation for this strand in naturalism as stemming from more recent critiques of *a priori* epistemology.⁴⁸

These critiques take two forms, which appeal to ideas of W. V. Quine and T. S. Kuhn, respectively.⁴⁹ Quine's attack on the notion of analyticity, in particular his explicit suggestion that any state-

⁴⁶In my judgment, Goldman's detailed account in the first part of *Epistemology and Cognition* exhibits the discomfort. Although he is clearly interested in recasting the meliorative project within a naturalistic framework, Goldman also takes up the traditional project of analytic epistemology. This can readily be understood in terms of the historical development of his ideas from an initial response to the Gettier problem. However, in his most recent writings, Goldman does separate two enterprises, that of elucidating our "epistemic folkways" and that of doing "scientific epistemology" (see "Epistemic Folkways and Scientific Epistemology"). In my judgment, he continues to place too much weight on the former.

⁴⁷Prominent defenders of this type of naturalism include Larry Laudan, Dudley Shapere, and, in a somewhat different way, Nicholas Rescher. For a clear statement of the main contentions, see Laudan's "Progress or Rationality: The Prospects for Normative Naturalism," *American Philosophical Quarterly* 24 (1987): 19–31, and "Normative Naturalism," *Philosophy of Science* 57 (1990): 44–59.

⁴⁸I shall not try to trace the ways in which ideas of Peirce and Dewey are recast in the writings of contemporary epistemologists such as Laudan and Rescher, nor shall I explore the influence of Peirce and Dewey on Quine's own critique of apriority.

⁴⁹Although the ideas of Kuhn and Quine are of first importance in the

ment is revisable, implies directly that epistemological principles specifying methods for forming or revising beliefs are themselves vulnerable to revision in the light of experience. The argument from Kuhn is less straightforward, using the mismatch between the deliverances of methodologies for science and scientific practice to undermine our confidence in *a priori* pronouncements about how science ought to be done.⁵⁰

Within the framework of post-Fregean epistemology, Quine's critique of analyticity automatically constituted an attack on the *a priori*. Many of the attempts to respond to Quine and to display the possibility of redrawing the analytic-synthetic distinction miss the crucial point that the resultant notion of analyticity should do the epistemological work that the tradition assigned to it—that is, it should solve the problem of the *a priori*.⁵¹ The distinction that Frege's successors celebrated was between those statements whose truth values can be known only on the basis of experience and those whose truth values are determined by the meanings of the constituent terms and *therefore* can be known independently of experience.

Quine's fundamental arguments against analyticity are distributed among three seminal papers, "Truth by Convention," "Two Dogmas of Empiricism," and "Carnap and Logical Truth." The first of these consists, for the most part, of a demonstration that, if

revival of naturalism, their own positions do not emerge in the way I suggest here. As "Epistemology Naturalized" (in *Ontological Relativity and Other Essays* [New York: Columbia, 1969], 69–90) makes clear, Quine's naturalism is a direct response to the failures of particular programs of logical empiricism. Kuhn's complex and sometimes elusive epistemological views have inspired numerous ventures in naturalistic approaches to science—for example, those of Ronald Giere and Barry Barnes—but it is far from clear whether Kuhn himself espouses naturalism.

⁵⁰Of course, the arguments of Kuhn and Quine were both made before the revivals of psychologism discussed in the previous section. To the best of my knowledge, doubts about the apriority of epistemological principles played little explicit role in the return of psychologistic epistemology. I have therefore not tried to tell a straightforwardly chronological story, but to present the challenges to Frege's two presuppositions—the ban on psychology and the belief in the apriority of epistemological principles—so as to bring out as clearly as possible the connections among aspects of the emerging naturalist position.

⁵¹This point is clearly made by Harman in "Quine on Meaning and Existence I."

we wished, all axiomatizable disciplines could be replaced by analytic surrogates. The recipe is to axiomatize any branch of inquiry that one chooses, and to institute linguistic conventions that fix the meanings of the terms so that the axioms are true. As Quine points out, this simple recipe works no epistemological magic.⁵² The resultant branch of inquiry can still be displaced by just the same inputs from nature that would have undermined its unreconstructed synthetic predecessor. We will now describe the replacement as an episode in which a particular kind of language was shown to be inappropriate for the description of reality. The point is pursued further in the final section of “Two Dogmas,” where Quine denies that the abandoning of linguistic conventions is, in principle, any different from the major shifts that have occurred in the history of science.⁵³ Later still, in “Carnap and Logical Truth,” Quine attends to the possibility of *a priori* knowledge that results from explicit postulation. The closing section of “Truth by Convention” had already scrutinized the idea of grounding all our logical knowledge in explicit conventional stipulation and found it wanting. Quine’s later treatment recognized the limited role that “legislative postulation” can play within an ongoing cognitive enterprise: we may introduce a statement by declaring that it is constitutive of the meanings of some constituent terms, but, by doing so, we do not ensure that it must be preserved within our system of beliefs, come what may.⁵⁴

Quinean arguments against the *a priori* can be extended by presenting them explicitly within the idiom of psychologistic epistemology.⁵⁵ If justification and knowledge turn on the characteristics of the processes that generate and sustain belief, then a *a priori* justification and a *a priori* knowledge result from the presence of processes that are, in some sense, independent of experience. If we now try to identify processes that could generate knowledge inde-

⁵²Quine, *The Ways of Paradox*, 94–95. For Kant’s anticipation, see Allison, *The Kant-Eberhard Controversy*, 175.

⁵³*From a Logical Point of View*, 43.

⁵⁴*The Ways of Paradox*, 124–25.

⁵⁵To say the least, this is a very un-Quinean way of developing the challenge to the *a priori*. As scores of commentators have noted, Quine’s writings are permeated by commitments to behaviorism, so that he would hardly be sympathetic to the psychologistic turn in epistemology as it has been developed by many of the authors cited in the previous section.

pendently of experience, we find only a few plausible candidates: the processes whereby logicians and mathematicians apprehend axioms and construct proofs, or those in which we defend conclusions by appealing to our understanding of our language.⁵⁶ In the former case, there are grounds for wondering if the processes could fulfill their knowledge-generating functions against a background of experiences that explicitly called their reliability into question; in the latter, we begin with the observation that experience might generate reasonable doubt about the utility of the concepts on which we rely, appeal to Quine's arguments for the view that there is no epistemological difference between changing one's doctrines and modifying parts of one's conceptual framework, and conclude that our current knowledge of the conclusions is dependent on the benign experiences that support the applicability of our conceptual framework. In fact, to use a recurrent Quinean formulation, all such knowledge-generating processes are dependent on our absorption of ancestral lore, so that we are always implicitly dependent on the struggles of our predecessors to fashion a language apt for the description of the world, and are thus, vicariously, dependent on their experiences.

If we think of ourselves as gaining knowledge by undergoing processes that are possible for us only through our absorption of the concepts and doctrines of predecessors, then the twentieth-century ideal of a synchronic reconstruction of the knowledge of individuals or groups—either by displaying the coherence of the statements that they espouse or by showing the chains of justification that lead back to foundational statements—appears absurd. Our knowledge is embedded in the history of human knowledge and not detachable from it.⁵⁷ Furthermore, our reliance on the authority of others is ineluctable. Surely, as a matter of fact, all of

⁵⁶This paragraph condenses a line of argument I have elaborated in "A Priori Knowledge," *Philosophical Review* 86 (1980): 3–23, and in chapters 1–4 of *The Nature of Mathematical Knowledge* (New York: Oxford University Press, 1983).

⁵⁷The denial of the *a priori* thus leads not to an opposing synchronic program that one might label "radical empiricism," but to a position whose emphasis on the growth of knowledge invites the title "historicism." My own terminology in *The Nature of Mathematical Knowledge* thus seems to me infelicitous.

us do depend on authorities from childhood up, but the Cartesian hope, re-enunciated by Frege in the context of mathematics, is that we could, in principle, retreat to a stove-heated room and take our cognitive lives into our own hands. If, however, the knowledge-generating power of the processes we can undergo is dependent on the endorsements and recommendations of others, if it could be subverted by the refusal of others to accept our conclusions, then there can be no Cartesian, or even Fregean, reconstruction. As Quine so frequently reminds us, we are all in Neurath's boat.⁵⁸

Similar conclusions can be drawn from one way of reading Kuhn's *The Structure of Scientific Revolutions*. Kuhn's passing remarks about the details of earlier accounts of scientific methodology stress the mismatch between the deliverances of methodology and the reasoning that scientists actually employ.⁵⁹ Unless one can show that attention to the historical record will close the gap between philosophers' methodologies and scientific practice (a course that few have pursued), methodologists are confronted with a dilemma. Either they can continue to insist that philosophers know *a priori* the principles of confirmation and evidence, concluding that the actual reasoning of scientists is cognitively deficient, or they can abandon the *a priori* status of methodological claims and use the performances of past and present scientists as a guide to formulating a *fallible* theory of confirmation and evidence. Since the first option has an uncomfortable air of arrogance, it is hardly surprising that most responses to Kuhn have followed the latter course.⁶⁰

However, Kuhn's reflections on the growth of science present positive views that undermine many facets of traditional epistemology and philosophy of science. For Kuhn, even more than for Quine, the idea of a synchronic reconstruction of human knowledge is absurd.⁶¹ We absorb ideas from our predecessors and are

⁵⁸See *Word and Object* (Cambridge: The MIT Press, 1960) and "Natural Kinds" (in *Ontological Relativity*, 114–38). For a different metaphorical expression of our dependence on the past, see the closing lines of "Carnap and Logical Truth" (in *The Ways of Paradox*, 125).

⁵⁹See the discussions of Popper and falsification, *The Structure of Scientific Revolutions* (Chicago: University of Chicago Press, 1962; 2d. ed., 1970), 146–47.

⁶⁰Witness the responses of such philosophers as Imre Lakatos, Larry Laudan, and Dudley Shapere.

⁶¹While both Kuhn and Quine have been concerned with the growth of

epistemically dependent upon them. Scientists endeavor to improve their epistemic states by appealing to standards that embody broadly shared values, values that are themselves articulated using prevailing beliefs about nature. Consequently, on Kuhn's picture of the growth of science, the methodological canons that function in acceptance or rejection of new ideas are intertwined with substantive beliefs.⁶² For Kuhnian, as well as Quinean, reasons we should reject the idea of *a priori* epistemology as a myth.

5.

I have traced two different lines of argument for the reform of post-Fregean epistemology, trying to show connections that have not always been acknowledged. I conclude that the separate routes converge on a single position, and the next task is to make that position as clear and explicit as possible. Traditional⁶³ naturalism (as I shall call it) contains theses about epistemology as well as theses within epistemology. First, the basic conception of epistemology.⁶⁴

- (1) The central problem of epistemology is to understand the epistemic quality of human cognitive performance, and to

knowledge, Kuhn's interest has focused on phylogeny, Quine's on ontology. Contrast *The Structure of Scientific Revolutions* and *The Essential Tension* (Chicago: University of Chicago Press, 1977) with *Word and Object* and *The Roots of Reference* (La Salle, Ill.: Open Court, 1973).

⁶²See "Objectivity, Value-Judgments and Theory Choice," in *The Essential Tension*, 322–25.

⁶³By choosing this term I intend to suggest that important elements in the epistemological tradition have been preserved in the contemporary position. Quite clearly, not all claims common to Bacon, Descartes, Locke, Hume, Kant, and Mill are honored in traditional naturalism. But traditional naturalism does carry on the classical normative project and does embed epistemology within a general view of knowing subjects and the world with which they interact.

⁶⁴This conception carries forward an important view common to thinkers as diverse as Bacon, Descartes, and Frege. Jaegwon Kim argues eloquently for this as the traditional conception of epistemology: see his "What Is Naturalized Epistemology?" *Philosophical Perspectives*, vol. 2, 381–405.

specify strategies through whose use human beings can improve their cognitive states.⁶⁵

Two points about (1) deserve explicit mention. First, the strategies considered might include both ways of forming new representations (intuitively, discovery) and ways of appraising representations already presented (justification), or simply the latter. For purposes of this essay, I shall focus on the problem of appraisal, although it is not clear that this can be kept distinct from issues about formulation.⁶⁶ Second, the problem is couched in terms of improvement of an existing state rather than the construction of a "proper" belief corpus *ab initio*. Many older epistemologies would present the problem in the latter way, regarding the improvement of cognitive state as consisting in a scorched-earth policy of intellectual slum clearance at the first step. (1) speaks of improvement because it absorbs the Quine-Kuhn idea that we are ineluctably dependent on the past.

Traditional naturalism adds to (1) some epistemological theses that have occupied us in the last two sections.

- (2) The epistemic status of a state is dependent on the processes that generate and sustain it.
- (3) The central epistemological project is to be carried out by describing processes that are reliable, in the sense that they

⁶⁵In the remainder of this essay I shall be primarily concerned with the more practical project of improving cognitive performance, but, as Peter Godfrey-Smith pointed out to me, the issue of identifying how successful we actually are and of accounting for our level of success should not be slighted. Many of the points that occur later can be adapted to discussion of that issue. I shall omit making the connections explicit for reasons of space.

⁶⁶Hans Reichenbach proposed distinguishing between the context of discovery and the context of justification, and confining epistemological analysis to the latter. For a characteristically lucid formulation and motivation, see Hempel, *Philosophy of Natural Science* (Englewood Cliffs, N.J.: Prentice-Hall, 1966), chap. 2. For reconsideration of the distinction, see Laudan, *Progress and its Problems* (Berkeley: University of California Press, 1977), and the essays in *Scientific Discovery: Logic and Rationality*, ed. Thomas Nickles (Dordrecht, The Netherlands: Reidel, 1980). A recent attempt to invade the territory banned by Reichenbach is Clark Glymour et al., *Discovering Causal Structure* (New York: Academic Press, 1987).

would have a high frequency of generating epistemically virtuous states in human beings in our world.⁶⁷

- (4) Virtually nothing is knowable *a priori*, and, in particular, no epistemological principle is knowable *a priori*.

As noted in section 3, the addition of (2) to (1) involves only minimal departures from the twentieth-century epistemological mainstream, simply denying the extreme apychologism of post-Fregean epistemology. (3) and (4) are far more substantive.⁶⁸ (4) depends on arguments of Quine and Kuhn. (3) claims that the standard of reliability is offered not as an analysis of our ordinary notions of rationality, justification, and/or knowledge, but as an ideal with which we hope to bring ourselves into conformity. One task for traditional naturalism is to articulate the notion of “epistemic virtue” against which potential cognitive improvements are to be judged.⁶⁹ The high frequency requirement is to be construed as follows: candidate processes will have conditions of application which are sometimes, though not always, satisfied in our world; within a representative sample of occasions on which conditions for application of a class of processes are satisfied, a correct naturalistic epistemology should specify those which maximize expected cognitive virtue. I shall refer to such processes as *cognitively optimal*. The ultimate goal of (traditional naturalistic) epistemology is to present a compendium of cognitively optimal processes for all those contexts in which human subjects find themselves.

⁶⁷In a more refined version, this simple reliabilist idea would need to be accompanied by an explicit discussion of the epistemic risks involved in various rival strategies. Here again I am indebted to Godfrey-Smith.

⁶⁸Although some apychologistic epistemologists, such as Chisholm, have considered the attainment of epistemic goals, they have not formulated a touchstone for epistemic principles in terms of the maximization of epistemic utility for cognitively limited creatures in the actual world. Thus, (3) seems to me to involve a marked departure from the twentieth-century epistemological mainstream.

⁶⁹As noted in section 3, conceptual clarification can play a valuable role within the naturalist enterprise, even though it is clearly understood that the concepts in question might be superseded. Traditional naturalists must undertake the project of formulating a *corrigible* epistemic ideal, and, as I shall note below, some of the tools of post-Fregean epistemology may be valuable in this endeavor (see the discussion immediately following the statement of objection (A), pp. 78–79).

Traditional naturalists⁷⁰ occupy an uncomfortable middle ground between earlier epistemologists and those who campaign for abandoning (or relativizing) normative projects. In the ensuing debates, each of the extreme positions can use its counterpart as a foil for denying the possibility of an intermediate position.⁷¹ Thus, post-Fregean epistemologists can attempt to oppose traditional naturalism by contending that it leads immediately to forms of relativism or skepticism that are unacceptable. Radical naturalists, by contrast, portray traditional naturalism as failing to break free from the errors of post-Fregean epistemology. If we are to have a clear view of forms of naturalism, we must understand what difference, if any, the acceptance of (1)–(4) makes, and what revisions in these theses might be made to produce more radical versions.

Although acceptance of (4) would alter one's assessment of the deliverances of post-Fregean epistemology, this might make little difference to epistemological practice. The mere fact that the apparent sources of epistemological recommendations (logic, conceptual analysis, probability theory) are no longer judged as *a priori*,⁷²

⁷⁰So far as I know, nobody has formulated the naturalist epistemological position precisely as I have done. Elements of the position can readily be found in Armstrong, Goldman, Dretske, Laudan, Shapere, Rescher, Kornblith and others. If I am right about the motivation for accepting those elements (outlined in sections 3 and 4), then these authors ought to adopt those parts of traditional naturalism that they have not explicitly affirmed in their writings. So I *hope* that traditional naturalism has a number of proponents.

⁷¹This strategy can variously be found in Richard Rorty, *Philosophy and the Mirror of Nature* (Princeton: Princeton University Press, 1979); David Bloor, *Knowledge and Social Imagery* (London: Routledge and Kegan Paul, 1974); and Hilary Putnam, *Realism and Reason*.

⁷²It may be held that this is defensible only in a strained and idiosyncratic sense of '*a priori*'. I claim, however, that the traditional notion of apriority rests on the idea that *a priori* knowledge is independent of experience, and that, when this is articulated, it is seen that claims that *p* is *a priori* imply that *p* can be known in such a way that the knowledge could not be undermined by any experience. Hence, for reasons analogous to those I give in the case of mathematics in *The Nature of Mathematical Knowledge*, I hold that the usual sources of philosophical knowledge do not yield *a priori* knowledge. I sympathize with those who think that my analysis of apriority somehow stacks the deck, and invite them to articulate a notion of apriority that will both cleave to the venerable idea of independence from experience and also avoid the negative conclusions which I reach,

does not mean that those sources need to be replaced with or extended by empirical investigations, say in psychology or in biology. I shall sum this up as the complaint

- (A) Empirical studies of our actual cognitive practices, whether they be psychological, biological, or historical, play only a minor role in the normative project of epistemology. The usual philosophical sources of normative principles are not displaced by traditional naturalism, which offers only the metaepistemological principle that the deliverances of these sources are not *a priori*.

(A) may be used by the post-Fregean epistemologist to suggest that traditional naturalism is much ado about very little, or by the radical naturalist to propose that the fundamental problems with post-Fregean epistemology are untouched by traditional naturalism.⁷³

Traditional naturalists ought to concede that there is a legitimate activity of using the arsenal of philosophical techniques (appealing to formal logic or probability theory, say) to articulate ideas about knowledge. The development of an account of epistemic value might well draw on such resources. Naturalists are not concerned to throw away useful tools but to deny some advertisements about what the tools can accomplish (i.e., generate *a priori* knowledge) and to insist that other instruments (findings from psychology or from the history of science, for example) may also be profoundly relevant to epistemology. (A)'s concern is with the relevance of the proposed additions.

One response to (A) notes that traditional naturalism does include a thesis, (3), amending post-Fregean standards for normative principles. Traditional naturalists aim to produce principles that

and which Quine and Putnam earlier defended in somewhat different ways.

⁷³Thus Kim closes "What Is Naturalistic Epistemology?" by suggesting that scientific findings have the same bearing on epistemology that they do on ethics. Conversely, Bloor finds no room for normative considerations within an account of the actual production of belief, and in this he is followed by many sociologists of science (see, especially, Bloor, *Knowledge and Social Imagery*, and Barnes and Bloor, "Relativism, Rationalism, and the Sociology of Scientific Knowledge," in *Rationality and Relativism*, ed. M. Hollis and S. Lukes [Cambridge: The MIT Press, 1982], 21–47).

can be deployed to promote cognitive success in the *actual* world, recommending that we use our current beliefs about the character of the world to formulate such principles. Empirical information about nature and our relation to the rest of nature must be relevant to the normative project. Now a second objection arises. If proper epistemic recommendations are crucially dependent on contingent information about the world, how could we acquire the information on which those recommendations depend? Because adequate epistemological principles emerge only late in inquiry—if at all—they must be based on a picture of nature obtained by using error-prone strategies. Consequently, the *apparent* information used in formulating our epistemic recommendations is likely to be misleading, with the result that *what we take to be* correct epistemic recommendations are infected with mistakes. The complaint:

- (B) Only if we can arrive at principles that would properly guide inquiry in any world and which can be validated *a priori* will the problem of normative epistemology be solved. For otherwise the dependence of epistemology on information that had to be obtained using admittedly error-prone methods will lead to an unresolvable form of skepticism.

Post-Fregean epistemologists can appeal to (B) to contend that the normative enterprise cannot be sustained without rejecting (3) and (4). By the same token, radicals are likely to insist that (4), taken in conjunction with (B), dooms traditional naturalism's attempt to accommodate the idea of normative epistemology.⁷⁴ I hope that it is now clear why traditional naturalism is threatened with instability.

The case for one form of radical naturalism can be made by focusing directly on (1). If epistemology is to be a normative discipline then, as we have seen already, its task is to specify those strategies which promote attainment of cognitive goals. But whose cognitive goals are these? Is there a single conception of the aims of

⁷⁴This argument is latent in much recent sociology of science (e.g., Harry Collins's *Changing Order* [London: Sage, 1985]). It can also be found in philosophical critiques of ventures in naturalistic epistemology (e.g., Gerald Doppelt's "The Naturalist Conception of Methodological Standards in Science: A Critique," *Philosophy of Science* 57 [1990]: 1–19).

inquiry that holds across all periods and all contexts? If not, then any formulation of epistemic recommendations must be explicitly relativized to some context within which cognitive goals are shared. Faced with a mismatch between actual strategies and epistemic recommendations, it would always be possible to respond by adjusting one's epistemic goals. Moreover, if there were normative principles specifying the legitimacy of this type of maneuver, these too would presuppose a conception of the goals of inquiry, which the participants in the context would be free to accept or reject. So there arises the complaint

- (C) The history of science reveals that the goals attributed to inquiry vary widely from field to field and from epoch to epoch. There can thus be no universal normative epistemology, and we must settle either for description of the ways in which people actually form their beliefs or for local recommendations about how those working within a particular context should operate to advance their goals.

One prominent form of contemporary naturalism, popular among many historians and sociologists of science, appeals to (C) to support the conclusion that normative epistemology is an exercise in empty moralizing.⁷⁵

Yet even those who hope to preserve some type of normative epistemology may want to defend a more radical view than that espoused by (1)–(4). Post-Fregean epistemology is relentlessly propositional. It deviates from older normative epistemologies (such as classical empiricism) with their discussions of the correct dynamic of ideas. Traditional naturalism uses the contemporary

⁷⁵Bloor, *Knowledge and Social Imagery*; Steven Shapin and Simon Schaffer, *Leviathan and the Air-Pump* (Princeton: Princeton University Press, 1985); Collins, *Changing Order*; Bruno Latour, *Science in Action* (Cambridge: Harvard University Press, 1987); Feyerabend, *Against Method*. As (C) makes plain, variation in goals can lead either to a thesis that all normative appraisals are relativized or to the dismissal of the enterprise of normative appraisal. Passages in Bloor, Latour, and Feyerabend suggest that variation in goals is so extensive that the history of science could have been radically different had those with different goals triumphed, and those authors draw the moral that there is no normative perspective from which the scientific tradition can be defended as privileged.

term 'representations' to formulate (1)–(4), conceiving representations as states having something like propositional content, so that the transition to traditional naturalism would take over much of the well-entrenched technical vocabulary for discussing psychological issues that is used in epistemological discussions. This, however, can be attacked on one of two grounds. First, recent studies of scientific knowledge have emphasized the role of nonverbal skills in processes that post-Fregean epistemologies have seen in terms of the relatively unproblematic adoption of statements.⁷⁶ Renewed interest in the character of experiments has shattered preconceptions about the acquisition of "empirical evidence."⁷⁷ Second, naturalism's emphasis on using our best scientific knowledge within epistemology should extend to the reformulation of the basic epistemological issues. What the cognitive sciences inform us about the character of our representations should be used in posing those issues. To the extent that some, much, or all, human knowledge emerges as fundamentally nonpropositional, the usual framework of epistemology will be superseded.⁷⁸ Thus,

- (D) Traditional epistemology formulates its problems and answers by thinking of knowledge as primarily propositional. This presupposition should be scrutinized in the light of historical and sociological analyses of cognitive performance and in the light of contemporary theories of human

⁷⁶See Michael Polanyi, *Personal Knowledge* (Chicago: University of Chicago Press, 1957); J. R. Ravetz, *Science and its Social Problems* (Oxford: Oxford University Press, 1971); Martin Rudwick, *The Great Devonian Controversy* (Chicago: University of Chicago Press, 1985); Harry Collins, *Changing Order*.

⁷⁷See Ian Hacking, *Representing and Intervening* (Cambridge: Cambridge University Press, 1983); Peter Galison, *How Experiments End* (Chicago: University of Chicago Press, 1988); Andrew Pickering, *Constructing Quarks* (Chicago: University of Chicago Press, 1984); James Bogen and James Woodward, "Saving the Phenomena," *Philosophical Review* 97 (1988): 303–52; and David Gooding et al., eds., *The Uses of Experiment*.

⁷⁸See Patricia Churchland, *Neurophilosophy* (Cambridge: The MIT Press, Bradford Books, 1986) and "Epistemology in the Age of Neuroscience," *Journal of Philosophy* 84 (1987): 544–52; Paul Churchland, *A Neurocomputational Perspective* (Cambridge: The MIT Press, Bradford Books, 1989); Stephen Stich, *The Case Against Belief* (Cambridge: The MIT Press, Bradford Books, 1983) and *The Fragmentation of Reason* (Cambridge: The MIT Press, Bradford Books, 1990).

cognition. Where necessary, the standard epistemological idioms of belief, justification, and so forth should be absorbed within a broader vocabulary or, perhaps, discarded entirely.

A second suggestion for modifying traditional naturalism, while retaining a normative focus, opposes the ruthlessly individualistic emphasis of both traditional naturalism and its pre-Fregean counterparts.⁷⁹ Although Bacon gave passing attention to the social character of knowledge, the primary epistemological problem, from the seventeenth century to today, has been how an *individual* may advance his epistemic ends.⁸⁰ Once we understand that the cognitive states achieved by individuals are dependent on those of their contemporaries and predecessors, the idea of studying the knowledge of Robinson Crusoes seems artificial and misleading. Strategies of improving the cognitive states of individuals cannot be specified without considering the communities to which those individuals belong.⁸¹

- (E) Epistemology must examine the attainment of knowledge by communities as well as by individuals, and should investigate strategies through which communities could advance their epistemic ends. The appropriate strategies for

⁷⁹(1)–(4) are focused on what individuals can do to improve their cognitive lives. Yet, as I understand the primary motivations for traditional naturalism, they consist, in part, of recognizing the social character of knowledge. (See the discussions in sections 3 and 4.) Thus, to be true to the underlying motivations for traditional naturalism, we need to draw out the full implications of Kuhn's and Quine's challenges to post-Fregean epistemology. See section 9 below.

⁸⁰See Bacon, *New Atlantis* (Oxford: Oxford University Press, 1961), and Thomas Sprat, *History of the Royal Society* (London, 1667). The individualistic emphasis is evident in Descartes, Locke, Hume, Kant, Mill, Frege, and virtually all twentieth-century epistemology. For concerns about the social organization of knowledge, the most important twentieth-century work is Robert Merton's, *The Sociology of Science* (Chicago: University of Chicago Press, 1973).

⁸¹This point has been emphasized by David Hull, in *Science as a Process* (Chicago: University of Chicago Press, 1988); by Goldman, in "Foundations of Social Epistemics," *Synthese* 73 (1987): 109–44; and in my "The Division of Cognitive Labor," *Journal of Philosophy* 87 (1990): 5–22.

individuals to follow cannot be identified without considering the communities to which they belong.

In the remainder of this essay, I shall use these complaints to try to assess the tenability of traditional naturalism and to explore its relations both to older epistemologies and to more radical versions of naturalistic epistemology and philosophy of science. In doing so, I shall make explicit some themes in contemporary naturalism that have so far been slighted—for example, the interest in evolutionary epistemology (including study of the implications of evolutionary biology) and debates about scientific realism.

6.

I begin with (A), offered by traditionalists as an indictment of the importance attributed to the naturalistic turn. Assume, for the purposes of this section, that the language of belief and justification is to be retained, and that our epistemic ends include the attainment of truth by individuals.⁸² Consider now the following elaboration of (A).

The principles of deductive logic, of probability theory, and of the methodology of science, can be understood as specifying norms for belief formation. Recognizing that q is a consequence of p , we see that if we believe that p then we should also believe that q . Similarly Bayes's theorem tells us that people who accept certain judgments about the frequency of diseases and the rates of errors on diagnostic tests ought to adopt particular values for the conditional probabilities of having a particular disease given a positive response on a particular test. If it is objected that psychological studies reveal that these are not in fact the ways in which people adjust their beliefs, then the right response should be that actual practice stands in need of correction.⁸³ Only if it could be argued

⁸²Discussion of these questions is postponed to sections 9 and 8 respectively.

⁸³See the essays in Daniel Kahneman, Paul Slovic, and Amos Tversky, eds., *Judgment under Uncertainty: Heuristics and Biases* (Cambridge: Cambridge University Press, 1982), especially those by Kahneman and Tversky; and Richard Nisbett and Lee Ross, *Human Inference: Strategies and*

that psychology shows that people *cannot* reason in the preferred ways would there be a *prima facie* case for revising normative judgments in the light of empirical evidence.⁸⁴ In the present instance, it is very hard to see how empirical disciplines could show any such thing, for the specification of the ideal strategy is itself evidence that this strategy can be recognized by human beings and that its epistemic merits can be appreciated. Hence empirical findings will be irrelevant to the normative projects of epistemology or methodology of science.⁸⁵

A weak response to this line of argument distinguishes between our ability to formulate a strategy and the possibility of integrating it into our cognitive lives. We may be able to reflect on the merits of Bayesian methods of reasoning, even though those methods are so unnatural for us that we cannot devise an educational regime that would make us fluent in their use. Post-Fregeans may legitimately reply that more imaginative attempts to devise programs of training would make the ideal processes accessible. As generations of teachers have found out for themselves, and as recent psychological investigations have confirmed, people can be taught to use Bayesian methods on appropriate problems.⁸⁶ So the weak response looks like the false counsel of laziness.

A stronger reply focuses on the normative implications of results in deductive logic.⁸⁷ Learning that *q* is a consequence of *p*, we are supposed to recognize that we should believe *q* if we believe *p*. So, the traditionalist concludes, we can frame a disjunctive constraint

Shortcomings of Social Judgment (Englewood Cliffs, N.J.: Prentice-Hall, 1980).

⁸⁴As Goldman points out ("Epistemics: The Regulative Theory of Cognition"), in epistemology as in ethics, "ought" implies "can."

⁸⁵Throughout this argument, 'empirical findings' should be read as referring to the deliverances of the sciences. The argument might concede the point that the traditional sources of epistemological principles are not *a priori*, simply concluding that this amounts to a reclassification of the status of those sources, not a demand for different sources.

⁸⁶See R. Holland et al., *Induction* (Cambridge: The MIT Press, Bradford Books, 1986), chap. 9.

⁸⁷For important arguments about the normative power of deductive principles, see Harman, *Thought and Change in View* (Cambridge: The MIT Press, Bradford Books, 1987); Stich, *The Fragmentation of Reason*; and Christopher Cherniak, *Minimal Rationality* (Cambridge: The MIT Press, Bradford Books, 1986).

on our belief-generating processes, debarring those that allow for the retention of inconsistent beliefs, and this constraint is not subvertible by empirical findings. However, on numerous occasions in the history of science, investigators have found themselves inclined to accept the members of a set of statements that they could recognize as jointly inconsistent, without knowing immediately what should be abandoned: Darwinian evolutionary theory survived Lord Kelvin's estimates of the age of the earth, Bohr's theory of the atom was retained and developed even though it was at odds with classical electromagnetic theory.⁸⁸ The phenomenon should be apparent from humbler situations, in which people know that they are inconsistent but do not yet see the right way to achieve consistency. It may even be universal, if each of us is modest enough to believe that one of our beliefs is false.

Deductive logic gives no normative advice here, beyond the admonition to be careful. Ultimately, something within the inconsistent set of beliefs will have to be changed if all members of the resultant set are to be true. How we should proceed in belief revision is a matter for nondeductive methodology to determine.⁸⁹

The argument for thinking that methodological recommendations are immune to scientific investigations turns on the notion that we can formulate, using logic and probability theory, belief-forming strategies that can be shown to be cognitively optimal. If we can formulate these strategies, then, apparently, we would be able to use them, and if we could use them, then, since they are cognitively optimal, we ought to use them. Two problems emerge here. One, which we shall consider in more detail in the next section, turns on the possibility that substantive methodology requires formulating strategies that are likely to yield good results, given the way *the world actually is*, and, consequently, identification of these strategies must draw on empirical information about the world. The second attacks directly the notion that if we can formulate a strategy then it follows that we can implement that strategy.

⁸⁸See Joe Burchfield, *Lord Kelvin and the Age of the Earth* (New York: Science History Publications, 1975), and Imre Lakatos's discussion of Bohr in "Falsification and the Methodology of Scientific Research Programmes," in *Criticism and the Growth of Knowledge*, ed. I. Lakatos and A. Musgrave (Cambridge: Cambridge University Press, 1970), 140–54.

⁸⁹See Harman, *Thought and Change in View*.

Consideration of simple algorithms, rules for checking deductive proofs and for solving typical Bayesian problems, favors a tight connection between formulation and implementation. We think of epistemic rules as specifying the steps that ought to be followed in arriving at beliefs, and, because those steps can be identified, there appears to be no bar to running through them. We forget that *application* of the rules usually involves procedures that cannot readily be reduced to algorithmic treatment. Results from psychology (and other empirical disciplines) become relevant when we scrutinize how formal procedures are adapted to belief generation in realistic contexts.

Consider a hallowed principle of methodology, the requirement of total evidence.⁹⁰ To a first approximation, this requirement demands that all available evidence be considered in forming belief. So, in fixing judgments of probability, one must take into account all the available information about frequencies of occurrence within relevant classes.

I ignore the worry that canvassing one's entire belief corpus might not be enough, that sometimes a cognitively optimal (or even an epistemically responsible) strategy is to seek *more* evidence. The real trouble is that, at least in complex scientific debates, and probably on all kinds of mundane occasions as well, people are incapable of bringing together all the relevant information that they have, let alone the entire stock of information that is available within their community. All kinds of things about plants, animals, continents, islands, rock strata, and geography were highly relevant to the evaluation of Darwin's "long argument" in the *Origin of Species*. None of the readers of that work could have followed the directive to consider "all the evidence": psychological research already tells us that this idealization is unrealistic. Yet Darwin's contemporaries, undergoing successive searches of parts of the information stored in their memories and combining their tentative conclusions in interactions with one another, were able to make the mass of potential evidence more manageable, to arrive at a conception of what the *relevant* evidence is and to use this to appraise and accept (some

⁹⁰The points of this paragraph are articulated by Goldman ("Epistemics: The Regulative Theory of Cognition," *Epistemology and Cognition*) and Harman (*Change in View*).

of) Darwin's proposals. To achieve specifications of cognitively superior strategies that human beings can use, the idealization of the requirement of total evidence should give way to descriptions of the processes through which evidence is rendered appreciable.

The need for psychology is hidden when we detach subjects from decision-making contexts. It is possible to specify cognitively optimal processes for explicitly circumscribed problems without considering psychological limitations. Post-Fregean ways of turning complex problems of belief formation into classroom exercises presuppose a highly idealized conception of the subject.⁹¹ Empirical studies are relevant to epistemology because we need to understand the cognitively superior ways in which creatures like ourselves could achieve neatly circumscribed problems in the first place.

7.

I have responded to (A) by considering the need to take into account the limitations of the knowing subject. But, as noted, the most obvious response to the complaint emphasizes the importance of information about the world with which the subject is attempting to come to grips, information which could be employed to identify those cognitive strategies that are likely to work, given the actual state of affairs. Before addressing complaint (B) directly, I want to recognize the attractions of restricting epistemology to the task of specifying cognitive strategies that are actually successful.

Consider the practice of inductive generalization. All human beings (even Popperians) need to form expectations about the future. One task of methodology is to explain how best to do this. A simple suggestion, obviously in need of qualification to accommodate the possibility of interfering background evidence, is that we should follow the straight rule: If the relative frequency with which the property *B* is found among the observed *As* is *r*, then predict that *r* of all *As* are *B*. There are well-known attempts to show that following this rule must succeed, whatever the world is like. Thus, Hans Reichenbach and Wesley Salmon have argued that, in the

⁹¹Hence the recurrent historians' complaint that the methodological principles urged by philosophers have virtually no bearing on the debates that have actually occurred in the history of science.

long run, use of the straight rule will converge on the limiting relative frequency if there is a limiting relative frequency to be found.⁹² Two familiar problems arise for this argument. Keynes's ironic reminder that "in the long run we are all dead" effectively points out that the notion of success in which we are interested is short-term convergence, and that cannot be guaranteed if we follow the straight rule. Nelson Goodman's celebrated "new riddle of induction" exposes the vacuity of the straight rule as stated: *Any policy of generalization whatsoever can be viewed as use of the straight rule if we choose appropriate predicates for describing what is observed in our sample.* Combining these points, we understand that a rule which appears to offer guidance whatever nature is like provides only the form in which rules with genuine content might be embedded. If counsel is to be given on matters of inductive belief generation, we must specify what kinds of predicates should be used in applying the straight rule so as to achieve convergence to actual relative frequencies in the short run.

The enterprise of articulating *a priori* substantive inductive policies that will be preferable to others *whatever the world is like* is obviously doomed. There are possible worlds in which use of rival predicates would bring rapid convergence to truth. We hope, but cannot demonstrate, that the system of predicates we actually use will lead to success in the actual world.

More exactly, our inductive methodology can be conceived as revising the system of predicates with which we articulate the straight rule, to favor those predicates and families of predicates that appear to give rise to stable generalizations. We might envisage that a higher-order analogue of the Reichenbach-Salmon argument would show that iterated pursuit of this policy will exhibit some desirable convergence. But such hopes will be dashed for exactly parallel reasons. Even the success of our policy for inductive revision depends on the character of the actual world.

Similar points can be gleaned from reflection on other epistemological issues. Since the late 1950s it has become a commonplace

⁹²See Hans Reichenbach, "On the Justification of Induction," in *Readings in Philosophical Analysis*, ed. H. Feigl and W. Sellars (New York: Appleton-Century-Crofts, 1949), 324–29, and Wesley Salmon, "On Vindicating Induction," in *Induction: Some Current Issues*, ed. H. E. Kyburg and Ernest Nagel (Middletown, Conn.: Wesleyan University Press, 1963).

that high-level theory can sometimes be preserved in the face of recalcitrant observations by jettisoning the supposed observational evidence.⁹³ Close study of scientific practice reveals, however, what Duhem pointed out long ago, that scientific judgment—*bon sens*—guides the acceptance and rejection of data.⁹⁴ Methodological counsels about how to respond to recalcitrant experience are grounded in empirical knowledge of how actual observers, their actual instruments, and their actual methods of formulating data are likely to err.

Doesn't this reformulation of epistemological projects lead directly to skepticism? Framing inductive policies in terms of the regularities we have observed in the world *so far* plays directly into Hume's hands. Without taking it on faith that those regularities will endure, we have no basis for supposing that the particular policies we employ will prove successful.

To see why traditional naturalism is especially vulnerable to skeptical challenges, consider the following development of (B). A traditional naturalistic epistemology would include principles of the form "*P* is cognitively optimal,"⁹⁵ where the notion of cognitive optimality is understood to depend on generation of cognitive success in the actual world and where the defense of the principle will involve appealing to empirical information about human beings and the rest of nature. Faced with a skeptic who denies that the principle should be accepted, traditional naturalists will thus appeal to some body of empirical statements. Persistent skeptics will then inquire why these statements should themselves be accepted. If the original skeptical challenge has been cleverly posed, the attack will have been directed at a process of belief formation that is widely implicated in the proper generation of beliefs about nature (at least on the naturalist's view). In consequence, defense of the

⁹³See Quine, "Two Dogmas of Empiricism," in *From a Logical Point of View* (New York: Harper, 1953), and Kuhn, *The Structure of Scientific Revolutions*.

⁹⁴Pierre Duhem, *The Aim and Structure of Physical Theory* (New York: Atheneum, 1962), 216–18.

⁹⁵Perhaps, also, comparative claims about the relative epistemic merits of different strategies. In striving for cognitively optimal processes, it may be helpful, at least initially, to be able to say which processes are superior to those that are actually employed.

body of empirical statements will require the use of *P*, at which point the skeptic will complain that the question has been begged.

The skeptic's demand is for synchronic reconstruction of beliefs: take the totality of things you believe, subtract this claim and everything that you cannot defend without assuming it, and now show that the claim is correct. With respect to some claims, synchronic reconstruction is possible. For theorems in mathematics and small parts of the sciences, successful synchronic reconstructions can be given and can serve useful purposes. Frege set his sights higher, aiming to show how the whole of mathematics could be reconstructed, while Descartes was yet more ambitious. Traditional naturalists, however, cognizant of the history of mathematics, science, and methodology, should know in advance of skeptical embarrassments that some forms of the problem of synchronic reconstruction are solvable and others are not. On their account there is no substantial body of *a priori* knowledge, so that successful synchronic reconstructions must always appeal to empirical findings. In consequence, one can produce unanswerable challenges by calling into question single claims or bodies of doctrine that are presupposed in all empirical investigations. On naturalism's own grounds, there are bound to be unanswerable forms of skepticism.

Traditional naturalists should therefore decline blanket invitations to play the game of synchronic reconstruction. Each of us absorbs information from our predecessors, and, through our own interactions with nature and with one another, we modify our collective picture of the world and of the proper ways to investigate it. Naturalists think of this process as leading to improvements, although there will be no way of showing that we are doing better without relying on some of our beliefs. Can matters be left with this diachronic picture of human knowledge? Or do the skeptical questions recur?

Naturalism offers the optimistic picture of a particular type of organism, beginning with rudimentary representations of nature and primitive notions of how to modify those representations, and gradually replacing these with cognitively superior representations and strategies. The skeptical questions concern (i) the possibility that we began in so primitive a state that we are incapable of working ourselves into any accurate representation of nature, and (ii) the possibility that there are constraints on the processes of modification that prevent us from making significant improvements. *If*

one believed that the historical process out of which contemporary beliefs have emerged was reliable (in the sense of having a high chance of generating truth), then the naturalistic reply to skeptics would succeed. However, the claim of reliability invites two types of worries: perhaps our initial state was so bad that there is nothing that we could have done to escape our misrepresentations, or perhaps the means of modification are so feeble or so confused that we have no chance of making substantive improvements. These concerns are reflected in (i) and (ii) respectively.⁹⁶

A promising and popular defense against (i) is to find “encouragement in Darwin.”⁹⁷ If our initial cognitive equipment were as unfortunate as the skeptic portrays it as being, then, the suggestion runs, our ancestors would have been eliminated by natural selection. They weren’t, so it wasn’t. In this way, we can appeal to Darwinian evolutionary theory to support the idea that our initial ways of classifying stimuli must correspond to objective regularities in nature, and our modes of reasoning must work reliably in producing accurate representations.

One complaint against the appeal to Darwin is rightly dismissed. If skeptics protest that a part of contemporary science is being taken for granted in evaluating aspects of the historical process out of which that science emerged, the appropriate naturalist reply is, “Of course. What else?” As I hope to have made clear, a central naturalist thesis is that some parts of our current scientific beliefs must be assumed in criticizing or endorsing others.

Unfortunately the invocation of natural selection will not do the intended job. Human brains have been assembled, over evolution-

⁹⁶Goldman discusses related issues in section 5.8 of *Epistemology and Cognition*.

⁹⁷The phrase is from Quine’s essay “Natural Kinds.” Similar thoughts were voiced by C. S. Peirce, and have recently been developed by Michael Ruse, *Taking Darwin Seriously* (Oxford: Basil Blackwell, 1986), and Nicholas Rescher, *A Useful Inheritance: Evolutionary Aspects of the Theory of Knowledge* (New York: Rowman and Allanheld, 1989). It is worth noting that there are *two* ways in which evolutionary concepts can be introduced into epistemology: one can try to appeal to selection pressures on hominids to tackle skeptical problems, or one can try to frame epistemological theories that have formal analogies to the theory of evolution by natural selection. Here I am concerned with the former project. For an illuminating overview of attempts to carry out the latter, see Michael Bradie, “Assessing Evolutionary Epistemology,” *Biology and Philosophy* 1 (1986): 401–59.

ary time, out of structures that were originally selected for properties far removed from any capacity for pursuing scientific investigations. Hence there may be severe phylogenetic constraints that prevent us from ever attaining accurate representations of nature. To the counter that our brains must be minimally competent at representing nature, that they must be able to alert us to the threats and possibilities of the environment, there are two obvious replies. One, in the spirit of the idea that the brain is the product of evolutionary tinkering, simply denies that this minimal competence will take us very far in establishing the reliability of the historical process out of which contemporary scientific beliefs have emerged. The second, scrutinizing the argument from selection in the style of contemporary neo-Darwinism, notes that the selection pressures felt by organisms are dependent on the costs and benefits of certain consequences. We think of hominids on the savannah requiring an accurate way to discern leopards, and conclude that parts of our ancestral schemes of representation, having evolved under strong selection, must accurately depict the environment. Yet, where selection is intense in the way it is here, the penalties are severe only for failures to recognize present predators. The hominid representations can be quite at odds with natural regularities, lumping all kinds of harmless things with potential dangers, provided that false positives are evolutionarily inconsequential and provided that the representation always cues the subject to danger.⁹⁸ There are easy ways of extending this line of thought to conclude that selection will favor organisms who play for safety and who consequently develop inaccurate systems of representation.

Initially, the appeal to Darwin promises to resolve skepticism of form (i). The tendency of the last paragraph is that contemporary evolutionary thinking actually reinforces (i). But, in my judgment, the exchange pits speculation against speculation. Nobody who has probed the subtle and rigorous analyses of the best evolutionary studies should have much patience with casual descriptions that seem to have more kinship with Kipling than with Darwin. The right response is neither optimism nor pessimism, but agnosticism.

⁹⁸This point is made by Stich in "Could Man be an Irrational Animal?" *Synthese* 64 (1984) (reprinted in *Naturalistic Epistemology*, ed. Hilary Kornblith [Cambridge: The MIT Press, Bradford Books, 1985], 249–67) and is further developed in *The Fragmentation of Reason*.

However, our agnosticism is remediable, at least in principle. More detailed accounts of human cognitive abilities and comparative studies of related organisms might enable us to adjudicate the skeptical question raised in (i). For the moment, let us simply note that success or failure with (i) is not decisive for the viability of traditional naturalism. Even if it were conceded that our ancestral epistemic situation was unfortunate, the reliability of products of the historical process could still be defended if we could show that there are unambiguous possibilities of continual correction. Even with a bad start, we might have gained improved representations of nature, perhaps accurate representations of nature, if our corrective capacities are sufficiently powerful.

At this point, we encounter (ii), a skeptical objection that directly denies our right to believe in self-correction. In fact it is wrong to speak of a single objection here, for there are several versions of a general difficulty. They are united in maintaining that study of decision making in science reveals that at many times there are alternative possibilities for modifying scientific belief, each of which might with equal justice have been adopted. We can imagine numerous possible histories of science, yielding divergent conceptions of nature and rival sets of epistemological principles, in each of which the protagonists retrospectively praise past decisions as exercises in self-correction. Because nothing distinguishes the actual course of events from these potential histories, there is no basis for concluding that the actual evolution of science is self-correcting while the others are not. The history of science comes to resemble a random walk, not a unidirectional process.⁹⁹

The root of this form of skepticism lies in arguments about the underdetermination of belief by encounters with nature. Such arguments begin with Duhem's point about the role of auxiliary assumptions in scientific testing.¹⁰⁰ They obtain prominence in the writings of Quine, whose treatment of the issue oscillates between two versions, one stressing the *logical consistency* of certain modifi-

⁹⁹This becomes explicit in many commentaries on Kuhn's *The Structure of Scientific Revolutions*, and is suggested by Kuhn's own efforts to provide an account of scientific progress in chapter 13 of that book.

¹⁰⁰As mentioned above, Duhem believed that underdetermination was typically resolved by the good sense of the scientist, and many practicing scientists react to formulations of the problem by dismissing it as an exer-

cations of belief and a stronger thesis asserting the *equal epistemic status* of those modifications.¹⁰¹ Attacks on the Duhem-Quine thesis and its disturbing implications typically exploit canons of methodology beyond those of deductive logic to urge the superiority of one alternative system of belief.¹⁰²

Kuhn's reflections on the growth of science complicate the issue. Where Quine had envisaged the possibility of alternative systems, each fully consistent with all possible stimuli, Kuhn portrayed scientists as deciding between systems, neither of which was fully concordant with the available stimuli. Alternatives might be defended by insisting on the significance of certain problems and/or achievements and downplaying others as unimportant. Moreover, there may be powerful forces that muffle the impact of our encounters with nature. Kuhn appealed to results in psychology to suggest serious constraints on our ability to perceive anomalous novelties, and he hinted at the power of authority to prevent the adoption of new ideas.

In the writings of those whom Kuhn has influenced, there is a battery of arguments that articulate the bare appeal to underdetermination.¹⁰³ The best of these draw on detailed studies of historical or contemporary science to expose one or more of the following features:

Shifting standards. Alternatives are defended by weighing problems and accomplishments differently.¹⁰⁴

cise in logic chopping. For a philosophical articulation of this line of response, see Larry Laudan, "Demystifying Underdetermination," in *Scientific Theories*, ed. C. Wade Savage, vol. 14 of *Minnesota Studies in the Philosophy of Science* (Minneapolis: University of Minnesota Press, 1990).

¹⁰¹Especially apparent in *Word and Object*, 22–24, 78. However, Quine uses the underdetermination thesis to make general points about problems of language and meaning, and it is something of a distortion to treat his arguments as bearing directly on the methodology of the sciences.

¹⁰²See Richard Boyd, "Realism, Underdetermination and the Causal Theory of Evidence," *Noûs* 8 (1973): 1–12; Laudan, "Demystifying Underdetermination"; and Clark Glymour, *Theory and Evidence* (Princeton: Princeton University Press, 1980).

¹⁰³Although these arguments often acknowledge a debt to Kuhn, it should be recognized that he disavows some of them.

¹⁰⁴See, for example, Gerald Doppelt, "Kuhn's Epistemological Relativism: An Interpretation and Defense," *Inquiry* 21 (1978): 33–86.

Theory-ladenness of observation. Theoretical disagreements are unresolvable because the parties report the evidence in the vocabulary of their preferred theories.¹⁰⁵

Assessment of experiments. The criterion for adequate performance of an experiment must ultimately be that it yields a conclusion that one is prepared to accept. Consequently, different theoretical conclusions can be sustained by judging some experiments to have been adequately performed, others badly done.¹⁰⁶

Social embedding. Conclusions about nature and about how to investigate nature are at least partly shaped by views about the proper social order.¹⁰⁷

Effects of authority. Heterodox views have only a small chance of acceptance or transmission because of the importance of reliance on authority within scientific communities.¹⁰⁸

Skepticism of type (ii) rests on detailed analyses of individual episodes in the history of science, purporting to show that these features are present and also that mechanisms that might more readily be understood as constituents of a self-correcting process are either absent or too weak to generate decisions.

Only a serious examination of the historical, sociological, and psychological material can resolve the issues that are raised by these versions of skepticism. But prior to any such examination, some preliminary points need to be appreciated. First, unlike some pre-

¹⁰⁵See N. R. Hanson, *Patterns of Discovery* (Cambridge: Cambridge University Press, 1958), and Paul Feyerabend, *Against Method* (London: Verso, 1975), both of whom formulated this theme independently of Kuhn. Other versions are due to Mary Hesse, "Is There an Independent Observation-Language?" (in *The Nature and Function of Scientific Theories*, ed. R. Colodny [Pittsburgh: University of Pittsburgh Press, 1970], 35–77) and Paul Churchland, *Scientific Realism and the Plasticity of Mind* (Cambridge: Cambridge University Press, 1979). For denial that the theory-ladenness of observation interferes with the objective comparison of theories, see Larry Laudan, *Progress and its Problems*.

¹⁰⁶See, H. M. Collins, *Changing Order*.

¹⁰⁷This theme is developed in many writings of David Bloor, Barry Barnes, and Steven Shapin, receiving its most detailed elaboration in Shapin and Simon Schaffer, *Leviathan and the Air-Pump*.

¹⁰⁸See Feyerabend, *Science in a Free Society* (London: New Left Books, 1978) and *Farewell to Reason* (London: Verso, 1987).

vious forms of skepticism that we have considered, the considerations just rehearsed attack traditional naturalism on its own ground: they recognize the diachronic defense of the joint progressiveness of science and epistemology, and they advance empirical evidence against it. Second, they cannot simply be dismissed by suggesting that it is inconsistent to appeal to empirical findings to undermine the accuracy of our empirical knowledge. For, like other skeptical arguments, these can reasonably be viewed as attempts at *reductio*. Third, although I have framed the issue in terms of skepticism, most proponents of the arguments under consideration advance positive views about the character of the history of science and about a naturalistic conception of science. From the perspective of traditional naturalism, those arguments are seen as skeptical challenges, but they might equally be viewed as steps towards radical forms of naturalism that either abandon the normative project of epistemology or honor it only insofar as it is explicitly relativized to frameworks, paradigms, sets of conventions, or “forms of life.” For radical naturalists of these stripes, the last link between naturalism and traditional epistemology, (1), is broken, replaced with the contention that naturalism’s task is simply to describe the ways in which beliefs are generated, in ontogeny and phylogeny.

Radical naturalism thus abandons the meliorative venture of Bacon and Descartes, letting epistemology fall into place as chapters of psychology, sociology, history of science.¹⁰⁹ Whether the collapse of traditional epistemology into radical naturalism is inevitable depends on whether the skeptical arguments that articulate (ii) can be resisted. To decide this we need both to do substantive work in psychology, sociology, and history of science, and to be clear about what is required for a successful defense of traditional naturalism.

Nobody should deny that there are instances of transient underdetermination in the history of science and in the practice of con-

¹⁰⁹The phrase is Quine’s (see “Epistemology Naturalized”), although contrary views are suggested in other places (such as Quine’s reply to Morton White in *The Philosophy of W. V. Quine*, ed. P. Schilpp [La Salle: Open Court, 1986]). In “What is Naturalistic Epistemology?” Kim urges that abandoning the normative constitutes changing the subject. See, also, Stroud, *The Significance of Philosophical Scepticism*.

temporary science, occasions on which what we now regard as optimal cognitive strategies are unable to decide between alternative systems of belief. From the perspective of the scientific community, it is a good thing if, under such circumstances, both rivals are developed and given the opportunity to prove their worth. The hope is that further encounters with nature yield a situation in which optimal cognitive strategies (as currently conceived) dictate a choice, either in favor of one of the original alternatives or for some new version that has emerged in the process of debate. Even if a decision is taken prematurely, there are still no grounds for thinking that science is not a self-correcting enterprise if there is subsequent convergence, with the triumphant system taking up the successes of its defeated rival. The real threat to traditional naturalism would arise from cases of transient underdetermination meeting two further conditions:

Continued Divergence. Developed versions of the current rivals continue to be incompatible with one another.

Indefinite Underdetermination. Those developed versions continue indefinitely to be underdetermined by the applications of cognitively superior strategies to inputs from nature.

Unless these conditions are met, it will not be possible to generate the skeptical conclusion that current conceptions of cognitively optimal strategies would allow for the acceptance of alternative views that would undermine them.

At this point, a strategy for defending traditional naturalism should become evident. With respect to the historical and contemporary cases, the aim will be to show that the case for continued divergence and indefinite underdetermination has not been made out. Instances of transient underdetermination can be expected (although even here the issues are often complicated by the fact that historical and sociological studies deploy an impoverished conception of the forms of reasoning available to the scientists under study).¹¹⁰ Resolution of scientific controversies often takes a long

¹¹⁰So, for example, Collins's major argument in *Changing Order* seems to presuppose that any consistent position is defensible. Similarly, in *Levia-*

time—for example, almost a century in the case of Copernicanism and about a decade for “the great Devonian controversy”—and it is folly to insist that there *must* have been available evidence to favor the final outcome from the beginning. All that traditional naturalism needs to show is that resolution is ultimately achieved, in favor either of one of the originally contending parties or of some emerging alternative that somehow combines their merits.¹¹¹

Implementing this strategy involves tackling some general questions. Those sensitive to the theory-ladenness of observation should be concerned that a premature triumph for one alternative will inevitably be sustained because subsequent stimuli will be accommodated in the victorious perspective. Similar anxieties arise with respect to the cramping effect of authority. Both questions submit to empirical investigation. Kuhn’s appeal to “new look” psychology can be updated, and one can try to ascertain the extent to which modifications of belief and/or changes in cognitive propensities (the “tacit knowledge” or *Fingerspitzengefühl* that some writers view as central to the practice of science) affect the processing of perceptual information.¹¹² Similarly, the effects of authority within a scientific community can be investigated by attempting to construct models of the dynamics of belief transmission within communities with specified authority structures.¹¹³ In both instances,

than and the Air-Pump, Shapin and Schaffer seem content with elaborating the consistency of Hobbes’s critique of pneumatic experiments.

¹¹¹For the reception of Copernicanism see Robert Westman, “The Melanchthon Circle, Rheticus, and the Wittenberg Interpretation of the Copernican Theory,” *Isis* 66 (1975): 165–93; for the Devonian controversy, Martin Rudwick, *The Great Devonian Controversy*. Rudwick’s study is especially useful for reminding philosophers that many scientific debates do not simply result in the victory of one of the originally expressed points of view, but that consensus is reached on a position that takes over elements from each of the original contenders.

¹¹²See, for example, Jerry Fodor, “Observation Reconsidered,” *Philosophy of Science* 51 (1984): 23–43; Paul Churchland, “Perceptual Plasticity and Theoretical Neutrality: A Reply to Jerry Fodor,” *Philosophy of Science* 55 (1988): 169–87; and Fodor, “A Reply to Paul Churchland’s ‘Perceptual Plasticity and Theoretical Neutrality,’” *Philosophy of Science* 55 (1988): 188–98.

¹¹³I consider this problem in “Authority, Deference, and the Role of Individual Reason,” forthcoming in *The Social Dimension of Scientific Knowledge*, ed. Ernan McMullin (Notre Dame: University of Notre Dame, 1992).

preliminary results reveal that the contamination of evidence is by no means inevitable.

Nor are there cogent *general* arguments from the possibility of shifting methodological standards or the omnipresence of the “experimenter’s regress.” Just as dismissals of unaided observations typically have costs—pleading hallucination is likely to flout background beliefs about the circumstances under which perceptions of various kinds of objects are reliable—so too dismissal of the problems one finds uncomfortable is liable to be constrained by background commitments.¹¹⁴ Regresses or circles in the appraisal of experiments arise when the setup or apparatus involved is detached from other types of investigation, when it is employed only in connection with some body of controversial findings.¹¹⁵ When an experimental design is used to generate a range of results, some of which are concordant with those obtained by other techniques, then the dismissal of its findings in a controversial area can be sustained only by defending a distinction in the conditions of application (as the Aristotelians tried, ultimately unsuccessfully, to limit the range of reliable application of the Galilean telescope).¹¹⁶

¹¹⁴This is readily recognizable in major episodes in the history of science: see, for example, the responses of Simplicio in Galileo’s *Dialogue*, Priestley’s continued efforts to respond to Lavoisier’s critique of the phlogiston theory, and the reactions of Poisson (and others) to Fresnel’s mathematization of the wave-front. (In this latter instance, there is no evidence of dismissal of problems, but much focus on whether Fresnel’s mathematics is correct.)

¹¹⁵Richard Miller and Ian Hacking have provided illuminating discussion of the justification of microscopic procedures: in *Fact and Method*, Miller shows how Leeuwenhoek and Hooke defended the power of the microscope to reveal a “new world”; Hacking offers less detailed coverage of a broader range of examples in *Representing and Intervening*. Collins’s contrary examples in *Changing Order* are particularly compelling precisely because they involve the use of apparatus which cannot be calibrated in unproblematic contexts.

¹¹⁶The inadequacy of Galileo’s theoretical explanation of the workings of the telescope (stressed by Feyerabend in *Against Method*) needs to be set in the context of (a) Galileo’s practical demonstrations of the reliability of the telescope on earth, (b) Galileo’s attack on the possibility of drawing a coherent distinction between the terrestrial and celestial spheres, and (c) Galileo’s linkage of his new celestial telescopic observations to phenomena that could be identified with the naked eye. So far as I know, it is an interesting, unanswered historical question whether celestial telescopic

Finally, the recognition of the social embeddedness of scientific debates, with consequent acknowledgement that at least *some* such debates are intertwined with broader questions about social order, should not lead us to replace a simplistic image of the noble scientist, moved solely by the desire for knowledge, with an equally simplistic vision of scientists as creatures of social or personal interests.¹¹⁷ There is no general move from a more realistic picture of science as practiced in social contexts to the refutation of traditional naturalism.

These remarks only clear a space within which traditional naturalism may try to take on the most important skeptical challenge to it. If I am right, (ii) is the dangerous form of skepticism, threatening to collapse traditional naturalism into a radical position that abandons or relativizes normative epistemology. The general arguments for that collapse are not cogent, and assessing the viability of traditional naturalism turns on a number of interesting questions: How penetrable is perception by cognition? What kinds of systems of authority inhibit or promote change? How is the significance of a problem or an accomplishment appraised? How are instruments and experimental designs assessed? How do social and cognitive interests combine in scientific decision making? Answers to these questions must be sought in the context of detailed studies of historical and contemporary scientific practice, if we are to determine whether science is an instrument of self-correction (as the traditional naturalists would have it) or whether it is simply a vehicle for the expression of different, incommensurable, forms of life.

8.

Further deep difficulties await traditional naturalism. Preservation of the normative project of epistemology requires develop-

observations became credible in part because of the possibility of disclosing a *sequence* of phenomena with a *sequence* of telescopes of increasing power. (This would be analogous to part of the defense of the microscope; see Miller, *Fact and Method*, 468.)

¹¹⁷There is a celebrated argument, due to Robert Merton and currently out of favor with many sociologists of science, to the effect that apparent scientific dedication to the truth is enforced by the social systems in which science is practiced. Even the most venal scientists would do well to behave as if concerned to attain the truth. See *The Sociology of Science*, 276–77.

ment of a conception of cognitive value. One complaint, (C), appeals to the history of inquiry to deny that there is any universal account that can be offered.

Before exploring potential responses to (C), it is worth asking whether traditional naturalism could preserve its normative character while conceding the objection. Some philosophers might regard (C) as a salutary counter to universalistic yearnings, which traditional naturalism shares with Fregean epistemological ventures.¹¹⁸ Traditional naturalism finds an objective standard for epistemological principles by seeing the project of inquiry as one in which cognitively limited beings, set in the actual world, seek a particular kind of representation of that world. Given the nature of the world, of the beings in question, and the kind of representation that is sought, there will be determinate answers to questions about how it is best to proceed, and hence an objective epistemological standard.

Recognizing alternative conceptions of cognitive virtue, different desirable characteristics that our representations should have, would not by itself doom the normative naturalistic enterprise. Perhaps these alternatives could be fused into a single, embracing view of the epistemic good, one that would serve as the basis of the objective standard. Or, even if the conceptions were genuinely incompatible, they might allow for a substantive practice of common normative appraisals, whichever conception was adopted. How-

¹¹⁸So, for example, a central theme of Miller's *Fact and Method* is that positivism's error lay in the search for universal principles of methodology. On Miller's account, scientific decisions are properly taken by appeal to topic-specific principles: microscopes are defended by showing that things that appear blurry to the naked eye (or to some cruder, but creditable, microscopic instrument) come into sharp focus. I believe that Miller's descriptions of the resolution of scientific debates are insightful. But it seems to me that they invite the universalist response that the topic-specific principles owe their epistemic force to the fact that they can be derived from more general methodological directives by setting particular parameters in particular contexts. The commonsense wisdom that is used to defend microscopy instantiates a principle roughly to the effect that a good way of forming beliefs is to use devices that generate a domain of results which includes a subdomain that is subject to independent check, all of whose members are recognized as correct. As will become apparent in the text, I shall suggest a similar strategy for coping with apparent variation in cognitive goals, viewing the diversity of aims found in the history of science as the elaboration of a common theme in different contexts.

ever, the historical evidence cited by the proponents of (C) is intended to disallow any such benign results. For any significant debate in the history of science, we can find, it is alleged, alternative visions of the cognitive good that will support rival systems of methodological principles, in such a way that each of the participants in the debate can appeal to one such set of principles to defend his preferred conclusion. Any practice of normative appraisal in which we engage presupposes a conception of cognitive virtue that can be challenged by those who disagree with us. Normative appraisal thus becomes a pointless kind of tub-thumping.

I shall thus suppose that (C) presents a serious challenge for traditional naturalism, and consider possible ways of responding by developing a single, compelling, conception of cognitive value. The most obvious cognitive value is truth, and perhaps the most obvious specification of the end of science is to say that we aim to achieve the truth, the whole truth, and nothing but the truth about nature. But even if we could give a clear sense to the notion of the whole truth about nature, I see no reason to think that this is anything that human beings could either attain or value. Innumerable questions about the world have no interest for us: for example, there are numerical relations among the entities in different domains, so that the number of *As* in *B* is greater than, less than, or equal to, the number of *Cs* in *D*; instantiating *A*, *B*, *C*, and *D* with certain kinds of properties quickly generates bizarre questions of crushing unimportance (“Are there more species of *Drosophila* than there are kangaroos in California?”). The aim of inquiry, I suggest, is to obtain *significant* truth. Significance is generated from our practical concerns, or from our epistemic interests.

Appeal to epistemic interests besides truth seems sufficiently nebulous that we might prefer a robust pragmatic line. Say that a question is significant for a person in a context just in case the person in that context has a goal that he is unable to achieve, which a true answer to the question would enable him to achieve. True beliefs are regarded as vehicles for obtaining practical ends, and the point of epistemology is to tell people how to appraise statements to form true beliefs and thus obtain their ends. At this point, however, we might wonder why we need take the detour through true beliefs at all. For large classes of practical ends, other kinds of representations will serve equally well—even better. Those who like to hike and hope to avoid poisonous snakes may be better

served by a principle that classifies snakes incorrectly, but that is easy to apply, than by an accurate distinction that is too difficult to apply speedily. If what we really care about are our practical goals, and if beliefs are only means to these goals, then we should not care whether those beliefs are true, only that they are effective.¹¹⁹

Dissatisfaction with a purely pragmatic reduction of the problem can be generated from thinking about a different kind of example. Some beliefs, especially scientific beliefs about human beings, make some people profoundly uncomfortable. Acceptance of these beliefs detracts from their happiness, without offering them any practical returns. From Socrates on, philosophers have encouraged people to value the truth even when it has proved uncomfortable, and it is hard not to respond to the intellectual honesty of those who do not waver.¹²⁰

An adequate account of cognitive virtue should not be simply pragmatic. The aim of inquiry is not simply to enable us to anticipate the deliverances of experience (prediction) or to shape nature to our own ends (control). Mention of two members of Claude Bernard's famous trinity recalls the third goal of science, understanding.¹²¹

¹¹⁹I here present a crude version of an argument that is elaborated with great sophistication by Stephen Stich. See his *The Fragmentation of Reason*.

¹²⁰A moving example of this is T. H. Huxley's response to Charles Kingsley's suggestion that Huxley might find comfort for his grief on the death of his beloved eldest son by accepting the consolations of Christianity. Huxley thanked his friend for his sympathy, but he was firm in his conviction that the hope he was being offered was illusory. Thus, Huxley writes: ". . . had I lived a couple of centuries earlier I could have fancied a devil scoffing at me . . . and asking what profit it was to have stripped myself of the hopes and consolations of the mass of mankind. To which my only reply was and is—Oh devil! truth is better than much profit."

Later in the same letter, Huxley elaborates on his scientific credo:

My business is to teach my aspirations to conform themselves to fact, not to try and make facts harmonise with my aspirations.

Science seems to me to teach in the highest and strongest manner the great truth which is embodied in the Christian conception of entire surrender to the will of God. Sit down before fact as a little child, be prepared to give up every preconceived notion, follow humbly wherever and to whatever abysses nature leads, or you shall learn nothing.

See Leonard Huxley, *Life and Letters of T. H. Huxley*, vol. 1 (New York: Appleton, 1913), 233, 235.

¹²¹See Bernard, *An Introduction to the Study of Experimental Medicine* (reprint, New York: Dover, 1957).

Cognitive value derives from the project of trying to understand nature. Some truths are worthless because they play no role in that project. Some falsehoods are valuable because they do play such a role. A primary task of traditional naturalism is to make these vague suggestions more precise.

One way to do so is to adopt a strongly realist approach to explanation. Minimal realism holds that there are objects independent of human cognition.¹²² Strong realism adds the thesis that, independently of us, these objects are assorted into natural kinds and that there are causal processes in which they participate. The task of science is to expose the causal structure of the world, by delineating the pre-existent natural kinds and uncovering the mechanisms that underlie causal dependencies.¹²³ Completed science would therefore not be an exhaustive description of all objects and events, but a presentation of the kinds of things there are (as in a hierarchical classification) and a delineation of the ordering of phenomena, a summation of how various kinds of effects are dependent on various types of mechanisms. Cognitively valuable statements are those that provide vivid exemplifications of these structures, as, for example, descriptions of certain types of crossings in *Drosophila* serve to expose particular genetic mechanisms.

¹²²Minimal realism also shares with strong realism the view that we are able to identify the characteristics of these objects. As Michael Devitt pithily remarks, the bare thesis that there are independent objects to which we sometimes succeed in referring is “anti-realism with a fig-leaf” (*Realism and Truth* [Princeton: Princeton University Press, 1984], 15). In formulating both versions of realism I have committed it to adoption of a correspondence theory of truth. It would be interesting to see if a similar account of cognitive value could be generated by making less substantial commitments, say by founding it on Arthur Fine’s “natural ontological attitude” (see his *The Shaky Game* [Chicago: University of Chicago Press, 1986]).

¹²³Positions of this general kind have been defended by Wesley Salmon, *Scientific Explanation and the Causal Structure of the World* (Princeton: Princeton University Press, 1984); Paul Humphreys, *The Chances of Explanation* (Princeton: Princeton University Press, 1989); Richard Boyd, *Realism* (Cambridge: Cambridge University Press, forthcoming); Nancy Cartwright, *How the Laws of Physics Lie* (Oxford: Oxford University Press, 1983); and Peter Railton, “A Deductive-Nomological Model of Probabilistic Explanation,” *Philosophy of Science* 45 (1978): 206–26. For a lucid overview, see Salmon, *Four Decades of Scientific Explanation* (Minneapolis: University of Minnesota Press, 1990).

An alternative approach avoids strong realism's commitment to independent natural kinds and causal mechanisms. Minimal realism can be developed by supposing that the assortment of objects and the ordering of the phenomena is driven not by an attempt to delineate the pre-existing causal structure of the world—for there is none—but by a search for unification. On this view, we endeavor to complete science by reducing the number of types of phenomena that we must take as basic.¹²⁴ Understanding nature consists in achieving a unified vision of it, and our only purchase on such notions as 'natural kind' and 'causal mechanism' is through recognizing the role that categories and relations of dependence play in our ordering of the world.

Plainly these views differ in their metaphysics, but they agree in taking as the goal of inquiry the production of a certain type of structured account.¹²⁵ For present purposes, we can overlook the important divergence concerning what makes that account correct, focusing on the similarities. If we are interested in cognitive virtue, then, in either case, we can see cognitive virtue as determined in terms of contributions to the structured account of nature.¹²⁶ Large numbers of truths will be quite irrelevant to the provision of the account and so will be cognitively worthless. Conversely, there will be falsehoods that are relevant, and thus count as cognitively valuable. When the phenomena concern the behavior of complex systems, delineating how basic mechanisms combine or how fundamental patterns of explanation are synthesized will sometimes hide the main lines of dependence. In such cases, we expose the

¹²⁴A characteristically lucid expression of this ideal is in T. H. Huxley: "In the end, the fundamental phenomena are incomprehensible and the aim of science is to reduce the fundamental incomprehensibilities to the smallest possible number" (*Darwiniana* [New York: Appleton, 1886], 165). For modern versions of the idea, see Michael Friedman, "Explanation and Scientific Understanding," *Journal of Philosophy* 71 (1974): 5–19; and my own "Explanatory Unification," *Philosophy of Science* 48 (1981): 507–31, and "Explanatory Unification and the Causal Structure of the World," in *Scientific Explanation*, ed. P. Kitcher and W. Salmon, vol. 13 of *Minnesota Studies in the Philosophy of Science* (Minneapolis: University of Minnesota Press, 1989), 410–505.

¹²⁵Differing from views like that offered by Bas van Fraassen in *The Scientific Image* (Oxford: Oxford University Press, 1980).

¹²⁶This idea is lucidly presented by Railton in "Probability, Explanation, and Information," *Synthese* 48 (1981): 233–56.

structure of nature through idealization, by pretending that the world is simpler than it is and offering an account that would be true only if some complications were removed.¹²⁷

The account I have sketched may easily be summarized in Aristotelian terms: inquiry aims to uncover the order of being. Strong realists think that the order of being is found, minimal realists that it is made. Assuming that one of these versions can be worked out in detail, the idea of autonomous cognitive virtue can be made less nebulous. However, this is only the beginning. Traditional naturalists are still vulnerable to the charge that their account of the epistemic good is parochial, consonant only with the views of a particular group at a particular epoch.¹²⁸

Complaint (C) denies the possibility of demarcating a single conception of cognitive value, shared by all those who regard themselves as pursuing projects of pure inquiry. Critics will contest the view that science from Aristotle to the present aims to uncover the order of being. They can point to the declarations of prominent scientists: to the advocates of science in the seventeenth century who regarded their project as that of revealing “the wisdom of God in the Creation,” to those who contended that the aim is to “save the phenomena,” to recurrent conflicts between realists and instrumentalists.¹²⁹ Such declarations seem to force traditional naturalism either to discount the claims of some eminent practitioners or to abandon the idea of a single conception of cognitive value.

Matters are not so desperate. Pious seventeenth-century campaigners for the pursuit and institutionalization of science had to convince their contemporaries of the value of the project of fathoming the order of being *relative to other human concerns*. They did so by suggesting that pure understanding would issue in greater

¹²⁷See Nancy Cartwright, “The Simulacrum Account of Explanation,” in *How the Laws of Physics Lie, and Nature’s Capacities and their Measurement* (Oxford: Oxford University Press, 1989).

¹²⁸There are also questions, which I shall not take up here, concerning how human cognitive ends (if understood along the lines just suggested) should be balanced against our practical concerns. For purposes of idealization, we can focus simply on the cognitive goals of inquiry, but we should acknowledge that this is only an idealization.

¹²⁹See, for example, Laudan, *Science and Values* (Berkeley: University of California Press, 1984); Shapin and Schaffer, *Leviathan and the Air-Pump*; and numerous historical studies.

mastery of nature (an argument renewed in nineteenth-century America,¹³⁰ and popular in testimony before Congress to this day), and, given the religious sentiments manifest in public discussions, that it would constitute a mode of worship. The fact that the purely epistemic goal can *also* be seen as a means to other desirable ends is compatible with its *intrinsic* desirability.

The harder cases are those in which scientists seem prepared to settle for less than the delineation of the order of being. But here, I think, we can avail ourselves of the distinction between fundamental goals and those that are derived from them within particular contexts.¹³¹ Discussions of anthropological relativism make plain how easy it is to impute variable ends by failing to allow for the possibility that common goals are articulated differently in different circumstances. In the history of science, too, debates as to whether a discipline should expose real causal dependencies or merely “save the phenomena” do not call into question the intrinsic value of the more ambitious goal. The issues concern attainability. The medieval astronomers whose criterion for success was the prediction and retrodiction of planetary motions despaired of combining their predictive models of the solar system with any account of the cosmological mechanisms. The nineteenth-century chemists who viewed molecular formulae as tools for the generation of accurate results about reactants and products in chemical combination doubted whether we could ever have access to the underlying processes.¹³²

Reflection on such examples should lead to an amendment of the account of cognitive value. The goal of pure inquiry is to produce a structured account of nature *insofar as that is possible for limited beings like ourselves*. I claim that the history of science is the history of attempts to achieve this goal—as well, of course, as to advance projects that are counted as practically significant—and that the explicit differences of aim that are found result from al-

¹³⁰See Daniel Kevles, *The Physicists* (New York: Knopf, 1971).

¹³¹This point is made by Alexander Rosenberg, “Normative Naturalism and the Role of Philosophy,” *Philosophy of Science* 57 (1990): 34–43.

¹³²For useful insights into arguments for and against instrumentalist conceptions of particular scientific theories, see Michael Gardner, “Realism and Instrumentalism in Nineteenth-Century Atomism,” *Philosophy of Science* 46 (1979): 1–34.

ternative conceptions of what it is possible for us to do. All this is concordant with the traditional naturalist's picture of limited agents struggling to attain a fixed goal, and finding, as they go, what they can expect to achieve and how they should plan their pursuit.

9.

I shall deal more briefly with the last two complaints, both of which prod traditional naturalism further away from traditional epistemology, but which do not challenge the retention of the normative dimension. The first of these, (D), questions the idiom in which epistemological issues are formulated. Epistemologists and philosophers of science routinely treat cognition in terms of the adoption or rejection of statements, in terms of belief, truth, justification, and evidence. These categories, it is suggested, derive from a venerable epistemological lineage—why should we think that its terms are sacrosanct? Sometimes the query is posed by suggesting that we have a “folk theory” that is just as vulnerable to displacement as were folk theories of heat, vision, or motion.¹³³ Moreover, the track record of our folk theory of mental life is not very impressive: psychology, we are told, has managed to accomplish very little by articulating propositional models of cognitive activity. Such propositional models also appear to separate our own species from other cognitive creatures (and human infants) who lack our impressive linguistic performances. Finally, recent studies of the brain, the medium of cognitive activity, reveal that it has an architecture seemingly designed for a different type of processing than the symbol manipulation beloved of folk psychology.

There are numerous well-known elaborations of and responses to these arguments.¹³⁴ My interest here is solely in how they should affect the *current* practice of naturalistic epistemology and philosophy of science. No doubt the central categories of epistemology

¹³³However, it is a bit hard to accuse the “folk” of subscribing to the full philosophical apparatus of belief, propositions, justification, and all the rest.

¹³⁴See Stich, *The Case Against Belief* and *The Fragmentation of Reason*; Patricia Churchland, “A Perspective on Mind-Brain Research,” *Journal of Philosophy* 77 (1980): 185–207, and *Neurophilosophy*; and Paul Churchland, *A Neurocomputational Perspective*. Among the replies are Barbara Von

might be displaced in the further course of scientific research. The cognitive life of human beings might be described in quite different terms—terms from the language of neural nets, say—and cognitive virtue might also be identified in these terms. However, in advance of developing this language in sufficient detail to account for the sophisticated reasoning that appears to occur in human inquiry, there is no way of formulating naturalistic claims about cognitively optimal strategies. The very advantage on which eliminativists sometimes insist—to wit, the display of kinship between human beings and other cognizers—is also a bar to the adumbration of naturalistic epistemology along eliminativist lines. For the goal of naturalistic epistemology and philosophy of science is to understand and improve our most sophisticated performances, and about these eliminativists have presently very little to say.

If we accepted the eliminativist indictment of traditional propositional approaches to cognition, then prospects for naturalism would be discouraging. In effect, we would be confronted with the choice between an inadequate idiom and one not yet developed. However, attacks on traditional ventures in cognitive psychology and cognitive science generally are often overstated. Furthermore, recent developments in these areas enlarge our vision of cognition in ways that are directly applicable to the study of historical and contemporary science. Although historians and sociologists have demonstrated forcefully that much of the cognitive activity of scientists is not obviously propositional, their focus has been primarily on scientific use of images and on the unarticulable skills (“tacit knowledge”) of scientists.¹³⁵ Both of these kinds of cognition are explored in work in contemporary cognitive science that retains the categories of traditional epistemology.¹³⁶ Naturalistic projects can be formulated and pursued within the framework suggested by

Eckhardt, “Cognitive Psychology and Principled Skepticism,” *Journal of Philosophy* 81 (1984): 67–88; Patricia Kitcher, “In Defense of Intentional Psychology,” *Journal of Philosophy* 81 (1984): 89–106; and Terry Horgan and James Woodward, “Folk Psychology is Here to Stay,” *Philosophical Review* 94 (1985): 197–226.

¹³⁵See Polanyi, *Personal Knowledge*; Rudwick, *The Great Devonian Controversy*; Ravetz, *Science and its Social Problems*; Collins, *Changing Order*; Gould, *Wonderful Life* (New York: Norton, 1989).

¹³⁶See, for example, John Anderson, *The Architecture of Cognition* (Cambridge: The MIT Press, 1983); and Holland et al., *Induction*. For applications to philosophical issues, see Goldman, *Epistemology and Cognition*.

such developments, enriching our epistemological vocabulary rather than eliminating it.

Two generic research strategies are obvious: (i) aim to develop the preferred rival idiom and defer projects of epistemic appraisal until they can be reformulated in these terms, and (ii) continue to use whatever resources from empirical studies of cognition can be used to formulate and address normative epistemological enterprises. It seems to me wrong for the philosophical community to pursue either of these strategies to the exclusion of the other.

My conclusion exemplifies an attitude to the growth of knowledge that is encapsulated in (E), a complaint which, I believe, should lead to an enrichment of traditional naturalism.¹³⁷ In rejecting the *a priori*, traditional naturalists recognize both the historical and social embeddedness of knowledge. Consequently, the enterprise of providing counsel about belief formation needs to be modified, to take into account the fact that the individual's own efforts at extending or revising belief may signify rather little in comparison with what is inherited from others. From the perspective of any individual, then, it is a matter of some importance that the achievement and transmission of consensus within a community promote cognitive ends. Instead of simply asking how individuals should adjust their beliefs, we should also investigate the distribution and dissemination of belief within communities.

Distinguish two kinds of cognitive goals that people have. On the one hand, there are *personal epistemic ends*: the subject aims to attain true beliefs (let us say). On the other, there are *impersonal epistemic ends*: the subject intends that the community to which she belongs attain the cognitive good. The pursuit of normative epistemology, taken over within traditional naturalism, approaches the question of specifying how individuals should achieve their personal epistemic ends by treating them in isolation from the communities to which they belong, effectively ignoring the possibility that cognitively optimal strategies may involve coordination of effort with

¹³⁷The enrichment consists in taking seriously the social aspects of knowledge. A few writers have recently begun to emphasize the importance of the issues here: see, for example, the final chapter of Husain Sarkar, *A Theory of Method* (Berkeley: University of California Press, 1983); John Hardwig, "Epistemic Dependence," *Journal of Philosophy* 82 (1985): 335–49; Alvin Goldman, "Foundations of Social Epistemics"; and David Hull, *Science as a Process*.

others. As I have remarked above, this is at odds with the traditional naturalist's reasons for rejecting the *a priori*.

Two main issues arise here: How do communities of individuals best carry out their distributed cognitive activity for the attainment of impersonal epistemic ends? How does the possibility of coordinating their efforts with those of others affect individuals' optimal attainment of their personal epistemic ends? We can approach the former by thinking of a community epistemic predicament as defined by (a) a set of individuals, (b) a distribution of beliefs among these individuals, and (c) a set of rival cognitive objects (hypotheses, theories, research programs, methods, and so forth). The task is to specify a distribution of decision rules for individuals and a distribution of activities by individuals that will yield the highest expected cognitive return (most effectively promote community attainment of the cognitive good). For some community epistemic predicaments, cognitive uniformity is the best solution. In other instances, we want to promote cognitive diversity, allowing for a genuine division of cognitive labor.¹³⁸

Actual communities are not planned in the light of analyses of what is cognitively optimal. Distributions of cognitive effort come about through the presence of cognitive characteristics of individuals and of social relations among them. Plainly, people differ in all kinds of cognitive propensities: some are attracted to novel ideas, others prefer to defend orthodoxy; because of differences in prior experience, people find different kinds of phenomena salient.¹³⁹ Moreover, human beings stand in complex social interrelations and are motivated by nonepistemic ends. Given a community epistemic predicament and the specification of an optimal solution, we should ask how various combinations of psychological distributions and social forces bear on the attainment of the optimum. As I have argued elsewhere, forces that have often been viewed as interfering with the advance of knowledge may sometimes play a productive role in the cognitive behavior of the community.¹⁴⁰

¹³⁸See my "The Division of Cognitive Labor."

¹³⁹This general point has been made in various ways by Kuhn ("Objectivity, Value-Judgment, and Theory Choice"), Ronald Giere (*Explaining Science* [Chicago: University of Chicago Press, 1988]), Howard Margolis (*Patterns, Judgment and Cognition* [Chicago: University of Chicago Press, 1987]), and Frank Sulloway ("Orthodoxy and Innovation in Science" [forthcoming]).

¹⁴⁰See "The Division of Cognitive Labor." This is not to deny that wel-

For the first problem, then, the enterprise of social epistemology consists in identifying community epistemic predicaments, solving the resultant optimization problems, and considering the ways in which various kinds of psychological and social factors contribute to or detract from the community pursuit of a good strategy.¹⁴¹ The second problem raises fundamental questions about authority, trust, and cooperation. Individual subjects often have the option of relying on others or trying to work matters out for themselves. When is it better to defer to others? How does one decide whom to trust? How is the individual's pursuit of personal epistemic ends to be reconciled with the role which each person ought to play within a well-functioning cognitive community?

Questions about authority and trust naturally arise when scientists announce heterodox findings. The decision to investigate such claimed results presupposes that one does not simply take over the prior probability of the claim (approximately zero, since it conflicts with firmly entrenched beliefs) but allows for the authority of those who report it. How exactly should one assess the claim in the light both of the implausibility of its content and of the authority of the announcers? How should their authority be evaluated?

Issues about cooperation result from recognizing the possibility of conflict between personal epistemic ends and impersonal epistemic ends (the ends of the community). From my perspective, the chances of attaining true belief may be enhanced (at least in the short run, the time scale that matters to me) if I pursue a different course of cognitive activity from that which would advance the community epistemic project. Is it possible that a cognitive community can degenerate into a set of consumers who exploit available resources rather than improving them? How can cognitive cooperation be sustained?

These questions deserve more careful formulations than I can give them here, and they deserve more thorough answers than

come distributions of effort might come about in other ways, for instance, through differences in cognitive style or cognitive "resources" (in Giere's useful term [*Explaining Science*, 213–14]). Miriam Solomon and Alvin Goldman have pursued related themes in currently unpublished work.

¹⁴¹Ideas from explicitly evolutionary accounts of the growth of knowledge (or, more generally, culture) are relevant here. For prime examples, see Robert Boyd and Peter Richerson, *Culture and the Evolutionary Process* (Chicago: University of Chicago Press, 1985), and David Hull, *Science as a Process*.

anyone has yet attempted.¹⁴² I hope that even a brief outline of them will exhibit the possibility of serious social epistemology and, in consequence, a legitimate extension of traditional naturalism's normative enterprise.

10.

My aim in this essay has been to chart the ways in which naturalism emerged from post-Fregean epistemology, to identify major naturalist theses, and to explore the possibility of pursuing the meliorative project of Bacon and Descartes within a naturalistic idiom. The route to naturalism involves at least a minimal reintroduction of psychology into epistemology (acceptance of the idea that the epistemic status of a belief state depends on the psychological processes that generate and sustain it) and the rejection of the *a priori*. In my judgment, these modifications of normative epistemology are difficult to resist. The most promising way of salvaging post-Fregean epistemology seems to me to be to minimize the importance of the changes. Hence my discussion of the complaint (A) with its suggestion that the shift to naturalism permits business as usual.

If, as I believe, (A) should be rejected, then the central question is whether naturalism allows any way to save the traditional meliorative project of epistemology. The two main issues here are the possibility of our sustaining the reliability of the historical process through which human knowledge has emerged, given a naturalistic perspective, and the availability of an account of cognitive value. These issues are complex and ramify beyond the bounds of this paper. In discussing them, I am conscious that I have only indicated lines of argument, suggesting possibilities for naturalistic epistemology to pursue.

Finally, I have considered two ways in which naturalistic epistemological projects might be recast. One would reform the basic language in which we discuss cognition; the other would take account of the social dimensions of human knowledge. As I have

¹⁴²For preliminary views about the possible roles of authority, see my "Authority, Deference, and the Role of Individual Reason."

indicated, I see no way at present to undertake the former venture, while pursuit of the latter seems long overdue.

Pre-Fregean epistemology studied human knowledge as part of the natural order, using psychological language which seemed quaint and repugnant to early-twentieth-century ears. But, after almost a century of eclipse, the naturalists have returned, campaigning for the relevance of psychology and biology to epistemology and denying (in contrast to most of their pre-Fregean ancestors) the possibility of *a priori* knowledge. Traditional naturalism is the obvious heir of pre-Fregean epistemology, because it preserves the ideal of a meliorative project. While academic philosophers tend to see the normative character of epistemology as a commonplace, the more general intellectual tendency of the age is to regard affirmations of normativity as antediluvian. A disturbingly large number of contemporary intellectuals perceive post-Fregean, “analytic,” “pure” philosophy as having collapsed. They conclude that this is the death of philosophy, and that the succession passes variously to history, sociology, or literary theory.¹⁴³ By focusing on the viability of traditional naturalism, I have tried to show that there are relatively unappreciated possibilities of extending the tradition. I cannot claim to have demonstrated that normative epistemology can survive a naturalistic metamorphosis, but I hope to have shown that reports of its demise are greatly exaggerated.¹⁴⁴

University of California, San Diego

¹⁴³See, for example, Rorty, *Philosophy and the Mirror of Nature*; Feysabend, *Farewell to Reason*; and numerous literary theorists.

¹⁴⁴I am grateful to the John Simon Guggenheim Foundation for support that enabled me to work on a longer study that pursues some of the themes I develop here, and to the University of California President's Research Initiative in the Humanities, which provided funds for a workshop on Naturalizing the Philosophy of Science (which I directed jointly with Stephen Stich). The participants in that workshop have influenced my ideas and formulations in important ways, and I am particularly indebted to Stich for numerous conversations about the topics discussed in this paper. I would also like to thank Michael Bishop, Fred Dretske, Peter Godfrey-Smith, Alvin Goldman, Patricia Kitcher, Larry Laudan, Gila Sher, and Elliott Sober, all of whom provided valuable comments on an earlier draft, and the editors of the *Philosophical Review*, who offered extensive, thoughtful suggestions. It should not, of course, be assumed that anyone agrees with the final version.