the Daltonian approach, many nonetheless were prepared to take it seriously, claiming that the serendipity of the Daltonian system made it at least sufficiently promising to be worthy of further development and refinement.

Whether the approach taken here to the problem of "rational pursuit" will eventually prevail is doubtful, for we have only begun to explore some of the complex problems in this area; what I would claim is that the linkage between progress and pursuit outlined above offers us a healthy middle ground between (on the one side) the insistence of Kuhn and the inductivists that the pursuit of alternatives to the dominant paradigm is *never rational* (except in times of crisis) and the anarchistic claim of Feyerabend and Lakatos that the pursuit of *any* research tradition—no matter how regressive it is—*can always be rational.*

Ronald N. Giere*

PHILOSOPHY OF SCIENCE NATURALIZED

In arguing a "role for history," Kuhn was proposing a naturalized philosophy of science. That, I argue, is the only viable approach to the philosophy of science. I begin by exhibiting the main general objections to a naturalistic approach. These objections, I suggest, are equally powerful against nonnaturalistic accounts. I review the failure of two nonnaturalistic approaches, methodological foundationism (Carnap, Reichenbach, and Popper) and metamethodology (Lakatos and Laudan). The correct response, I suggest, is to adopt an "evolutionary perspective." This perspective is defended against one recent critic (Putnam). To argue the plausibility of a naturalistic approach, I next sketch a naturalistic account of theories and of theory choice. This account is then illustrated by the recent revolution in geology. In conclusion I return to Kuhn's question about the role of history in developing a naturalistic theory of science.

1. KUHN'S NATURALISM

In the very first chapter of *The Structure of Scientific Revolutions,* Kuhn sought to establish "a role for history." Part of that role, he implied, is as data for "a theory of scientific inquiry." And by "theory" he meant something comparable to theories in the sciences themselves. Thus, referring to standard philosophical distinctions, such as that between "discovery" and "justification," he wrote:

> Rather than being elementary logical or methodological distinctions, which would thus be prior to the analysis of scientific knowledge, they now seem integral parts of a traditional set of substantive answers to the very questions upon which they have been deployed. That circularity does not at all invalidate them. But it does make them parts of a theory and, by doing so, subjects them to the same scrutiny regularly applied to theories in other fields. If they are to have more than pure abstraction as their content, then that content must be discovered by observing them in application to the data they are meant to eluci-

*The support of the National Science Foundation is hereby gratefully acknowledged. My colleagues at Indiana and a reviewer supplied many helpful suggestions.

Reprinted by permission from the author and *Philosophy of Science,* Vol. 52, September 1985.

date. How could history of science fail to be a source of phenomena to which theories about knowledge may legitimately be asked to apply? ([1962] 1970, p. 9)

Although he did not use exactly these words, Kuhn was advocating a *naturalized* philosophy of science.

The many philosophical criticisms of Kuhn's work focused mainly on the details of his naturalistic account. His account of revolutionary theory change, which invokes only naturalistic notions like gestalt switches and persuasion, was a frequent target. Few critics, however, raised the general question whether *any* purely naturalistic theory of science might be correct. No doubt this was due to the unquestioned presumption that no such account could be correct. It is precisely this presumption I wish now explicitly to challenge.

For some, I admit, no challenge is necessary. Some philosophers regard the philosophy of science as merely a branch of epistemology. And some of these philosophers follow Quine (1969) in the project of naturalizing epistemology. Others embrace a version of evolutionary epistemology. But these are still minorities. Naturalism in the philosophy of science is generally rejected not only by the successors to logical empiricism but also by most of those who agree with Kuhn in adopting a historical methodology. Thus Lakatos (1970), Toulmin (1972), Laudan (1977), and Shapere (1984) have each sought to show that the process of scientific inquiry is not only historical but *rational* as well. Rationality is not a concept that can appear in a naturalistic theory of science—unless reduced to naturalistic terms.[1]

My argument will be both negative and positive. I will first exhibit what seem to be the main general objections to a naturalistic approach. These objections, I suggest, are too strong. There seems no viable non-naturalistic way around them either. I review the failure of two such non-naturalistic approaches, methodological foundationism (Carnap, Reichenbach, and Popper) and metamethodology (Lakatos and Laudan). The correct response, I suggest, is to adopt an "evolutionary perspective." This perspective is defended against one recent critic (Putnam). To argue the plausibility of a naturalistic approach, I next sketch a naturalistic account of theories and of theory choice. This account is illustrated by the recent revolution in geology. In conclusion I return to Kuhn's question about the role of history in developing a naturalistic theory of science.

2. SOME ARGUMENTS AGAINST NATURALISM

The following are some general forms of argument that one would expect to be raised against any proposal to naturalize the philosophy of science.

The circle argument. The general idea behind the circle argument is that the use of scientific methods to investigate scientific methods must be circular, beg the question, or lead to a regress. A more explicit version of the argument might go something like this: One of the things any study of science must investigate is the methods (criteria, canons, etc.) scientists use in evaluating evidence. To pursue such an investigation *scientifically* requires using data about scientific practice to reach conclusions about scientific methods. Thus, any empirical investigation aimed at discovering the criteria that scientists use for evaluating evidence would necessarily presuppose at least some of the criteria it was supposedly setting out to discover. So not *all* the methods of science could be discovered by scientific investigation. At least *some* must be discoverable by other means.[2]

The circle argument is a version of classic arguments concerning the justification of induction. This relationship may partly explain the power of the argument within the philosophical community.

The argument from norms. This argument appeals to the distinction between facts and norms. A naturalistic study of science, it is claimed, could at most *describe* the methods scientists use in coming to adopt hypotheses or theories. The goal of the philosophy of science, however, is not merely to *describe* the methods scientists employ, but to *prescribe* what methods they *should* employ. We want to know not merely what criteria scientists in fact use in adopting theories; we want to know which are the *right* criteria. A naturalistic philosophy of science would be powerless to answer such questions.

The argument from relativism. This argument may be viewed as a corollary to the argument from norms, but relativism has been so much discussed of late that this form of argument deserves independent billing. The argument has the form of a *reductio.* A naturalistic philosophy of science, it is claimed, would be powerless to distinguish good from bad science. It would, for example, have to treat "creation theory" on a par with evolutionary theory. Such a philosophy of science would be at best worthless, at worst pernicious.

These, in brief, are some of the main arguments against a naturalized philosophy of science. A reply takes a bit longer.

3. METHODOLOGICAL FOUNDATIONISM.

Circle arguments have always been among the most powerful in the philosopher's arsenal. Their use, however, commits the user to constructing a defense against a similar attack. Regarding our particular circle argument, the traditional defense has been some form of *methodological foundationism*—the construction of a method whose correctness can be certified a priori.

One connection between naturalism and foundationism has been well charted by Quine (1969). For Quine it was the foundationist inability to reduce mathematics to logic (or semantics to behavior) that left us no alternative but to naturalize epistemology. It requires only a little elaboration to see that a similar connection exists within the philosophy of science.

When Kuhn's book first appeared, the methodologically foundationist programs of Carnap, Reichenbach, and Popper were among the most active areas of research in the philosophy of science. Carnap and Reichenbach, though not Popper, were also tempted by foundationism with regard to the data furnished by experience. Since this latter aspect of foundationist programs is not at issue here, I shall say no more about it.

Carnap was originally attracted to Russell's foundationism which utilized the method of logical construction developed in the context of the foundations of mathematics. Here the regress stops at a foundation of *logic.* Logic also provides a normative component for scientific reasoning and a bulwark against relativism.

Carnap and the early Logical Empiricists gave up Russell's strict form of methodological foundationism not so much for technical reasons as for broadly empirical reasons. They concluded that the methods of logical constructions could not yield the laws of physics as they understood them, and they were unwilling to reject the laws of physics as philosophically unsound. But few were willing to give up the idea that logic provides the foundation for scientific method.[3]

Even Popper, who otherwise was quite critical of many Logical Empiricist doctrines, rested his methodology on the simple logical rule of *modus tollens.* Here again, some of the most severe criticisms of Popper's methodology have been broadly empirical. Adopting Kuhn's claim that all theories have faced anomalies throughout their careers, Lakatos (1971), for example, argued that if we follow Popper's rules, we should have to regard all theories as falsified. Assuming falsified theories should be rejected, we reach the empirically unacceptable conclusion that all theories should be rejected.

Among the major Logical Empiricist figures, Reichenbach was unusual in seeking a methodological foundation not in logic but in a pragmatic rule of action. Assuming that the ultimate goal of science is to determine limiting relative frequencies in infinite sequences, he offered an *a priori* argument that his inductive rule of inference, the "straight rule," guaranteed success in reaching this goal so long as the goal was obtainable at all. Unfortunately, there is a continuum of "crooked rules" for which the same justification can be given. Hacking (1968) delivered the *coup de grâce* to the program by showing that any long-run justification sufficient to justify only the straight rule required the *empirical* assumption that the sequences in question be *random*. This exposed the program to the regress argument it was explicitly designed to elude. Reichenbach's program could also have been criticized on the *empirical* ground that science in fact has stronger goals than the long-run discovery of limiting relative frequencies. But this was not the argument that led to its demise.

During the 1940s, after moving to the United States, Carnap took up Keynes's program of developing an *inductive logic* that would be a formal generalization of deductive logic. His own semantic theories of the previous decade provided the formal background for this attempt. Being a logic, this inductive logic would stop the regress and provide the norms to defeat relativism. Carnap's inductive logic has been criticized on the empirical ground that it is too simple to tell us anything about the evaluation of actual scientific theories. The main reason most people gave up the program, however, was more technical. The logic requires a measure of the initial probability of all hypotheses. But the space of possible measures, like the space of Reichenbachian rules of inference, is so large that there seems no *a priori* and nonarbitrary way to justify a unique measure.

Recalling that the first chapter of Carnap's *Logical Foundations of Probability* (1950) was entitled "On Explication," we are reminded of another program for grounding a particular inductive logic, namely, as an explication of our prereflective concept of evidential support. But how do we know when we have correctly captured this concept, or even that "we" have a univocal concept of this sort? Carnap says only that the "explicatum" must be "similar to" the "explicandum." If this similarity is to be determined empirically, and there seems no other way, then Carnap too was caught in the circle argument.

Richard Jeffrey (1973) once argued that the correct inductive logic would be the one that eventually agrees with our inductive intuitions when the Carnapian program is sufficiently developed for more complex languages. To avoid the circle, this view must assume that our intuitions are given directly without empirical investigation. Moreover, this interpretation leaves it open whether the program ever will be sufficiently developed—which makes it impossible for us currently to say whether science ever has been or is now a rational enterprise. This is a rather weak foundation. And like any explication, it provides no protection from relativism. The logic is at best *descriptive* of *our* intuitions. It does not insure us that our intuitions themselves are correct.

The major remaining stronghold of methodological foundationism is "Bayesian inference" or one of its near relatives. Here the problem of picking out a unique initial probability measure is avoided by relativizing to an individual agent. Individuals supply their own initial measures. Rationality then consists in how one *revises* probability assignments in light of new evidence. It has often been objected that Bayesian inference goes too far in the direction of relativism. Moreover, the same type of problem that plagued both Carnap and Reichenbach arises here too because there are many different logically possible ways of "conditionalizing" on the evidence, and no *a priori* way of singling out one way as uniquely rational. One is reduced either to appealing to something like "explication," or to investigating actual reasoning, which reintroduces the circle.[4]

Adopting Quine's form of inference, I should like now to conclude that methodological foundationism is a hopeless program and thus that naturalism, in spite of the circle argument, is our only alternative. There is, however, a further line of inquiry that must be considered, if only because it has been so prominent in the post-Kuhnian literature.

4. METAMETHODOLOGY

Imre Lakatos (1971) introduced the term "metamethodology" to describe his method for investigating the relative superiority of any proposed theory of scientific method. Laudan (1977) adopted a similar strategy—though differing in detail. For brevity of exposition, I will concentrate on Laudan's approach.[5]

The connection between metamethodology and the circle argument arises as follows. If Lakatos and Laudan really had been taking a naturalistic approach to methodology, they would have adopted the reflexive strategy of applying their methodology to itself. This, however, is not their official doctrine. That they deliberately avoided a reflexive strategy *because* of its obvious circularity, I cannot say. Their metamethodologies, however, are not reflexive and thus not blatantly circular. Whether they can achieve their ends while still avoiding circularity is another question.

In Laudan's theory of scientific rationality, the measure of progress, and therefore of rationality, is problem-solving effectiveness. Roughly speaking, the problem-solving effectiveness of a research tradition is the weighted number of empirical problems solved by its latest theory minus the weighted number of outstanding anomalies and conceptual problems. Problems are weighted by their importance to the research tradition. The relative acceptability of one research tradition over another is determined by its relative problem-solving effectiveness. This measure, Laudan claims, provides a *rational* way of deciding the relative acceptability of two research traditions.

Laudan's metamethodological strategy is to seek first a set of "preferred pre-analytic intuitions about scientific rationality" (PIs). That is, looking at the history of science, we find

> a subclass of cases of theory-acceptance and theory rejection about which most scientifically educated persons have strong (and similar) intuitions. This class would include within it many (perhaps even all) of the following: (1) it was rational to accept Newtonian mechanics and to reject Aristotelian mechanics by, say, 1800; . . . ; (4) it was irrational after 1920 to believe that the chemical atom had no parts;. . . . (1977, p. 160)

The next step is to apply the methodology (Laudan's theory) to the PIs in order to determine the relative problem-solving effectiveness of the traditions in question. This assumes, of course, that we can indeed identify, count, and weigh the relevant problems. Comparison of computed problem-solving effectiveness will tell us which tradition was in fact most progressive and thus which should have been accepted according to Laudan's methodology. "The degree of adequacy of any theory of scientific appraisal is proportional to how many of the PIs it can do justice to" (1977, p. 161).

Assume, for the sake of argument, that Laudan's methodology agrees with *all* the PIs. Could we then be confident that it is "a sound explication of what we mean by rationality" (1977, p. 161)? I think not. The most one could conclude is that Laudan has identified a highly reliable *symptom* of the basis for our pre-analytic judgments of theory-acceptance and theory-rejection.

Suppose, contrary to Laudan, that our pre-analytic judgments are really based on an assessment of the approximate *truth* of the theories in question, and that we take problem-solving effectiveness as our best *evidence* for approximate truth. Laudan's method of assessment would then yield the same judgments of acceptance and rejection, but fail to capture the real basis of our judgments. The trouble is that comparison with our gross judgments of acceptance and rejection does not test the fine structure of the methodological theory. To test the fine structure, however, would require a more detailed empirical inquiry, and this would immediately raise the problem of circularity.

It may be, however, that Laudan would not be all that bothered by learning that he has not avoided circularity. His main concern, like that of Lakatos before him, is to have a *normative* theory of rationality. Let us, therefore, move on to the argument from norms.

Does the type of metamethodology advocated by Lakatos and Laudan yield methodological principles which are genuinely normative? Not really. At its most successful, the metamethodology would tell us only that we had discovered a general *description* of situations which we intuitively regard as clear cases of rational acceptance or pursuit. We might have correctly identified the descriptive component of the methodology, without capturing its normative force. To claim we had captured the normative component would require that we make the judgments we do *because* of considerations based on problem-solving effectiveness. In Kantian terms, Laudan's metamethodology could at most show only that we are acting in accord with his methodology, not that we are acting out of regard for that methodology. It cannot show that his methodology is actually embodied as a norm in our judgments.

This point is all the more pronounced if we consider not merely our own current preferred intuitions, but those of the historical actors in the episodes considered. Laudan does not attempt to show that actual scientists in historical contexts made the judgments they did because of considerations of problem-solving effectiveness. He is content to point out the correlation between their judgments and our calculations of actual problem-solving effectiveness. That is scant evidence that such considerations were normatively operative at the time.

Being at bottom a strategy for explication, not justification, Laudan's metamethodology also fails to provide a strong defense against relativism. Questions about the rationality of the whole western scientific tradition are ruled out because the metamethodology begins with the assumption that some judgments (the PIs) are rational. It is these we use to test the theory of rationality. This Laudan freely admits (p. 161). He fails to point out, however, that this leaves us defenseless against the cultural anthropologist who claims that the belief systems of non-Western cultures cannot rationally be judged by the standards of Western science.

5. AN EVOLUTIONARY PERSPECTIVE

I conclude that neither methodological foundationism nor metamethodology can break the circle or provide the norms needed to defeat relativism. This hardly proves that there is no way to achieve these ends. It does, however, provide some motivation for seeking to *understand* how a naturalized philosophy of science might fruitfully be pursued. I would suggest that evolutionary theory, together with recent work in cognitive science and the neurosciences, provides a basis for such an understanding. The following is the barest sketch of how the story might go.[6]

Human perceptual and other cognitive capacities have evolved along with human bodies. We share many of these capacities with other primates and even lower mammals. Indeed, those parts of our brains responsible for our more advanced linguistic abilities are built upon and linked to those parts that we share with other mammals. There can be no denying that these capacities are fairly well adapted to the environment in which they evolved. Without considerable adaptation, we would very likely not be here. Nor are these capacities trivial. The amount of perceptual and neural processing required just for a human to walk without falling or bumping into things is fantastically large and very complex.

The capacities evolution favors, of course, are just those that confer biological fitness, that is, the ability to survive and leave offspring. The ability to do modern science had nothing to do with the evolution of our perceptual and cognitive capacities—indeed, doing science may very well be detrimental to our survival as a species. The general problem faced by a naturalistic philosophy of science, then, is to explain how creatures with our natural endowments manage to learn so much about the detailed structure of the world—about atoms, stars and nebulae, entropy and genes. This problem calls for a *scientific* explanation.

Empiricist philosophers emphasized the role of immediate perceptual experience in their analyses of knowledge because of the high degree of subjective certainty attached to such experience. From an evolutionary perspective, the subjective certainty is indeed causally connected with the more direct source of the reliability of such judgments, which lies in our evolved capacities for interacting with our world. But the operation of these capacities is largely unrecorded in our conscious experience. Rationalist philosophers, on the other hand, focused on our more general subjective institutions, such as, that space has three dimensions and that time exhibits a linear structure. These judgments seem to be built into the way we think. And indeed they are, for the aspects of the world relevant to our biological fitness have roughly that structure.

Neither empiricists nor rationalists could see how to get beyond their subjective experience or intuitions. This led to the familiar philosophical views that the world is nothing more than the sum total of our sense experience or that it is totally unknowable. In fact, we possess built-in mechanisms for quite direct interaction with aspects of our environment. The operations of these mechanisms largely bypass our conscious experience and linguistic or conceptual abilities.[7]

Thinkers struggling to understand the nature of their own knowledge in the seventeenth and eighteenth centuries may be forgiven for not appreciating evolutionary theory or contemporary neurobiology. A century after Darwin a similar lack of appreciation is less forgivable.

The traditional philosophical skeptic would of course seek to reintroduce the circle argument. To invoke evolutionary theory to understand how we know about the world, he would say, simply begs the question. Evolutionary theory is a fairly advanced, and therefore problematic, form of scientific knowledge. Our problem is to *justify* that knowledge using something less problematic, such as, what one can "directly" experience or intuit.

At this point, however, the skeptic's reply is equally question-begging. Three hundred years of modern science and over a hundred years of biological investigation have led us to the firm conclusion that no humans have ever faced the world guided only by their own subjectively accessible experience and intuitions. Rather, we now know that our capacities for operating in the world are highly adapted to that world. The skeptic asks us to set all this aside in favor of a project that denies our conclusion. And he does so on the basis of what we claim to be an outmoded and mistaken theory about how knowledge is, in fact, acquired.

It should be noted that the above appeal to evolutionary theory is far more modest than that of numerous advocates of "evolutionary epistemology." It is limited to explaining why we need not worry about our failure to break the circle argument. Others have advanced the more extensive claim that evolutionary theory itself provides a good model for the overall development of scientific knowledge. I doubt that it is a very good model for this more ambitious purpose, and I shall suggest a far different account. We agree, however, that the issue is an empirical one, to be settled by scientific procedures, and not by philosophical arguments.[8]

Finally, an evolutionary perspective provides a program for dealing with norms and the problem of relativism. At some stage in the evolutionary process, the evolution of human organisms and human societies became coextensive. Even modestly complex societies require some social organization. Norms make it possible to maintain the requisite degree of social organization. Nor need the naturalist regard these as mere regularities in social behavior. Norms are taught and enforced by various means of social control. The regularity is a product of these social actions. What the naturalist denies is that there is any basis for the norms that transcends the society in its actual physical context. But does this view not leave us open to a radical form of relativism?

An evolutionary perspective places definite limits on how different a human society on earth could be. It is not physically possible that there should exist on earth a culture totally alien to us. Humans walk, talk, eat, sleep, and procreate. Correspondingly, they must acquire food and shelter. We could not fail to understand these activities. How a society goes about doing these things, on the other hand, is not uniquely determined by our biological nature, even if we include the physical circumstances of that society. There is always more than one way to skin a cat. Moreover, there is no supracultural basis for the norm that cats are to be skinned one particular way (or perhaps not at all). At this level, cultural relativism is correct. Does this imply that "creation theory" is as good as evolutionary theory? No more than it implies that prayer is as effective as penicillin for curing infections. Vindicating this reply, however, requires a positive theory of science.

6. REALISM, REFERENCE, AND RATIONALITY

Hilary Putnam (1982) has recently presented several arguments against the possibility of naturalistic or evolutionary epistemologies. One argument is that evolutionary epistemology presupposes metaphysical realism which, he claims to have shown, is incoherent. A second, more general argument is that naturalistic epistemologies attempt to eliminate normative reason. But reason, being both "immanent" and "transcendent," cannot be eliminated without committing "mental suicide" (1982, p. 22). Just explicating these arguments, let alone refuting them, would be a major undertaking. Here I can only attempt to locate some main points of disagreement and suggest where Putnam goes wrong.

Let us adopt Putnam's simple characterization of metaphysical realism as the view that "there is exactly one true and complete description of 'the way the world is'" (1981, p. 49). Must the evolutionary epistemologist or naturalistic philosopher of science make any such supposition? I don't see why. The naturalistic position is that our cognitive capacities are an evolutionary development of those possessed by lower primates and other animals. It is these same capacities the naturalistic philosopher of science employs in attempting to study the scientific activities of his fellow humans. Surely our primate ancestors could not be accused of being metaphysical realists. In so far as our cognitive abilities are continuous with theirs, why should we be any different? Perhaps some evolutionary epistemologists have indeed espoused metaphysical realism, maybe even claiming evolutionary support for such a position. But this is surely no necessary feature of an evolutionary perspective in epistemology.

In Putnam's terms, my naturalistic philosopher of science might be called an "internal realist." But naturalistic philosophers of science holding internal realism are, Putnam claims, no better off. He sees such a view as an attempt of *define* "rationality" in terms of the use of evolved capacities. The suggested formula is: Rational beliefs are those arrived at using evolved capacities for forming beliefs. But this formula is either obviously false or vacuous depending on whether we include all beliefs or only the rational ones. Thus, Putnam concludes, "The evolutionary epistemologist must either presuppose a 'realist' (i.e., metaphysical) notion of truth or see his formula collapse into vacuity" (1982, p. 5).

Here Putnam assumes that one of the tasks of a naturalistic epistemology would be to provide a *definition* of rationality. But one of the main points of an evolutionary perspective is that there is no sharp boundary between animals and humans, and thus between irrational and rational. From an evolutionary perspective, different organisms deal with aspects of their environments in more or less effective ways. Doing science is one of the ways we humans deal with aspects of our environment. Turning our attention to that process itself, we should expect to find that, in various respects, some people are more effective than others. And we would seek to explain why and how this comes about. Attempting to draw a fundamental distinction between rational and irrational activities is itself not an effective way to understand science, or any other human activity.

Of course I do not deny that providing a characterization of rationality is a well-entrenched feature of epistemology. By defining man as the rational animal, Aristotle bequeathed to philosophy the task of discovering the essence of rationality. We have given up essentialism in biology. It is about time we gave it up in epistemology, and for similar reasons.

As noted above, Putnam also has more general arguments purporting to show why reason (and by implication epistemology and the philosophy of science) cannot be naturalized. One line of argument is that reason requires language, which requires reference, which cannot be naturalized. Moreover, reason and language necessarily involve *values*, which also cannot be naturalized. I could not begin to untangle these arguments here. The most I can do is point out that if Putnam is correct, then there are genuinely *emergent* properties, for example, the property of being rational.[9] Somewhere along the line from fishes to philosophers there emerged fundamentally irreducible properties that science alone cannot explain.

Arguments against emergentism have been given by many philosophers, including, a generation ago, Putnam himself (Putnam and Oppenheim 1958). I shall not attempt to review them here. I only marvel that anyone could think these arguments refuted by an analysis of the possible reference of 'cat' and 'cherries' (Putnam 1981, chap. 2).

From a naturalistic perspective, the urge to find some essential difference between animals and humans is itself something to be explained. The evolutionary process produced a species of creatures that has spent much of its history denying its evolutionary origins. Why do humans keep insisting on their special (if not outright superior) place in nature? Psychologists, sociologists, and historians of religion have, in various guises, attempted to answer this question. What strikes me is how self-serving the emergentist program can be. Humans arguing that humans are a breed apart. One wonders if the rejection of a naturalistic approach to the philosophy of science (and philosophy generally) does not serve a much narrower self-interest. If the philosophy of science is naturalized, philosophers of science are on the same footing with historians, psychologists, sociologists, and others for whom the study of science is itself a scientific enterprise. The most philosophers of science could claim is to be the *theoreticians* of a developing science of science on the model of theoretical, as opposed to experimental, physics. Would that not be status enough?

7. MODELS AND THEORIES

As is clear from the form of the circle argument, a crucial test for any naturalistic theory of science is its account of *theory choice.* Since it would be impossible adequately to develop and defend a naturalistic account of theory choice in a short space, I will only present enough to show that such an account is both possible and at least somewhat plausible. Before one can discuss theory choice, however, it is necessary to say something about the objects of choice, namely, theories.

Since Euclid there has existed a more or less continuous tradition of representing theoretical knowledge in the form of an axiomatic system. Newton was part of this tradition, and so were the founders of modern logic. For most of this century, philosophers who have drawn their inspiration from logic and the foundations of mathematics have assumed that a theory is some type of formal, axiomatic system. The fact that scientists in the twentieth century rarely present theories in axiomatic form has not been very troubling because the philosopher's task has been seen as one of reconstruction, conceptual analysis, or justification—not description. If, however, one takes the descriptive task as fundamental, the axiomatic account clearly is not adequate. For the most part it is simply not true that theoretical scientists are engaged in developing axiomatic systems. This point is obvious for the major recent theoretical developments in sciences such as biology and geology, but it holds even for physics. Where are we to find a better account of scientific theories?

If we restrict ourselves to recent science (since 1900 or 1945), the task is easier because the transmission of theoretical knowledge has become quite uniform. It relies heavily on the advanced *textbook.* Until beginning dissertation research, most scientists in most fields learn what theory they know from textbooks (in conjunction with lectures, which also follow a textbook format). Thus, if we wish to learn what a theory is from the standpoint of scientists who use that theory, a good way to proceed is by examining the textbooks from which they learned much of what they know about that theory.

Classical mechanics provides a good example. Many sciences were modeled on mechanics and borrowed heavily from its mathematical techniques. And for many scientists and engineers today, classical mechanics provides their first experience with a real theory. In addition, classical mechanics has been a standard example for philosophers advocating an axiomatic account of theories. It thus allows a direct comparison of the merits of any rival account.

Looking at typical upper-division or graduate-level texts, what do we find? Often there is a chapter of mathematical preliminaries. The first substantive chapter, however, almost invariably presents Newton's three laws of motion. One needs a force function. The following chapters, therefore, are typically devoted to the use of Newton's laws of motion with various force-functions. A not too advanced text might devote a chapter to uniform forces—Galileo's problem of falling bodies. A typical next chapter takes up the case of a linear restoring force in one dimension, Hooke's Law—which yields a linear harmonic oscillator. Later one meets the Law of Universal Gravitation, the inverse-square force that yields orbital motion in two dimensions. And so on.

Within each chapter one finds, among other things, the following: (i) mathematical solutions to the equations of motion incorporating the specific force-function at issue, and (ii) examples of kinds of real systems to which these particular equations of motion might be applied. One learns, for example, that a linear restoring force yields a sinusoidal motion, and that the horizontal motion of a pendulum is approximately sinusoidal.

One of the most significant other things one learns is that none of the systems cited as examples *exactly* fit the equations. The horizontal restoring force of the pendulum, for example, is only linear in the limit as the angle of swing approaches zero. Regarding the equations as straight-forward statements which are then either true or false is, therefore, bound to misrepresent the situation. How, then, should we represent it?

I suggest we take the equations as characterizing an abstract, idealized system, for example, the simple harmonic oscillator. Calling such a system a "model" (or theoretical model) agrees pretty well with both scientific and philosophical usage. Claims about real systems, then, have the form: the real system *is similar to* the model. A pendulum with small amplitude, for example, is similar to a simple harmonic oscillator. I will call such claims "theoretical hypotheses." Implicit in any theoretical hypothesis is a specification of the respects and degrees in and to which the similarity is claimed to hold. At this point one could introduce truth and falsity for theoretical hypotheses, but a claim of truth here would be redundant, serving only to facilitate semantic assent.

The typical advanced text, then, presents the student with a *cluster of models* (really a cluster of clusters) together with a number of hypotheses about real things claimed to be similar to one or another of the models. For the purpose of developing a naturalistic theory of science, I suggest we understand the word 'theory' as including both the cluster or models and a broad range of hypotheses utilizing these models. Restricting 'theory' either to the models or to hypotheses produces too great a variance with how scientists use the term. For all sorts of reasons, it is best to stick as closely to scientific usage as is compatible with developing an overall, adequate theory of science.

In working through a standard text, students learn many things that are best not regarded as explicitly part of the theory, but that are very important nonetheless. They learn the accepted *interpretation* of general terms such as 'position', 'mass', and 'force'. They also learn how to *identify* particular positions, masses, and forces. Any theory of science must assume that scientists have the ability to make these sorts of interpretations and identifications. Securing a better understanding of how this is done, however, can safely be left to linguistics or, more generally, to the cognitive sciences.

It is evident that the above account of theories is realistic without going to the extreme of "metaphysical realism." Indeed, it is compatible with some recent forms of anti-realism. I would call it "constructive realism." It is "constructive" because models are humanly constructed abstract entities. It is realistic because it understands hypotheses as asserting a genuine similarity of structure between models and real systems without imposing any distinction between "theoretical" and "observational" aspects of reality. It is not "metaphysical" in that it makes no claim that there is one true and complete description of any real system. A constructive realist need not claim, for example, that there is a uniquely correct classical model for describing any actual pendulum. Nor must one claim similarity with the real world for *every* aspect of a model. One can be selective in choosing those respects in which the similarity is claimed to hold.[10]

From a naturalistic perspective, then, the theory of classical mechanics appears not to have the structure of an axiomatic system. At best an axiomatic structure could be imposed on one particular type of model, for example, systems of particles subject only to inverse-square forces. Nor, contrary to Popper's philosophy, do *universal* statements play a major role. No longer does one find sweeping Laplacian generalizations about all bodies in the universe. The typical hypothesis only asserts a similarity between a model and a more or less restricted class of real systems such as pendulums. There are many more lessons to be learned from a serious study of science textbooks, but these are sufficient to proceed to a sketch of naturalistic theory choice.

8. NATURALISTIC THEORY CHOICE

In the philosophical literature, the problem of theory choice has almost universally been understood as one of characterizing *rational* choice. Most philosophers have been willing to grant that it would be rational to choose theories that are true (or at least approximately true). The trouble is, of course, that we do not have an independent check on which ones are true.

Philosophical treatments of theory choice, therefore, have generally proceeded by focusing on properties other than truth, and then attempted to establish a general principle saying it is rational to choose theories with the specified properties. Among the many properties of theories suggested for this role have been: simplicity, falsifiability, high degree of logical probability, high degree of corroboration, predictive power, explanatory power, fruitfulness, and so on. The preferred way of establishing the required general principle is by demonstrating a connection between the specified properties and truth. Despairing of establishing any such connection with truth, however, many philosophers have argued for the rationality of theory choice in terms of these other properties themselves.

The post-Positivist switch to larger units of analysis (paradigms, research programmes, or research traditions) has not significantly changed the general strategy. The difference is that now one focuses on properties of the larger unit, such as progressiveness, and then argues that it is rational to choose a tradition with these properties. The choice of theories is subordinated to the choice of the corresponding larger unit.

All of these approaches assume the more general principle of rationality that scientists generally strive to make a *rational* choice, however this is defined. Other than this general principle, philosophical accounts of theory choice make scant reference to the actual flesh-and-blood scientists who do the choosing. The approach is almost totally "top down." A naturalistic approach to theory choice is explicitly "bottom up." It begins with real agents facing various choices in the course of their actual scientific lives. It assumes that choosing theories is not too dissimilar from choosing anything else, and then looks at how humans in fact make choices.

If our naturalistic theory of science is not to be *merely* historical, we need a *theory* of theory choice. I would suggest that decision theory includes some models of choice that can provide at least a start. Decision theory, however, has a split personality. Sometimes it operates as an account of *rational* choice; other times it is more descriptive. Here we want the descriptive mode, which may be viewed as a specialized part of ordinary belief-desire psychology.

Taken descriptively, decision-theoretic models begin with a *decision problem* that may be represented as a matrix defined by a set of possible options and a set of possible states of the world. The agent's *desires* (or values) are represented by a ranking or utility measure over the option-state pairs, the *outcomes* of the decision process. The result of adding the agent's values is a completed value (or "payoff") matrix. The role of the agent's beliefs in decision making is more complicated, as will be illustrated below.

The focus of rational decision-theory has always been on the *decision rule* (or decision strategy) that defines *the* rational choice as a function of the payoff matrix. The problem of rational decision-theory has been to establish a uniquely rational decision strategy. Descriptive decision-theory looks instead at the characteristics of the decision strategies that are actually used.

Among the most promising descriptive strategies is *satisficing.* Agents following a satisficing strategy must have a good idea of their minimum satisfactory payoff—their satisfaction level. They then survey their options to see whether any have at least a satisfactory payoff for each possible state of the world. If such an option exists, that is the one chosen. If there is no satisfactory option, agents must either lower their satisfaction level or otherwise change the decision problem. Following a satisficing strategy thus guarantees at least a satisfactory payoff—unless, perhaps, no decision is made.[11]

One could, of course, go on to argue that satisficing is rational. But there is no need to do so. Rather, we can take the satisficing strategy as part of our theoretical model of human decision making. We can then investigate the characteristics of the model and inquire of the circumstances, if any, in which humans fit this model. The fit need not be perfect. Like many theoretical models, this one is highly idealized. My hypothesis is that scientists typically follow something approximating a satisficing strategy when faced with the problem of choosing among scientific theories. If this is correct, we have a good scientific explanation of theory choice in science.

9. THE REVOLUTION IN GEOLOGY

A naturalistic approach to theories and theory choice can be nicely illustrated by the recent revolution in the earth sciences.[12]

In the early decades of this century, earth scientists, that is, geologists, geophysicists, climatologists, etc., described the earth as having originated as a much warmer body that since cooled and contracted. In the process, it was thought, the heavier material tended to collect at the core, leaving the lighter material in the mantle and crust. I would say that earth scientists had constructed a *cluster of models* built around the idea of a slowly rotating molten sphere suspended in space. The *theoretical hypotheses* implicit in the texts of the time asserted that the earth is similar in many respects to such models. Much scientific work consisted in working out the consequences of these hypotheses using lots and lots of further information such as the relative abundances of various elements in the earth's crust.

"Contractionist" models were the product of many individual decisions by earth scientists over a long period of time. Many scientists, of course, never made any explicit decisions about any features of these models. They just learned what they were taught. For them, these models formed the background of their practice as earth scientists. The models were part of what is now called a "paradigm" or "research tradition." Here we are not concerned with how this tradition came to be, but how it came to be abandoned in the 1960s.

It is a consequence of contractionist hypotheses that the oceans and continents are relatively fixed. There may be some vertical motion as the planet continues to contract, but little horizontal motion. Geographers, of course, had long noted the remarkable fit in the coastlines of Africa and South America. This immediately suggests that the two were once joined and later drifted apart. There was, however, no known way to reconcile such horizontal motions with hypotheses based on contractionist models. There were no known forces that could operate horizontally in the crust and, to make matters even worse, the ocean floors are made of harder material than the continents. These seem to have been among the main factors in the decision, by prominent earth scientists, not to take seriously hypotheses based on drift models.

Between 1910 and 1915, drift models of the earth were revived by the German meteorologist, Alfred Wegener. Rather than attempting to review the development of Wegener's thought, I will examine briefly the distinction (recently emphasized by Laudan [1977]) between *pursuit* and *acceptance*. In my view, pursuit focuses on models, acceptance on hypotheses.

In 1910, "under the direct impression produced by the coastlines on either side of the Atlantic," Wegener (1966, p. 1) thought about drift models, but did not *pursue* their elaboration. In 1911 he did pursue such models in connection with reports of evidence for the prior existence of "land bridges" between Brazil and Africa. By the beginning of 1912, he had pretty much *accepted* some drift hypotheses. That is, he had concluded that drift models pretty well fitted the actual history of the earth's development. The 1915 German edition of *The Origin of Continents and Oceans* elaborated one particular model, but was primarily designed to convince others that, contrary to the prevailing view, some corresponding drift hypotheses were correct. The book did not succeed in getting many other scientists to accept any of his drift hypotheses, but it did convince a few others that drift models were worth pursuing.

The distinction between pursuit and acceptance is a naturalistic counterpart to the Logical Empiricists' distinction between discovery and justification. The Logical Empiricists held that there is a "logic" of justification, but not of discovery. Their critics argued that there is also a logic of discovery. From a naturalistic viewpoint, there is neither a logic of discovery nor of justification. Both pursuit and acceptance are alike simply as examples of human decision making. Instead of asking, "Is there a logic of discovery?" the critics should have asked, "Is there is a logic of justification?"

There was, however, a germ of truth in the Logical Empiricists' position. Decisions to pursue a type of model (a "program"?) seem to be much more complex, and thus more difficult to study, than decisions to accept corresponding hypotheses. What leads people like Wegener to spend time developing models that neither they (at least initially) nor the vast majority of their professional colleagues believe to fit the real world? What, in decision-theoretic terms, is their payoff matrix? I have considered this question not only for Wegener, but also for later figures in the story such as Arthur Holmes, J. Tuzo Wilson, Harry Hess, and Fred Vine. It is difficult to find any general patterns. The decisions seem about as varied as the individuals.

Not that acceptance decisions are always transparent. Wegener's decision to accept drift hypotheses had to have been idiosyncratic since few of the readers of his book made similar decisions. By contrast, the decision to accept a modified drift hypothesis in the 1960s was widely shared. This, I think, was because the scientific context in the 1960s so strongly structured everyone's payoff matrix that individual differences tended not to matter. The decision was robust under exchange of professionally competent individuals.

Most commentators agree that the crucial episode in the 1960s revolution was the verification of the Vine-Matthews-Morley hypothesis (VMM) and of Wilson's related hypothesis regarding transform faults. For simplicity I will concentrate on VMM.[13]

In retrospect, it is possible to identify two relatively independent lines of development during the 1950s that made possible the revolution in the 1960s. One, in oceanography, was the discovery of large systems of ocean ridges running roughly in a north-south direction. The mid-Atlantic ridge and the eastern Pacific ridges were among the first to be explored. The second line of development, in paleomagnetism, was the discovery of several apparently global reversals in the earth's magnetic field.

In 1960, Harry Hess, reviving an idea of Arthur Holmes, suggested that ocean ridges were formed by currents of molten material rising from the core and spreading laterally, east and west, from the center of the ridge. This hypothesis, named "sea floor spreading" by Dietz a year later, immediately suggested a mechanism for continental drift. The continents are carried along on top of the spreading sea floor material.

In 1963, Fred Vine and Drummond Matthews (and, independently, Lawrence Morley) put together Hess's model of sea floor spreading with the possibility of magnetic-field reversals. The resulting hypothesis implies that there should be stripes of oppositely directed remnant magnetism in the material of the ocean floor parallel to the ridges. The pattern of magnetic reversals found on land should be duplicated on the sea floor in the pattern of "stripes" parallel to and symmetrical with the ridges. Since the reversals seemed not to occur at regular intervals, the pattern carries a very distinctive "signature."

It is not clear exactly when people like Hess and Vine switched from pursuit to acceptance of VMM and a generalized drift hypothesis. It is clear that few people in the larger earth sciences community took up pursuit of drift models, let alone acceptance of drift hypotheses. This changed dramatically in 1966–67.

In 1966, a research vessel taking magnetic soundings across the Pacific-Antarctic Ridge brought back clear evidence of the magnetic pattern predicted by VMM. About the same time, a similar pattern of reversals was observed in cores of sea floor sediment. The impact on the earth sciences community was swift and complete. Within a year just about everyone with professionally competent knowledge of the situation accepted some sort of drift hypothesis.

Why did the community of earth scientists rush to accept drift hypotheses? Part of the explanation, I suggest, is that the payoff matrix for the decision to accept drift was clear and simple to just about everyone. First, the options were clear. One option was to accept the hypothesis that some drift model fits the earth better than a static model. That is, accept the hypothesis that there are large-scale horizontal movements involving both the sea floor and the continents. The single alternative was to retain the hypothesis that a static, contractionist model is basically correct. This implies rejecting the hypothesis of large-scale horizontal movement.

Among the relevant states of the world are: (i) that the tectonic structure of the world is more similar (in the relevant respects) to drift models than to static models; and (ii) the reverse. Are these the only relevant states? In general, clearly not. It would be relevant to almost anyone whether or not a majority of their professional peers made the same decision. Few people place great value on being proven correct, as Wegener was, only after they are dead. This factor can be neutralized if we restrict our attention to scientists whose professional work was directly affected by which type of model they chose. Such people cannot comfortably wait to see which way the wind is blowing. They have a strong interest in being right at the right time.

For our suitably restricted class of earth scientists, then, the basic structure of the decision problem is as shown in figure 1. The value ranking depends on only two fairly weak, and very plausible, assumptions. One is that being objectively right is regarded as satisfactory even if one would prefer that the world were otherwise. This does not mean assuming that scientists place an *intrinsic* positive value on being objectively right. They may in fact believe that in this case being right will yield valuable short-term professional payoffs. The second assumption is that most earth scientists who were not directly involved in the research on oceanography or paleomagnetism would have preferred that contractionist hypotheses had been correct. This is simply because their training and skills were developed in the context of such models. Switching carries the cost of acquiring new knowledge and new skills. These assumptions yield the ranking of outcomes exhibited in figure 1.

	DRIFT MODELS APPROXIMATELY CORRECT	STATIC MODELS APPROXIMATELY CORRECT
ADOPT DRIFT MODELS	SATISFACTORY	TERRIBLE
RETAIN STATIC MODELS	BAD	EXCELLENT

FIGURE 1. Decision Problem for Geologists in 1966.

As it stands, the matrix does not make obvious which choice one would prefer. But there is some vital information that has not yet been factored into the decision problem. This is in two parts. The first is that finding the magnetic profiles predicted by VMM was thought to be quite likely if Hess's model were roughly correct. The second is that finding such profiles on any remotely plausible static model was thought to be quite unlikely. Static models had few known resources for accommodating the existence of such a pattern on so large a scale. These judgments were widely shared by people pursuing, or otherwise developing, either type of model. And they support a clearly satisfactory decision rule: If VMM is verified, accept drift models; if not, continue accepting static models. What makes this rule satisfactory is that, given the above judgments, a satisfactory or excellent outcome was very likely while a bad or terrible outcome was very unlikely. A satisficer would ask for no more. If earth scientists are satisficers, we have a plausible explanation of their choice.

Notice how far from methodological foundationism this account is. It assumes agreement that the technology for measuring magnetic profiles is reliable. The Duhem-Quine problem is set aside by the fact that one can build, or often purchase commercially, the relevant measuring technology. The background knowledge (or auxiliary hypotheses) are embodied in proven technology.

The above account also assumes agreement on what is likely or not depending on which type of hypothesis is correct. The fact that there were *logically* possible contractionist models that could yield VMM was irrelevant. What mattered was whether anyone in the opposing camp thought they could come up with rival hypotheses that fitted the data. In this case, most of those who had been developing static models simply gave up. It was only in the context of this vast background of shared judgments that the data was able decisively to force the decision on a typical satisficer.

The above sketch must suffice. It is hardly enough to convince anyone to accept my particular naturalistic hypotheses regarding even just this one case. I hope it is enough to convince some that such models are worth pursuing.

10. A ROLE FOR HISTORY

Kuhn was of course correct in thinking that a naturalized philosophy of science would provide a role for history. The role, he suggested, was as *evidence* for theoretical claims about science. Yet the use of history by philosophers of science (recall the metamethodology of Lakatos and Laudan) suggests that this evidential relationship is more complex than it might seem.

It is useful to consider how some other sciences use the historical record as evidence for their theories. Evolutionary biology and economics provide appropriate models because they seem nicely to bracket a proposed theoretical science of science.

Turning first to evolutionary biology, it is generally thought that the fossil record provides historical evidence for evolutionary theory. I am far from convinced that this record, by itself, provides a satisfactory basis for deciding that any evolutionary theory is correct. Here, however, I am concerned with a narrower issue. Those who have used the fossil record as evidence for evolutionary theory have generally assumed that the underlying mechanisms of evolution, whatever they might be, are relatively stable. Few biologists have ever argued that we might need different models of evolution for different epochs. The major recent controversies over punctuated equilibria or mass extinctions concern the nature and rate of changes in the environment—not our models of the underlying evolutionary mechanisms.

In contrast to evolutionary biology, the most successful theoretical models in economics, whether macro or micro, are *equilibrium* models. The data for such models are, therefore, not historical in the sense that they follow economic developments over time. For models that do use genuinely historical data, one must turn to theories of economic *development*. These, however, are generally thought to provide the least successful models in all of economics, Marxism being the most obvious example. The generally accepted reason for the poor record of models of development is that the economic mechanisms themselves change over time. It is not simply a matter of looking at the same mechanisms operating in a different environment. A rural, agrarian society, for example, seems to embody different economic mechanisms than an urban, industrial society.

Following Kuhn, historically minded philosophers of science have argued, using historical examples, that not only the content of science changes with time. Its aims and methods change as well. This seems to imply that the relation between theories of science and the history of science follows the economic rather than the biological pattern. Indeed, the Kuhnian model of development—normal science, crisis, revolution, new normal science—seems to have as much, or as little, theoretical content as the Marxian stages of economic development. One wonders whether any theory of scientific development that includes changes in aims and methods could do much better. Yet most historically oriented philosophers of science since Kuhn seem to be aiming at a similarly grand theory of development. Few would describe their own aims in these terms, of course. But illustrating the same point using historical cases ranging from Newton, through Lavoisier, to Einstein, and even J. D. Watson, betrays the intent.

The options for a naturalistic theory of science, then, are these. The first is an ambitious strategy that seeks mechanisms of scientific development that can explain not only changes in content but also changes in aims and methods. One could then claim to have similar mechanisms operating over long periods, say from the seventeenth century to the present. The danger in this strategy is ending up with only vaguely defined models of science. A second, much less ambitious strategy would be to restrict attention to shorter epochs such as science in the seventeenth century or since World War II. Here the danger is ending up with models of only very restricted applicability.[14]

Following my own theory of science, I would suggest a third, hybrid strategy. Begin with the less ambitious strategy, and then try to link up the various models so as to obtain a cluster of partially overlapping models covering several epochs, perhaps, even, most of science since Newton. That, I suspect, is the most that can be done.

The suggested model of science provides some hope for thinking that the third strategy can be successful. The activities of model construction and model choice abstract from the scientific context in much the same way as models of mutation and selection abstract from the biological environment. These activities may take place in many different social and economic settings. Different aims, or values, may be reflected in the structure of decisions concerning specific hypotheses. So may the information yielded by new methodologies. Whether this is enough to provide informative similarities among widely separated epochs remains to be seen.

My aim in this paper, however, has not been to argue for a particular strategy or a particular model of science. These have been noted only to illustrate the possibilities opened up by a naturalistic approach. The main thesis is that the study of science must itself be a science. The only viable philosophy of science is a naturalized philosophy of science.

NOTES

1. Among prominent evolutionary or naturalistic epistemologists, I count Donald Campbell (1974) and Abner Shimony (1971, 1981). Toulmin's (1972) view is evolutionary but perhaps not naturalistic. Popper (1972) appropriates the title "evolutionary" without adequate justification. Friedman (1979) reaches conclusions that are naturalistic but not evolutionary. Arthur Fine's (1984) recently espoused "natural ontological attitude" encompasses a natural epistemological attitude as well. Outside the philosophy of science, advocates of naturalistic or evolutionary epistemologies are too numerous even to begin mentioning.

2. I first formulated this argument (in Giere 1973) as an expression of what I then took to be a majority view among philosophers of science. I did not intend to argue that it was impossible to establish a connection between the philosophy of science and the historical practice of science; only that the authors under review had failed adequately to address the most serious difficulties. Indeed, my own solution at the time (Giere 1975) was basically naturalistic, though not evolutionary.

3. The importance of broadly empirical considerations in the Logical Empiricists' rejection of Russellian foundationism has been emphasized recently by Hempel (1983). In this paper, Hempel distinguishes "normative" from "descriptive-naturalistic" methodologies, and argues for a mixed approach.

4. For a summary of the relevant literature, and many references, see Giere (1979). Advocates of "Bayesian inference" seem to assume that their reconstruction of scientific inference, while not strictly reducible to deductive logic, nevertheless somehow carries the normative force associated with deductive logic. I do not understand the basis for this assumption. I am not even convinced that deductive logic possesses the normative powers commonly ascribed to it.

5. The following discussion applies only to the Laudan of *Progress and Its Problems*. I understand he no longer subscribes to this metamethodology. Lakatos is somewhat ambiguous as to whether his metamethodology is just his ordinary methodology applied at the meta-level or a different methodology altogether.

6. What follows owes at least part of its inspiration to Paul Churchland's (1979) notion of an "epistemic engine." See also Patricia Smith Churchland and Paul M. Churchland (1983).

7. For some recent neurobiological findings relevant to the mechanisms underlying spatial coordination among mammals, see O'Keefe and Nadel (1978), and Pellionisz and Llinas (1982).

8. Here I am thinking particularly of Donald Campbell (1974) and his followers.

9. From informal comments at a conference in May, 1984, I infer that Putnam himself might agree that his view is a variety of emergentism.

10. The label "constructive realism" was originally intended as a direct contrast to van Fraassen's (1980) "constructive empiricism." See Giere (1984 and forthcoming). My view of theories is a liberal version of his "semantic" conception of theories, and similar to Frederick Suppe's (1973) conception. Van Fraassen's distinction between "observable" and "theoretical" seems to me a philosophical imposition. It is very difficult to interpret the actual practice of scientists as honoring such a distinction. I find Nancy Cartwright's (1983) anti-realism much more congenial, perhaps even compatible with a constructive realism. I could also agree with much of what Putnam (1978) says about "internal realism." There are some general similarities between my view and the "structuralist" approach of Sneed (1971) or Stegmüller (1979). This school, however, seems primarily interested in reconstruction and philosophical vindication, and little concerned with description.

11. Satisficing has been developed primarily by Herbert Simon. See his 1972 for further details and references.

12. Historians and philosophers of science have recently begun to give the revolution in geology the attention it deserves. See, for example, Frankel (1982), Rachel Laudan (1981), or Ruse (1981), and the references cited therein.

13. For technical details, see Frankel (1982). The reader is invited to compare Frankel's Laudanistic interpretation of this episode with my decision-theoretic account.

14. For an example of the ambitious strategy, see L. Laudan (1984).

REFERENCES

Campbell, D. R. (1974), "Evolutionary Epistemology", in *The Philosophy of Karl Popper,* P. A. Schilpp (ed.). La Salle: Open Court, pp. 413–63.

Carnap, R. [1950] (1962), *Logical Foundations of Probability.* 2nd edition. Chicago: University of Chicago Press.

Cartwright, N. D. (1983), *How the Laws of Physics Lie.* Cambridge: Cambridge University Press.

Churchland, P. M. (1979), *Scientific Realism and the Plasticity of Mind.* Cambridge: Cambridge University Press.

Churchland, P. S., and Churchland, P. M. (1983), "Stalking the Wild Epistemic Engine", *Noûs 17:* 5–18.

Fine, A. (1984), "And Not Anti-Realism Either", *Noûs 18:* 51–65.

Frankel, H. (1982), "The Development, Reception and Acceptance of the Vine-Matthews-Morley Hypothesis", *Historical Studies in the Physical Sciences 13:* 1–39.

Friedman, M. (1979), "Truth and Confirmation", *The Journal of Philosophy 76:*361–82.

Giere, R. N. (1973), "History and Philosophy of Science: Intimate Relationship or Marriage of Convenience?", *British Journal for the Philosophy of Science 24:* 282–97.

———. (1975), "The Epistemological Roots of Scientific Knowledge", in *Induction, Probability, and Confirmation,* G. Maxwell and R. M. Anderson, Jr. (eds.). Minnesota Studies in the Philosophy of Science, vol. 6. Minneapolis: University of Minnesota Press, pp. 212–61.

———. (1979), "Foundations of Probability and Statistical Inference", in *Current Research in Philosophy of Science,* P. D. Asquith and Henry E. Kyburg, Jr. (eds.). East Lansing: Philosophy of Science Association, pp. 503–33.

———. (1984), "Toward a Unified Theory of Science", in *Science and Reality,* J. T. Cushing, C. F. Delaney, and G. Gutting (eds.). Notre Dame: University of Notre Dame Press.

———. (forthcoming), "Constructive Realism", in *Images of Science,* P. M. Churchland and C. Hooker (eds.). Chicago: University of Chicago Press.

Hacking, I. (1968), "One Problem about Induction", in *The Problem of Inductive Logic,* I. Lakatos (ed.). Amsterdam: North-Holland, pp. 44–58.

Hempel, C. G. (1983), "Valuation and Objectivity in Science", in *Physics, Philosophy and Psychoanalysis,* R. S. Cohen and L. Laudan (eds.). Dordrecht: D. Reidel, pp. 73–100.

Kuhn, T. S. [1962] (1970), *The Structure of Scientific Revolutions.* 2nd edition. Chicago: University of Chicago Press.

Jeffrey, R. (1973), "Carnap's Inductive Logic", *Synthese 25:* 299–306.

Lakatos, I. (1970), "Falsification and the Methodology of Scientific Research Programmes", in *Criticism and the Growth of Knowledge,* I. Lakatos and A. Musgrave (eds.). Cambridge: Cambridge University Press.

———. (1971), "History of Science and Its Rational Reconstructions", in *PSA 1970,* R. S. Cohen and R. C. Buck (eds.). Boston Studies in the Philosophy of Science, vol. 8. Dordrecht: D. Reidel, pp. 91–135.

Laudan, L. (1977), *Progress and Its Problems.* Berkeley and Los Angeles: University of California Press.

———. (1984), *Science and Values: The Aims of Science and Their Role in Scientific Debate.* Berkeley and Los Angeles: University of California Press.

Laudan, R. (1981), "The Recent Revolution in Geology and Kuhn's Theory of Scientific Change", in *PSA 1978,* P. D. Asquith and I. Hacking (eds.). Vol. 2. East Lansing: Philosophy of Science Association, pp. 227–39.

O'Keefe, J., and Nadel, L. (1978), *The Hippocampus as a Cognitive Map.* Oxford: Clarendon Press.

Pellionisz, A., and Llinas, R. (1982), "Space-Time Representation in the Brain: The Cerebellum as a Predictive Space-Time Metric Tensor", *Neuroscience 7:* 2949–70.

Popper, K. R. (1972), *Objective Knowledge.* Oxford: Clarendon Press.

Putnam, H. (1978), *Meaning and the Moral Sciences.* London: Routledge and Kegan Paul.

———. (1981), *Reason, Truth and History.* Cambridge: Cambridge University Press.

———. (1982), "Why Reason Can't Be Naturalized", *Synthese 52:* 3–23.

Putnam, H., and Oppenheim, R. (1958), "Unity of Science as a Working Hypothesis", in *Concepts, Theories and the Mind-Body Problem,* H. Feigl, M. Scriven, and G. Maxwell (eds.). Minnesota Studies in the Philosophy of Science, vol.

2. Minneapolis: University of Minnesota Press, pp. 3–36.

Quine, W. V. O. (1969), "Epistemology Naturalized", in *Ontological Relativity and Other Essays.* New York: Columbia University Press.

Ruse, M. (1981), "What Kind of a Revolution Occurred in Geology?", in *PSA 1978,* P. D. Asquith and I. Hacking (eds.). Vol. 2. East Lansing: Philosophy of Science Association, pp. 240–73.

Shapere, D. (1984), *Reason and the Search for Knowledge.* Dordrecht: D. Reidel.

Shimony, A. (1971), "Perception from an Evolutionary Point of View", *The Journal of Philosophy 67:* 571–83.

———. (1981), "Integral Epistemology", in *Scientific Inquiry and the Social Sciences,* M. B. Brewer and B. E. Collins (eds.). San Francisco: Jossey-Bass, pp. 98–123.

Simon, H. A. (1972), "Theories of Bounded Rationality", in *Decision and Organization,* R. Radner and C. B. McGuire (eds.). Amsterdam: North-Holland, pp. 161–76.

Sneed, J. D. (1971), *The Logical Structure of Mathematical Physics.* Dordrecht: D. Reidel.

Stegmüller, W. (1979), *The Structuralist View of Theories.* New York: Springer.

Suppe, F. (1973), "Theories, Their Formulations, and the Operational Imperative", *Synthese 25:* 129–64.

Toulmin, S. (1972), *Human Knowledge.* Princeton: Princeton University Press.

van Fraassen, B. C. (1980), *The Scientific Image.* Oxford: Oxford University Press.

Wegener, A. (1966), *The Origin of Continents and Oceans.* New York: Dover.

The Logic of Discovery

Norwood Russell Hanson

IS THERE A LOGIC OF SCIENTIFIC DISCOVERY?

Is there a logic of scientific discovery? The approved answer to this is "No." Thus Popper argues:[1] "The initial stage, the act of conceiving or inventing a theory, seems to me neither to call for logical analysis nor to be susceptible of it." Again, "There is no such thing as a logical method of having new ideas, or a logical reconstruction of this process." Reichenbach writes that philosophy of science "cannot be concerned with [reasons for suggesting hypotheses], but only with [reasons for accepting hypotheses]."[2] Braithwaite elaborates: "The solution of these historical problems involves the individual psychology of thinking and the sociology of thought. None of these questions are our business here."[3]

Against this negative chorus, the "Ayes" have *not* had it. Aristotle (*Prior Analytics* II, 25) and Peirce[4] hinted that in science there may be more problems for the logician than just analyzing the arguments supporting already invented hypotheses. But contem-