Chapter 1
Something between a Preface and an Introduction

One good way to begin a philosophical essay is to locate the project in philosophical space by saying how the questions to be asked and the answers to be defended relate to the questions and answers that have occupied other writers. Another way to get started, particularly when the views to be defended are at odds with much of received opinion—and mine certainly are—is to say something about the evolution of one's own thinking, in an effort to explain how one arrived at such unorthodox views. In this introductory chapter, I propose to do a bit of both. The autobiographical approach will predominate, not because I imagine the reader will find the details of my intellectual autobiography to be of any intrinsic interest, but simply because it provides a perspicuous way to give an overview of the book and to explain how its various themes come to be linked together. Though issues in the philosophy of mind, the philosophy of language, and the philosophy of psychology will all be center stage at one point or another in the pages to follow, the cluster of questions that motivate the volume fall squarely within the domain of epistemology. So let me start by saying how I view that domain.

1.1 *Three Traditional Epistemological Projects*

There are, as I see it, at least three interrelated projects that traditionally have been pursued in epistemology, with different authors naturally enough emphasizing different ones. The first of these projects focuses on the evaluation of methods of inquiry. It tries to say which ways of going about the quest for knowledge—which ways of building and rebuilding one's doxastic house—are the good ones, which are the bad ones, and why. Since reasoning is central to the quest for knowledge, the evaluation of various strategies of reasoning often plays a major role in the assessment of inquiry.

There is no shortage of historical figures who have pursued this sort of epistemological investigation. Much of Francis Bacon's epistemological writing is devoted to the project of evaluating and criticizing strategies of inquiry, as is a good deal of Descartes's. Among more modern epistemological writers, those like Mill, Carnap, and Popper, who are concerned with the logic and methodology of science, have tended to emphasize this aspect of epistemological theory. From Bacon's time to Popper's, it has frequently been the case that those who work in this branch of epistemology are motivated, at least in part, by very practical concerns. They are convinced that defective reasoning and bad strategies of inquiry are widespread, and that these cognitive shortcomings are the cause of much mischief and misery. By developing their accounts of good reasoning and proper strategies of inquiry, and by explaining why these are better than the alternatives, they hope others will come to see the error of their cognitive ways. And, indeed, many of these philosophers have had a noticeable impact on the thinking of their contemporaries.[1]

A second traditional epistemological project aims to understand what knowledge is, and how it is to be distinguished from other cognitive states like mere opinion or false belief. For Plato, and for many other philosophers as well, the effort to understand what knowledge is was taken to be an inquiry into the nature of a natural kind. It was the form or essence of the natural kind that the inquiry sought to uncover. With the "linguistic turn" in twentieth-century philosophy, this project has been reconstrued as a quest for the correct definition of the word 'knowledge' or for the correct analysis of the concept of knowledge. Since the publication, in 1963, of Gettier's brief and enormously influential attack on the venerable view that 'knowledge' could be defined as 'justified true belief', the analytic enterprise has grown into a thriving cottage industry.[2]

A third project that has loomed large in epistemology has been elaborating replies to the arguments of those skeptics—real or more often imaginary—who deny that we have knowledge, or certainty, or some other epistemologically valuable commodity, and who often go on to claim that knowledge, certainty, or what have you is impossible to obtain. Answers to the skeptic have been a persistent motif in epistemology from Descartes to G. E. Moore, and right down to the present.[3]

Clearly, these three projects are linked together in a variety of ways. To answer the skeptic, a natural first step might be to develop an analysis of knowledge or certainty so that we can be clear on exactly what it is the skeptic is claiming we don't have or can't get. Moreover, in attempting to say what knowledge is, an epistemological theorist will often find it necessary to give some account of good reasoning or good strategies of inquiry, since whether a given belief will count as an instance of knowledge is often said to depend, in part, on whether the belief was arrived at in an appropriate way.

My own interests are not distributed equally among these three projects. Indeed, for as long as I can remember, I have found the latter two projects to be somewhat dreary corners of philosophy. On the few occasions when I have taught the "analysis of knowledge" literature to undergraduates, it has been painfully clear that most of my students had a hard time taking the project seriously. The better students were clever enough to play fill-in-the-blank with 'S knows that p if and only ____'. They could recognize the force of the increasingly arcane counterexamples that fill the literature, and they occasionally produced new counterexamples of their own. But they could not, for the life of them, see why anybody would want to do this. It was a source of ill-concealed amazement to these students that grown men and women would indulge in this exercise and think it important—and of still greater amazement that others would pay them to do it! This sort of discontent was all the more disquieting because deep down I agreed with my students. Surely something had gone very wrong somewhere when clever philosophers, the heirs to the tradition of Hume and Kant, devoted their time to constructing baroque counterexamples about the weird ways in which a man might fail to own a Ford, or about strange lands that abound in trompe l'oeil barns.[4] But just what had gone wrong I was, at that time, quite unable to say. Though the arguments developed in this book began with concerns far removed from the analysis of knowledge and other epistemic notions, as my position evolved I began to see with increasing clarity what it was that made the project of analyzing epistemic terms seem so wrongheaded. I'll say more on this a bit later in this chapter; my full brief against "analytic epistemology" will be set out in chapter 4.[5]

I have confessed that for about as long as I can remember I have had deep, though largely inarticulate, misgivings about the project of analyzing epistemic notions. I have also long harbored similar concerns about the effort to construct responses to the epistemic skeptic. As I got clearer on what I thought was wrong with the analytic project, I also got clearer about why responding to the skeptic often seems a waste of time. I'll elaborate on this theme, in 1.4.1.1.

Before beginning the work that led to this book, my attitude toward the remaining entry on my list of epistemic projects—the evaluation of strategies of reasoning and inquiry—was much more conventional and benign. The part of the literature I knew seemed far less frivolous than the literature on the analysis of knowledge and far more persuasive than the attempts to reply to the skeptic. Moreover, the issues themselves struck me as important ones with real, practical implications both for the conduct of science and for the governing of one's cognitive affairs in everyday life. But my own philosophical interests had long been centered on the philosophy of language and the philosophy of psychology, and this branch of epistemology seemed quite unrelated to those interests.

1.2 Finding Connections: The Psychology of Reasoning, the Evaluation of Inquiry, and the Analysis of Intentional Content

I began to see that these domains might be more closely related than I had thought when, a bit over a decade ago, my friend and former colleague Richard Nisbett posed an intriguing problem to me. To explain Nisbett's problem, I'll have to back up a bit and fill in some background.

1.2.1 The Empirical Exploration of Reasoning

Nisbett, along with a number of other experimental social psychologists, was exploring the ways in which normal human subjects (well, undergraduates actually) go about the business of reasoning, on quite ordinary problems, in relaxed and unthreatening surroundings. What they found was both fascinating and more than a bit unsettling. On many sorts of problems their subjects, despite being fairly bright, seemed to reason very badly and to do so in more or less predictable ways. Indeed, in some domains the reasoning was so strikingly bad that Nisbett and his colleagues were led to describe the implications of their research as "bleak."[6] In the intervening years, much of this work has become well known, and there are several excellent surveys available.[7] But since the implications of these findings will be a recurrent theme in the pages to follow, I had best set out a few examples for readers who may not be familiar with the literature. Those for whom this is all old hat may wish to scoot ahead to 1.2.2.

1.2.1.1 The Selection Task One of the most extensively investigated examples of *prima facie* failure in reasoning is the so-called selection task first studied by Wason and Johnson-Laird.[8] In a typical selection task experiment, subjects are presented with four cards like those in figure 1. Half of each card is masked. Subjects are then given the following instructions:

> Your job is to determine which of the hidden parts of these cards you need to see in order to answer the following question deci-

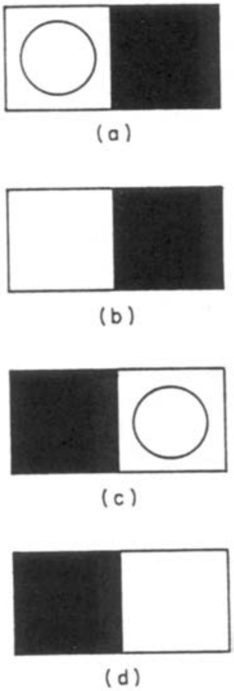

Figure 1

sively: FOR THESE CARDS IS IT TRUE THAT IF THERE IS A CIRCLE ON THE LEFT THERE IS A CIRCLE ON THE RIGHT?

You have only one opportunity to make this decision; you must not assume that you can inspect the cards one at a time. Name those cards which it is absolutely essential to see.

Wason and Johnson-Laird found that subjects, including very intelligent subjects, typically do very badly on this question. In one group of 128 university students, only *five* got the right answer. Moreover, the mistakes turn out not to be randomly distributed. The two most common wrong answers are that one must see both (a) and (c), and that one need only see (a); subjects find it particularly difficult to understand why (d) must be removed. In the years since Wason and Johnson-Laird's first studies of the selection task, there have been many further studies looking at related tasks that vary from this one in a number of dimensions. Some of those studies have indicated that subjects do a much better job on structurally analogous problems if the subject matter of the problem is more realistic, or more familiar to them, or if it can be fitted into one or another preexisting schema for reasoning. These results have provided a rich data base for theorists trying to understand the cognitive mechanisms underlying this sort of reasoning. Though at present there is no consensus at all about what those mechanisms are.[9]

1.2.1.2 The Conjunction Fallacy A second example of research revealing apparent deviations from normative standards of inference focuses on the way people assess the probability of logically compound events or states of affairs. It is a truism of probability theory that the likelihood of a compound event or state of affairs must be less than or equal to the likelihood of the component events or states of affairs. If the components are probabilistically independent, the probability of the compound is equal to the product of the probabilities of the components. If the components are not probabilistically independent, matters are more complicated. But in no case will the probability of the compound be *greater* than the probability of any component. There are, however, a number of experiments which demonstrate that people regularly violate this basic tenet of probabilistic reasoning and fall into what has been called "the conjunction fallacy."

In one such experiment Tversky and Kahneman posed a number of questions like the following:[10]

> Linda is 31 years old, single, outspoken and very bright. She majored in philosophy. As a student she was deeply concerned with issues of discrimination and social justice, and also participated in antinuclear demonstrations.
>
> Please rank the following statements by their probability, using 1 for the most probable and 8 for the least probable.
>
> (i) Linda is a teacher in an elementary school.
> (ii) Linda works in a bookstore and takes Yoga classes.
> (iii) Linda is active in the feminist movement.
> (iv) Linda is a psychiatric social worker.
> (v) Linda is a member of the League of Women Voters.
> (vi) Linda is a bank teller.
> (vii) Linda is an insurance salesperson.
> (viii) Linda is a bank teller and is active in the feminist movement.

In this experiment, 89 percent of the subjects ranked (viii) as more likely than (vi). Moreover, the result turns out to be very robust. Concerned that subjects might tacitly suppose that (vi) really meant

Linda is a bank teller and is *not* active in the feminist movement

Tversky and Kahneman replaced (vi) with

(vi′) Linda is a bank teller whether or not she is active in the feminist movement

and tried the new material on a second set of subjects. The results were essentially the same. But perhaps subjects were distracted by all the other options and failed to notice the relationship between (vi) and (viii). To test this, 142 subjects were given the original problem with all the alternatives except (vi) and (viii) deleted and asked to indicate which of the two alternatives was more likely. Eighty-five percent said that the conjunction was more likely than the conjunct.

1.2.1.3 Pseudodiagnosticity Suppose we were interested in the effectiveness of a new drug in treating a certain disease, and we have evidence that 79 percent of patients suffering from the disease who have taken the drug recover completely within a month. What should we conclude about the efficacy of the drug? The answer, of course, is that we do not yet have enough information, since we must also know the spontaneous recovery rate—the rate of recovery of people who have not taken the drug. Without such information, there is no saying whether the drug fosters recovery or impedes it. However, there are a number of studies indicating that people often do draw conclusions in cases like this, even when they have no information at all about relevant base rates.[11] Moreover, in one striking study, Doherty, Mynatt, Tweney, and Schiavo showed that subjects are reluctant to seek out diagnostically relevant base rate information, even when it is readily available.[12]

These investigators set subjects the task of determining whether a certain archaeological find had come from Coral Island or from Shell Island. The find was a clay pot, and the subjects were given a list of the pot's characteristics (smooth clay—not rough; curved handles—not straight) and so on for a number of other binary characteristics. The subjects were then given a booklet from which they could get some data about the kinds of pots that had been produced on the two islands. The data were arrayed as follows, with each pair of percentages covered by an opaque sticker as indicated.

On one page of the booklet there were a total of 12 stickers, and in order to get the data they needed, subjects were permitted to peel off any six of them. The most useful or "diagnostic" information would be gleaned only if subjects removed both stickers in a given row, and therefore the optimal strategy would be to select three row-pairs.

	Coral Island	Shell Island
Curved handles	21%	87%
Straight handles	79%	13%
Smooth clay	19%	91%
Rough clay	81%	9%
.	.	.
.	.	.
.	.	.

However, only 11 of 121 subjects removed three pairs; nine removed two pairs, and thirty removed one pair. The remaining 71 subjects (that's 59 percent) removed no pairs at all. Thus, the majority of subjects formed their belief about which island the pot had come from on the basis of "pseudodiagnostic" information. Though it was readily available, they chose not to seek out the information that would be of most use to them.

There might be some temptation to suppose that results like these are artifacts of the rather artificial experimental format. However, as Nisbett and Ross point out, the logic exhibited by these experimental subjects "is suspiciously similar to the logic shown by poorly educated laypeople in discussing a proposition such as: Does God answer prayers? Yes, such a person may say, because many times I've asked God for something and He's given it to me."[13]

1.2.1.4 Belief Perseverance My final example of a research program that has uncovered apparent irrationality is the work Ross and his colleagues, exploring how people modify their beliefs when the evidence for those beliefs is no longer accepted.[14] One of the experimental strategies used in this work is the so-called debriefing paradigm in which subjects are given evidence that is later completely discredited. But despite being debriefed and told exactly how they had been duped, subjects tend to retain to a substantial degree the beliefs they formed on the basis of the discredited evidence.

In one such experiment, subjects were presented with the task of distinguishing between authentic and inauthentic suicide notes. Some of the notes, they were told, had been found by the police, while others were written by students as an exercise. As the subjects worked on the task, they were provided with false feedback indicating that overall they were performing close to the average level, or (for another group of subjects) much above the average level, or (for a third group of subjects) much below the average level. Following this, each subject was debriefed, and the predetermined nature of the feedback was clearly explained. They were not only told that their feedback had been false but were also shown the experimenter's instruction sheet assigning them to the success, failure, or average group and specifying the details of the feedback that they had been given. Subsequent to this, and allegedly for a quite different reason, subjects were asked to fill out a questionnaire on which they were asked to estimate their actual performance at the suicide note task they had completed, to predict how well they would do on related tasks, and to rate their ability at suicide note discrimination and similar tasks. The striking finding was that, even after debriefing, subjects who had initially been assigned to the success group continued to rate their performance and abilities far more favorably than did subjects in the average group. Subjects initially assigned to the failure group showed the opposite pattern of results. Once again, further experiments suggested that these results reflect a robust phenomenon that manifests itself in many variations on the experimental theme, including some conducted outside the laboratory setting. The phenomenon has been labeled "belief perseverance."

1.2.2 Nisbett's Problems and Goodman's Solution

These were the sorts of experimental findings Nisbett had in mind when he posed a problem to me that went something like this: *When I present these experimental results to various professional audiences and draw the obvious, pessimistic conclusions about the reasoning abilities of the man or woman in the street, people raise various sorts of objections. Some of the objections are about experimental design, "ecological validity" and similar issues. And these I know how to handle. But from time to time someone will challenge my claim that in a particular experiment, subjects who give a certain answer are in fact reasoning badly. These critics demand to know why* I *get to say which inferences are the good ones and which are the bad ones. They want to know what it is that makes the subject's inference bad and the inference I think they should draw good.* Like Nisbett, I was at the time strongly inclined to say that the subjects, not the experimenters, were the ones who were reasoning badly. But if nothing more than this could be said, the debate between Nisbett and his critics would degenerate into an exchange of raw intuitions about which inferences are good ones. What more could be done, Nisbett asked? How could it be *shown* that the subjects were reasoning badly?

When Nisbett first posed the question to me, I thought I knew the answer. Having been, as an undergraduate, a student of Nelson Goodman's, I had cut my philosophical teeth on what I took to be an elegant, powerful, and entirely persuasive answer to the question of how inferences and rules of inference are to be justified. The way to do it, Goodman had argued, is via a process of mutual adjustment in which judgments about particular inferences and judgments about inferential rules are brought into accord with one another.[15] The justification for rules of inference lies in the accord thus achieved. But as soon as I proposed Goodman's process as a solution to Nisbett's problem, it became clear that this very influential account of inferential justification could not be quite right; there must be a bug somewhere. For, read literally, Goodman's account of what it is for an inference to be justified, when conjoined with a plausible extrapolation of the empirical data about actual inferential practice, seems to entail that some very strange inferences are justified. At the time, the bug did not strike me as a major one. All that was needed, it seemed, was a bit of fine tuning of Goodman's picture of inferential justification. And after a few months of discussion, Nisbett and I thought we knew just how that fine tuning should go. We published our proposal in Stich and Nisbett (1980).

Even before that article was in print, I had come to realize that our fine tuning was going to need some fine tuning. For our story, no less than Goodman's, had some very curious counterintuitive consequences. Further infelicities were noted by Conee and Feldman in their critique of our paper.[16] Still, I was reasonably confident that, with a bit more effort and a bit more thought, I would find a way to patch up Goodman's account so that it would not sanction what seemed to be patently unjustified inferences. Indeed, I even supposed I had an argument of sorts demonstrating that there *must* be a way to repair Goodman's account. (This is the argument I attribute to the "neo-Goodmanian" in 4.4.) It took about four years of trying intermittently and failing consistently before I came to suspect that my lack of success might be symptomatic of something besides my own limited intellectual endowments. But here I'm getting ahead of myself. For while hunting for some variation on Goodman's theme that would avoid counterintuitive consequences, I had a number of other, warmer irons in the fire.

1.2.3 Good Reasoning and Intentional Content: The Davidson/Dennett Argument That Bad Reasoning Is Impossible

As I became better acquainted with the data on human inference collected by psychologists like Tversky, Kahneman, Ross, Nisbett, and others and with the conclusions these authors wanted to draw from those data about the often questionable quality of everyday reasoning, I began to realize that these conclusions sit very uncomfortably with certain aspects of Donald Davidson's much-discussed theories in the philosophy of language and with some kindred ideas in the philosophy of mind developed by Daniel Dennett—ideas that were just then attracting a great deal of attention. Davidson and Dennett, both inspired by Quine, have offered accounts of how we go about interpreting, or assigning "intentional content," to a person's utterances and to the mental states that they presumably express. While differing in a variety of ways, both accounts require a high degree of rationality as a prerequisite for intentional interpretation. People's beliefs must be mostly true, and the inferences they draw must be mainly the right or normatively appropriate ones. If they are not, Davidson and Dennett maintain, it will be impossible to assign any interpretation to their verbal output or ascribe any content to their mental states. But acoustic output that admits of no interpretation is not language at all, and mental states without content can not be beliefs or thoughts. It follows that inference which is seriously and systematically irrational is a conceptual impossibility. For inference is a process in which beliefs are generated or transformed. But without a high level of rationality and truth there can be no belief, and without belief there can be no inference. Thus it is simply incoherent to suggest that people reason in ways that depart seriously and systematically from what is rational or normatively appropriate. If Davidson and Dennett are right, the psychologists who claim to have evidence for extensive irrationality in human inference must be mistaken. No evidence could possibly support such a conclusion, since the conclusion is conceptually incoherent.

This is a type of doctrine for which I have never had much sympathy. Philosophy has a long history of trying to issue a priori ultimatums to science, decreeing what must be the case or what could not possibly be the case. And those a priori decrees have a dismal track record. *Pace* Kant, space is not Euclidean, nor are the laws of physics Newtonian. *Pace* Hegel, there are nine planets, not seven. But underlying my distrust of the a priori arguments against the possibility of systematically defective reasoning, there was more than a general skepticism about philosophy's attempts to constrain science. For at the time I first saw the conflict between the theories advanced by Davidson and Dennett and the thesis about widespread human irrationality defended by various psychologists, I was working out the details of my own account of intentional interpretation or "content ascription."[17]

My account agreed with Davidson and Dennett that there is indeed a link of sorts between rationality and content and thus that significant departures from rationality make the ascription of content difficult or impossible. Moreover, unlike Dennett's story, or Davidson's, my account of content ascription tries to make it clear *why* content and good reasoning are linked; it offers an *explanation* of the connection. But it seemed to me that if my account of content ascription was on the right track, and my explanation of the link between rationality and content was even roughly right, it would undermine the conclusion of the Davidson/Dennett argument. For it follows from my account that the distinction between those mental states to which content can be comfortably ascribed, and those to which it can be ascribed either tenuously or not at all, is a distinction entirely bereft of theoretical interest. It is a parochial, observer-relative, context-sensitive distinction that marks no significant psychological boundary.

If this is right, then the sort of impossibility that the Davidson/Dennett argument might establish is simply not worth worrying about. For at best, what that argument shows is that systematically irrational people cannot engage in "real" inference at all but only in "inferencelike" mental processes. These processes don't count as inference, properly so-called, because they generate and transform mental states that cannot be intentionally described and thus do not count as "real" beliefs. But if, as I maintain, the distinction between "real" beliefs and those "belieflike" mental states that are not intentionally characterizable is a vague and parochial one that marks no significant psychological boundary, then the same is true of the distinction between "real" inference and contentless "inferencelike" processes. And we hardly need worry about the fact that seriously irrational people don't infer at all, if they do something that is like inference except in ways that are of no psychological interest.

Setting out a detailed defense of all of this was my first major project after finishing the book in which my account of content ascription was developed. Most of the work was done at the Center for Advanced Study in the Behavioral Sciences in Stanford. The Center provided me with a year free of teaching and administrative obligations and a study whose floor-to-ceiling window offered a spectacular view of San Francisco Bay. For both, I remain deeply grateful.[18] A first pass at defending my views on rationality, content, and the Davidson/Dennett argument was published in Stich (1984a), though I have since come to think that parts of that article are less than satisfactory. Chapter 2 of this volume is a revised—and I hope improved—version of that material.

1.2.4 *The Varieties of Cognitive Pluralism*

While working on the original version of that paper, I began to see, if only dimly, the outlines of what would ultimately become one of the more radical theses of this book. To explain the thesis, let me introduce a pair of claims, each of which might plausibly be labeled *cognitive pluralism.* To set them apart, I'll call one *descriptive cognitive pluralism* and the other *normative cognitive pluralism.*[19] The descriptive claim is one that has been much debated by social scientists and, more recently, by historians of science. What it asserts is that different people go about the business of cognition—the forming and revising of beliefs and other cognitive states—in significantly different ways. For example, it has been urged that people in certain "primitive" or preliterate societies think or reason very differently from the way modern, western, scientifically educated people do.[20] Closer to home, it has been suggested that different individuals in our own society solve cognitive problems in markedly different ways—ways that indicate differences in underlying cognitive processes.[21] These claims are, or at least they appear to be, empirical claims—the sort of claims that might be supported by various sorts of observations, experiments, and historical research. The denial of descriptive pluralism about cognition is *descriptive monism,* the thesis that all people exploit much the same cognitive processes. Clearly, the distinction between descriptive monism and descriptive pluralism is best viewed not as a hard and fast one but as a matter of degree. No one would deny that people differ from one another to some extent in the speed and cleverness of their inferences, nor would it be denied that in attempting to solve cognitive problems, different people try different strategies first. But if these are the only sorts of cognitive differences to be found among people, descriptive monism will be vindicated. If, on the other hand, it should turn out that different people or different cultures use radically different "psycho-logics," or that the revising and updating of their cognitive states is governed by substantially different principles, pluralism will have a firm foot in the door. The more radical the differences, the further we will be toward the pluralistic end of the spectrum.

Normative cognitive pluralism is not a claim about the cognitive processes people do use; rather it is a claim about *good cognitive processes*—the cognitive processes that people *ought to* use. What it asserts is that there is no unique system of cognitive processes that people should use, because various systems of cognitive processes that are very different from each other may all be equally good. The distinction between normative pluralism and normative monism, like the parallel distinction between descriptive notions, is best viewed as a matter of degree, with the monist end of the spectrum urging that all normatively sanctioned systems of cognitive processing are minor variations of one another. The more substantial the differences among normatively sanctioned systems, the further we move in the direction of pluralism.

Historically, it is probably true that much of the support for normative pluralism among social scientists derived from the discovery (or putative discovery) of descriptive pluralism, along with a certain ideologically inspired reluctance to pass negative judgments on the traditions or practices of other cultures. But normative pluralism was certainly not the only response to descriptive pluralism among social scientists. Many reacted to the alleged discovery of odd reasoning patterns among premodern peoples by insisting on monism at the normative level and concluding that the reasoning of premodern folk was "primitive," "prelogical," or otherwise normatively substandard.[22] I don't have a good guess as to whether normative monism or normative pluralism is more widespread among contemporary social scientists. But among philosophers, both historical and contemporary, normative cognitive pluralism is very clearly a minority view. The dominant philosophical view is that there is only one good way to go about the business of reasoning or, at most, a small cluster of similar ways. Good reasoning, philosophers typically maintain, is rational reasoning, and in the view of most philosophers, it is just not the case that there are alternative systems of reasoning differing from one another in important ways, all of which are rational.

When, in response to Nisbett's query, I began thinking seriously about the issue of cognitive or epistemic virtue—about what it is that makes a strategy of inference or reasoning a good one—I unquestioningly fell in with the prevailing prejudice and assumed that normative monism was correct. It was only as I repeatedly tried to give a monistic normative account, and repeatedly failed, that I began to suspect normative pluralism might be the better view. The account of cognitive virtue I have come to defend is floridly pluralistic. Moreover, it is relativistic as well, since it entails that different systems of reasoning may be normatively appropriate for different people. For a long time I was rather embarrassed to hold such a view, since it aligns me with a small minority among philosophers and an even smaller minority among philosophers whose work I most respect. By and large, I'm afraid, the writings of my fellow relativists are more than a bit obscure and scruffy. But I am convinced that it is possible to be a cognitive relativist without being muddleheaded, which is a good thing, since that is where the arguments lead.[23]

I said earlier that I first began to worry about cognitive pluralism while at work on my response to the Davidson/Dennett argument for the impossibility of irrationality. Let me take a moment to note how those issues are interwoven.[24] Suppose that normative cognitive monism is right—that in matters cognitive, the good is one rather than many. Then, if the Davidson/Dennett line is defensible, and serious departures from good reasoning are indeed conceptually impossible, it follows that *descriptive* cognitive monism is true as well. For if all genuine cognitive systems must be largely rational, and all rational systems are minor variations on a common theme, then all actual cognitive systems must be very similar to one another. This is, on the face of it, a quite astounding result since both descriptive monism and its denial, descriptive pluralism, *appear* to be empirical theses. But the argument just sketched for descriptive monism and against descriptive pluralism is not an empirical argument at all. One of its premises (the Davidson/Dennett thesis) purports to be a conceptual claim, while the other is a normative claim. So if the argument works, I suppose we would have to conclude that, contrary to appearances, descriptive monism and pluralism are not genuine empirical theses. This is, near enough, the view urged by Davidson in his argument against "the very idea of [alternative] conceptual schemes," and his argument, though characteristically illusive, seems to have some points in common with the one I have sketched.[25] However, the argument I've sketched is not one with which I have sympathy, and I shall argue that both of its premises are mistaken. The argument against the Davidson/Dennett thesis developed in chapter 2, and while there is a sense in which the entire book mounts an argument against normative monism, the issue is center stage in chapter 6.

The initial motive for my concern about the Davidson/Dennett thesis that rationality is a prerequisite for cognition was that it threatened to undermine the empirical explorations of irrationality that were producing, and have continued to produce, surprising and unsettling insight into human cognition. A second concern, one that became increasingly important as work on this volume proceeded, was that if the thesis were true, then much of the urgency would be drained from the project of assessing strategies of reasoning and inquiry. The interest and vitality of this branch of epistemological research can be traced, in significant measure, to the practical worries it addresses: People out there are reasoning badly, and this bad reasoning is giving rise to bad theories, many of which have nasty consequences for people's lives. But if Davidson and Dennett are right, then these concerns are overblown. Cognition *can't* be all that bad. Perhaps the reasoning of the man and woman in the street (or the jury box, or the legislature) is not quite normatively impeccable, but we need not worry about them departing in major ways from the normative ideal. This Panglossian doctrine reduces the normative evaluation of inquiry to a rather bloodless, scholastic preoccupation. We can still, if we wish, try to say what it is that makes good reasoning good. But the project can hardly be infused with the reformer's zeal, since we know in advance that there is nothing much to reform.

1.2.5 The Evolutionary Argument that Bad Reasoning Is Impossible

The conceptual argument suggested by Davidson and Dennett is not the only route to Panglossian optimism about human cognition. There is another argument for much the same conclusion hinted at by Dennett and by many other authors as well. It maintains that *biological evolution* guarantees that all normal cognitive systems will be rational, or nearly so, since organisms whose cognitive systems depart too drastically from the normative standard will run a very high risk of becoming posthumous before they have had a chance to pass on their genes to offspring. In saying that the evolutionary argument is "hinted at" by many authors, I chose my words quite deliberately, since I have been able to find nothing in the literature that amounts to anything even close to a full-dress argument. So to explore the plausibility of the view, I set about trying to build an argument myself. A first attempt was included, and criticized, in Stich (1985), and subsequent versions were tried out in front of a number of audiences from Adelaide to Helsinki. From their many helpful suggestions I have cobbled together, in chapter 3, what I believe to be the most detailed and plausible version of the evolutionary argument yet offered.[26]

That argument divides into two parts, one of which maintains that evolution produces organisms with good approximations to optimally well-designed systems, while the other maintains that an optimally well-designed cognitive system is a rational one. But, as I try to establish in chapter 3, the second part is very dubious. The first part is worse. It simply can't get off the ground without the help of a cluster of serious though widespread misunderstandings about evolution and natural selection. Though I had long realized that there is something very wrong with that part of the argument, its problems came into sharp focus only after I joined the philosophy department of the University of California at San Diego in 1986. At UCSD, I had the singular good fortune to have Philip Kitcher as a colleague and mentor on these matters. My critique of the evolutionary argument, set out in detail in chapter 3, borrows frequently from Kitcher's work and has benefited enormously from his good advice.

1.2.6 Competence, Performance, and Reflective Equilibrium: Cohen's Argument That Bad Reasoning Is Impossible

The Davidson/Dennett argument and the evolutionary argument conclude that widespread departures from normative standards of reasoning are impossible or unlikely. Thus, they challenge the conviction, widely held by psychologists who study reasoning, that the cognitive processes of ordinary men and women could be significantly improved. But the empirical literature on reasoning was not center stage in the research agendas of Davidson, Dennett, or others who urged versions of these arguments. So the conflict between those arguments and the usual interpretation of the empirical results went largely unnoticed.[27] But to L. J. Cohen, those results were both salient and paradoxical. Many of the undergraduates who serve as subjects for the experiments, Cohen noted, will go on to become leading scientists, jurists, and civil servants. How could they be so successful, Cohen asked, if they do not know how to reason well?

The answer that Cohen urged was that the subjects *do* know how to reason well.[28] Indeed, he argued, the suggestion that they don't is demonstrably incoherent. Central to Cohen's argument is the distinction between *competence* and *performance* that has loomed large in recent linguistics. In the linguistic domain, a person's competence is typically identified with his tacit knowledge of the grammatical rules of his language. In the domain of reasoning, Cohen urged, a person's competence can be identified with his tacit knowledge of his "psychologic"—the rules he exploits as he goes about the business of reasoning. The crucial, and enormously clever, step in Cohen's argument is his demonstration that if we adopt something like Goodman's account of what it is for an inferential rule to be justified, it follows that the rules constituting a person's reasoning competence will inevitably be justified. Thus in the domain of inference, people's competence must be normatively impeccable. In 4.2, I'll set out a detailed sketch of Cohen's argument. At the time I first heard Cohen's argument, Nisbett and I had already become convinced that the version of Goodman's normative account needed for Cohen's argument to work would have to be rejected. Thus yet another argument for the inevitability of rationality came to grief.

1.3 The Search for a Theory of Cognitive Evaluation: Clearing the Ground

From time to time I have jokingly described my efforts in chapters 2 and 3, and in my critique of Cohen, as an attempt to make the world safe for irrationality. The point is not, of course, that irrationality is a good thing or that bad reasoning is to be encouraged. But if bad cognitive processing were either conceptually or biologically impossible, it would make nonsense of the empirical exploration of reasoning and its foibles. It would also turn the effort to articulate and defend a normative theory of cognition into an arcane academic exercise of no particular practical importance. This second consequence of the irrationality-is-impossible thesis came to seem doubly unwelcome to me since, while working on chapters 2 and 3, I was also investing considerable time and effort trying to rework Goodman's proposal into a defensible criterion for distinguishing good inferential strategies from bad ones.

1.3.1 Goodman's Project Might Turn Out to Be Impossible

This work was not going at all well, however. As time passed, I accumulated a substantial collection of variations on Goodman's idea and an even more substantial collection of arguments showing that none of them worked. Some of those variations, and the arguments against them, are assembled in chapter 4.

At about the time I was getting thoroughly discouraged with my neo-Goodmanian quest, two lines of argument began to take shape in my mind. The first of these suggested that the Goodmanian project might very well turn out to be impossible. Initially at least, that struck me as a most unwelcome conclusion, since I had long thought that Goodman's approach to building a normative theory of cognition was by far the most promising one available. The conclusion followed from a cluster of considerations all of which shared a common theme: The Goodmanian approach tacitly presupposes a number of empirical theses, and each of these stands in some serious risk of turning out to be false. To see these empirical presuppositions, it helps to back up a bit and get a broader view of what Goodman was up to.

Goodman had sketched a procedure or test, that a system of inferential rules should pass if it is to count as rational or justified. Others, including Nisbett and I, had argued that Goodman's test was inadequate and had proposed a variety of modifications. But what counts as getting the story right here? What is the relation between rationality and the right test supposed to be, and why is the fact that a system of inference passes some test or other supposed to show that the system is rational? I think the most plausible answer for a Goodmanian to give is that the right test, when we discover it, will be an *analysis* or *explication* of our ordinary concept of rationality (or some other commonsense concept of epistemic evaluation).[29] The test—which will be a tidied up version of the procedures we actually follow in evaluating the merits of an inferential system—provides necessary and sufficient conditions for rationality because it unpacks our concept of rationality; it tells us what that concept comes to.

Now for this sort of answer to be defensible, it must be the case that our commonsense concept of rationality is univocal and more or less coherent and that it is structured so as to admit of an analysis or explication in terms of necessary and sufficient conditions. It must also be the case that the procedures we use for deciding whether a system is rational exhausts the content of the concept. None of this can safely be supposed a priori. The conclusion to be drawn from these considerations is not that our commonsense concepts of epistemic evaluation *do not* admit of the sort of explication that Goodman's project seeks—the evidence needed to settle that is far from in—but merely that they *may* not. The feasibility of the Goodmanian project is very much hostage to the psychological facts. Moreover, as argued in chapter 4, what little we know about the mental representation of concepts gives scant reason for optimism.

1.3.2 The Irrelevance of Analytic Epistemology

These ominous thoughts had taken shape during the year I spent at the Center for Advanced Study in the Behavioral Sciences. At the end of that year I took my forebodings with me to Australia, where I was to spend a year as visiting professor at the University of Sydney. There, perhaps influenced by the refreshing Australian iconoclasm of my colleagues and students in the Department of Traditional and Modern Philosophy, a second line of argument began to come into focus, one that was both more radical and more liberating. This second line began as yet another critique of the Goodmanian project. But it soon became evident that its real target was much larger. If the argument is right, it undermines the entire analytic epistemology tradition, a tradition that has been dominant in the English-speaking world for the last quarter century or more. It is a hallmark of that tradition to seek criteria of cognitive evaluation in the analysis or explication of our ordinary concepts of epistemic evaluation. However, one of the conclusions that drops out of this second line of argument is that, for almost anyone who takes the project of evaluating cognitive processes seriously, analytic epistemology is going to be a hopeless nonstarter. If the analytic strategy were the only one around, this conclusion would be as disheartening as it is radical. However, the argument that undermines analytic epistemology also highlights the virtues of a very different strategy for cognitive evaluation.

The starting point for the argument against analytic epistemology is the observation that if descriptive cognitive pluralism is true—if different people go about the business of reasoning in significantly different ways, some of which may be substantially better than others—then much of this divergence is likely to be traceable to cultural differences, though genetic factors and idiosyncratic differences in individual experience may also play a role. In attempting to evaluate these divergent strategies and in deciding which of them we ought ourselves to use, we are trying to decide among a variety of cultural products. The analytic epistemologist proposes to evaluate these differing cognitive processes by explicating our intuitive notions of cognitive evaluation, and then exploring which inferential processes fall most comfortably within the extention of those notions. But these intuitive notions of cognitive evaluation are themselves local cultural products, and there is no reason to think that they won't exhibit just as much intercultural and interpersonal variation as the cognitive processes that they evaluate. In light of this, it is hard to see why most people would *care* very much whether a system of cognitive processes falls within the extension of some ordinary notion of epistemic evaluation—why, for example, they would care whether their reasoning falls within the boundaries of the intuitive notion of rationality—unless of course there is some reason to think that falling within the extension of one of these concepts correlates with something else we do care about.

To be sure, there may be some exceptions here. It is not unthinkable for a person to find intrinsic value in having cognitive processes sanctioned by our culturally inherited concepts of epistemic evaluation, just as a person might find intrinsic value in adhering to the traditional social practices of his ethnic group. In each case the person recognizes that the concepts or practices in question are just one set among many that people might or do exploit. Yet he values his own, and does so for no further reason. He values them for their own sake. At one time I thought it was clear that this was the position embraced by David Stove, the most outspoken of my Sydney colleagues and a man whom I came to regard as one of the most acute, and most conservative, cultural critics of our time. However, Stove has protested that this is a hopeless caricature of his view. The reader may wish to read Stove (1986) and judge for him- or herself.

Some writers have been tempted by the Wittgensteinian idea that epistemic assessments must come to an end with the criteria embedded in our ordinary concepts of cognitive evaluation. But surely this is nonsense. Both our notions of epistemic evaluation and (more important) our cognitive processes themselves can be evaluated *instrumentally*. That is, they can be evaluated by how well they do at bringing about states of affairs that people do typically value—states of affairs like being able to predict or control nature, or contributing to an interesting and fulfilling life. The idea of viewing cognitive processes as mental tools, to be evaluated as we might evaluate other sorts of tools, has its roots in the pragmatist tradition, and it came to play a quite central role in my thinking, once I had finally persuaded myself that the Goodmanian strategy, and the analytic tradition in which it was embedded, were of no use to the project of evaluating cognitive strategies.

1.3.3 An Attack on Truth

In rejecting the appeals to our ordinary notions of epistemic evaluation—appeals to rationality, justification, and the rest—as final arbiters in our efforts to choose among competing strategies of inquiry, I was, in effect, denying that rationality or justification have any intrinsic or ultimate value. At that juncture, a natural question to ask was whether there was any other paradigmatically epistemic feature of our cognitive lives that might be taken to be intrinsically valuable. And when the question was posed in this way, there was an obvious candidate: *truth.* However, from my earliest work in the philosophy of language, I had harbored a certain skepticism about the utility, indeed even the intelligibility, of the notion of truth.[30] And in the process of polishing the argument against analytic epistemology, I came to suspect that there was a largely parallel argument to be mounted against truth.[31] Thus I came to think that neither being rational nor generating truth would turn out to be an intrinsically valuable feature for cognitive processes to have. If the argument about the value of truth could be sustained, the natural upshot for the normative theory of cognition would be a thoroughgoing pragmatism which holds that all cognitive value is instrumental or pragmatic—that there are no intrinsic, uniquely cognitive values. And this, indeed, is the position I finally came to defend. But once again I have let my story get ahead of itself. So let me go back and retrace the steps that led me to reject truth as a cognitive virtue.

The path that led me to the surprising conclusion that truth is not to be taken seriously as a cognitive virtue begins with the observation that Goodmanian analyses of notions of cognitive evaluation are only one approach in the analytic epistemology tradition. Another well-explored idea is that the rationality or justifiedness of a set of cognitive processes can be explicated by appeal to the success or failure of the processes in producing true beliefs. There are lots of variations on this "reliabilist" theme, but all of them offer analyses on which the normative status of a cognitive process is at least in part a function of how well it does in producing true beliefs.[32] An advocate of reliabilism might be tempted to think that this sort of analysis blunts the conclusion of the argument against analytic epistemology. To be sure, the reliabilist might concede, the mere fact that a cognitive process is sanctioned by some socially transmitted notion of epistemic evaluation or other is no reason to favor that process. But if the evaluative notion can be explicated along reliabilist lines, the situation is very different. For then processes that fall within the extension of the notion do a good job of producing true beliefs. And for most people, having true beliefs *is* intrinsically valuable.

Now it surely is the case that many people, if asked, would profess to value having true beliefs. But most of these same people would be hard pressed to say anything coherent about what it is for a belief to be true and thus would be quite unable to explain what it is that they value. This is not to suggest that there is nothing to be said on that score, however. Quite the opposite. In recent years philosophers, particularly those concerned with issues in the philosophy of mind and the philosophy of language, have lavished considerable attention on theories of mental representation (or "psycho-semantics") whose central concern is to explain how psychological states like beliefs could come to have semantic properties—properties like being true or being false, being about (or referring to) a particular person, and so forth. The strategy I pursued in exploring whether we might undermine the conviction of those who think they take true belief to be intrinsically valuable was to extract, from some of the more plausible and well–worked out of these theories of mental representation, an account of just what it is, on those theories, for a belief to be true. With such an account in hand, we can have people ponder the following question: If *that* is indeed what it is for a belief to be true, do you really care whether your beliefs are true? Do people really value having true beliefs once they are offered a clear view of what having a true belief comes to? The result of this exercise, at least in my own case (and I don't think my values are idiosyncratic here), was a consistently negative answer. Moreover, the negative judgment was most firm in just those cases where the account of true belief was clearest and most explicit. The argument on this point is developed in chapter 5, and while it is not an argument that admits of a comfortable summary, the following brief remarks may help a bit in anticipating what is to come.

Let's start with beliefs. What sorts of things are they? On one widely held view, the so-called token identity theory, belief tokens (like my current belief that Princeton is south of New Brunswick) are brain states—no doubt fairly complex ones. If we accept this view, the next question to ask is how a brain state could possibly have semantic properties: how could a complex neural event have a truth value? A variety of plausible answers begin with the idea that there is a function

(I'll call it the "interpretation function") that maps belief-brain-states to a class of entities that are more comfortably semantically evaluable—entities like propositions, or truth conditions, or situations, or possible states of affairs. According to theories of this sort, a belief is true if and only if the proposition to which it is mapped is true (or if and only if the possible state of affairs obtains, etc.). But of course mappings or functions are easy to come by. If there is one function from belief-brain-states to propositions, then there are indefinitely many. Obviously, not just any mapping will do. Which function is the right one?

This last question may be read in two different ways. On one reading it is a request for a detailed story about the right interpretation function, a story that specifies which belief-brain-states the theorist would map to which propositions. In the literature on mental representation, there is much careful discussion, and no shortage of debate, about how the interpretation function is to be constructed. On another reading, the question is not a request for details but rather a request for a criterion of what it is to get those details right. Suppose a pair of theorists offer competing detailed stories about the function from brains states to propositions. How would we go about deciding which one is the right one? On this issue there is relatively little discussion in the literature. However, as I argue in chapter 5, it is clear from the sorts of arguments that theorists use in defense of their proposed interpretation functions that one strong constraint on getting the mapping right is that it must generally capture our intuitive judgments about the content or truth conditions of the mental states in its domain. Absent special circumstances, a function that assigns counterintuitive truth conditions is the wrong function.

But now what is so special about those intuitions? Why do they get to have such a significant say in deciding which interpretation function is the right one? Presumably the intuitions theorists exploit are a socially acquired set of judgments that may well vary from one individual or culture to another. Just as it is hard to see why the intuitive, socially shaped notions of epistemic evaluation that prevail in a culture should command such commitment, so too it is hard to see why anyone would take our intuitively sanctioned interpretation function to be special or important, unless of course there is some reason to believe it correlates with something else that is more generally valued. But recall that, on the accounts of mental representation in question, a belief is true if and only if it is mapped to a true proposition *by the intuitively sanctioned mapping function*. If it is granted that there is nothing uniquely special or important about that intuitive function—that it is simply one mapping among many—it would seem to follow that there is nothing special or important about having true beliefs. One could still find true beliefs to be intrinsically valuable, of course. But in light of these considerations, it seems a curious, culturally local value, on a par with finding intrinsic value in the cultural practices of one's ethnic group. Alternatively, one could give up on the intrinsic value of true beliefs and urge, instead, that true beliefs and the cognitive systems that tend to produce them are instrumentally valuable because they foster our pursuit of other goals. This last idea is considered at some length in chapter 5. Though I have no knockdown argument against it, I try to show that making a case even for the instrumental value of true beliefs is, to put it mildly, no easy matter.

1.4 Epistemic Pragmatism

On more than one occasion, when recounting my none too sanguine views about the notions of rationality and truth, I have been accused of intellectual vandalism—defacing established structures and offering nothing positive to replace them. And I must admit that from time to time, particularly when the shortcomings of various traditional doctrines were coming clearly into focus while the outlines of a defensible alternative were still very hazy, I began to wonder whether the charge might not be justified. But as work progressed, I became increasingly convinced that a pragmatic account of cognitive evaluation could avoid the difficulties that scuttle analytic and truth-linked accounts.

Central to a pragmatic account is the very Jamesian contention that *there are no intrinsic epistemic virtues.*[33] Rather, for the pragmatist, cognitive mechanisms or processes are to be viewed as tools or policies and evaluated in much the same way that we evaluate other tools or policies. One system of cognitive mechanisms is preferable to another if, in using it, we are more likely to achieve those things that we intrinsically value. At the beginning of chapter 6, I try to show why this sort of pragmatic account of cognitive evaluation is suggested by the shortcomings of other accounts. I then go on to explore some of the problems to be expected when we unpack the notion of one cognitive system being more likely than another to enable us to achieve those things that we intrinsically value.

1.4.1 Objections to Pragmatism

Though the attractions of a pragmatic account of cognitive evaluation were fairly easy to see, it took me a long time to take pragmatism seriously. For along with the virtues of the view, there were a pair of obvious objections, each of which initially seemed quite overwhelming. The first objection is that pragmatism leads to relativism. And relativism, I had long thought, is to be avoided at all costs in matters epistemic. The second objection is that pragmatism is viciously circular, since there is no way we could show that our cognitive system is pragmatically preferable without using the very system whose superiority we are trying to establish. The more I thought about these objections, however, the less objectionable they came to appear.

1.4.1.1 Relativism It can hardly be denied that a pragmatic account of cognitive evaluation is relativistic, since the pragmatic assessment of a cognitive system will be sensitive to both the values and the circumstances of the people using it. Thus it may well turn out that one cognitive system is pragmatically better than a second *for me*, while the second is pragmatically better than the first *for someone else*. But while it is obvious that pragmatism is relativistic, I gradually came to realize that it is not at all obvious why this is a bad thing. Indeed, despite the widespread prejudice against relativism, I found it surprisingly hard to find any plausible, published arguments aimed at showing why epistemic relativism would be unwelcome. Since serious arguments were in such short supply in the literature, I set about quizzing philosophical friends about the grounds of their antipathy toward the view.

The charge that I heard most often was that relativistic accounts of cognitive assessment play into the hands of epistemic nihilists because they abandon any serious attempt to separate good cognitive strategies from bad ones. For both the relativist and the nihilist, it was said, "anything goes." But this, as I argue in 6.2.2.1, is just a mistake as far as epistemic pragmatism is concerned. The pragmatic assessment of cognitive strategies is certainly relativistic, but it is no more nihilistic than the pragmatic assessment of investment strategies or engineering techniques.

A second accusation was that relativism, or indeed any version of normative pluralism, leads to skepticism by driving a wedge between good reasoning, on the one hand, and truth, on the other. If significantly different systems of reasoning may be preferable for different people, and if these systems generate significantly different beliefs on the basis of similar sensory input, then presumably different people may end up with different beliefs even though they have much the same evidence and are all invoking normatively sanctioned systems of reasoning. Under those circumstances, however, it is hard to see how we could ever defend the view that good reasoning is likely to lead to truth. And, of course, one of the classic worries of the skeptic is that no matter how good our reasoning may be, it will not lead us to the truth.

When I first began thinking about how an epistemic pragmatist might respond to this challenge, my goal was to find some weakness in the argument; and as we'll see in 6.2.2.2, they are not all that hard to find. After finding them, however, it occurred to me that I had a more powerful, and more radical, reply available, one which was not hostage to the success or failure of various possible strategies for shoring up the weaknesses in the argument. For even if we grant that a relativistic account of cognitive evaluation will make it impossible to show that good reasoning generally leads to truth, this should be no cause for concern unless we have some reason to *want* our cognitive systems to produce true beliefs. And the burden of my argument in chapter 5 is that most of us have no such reason. Epistemic skepticism generally simply assumes the importance or desirability of the commodity it claims to be beyond our reach—be it certainty, or truth, or justification. And for the most part those who have done battle with the skeptic have shared this assumption. But on my view, the best first response to the skeptic who maintains that we cannot achieve certainty, truth, knowledge, or what have you, is not to argue that we can. Rather, it is to ask, so what? There is no point in worrying about whether our cognitive system can generate beliefs with a given property unless we think such beliefs would be of some value. So before investing any effort in the skeptic's argument, we should demand some explanation of why it matters whether our beliefs are true. And if the arguments in chapter 5 are on the right track, that explanation is not going to be easy.

1.4.1.2 Circularity While working on relativism, I realized that truth-linked accounts of cognitive evaluation are also typically relativistic, though the critics of truth-linked accounts rarely make much of it.[34] By contrast, critics are forever protesting that truth-linked accounts are unacceptable because they are circular. And this, it turned out, was a great convenience for my project, since both the accusations and the responses are much the same when the targets of the circularity complaint are pragmatic accounts rather than truth-linked accounts. Even more convenient was the fact that Goldman's admirable, systematic defense of truth-linked theories had recently appeared (1986), and in that volume he set out some of those responses in a singularly clear and persuasive way. Thus, when it came time for me to defend pragmatic accounts against the charge of circularity, in 6.3, I simply adapted Goldman's arguments and added a few bells and whistles of my own.

1.4.2 *The Empirical Study of Human Reasoning: Pragmatism Applied*

The line of thought that ultimately led me to advocate epistemic pragmatism had begun, almost a decade earlier, with Nisbett's worry about replying to critics who insisted that in many experiments which allegedly demonstrated that people reason poorly, the subjects were in fact reasoning quite well. Having convinced myself of the merits of a pragmatic account of cognitive evaluation, the obvious next step was to apply that account to the disputed cases. But when I set out to do so, I was in for a pair of surprises. The first surprise was that the pragmatic story, as I had developed it to that point, did not really address the question of whether the subjects were reasoning well or badly. For, as I tell the story, the pragmatic account is a *comparative* account—it tells us whether one system is better than another (for a given person, at a given time). It does not tell us whether any given system is a good one (for a person, at a time) full stop. To address this question, it's tempting to say that a system is a good one if it is pragmatically at least as good as any possible alternative. But this leaves us with the problem of saying just how we are to understand the notion of a "possible alternative," and that, it turns out, is far from trivial. For, as Christopher Cherniak has argued, if we follow a venerable tradition in epistemology and take "possible" to mean "logically possible," then many of the "possible alternative strategies" are going to be *vastly* beyond anything that brains like ours could use. And it seems simply perverse to judge that subjects are doing a bad job of reasoning because they are not using a strategy that requires a brain the size of a blimp. It seems, then, that in deciding whether a person is doing a good job of reasoning, we should compare his cognitive system not to all logically possible alternatives but only to the *feasible* alternatives. But which ones are these?

One way to approach this question would be to adopt the standard strategy of analytic philosophy: We try to set down necessary and sufficient conditions for feasibility, and then test these conditions against our intuitions about a host of particular cases. But that strategy looked singularly unappealing. For, as I argue in 6.4, what we take to be feasible is going to depend on the purposes at hand and on the technologies available. Thus there is no clear sense to questions about the goodness or badness of a given inferential strategy when asked abstractly. Before trying to answer such questions, we have to get clear on what William James would call the "cash value" of the question; we have to ask what sorts of actions would be suggested by one answer or another.

While there are many projects that might involve an evaluation of people's cognitive processes, one that has played a central role in epistemology from Descartes to Popper and Goldman is the project of *improving* people's cognitive performance. If that is our goal in assessing existing cognitive systems, then the relevant class of feasible alternatives are those that we could actually get people to use. Which those are is not something we can discover without some serious empirical exploration. Though from the work of Cherniak, Harman, and others, we can be pretty certain that some very demanding cognitive strategies, like those that require the elimination of all inconsistencies or those that require us to keep track of the evidence for all of our beliefs, are well beyond the reach of brains like ours. The result of all this—and this was my second surprise—is that it is not at all clear that Nisbett's critics were wrong, since it is far from obvious that the subjects were reasoning badly. To show that they were, we would have to show that there is a pragmatically superior alternative and that people can actually be taught to use it. But surprisingly little is known about the effectiveness of various techniques aimed at altering people's cognitive processes. Until we know more, we often won't be able to say whether Nisbett's subjects, and those who reason like them, are reasoning well or badly.

It is a consequence of the sort of epistemic pragmatism I advocate that a pair of venerable epistemological concerns—the assessment of people's cognitive performance and the improvement of people's cognitive performance—become inextricably interwoven with empirical explorations of psychological feasibility. Those explorations are in turn dependent on the advancing state of the art in the various technologies which help to determine what is feasible. There is a long tradition in epistemology which would reject out of hand any proposal that makes epistemological questions dependent on empirical findings or technological developments. But that is a tradition which I, in the company of a growing number of philosophers, take to be sterile and moribund. Another, younger tradition in epistemology, tracing to James and Dewey, finds nothing untoward in the suggestion that epistemology is inseparable from science and technology. From the perspective of that tradition, the doctrines defended in the pages to follow hardly look radical at all.