

Studies in Epistemology and Cognitive Theory

Series Editor: *Paul K. Moser, Loyola University of Chicago*

Blind Realism: An Essay on Human Knowledge and Natural Science

by Robert Almeder, Georgia State University

Epistemic Virtue and Doxastic Responsibility

by James A. Montmarquet, Tennessee State University

Rationality, Morality, and Self-Interest: Essays Honoring Mark Carl Overvold

edited by John Heil, Davidson College

The Problem of the Criterion

by Robert P. Amico, Saint Bonaventure University

Socializing Epistemology: The Social Dimensions of Knowledge

edited by Frederick F. Schmitt, University of Illinois at Urbana-Champaign

Metaepistemology and Skepticism

by Richard Fumerton, University of Iowa

Warrant in Contemporary Epistemology: Essays in Honor of Plantinga's

Theory of Knowledge

edited by Jonathan L. Kvanvig, Texas A&M University

A Defense of the Given

by Evan Fales, University of Iowa

Pragmatism and Realism

by Frederick L. Will; edited by Kenneth R. Westphal, University of New Hampshire

Realism/Antirealism and Epistemology

edited by Christopher B. Kulp, Santa Clara University

Belief and Knowledge: Mapping the Cognitive Landscape

by Kenneth M. Sayre, University of Notre Dame

In Defense of Radical Empiricism: Essays and Lectures by Roderick Firth

edited by John Troyer, University of Connecticut—Storrs

Illusions of Paradox: A Feminist Epistemology Naturalized

by Richmond Campbell, Dalhousie University

Rethinking Intuition: The Psychology of Intuition and Its Role in Philosophical Inquiry

edited by Michael R. DePaul, University of Notre Dame, and William Ramsey, University of Notre Dame

STICK

Rethinking Intuition

The Psychology of Intuition and Its Role in Philosophical Inquiry

edited by

Michael R. DePaul
and William Ramsey

NOTICE

This material may be
protected by copyright
law (Title 17 U.S. Code.)

ROWMAN & LITTLEFIELD PUBLISHERS, INC.
Lanham • Boulder • New York • Oxford

Chapter 6

Reflective Equilibrium, Analytic Epistemology and the Problem of Cognitive Diversity

Stephen Stich

This chapter is about different ways of thinking—or cognitive diversity, as I shall sometimes say—and the problem of choosing among them. In the pages to follow I will defend a pair of claims. The first is that one influential proposal for solving the problem of cognitive diversity, a proposal that invokes the notion of reflective equilibrium, will not work. The second is much more radical. What I propose to argue is that although some of the objections to the reflective equilibrium solution turn on details of that idea, the most serious objection generalizes into an argument against an entire epistemological tradition—the tradition that I shall call “analytic epistemology.” Before attending to either of these claims, however, I will have to say something about how I conceive of cognition and cognitive diversity.

Cognition and Cognitive Diversity

Let me begin with a simplifying assumption that I hope you will not find wildly implausible. I shall assume that in humans and other higher animals there is a distinct category of mental states whose function it is to store information about the world. When the organisms in question are normal, adult humans in a culture not too remote from our own, folk psychology labels these states *beliefs*. Whether or not this folk label can be used appropriately for the belief-like states of animals, automata, young children and exotic folk is a question of considerable controversy. (See Davidson 1982; Routley 1981; Stich 1979, 1983: 89-106, 1984.) For present purposes, however, it is a controversy best avoided. Thus, I propose to adopt the term “cognitive state” as a broad cover term whose extension includes not only beliefs properly so-called, but also the belief-like information-storing mental states of animals, young children, and those adult humans, if any there be, whose cognitive lives differ substantially from our own.

Our beliefs, and the cognitive states of other creatures, are in a constant state of flux. New ones are added and old ones removed as the result of perception, and as a result of various processes in which cognitive states interact with each

other. In familiar cases, folk psychology provides us with labels like “thinking” and “reasoning” for these processes, though once again the propriety of these labels becomes controversial when the cognitive states being modified are those of children, animals, or exotic folk. So I will use the term “cognitive processes” as a cover term whose extension includes our own reasoning processes, the updating of our beliefs as the result of perception, and the more or less similar processes that occur in other organisms.

Cognitive processes are biological processes; they are something that brains do. And, like other biological processes, they have been shaped by natural selection. Thus, it is to be expected that our genes exert an important influence on the sorts of cognitive processes we have. It is also to be expected that the cognitive processes of other species with other needs and other natural environments will be in varying degrees different from those to be found among humans. But from the fact that genes inevitably exert a major influence on cognitive processes it does not follow that all of our cognitive processes are innate, or, indeed, that any of them are.

To see the point, we need only reflect on the case of language. My ability to speak English is a biological ability; processing English is something my brain does. Moreover, my genes are surely heavily implicated in the explanation of how I came to have a brain that could process English. Still, English is not innate. The ability to process English is an ability I acquired, and had I been raised in a different environment I might have acquired instead the ability to speak Korean or Lapp. This is not to deny that *something* relevant to language is innate. All normal human children have the ability to acquire the language spoken around them. And that is a very special ability. There is no serious evidence indicating that members of any other species can acquire human languages or anything much like them.

Now the point I want to stress is that, as far as we know, human cognitive processes may be like human language processing abilities. They may be acquired in ways that are deeply dependent on environmental variables, and they may differ quite radically from one individual or culture to another. Of course, it is also possible that human cognitive processes are much less plastic and much less under the influence of environmental variables. It is possible that cognition is more similar to digestion than to language. To make matters a bit messier, there is no reason a priori for all cognitive processes to be at the same point on this continuum. It may be that some of our cognitive processes are shared by all normal humans, while others are a part of our cultural heritage.¹ I am inclined to think that this last possibility is the most plausible one in the light of available evidence, and for the remainder of this chapter I will take it for granted. But it must be admitted that the evidence is both fragmentary and very difficult to interpret. (See Cole and Scribner 1974; Cole and Means 1981.)

If we suppose that there is a fair amount of acquired diversity in human cognitive processes, and that patterns of reasoning or cognitive processing are to some substantial degree molded by cultural influences, it adds a certain urgency to one of the more venerable questions of epistemology. For if there are lots of different ways in which the human mind/brain can go about ordering and reordering its cognitive states, if different cultures could or do go about the business of

reasoning in very different ways, *which of these ways should we use?* Which cognitive processes are the *good* ones? It is just here that the analogy with language breaks down in an illuminating way. Most of us are inclined to think that, at least to a first approximation, one language is as good as another. The one you should use is the one spoken and understood by the people around you.² By contrast, most of us are *not* inclined to accept this sort of thorough-going relativism about cognitive processes. If primitive tribesmen or premodern scientists or our own descendants think in ways that are quite different from the ways we think, few of us would be inclined to suggest that all of these ways are equally good. Some ways of going about the business of belief revision are better than others. But just what is it that makes one system of cognitive processes better than another, and how are we to tell which system of reasoning is best? In the remaining sections of this chapter I want to consider one influential answer to this question. I shall argue that both the answer itself and the philosophical tradition it grows out of should be rejected.

Reflective Equilibrium as a Criterion for Assessing Cognitive Processes

The answer I will disparage was first suggested about three decades ago when, in one of the more influential passages of twentieth-century philosophy, Nelson Goodman (1965) described a process of bringing judgments about particular inferences and about general principles of inference into accord with one another. In the accord thus achieved, Goodman maintained, lay all the justification needed, and all the justification possible for the inferential principles that emerged. Other writers, most notably John Rawls, have adopted a modified version of Goodman's process as a procedure for justifying moral principles and moral judgments. To Rawls, too, we owe the term ‘reflective equilibrium,’ which has been widely used to characterize a system of principles and judgments that have been brought into coherence with one another in the way that Goodman describes (Rawls 1971: 20ff).

It is hard to imagine the notion of reflective equilibrium explained more eloquently than Goodman himself explains it.

How do we justify a *deduction*? Plainly by showing that it conforms with the general rules of deductive inference. An argument that so conforms is justified or valid, even if its conclusion happens to be false. An argument that violates a rule is fallacious even if its conclusion happens to be true. . . . Analogously, the basic task in justifying an inductive inference is to show that it conforms to the general rules of *induction*. . . .

Yet, of course, the rules themselves must ultimately be justified. The validity of a deduction depends not upon conformity to any purely arbitrary rules we may contrive, but upon conformity with valid rules. When we speak of *the* rules of inference we mean the valid rules—or better, *some* valid rules, since there may be alternative sets of equally valid rules. But how is the validity of rules to be determined? Here . . . we encounter

philosophers who insist that these rules follow from some self-evident axiom, and others who try to show that the rules are grounded in the very nature of the human mind. I think the answer lies much nearer to the surface. Principles of deductive inference are justified by their conformity with accepted deductive practice. Their validity depends upon accordance with the particular deductive inferences we actually make and sanction. If a rule yields unacceptable inferences, we drop it as invalid. Justification of general rules thus derives from judgments rejecting or accepting particular deductive inferences.

This looks flagrantly circular. I have said that deductive inferences are justified by their conformity to valid general rules, and that general rules are justified by their conformity to valid inferences. But this circle is a virtuous one. *A rule is amended if it yields an inference we are unwilling to accept; an inference is rejected if it violates a rule we are unwilling to amend.* The process of justification is the delicate one of making mutual adjustments between rules and accepted inferences; and in the agreement thus achieved lies the only justification needed for either.

All this applies equally well to induction. An inductive inference, too, is justified by conformity to general rules, and a general rule by conformity to accepted inductive inferences. (Goodman 1965: 66-67; emphasis is Goodman's.)

There are three points in this passage that demand a bit of interpretation. First, Goodman claims to be explaining what justifies deductive and inductive inferences. However, it is not clear that, as he uses the term, *inference* is a cognitive process. It is possible to read Goodman as offering an account of the justification of principles of logic and of steps in logical derivations. Read in this way, Goodman's account of justification would be of no help in dealing with the problem of cognitive diversity unless it was supplemented with a suitable theory about the relation between logic and good reasoning. But as several authors have lately noted, that relation is much less obvious than one might suppose (Cherniak 1986: chap. 4; Harman 1986: chap. 2; Goldman 1986: section 5.1). It is also possible to read Goodman as speaking directly to the question of how we should go about the business of reasoning³ and offering a solution to the problem of cognitive diversity. This is the reading I propose to adopt.

A second point that needs some elaboration is just what status Goodman would claim for the reflective equilibrium test he describes. It is clear Goodman thinks we can conclude that a system of inferential rules is justified if it passes the reflective equilibrium test. But it is not clear *why* we can conclude this. Two different sorts of answers are possible. According to one answer, the reflective equilibrium test is *constitutive* of justification or validity. For a system of inferential rules to be justified just *is* for them to be in reflective equilibrium. Another sort of answer is that if a set of inferential principles passes the reflective equilibrium test, this counts as good *evidence* for them being valid or justified. But, on this second view, being in reflective equilibrium and being justified are quite different. One is not to be identified with the other. I am inclined to think that it is the former, constitutive, view that best captures Goodman's intentions.

But since my concern is to criticize a view and not an author, I do not propose to argue the point. Rather, I will simply stipulate that the constitutive reading is the one I am stalking.⁴

The third point of interpretation concerns the status of the claim that reflective equilibrium is constitutive of justification. On this point, there are at least three views worth mentioning. The first is that the claim is a *conceptual truth*—that it follows from the meaning of 'justification' or from the analysis of the concept of justification. Like other conceptual truths, it is both necessarily true and knowable a priori. If we adopt this view, the status of the claim that reflective equilibrium is constitutive of justification would be akin to the status of the claim that being a closed, three-sided plane figure is constitutive of being a triangle, though the claim about justification is, of course, a much less obvious conceptual truth. A second view is that the claim is a nonconceptual necessary truth that is knowable only a posteriori. This would accord it much the same status that some philosophers accord to the claim that water is H₂O. Finally, it might be urged that the claim is being offered as a stipulative proposal. It is not telling us what our preexisting concept of justification amounts to, nor what is essential to the referent of that concept. Rather, in a revisionary spirit, it is proposing a new notion of justification. Actually, the divide between the first and the last of these alternatives is not all that sharp, for one might start with an analysis of our ordinary notion and go on to propose modifications in an effort to tidy the notion up a bit here and there. As the changes proposed get bigger and bigger, this sort of "explication" gradually shades into pure stipulation. So long as the changes an explication urges in a preexisting concept are motivated by considerations of simplicity and do not result in any radical departures from the ordinary concept, I will count them as a kind of conceptual analysis. I think a good case can be made that Goodman took himself to be providing just such a conservative explication. But again, since it is a view rather than an author that I hope to refute, I will simply stipulate that the conceptual analysis or conservative explication interpretation is the one to be adopted here.

Does the Reflective Equilibrium Account Capture our Notion of Justification?

Goodman, as I propose to read him, offers us an account of what our concept of justified inference comes to. How can we determine whether his analysis is correct? One obvious strategy is to ask just what systems of inferential rules result from the process of mutual adjustment that Goodman advocates. If the inferential systems generated by the reflective equilibrium process strike us as systems that a rational person ought to invoke, this will count in favor of Goodman's analysis. If, on the other hand, the reflective equilibrium process generates what we take to be irrational or unjustified inferential rules or practices, this will cast doubt on Goodman's claim to have captured our concept of justification. Since we are viewing conceptual explication as a kind of analysis, we should not insist that Goodman's account coincide perfectly with our intuitive judgments. But if there are lots of cases in which Goodman's account entails that a system of in-

ferential rules is justified and intuition decrees that it is not, this is a symptom that the analysis is in serious trouble.

In an earlier paper, Nisbett and I exploited the strategy just described to argue that the reflective equilibrium account does not capture anything much like our ordinary notion of justification (Stich and Nisbett 1980). On the basis of both controlled studies and anecdotal evidence, we argued that patently unacceptable rules of inference would pass the reflective equilibrium test for many people. For example, it appears likely that many people infer in accordance with some version of the gambler's fallacy when dealing with games of chance. These people infer that the likelihood of throwing a seven in a game of craps increases each time a nonseven is thrown. What is more, there is every reason to think that the principle underlying their inference is in reflective equilibrium for them. When the principle is articulated and the subjects have had a chance to reflect upon it and upon their own inferential practice, they accept both. Indeed, one can even find some nineteenth-century logic texts in which versions of the gambler's fallacy are explicitly endorsed. (In a delightful irony, one of these books was written by a man who held the same chair Goodman held when he wrote *Fact, Fiction and Forecast*.)⁵ It can also be shown that many people systematically ignore the importance of base rates in their probabilistic reasoning, that many find the principle of regression to the mean to be highly counterintuitive, that many judge the probability of certain sequences of events to be higher than the probability of components in the sequence, and so forth.⁶ In each of these cases, and in many more that might be cited, it is very likely that, for some people at least, the principles that capture their inferential practice would pass the reflective equilibrium test. If this is right, it indicates there is something very wrong with the Goodmanian analysis of justification. For on that analysis, to be justified is to pass the reflective equilibrium test. But few of us are prepared to say that if the gambler's fallacy is in reflective equilibrium for a person, then his inferences that accord with that principle are justified.

Of course, each example of the infelicitous inferential principle that allegedly would pass the reflective equilibrium test is open to challenge. Whether or not the dubious principles that appear to guide many people's inferential practice would stand up to the reflective scrutiny Goodman's test demands is an empirical question. And for any given rule, a Goodmanian might protest that the empirical case has just not been made adequately. I am inclined to think that the Goodmanian who builds his defenses here is bound to be routed by a growing onslaught of empirical findings. But the issue need not turn on whether this empirical hunch is correct. For even the possibility that the facts will turn out as I suspect they will poses a serious problem for the Goodmanian story. It is surely not an a priori fact that strange inferential principles will always fail the reflective equilibrium test for all subjects. And if it is granted, as surely it must be, that the gambler's fallacy (or any of the other inferential oddities that have attracted the attention of psychologists in recent years) could possibly pass the reflective equilibrium test for some group of subjects, this is enough to cast doubt on the view that reflective equilibrium is constitutive of justification as that notion is ordinarily used. For surely we are not at all inclined to say that a person is justified in using any inferential principle—no matter how bizarre it may be—

simply because it accords with his reflective inferential practice.

Faced with this argument the friends of reflective equilibrium may offer a variety of responses. The one I have the hardest time understanding is simply to dig in one's heels and insist that if the gambler's fallacy (or some other curious principle) is in reflective equilibrium for a given person or group, then that principle is indeed justified for them. Although I have heard people advocate this line in conversation, I know of no one who has been bold enough to urge the view in print. Since no one else seems willing to take the view seriously, neither will I.

A very different sort of response is to urge that the notion of reflective equilibrium is itself in need of patching—that some bells and whistles must be added to the justificatory process Goodman describes. One idea along these lines is to shift from narrow Goodmanian reflective equilibrium to some analog of Rawls's "wide reflective equilibrium" (Rawls 1974). Roughly, the idea here is to broaden the scope of the judgments and convictions that are to be brought into coherence with one another. Instead of attending only to our assessments of inferential principles, wide reflective equilibrium also requires that our system of inferential rules is to cohere with our semantic, epistemological, metaphysical, or psychological views. Just how various philosophical or psychological convictions are supposed to constrain a person's inferential principles and practice has not been spelled out in much detail, though Norman Daniels, whose papers on wide reflective equilibrium are among the best around, gives us a hint when he suggests, by way of example, that Dummett's views on logic are constrained by his semantic views (Daniels 1979, 1980a, 1980b). It would also be plausible to suppose that the classical intuitionists in logic rejected certain inferential principles on epistemological grounds.

A rather different way of attempting to preserve a reflective equilibrium account of justification is to restrict the class of people whose reflective equilibrium is to count in assessing the justification of inferential principles. For example, Nisbett and I proposed that in saying an inferential principle is justified, what we are saying is that it would pass the (narrow) reflective equilibrium test for those people whom we regard as experts in the relevant inferential domain (Stich and Nisbett 1980).

A dubious virtue of both the wide reflective equilibrium and the expert reflective equilibrium accounts is that they make clear-cut counterexamples harder to generate. That is, they make it harder to produce actual examples of inferential rules, which the analysis counts as justified and intuition does not. In the case of wide reflective equilibrium, counterexamples are hard to come by just because it is so hard to show that anything is in wide reflective equilibrium for anyone. ("Would she continue to accept that rule if she thought through her epistemological and metaphysical views and came to some stable equilibrium view?" Well, God knows.) In the case of the expert reflective equilibrium account, the dubious but reflectively self-endorsed inferential practice of the experimental subject or the Las Vegas sucker just do not count as counterexamples, since these people do not count as experts.

But though clear-cut cases involving actual people may be harder to find, each of these elaborations of the reflective equilibrium story falls victim to the argument from possible cases offered earlier. Consider wide reflective equilibrium

first. No matter how the details of the wide reflective equilibrium test are spelled out, it is surely not going to turn out to be impossible for a person to reach wide reflective equilibrium on a set of principles and convictions that includes some quite daffy inferential rule. Indeed, one suspects that by allowing people's philosophical convictions to play a role in filtering their inferential principles, one is inviting such daffy principles, since many people are deeply attached to outlandish philosophical views. The expert reflective equilibrium move fares no better. For unless experts are picked out in a question-begging way (e.g., those people whose inferential practices are in fact justified), it seems entirely possible for the expert community, under the influence of ideology, recreational chemistry, or evil demons, to end up endorsing some quite nutty set of rules.⁷

A "Neo-Goodmanian" Project

At this point, if the friend of reflective equilibrium is as impressed by these arguments as I think he should be, he might head off to his study to work on some further variations on the reflective equilibrium theme that will do better at capturing our concept of justification. Despite a string of failures, he might be encouraged to pursue this project by a line of thought that runs something like the following. I will call it the *neo-Goodmanian* line.

It can hardly be denied that we do *something* to assess whether or not an inferential practice is justified. Our decisions on these matters are certainly not made at random. Moreover, if there is some established procedure that we invoke in assessing justification, then it must surely be possible to describe this procedure. When we have succeeded at this we will have an account of what it is for an inferential practice to be justified. For, as Goodman has urged, to be justified just *is* to pass the tests we invoke in assessing an inferential practice. Our procedures for assessing an inferential practice are constitutive of justification. Granted, neither Goodman's narrow reflective equilibrium story nor the more elaborate stories told by others has succeeded in capturing the procedure we actually use in assessing justification. But that just shows we must work harder. The rewards promise to repay our efforts, since once we have succeeded in describing our assessment procedure, we will have taken a giant step forward in epistemology. We will have explained what it is for a cognitive process to be justified. In so doing we will have at least begun to resolve the problem posed by cognitive diversity. For once we have a clear specification of what justification amounts to, we can go on to ask whether our own cognitive processes are justified or whether, perhaps, those of some other culture come closer to the mark.

There is no doubt that this neo-Goodmanian line can be very appealing. I was myself under its sway for some years. However, I am now persuaded that the research program it proposes for epistemology is a thoroughly wrong-headed one. In the pages that follow I will try to say why. My case against the neo-

Goodmanian project divides into two parts. First I shall raise some objections that are targeted more or less specifically on the details of the neo-Goodmanian program. Central to each of these objections is the fact that the neo-Goodmanian is helping himself to a healthy serving of empirical assumptions about the conceptual structures underlying our commonsense judgments of cognitive assessment, and each of these assumptions stands in some serious risk of turning out to be false. If one or more of them is false, then the project loses much of its initial attractiveness. In the following section I will set out a brief catalog of these dubious assumptions. The second part of my critique is much more general and I will be after much bigger game. What I propose to argue is that neither the neo-Goodmanian program nor any alternative program that proposes to analyze or explicate our presystematic notion of justification will be of any help at all in resolving the problem posed by cognitive diversity. But here I am getting ahead of myself. Let me get back to the neo-Goodmanian and his dubious empirical presuppositions.

Some Questionable Presuppositions of the Neo-Goodmanian Project

Let me begin with a fairly obvious point. The neo-Goodmanian, as I have portrayed him, retains his allegiance to the idea of reflective equilibrium. We last saw him heading back to his study to seek a more adequate elaboration of this notion. But nothing the neo-Goodmanian has said encourages us to expect that reflective equilibrium or anything much like it plays a role in our procedure for assessing the justification of a cognitive process. So even if it is granted that we have good reason to work hard at characterizing our justification-assessing procedure, we may find that the notion of reflective equilibrium is simply a non-starter. Confronted with this objection, I think the only move open to the neo-Goodmanian is to grant the point and concede that in trying to patch the notion of reflective equilibrium he is simply playing a hunch. Perhaps it will turn out that something like reflective equilibrium plays a central role in our assessments of justification. But until we have an accurate characterization of the assessment process, there can be no guarantees.

Two further assumptions of the neo-Goodmanian program are that we ordinarily invoke only *one* notion of justification for inferential processes, and that this is a *coherent* notion for which a set of necessary and sufficient conditions can be given. But once again these are not matters that can be known in advance. It might be that different people mean different things when they call a cognitive process 'justified' because there are different notions of justification in circulation. These different meanings might cluster around a central core. But then again they might not. There are lots of normatively loaded terms that seem to be used in very different ways by different individuals or groups in society. I would not be at all surprised to learn that what I mean by terms like 'morally right' and 'freedom' is very different from what the followers of the Rev. Falwell or admirers of Col. Khadafi mean. And I would not be much more surprised if terms of epistemic evaluation turned out to manifest similar interpersonal ambiguities.

Even discounting the possibility of systematic interpersonal differences, it might be that in assessing the justification of a cognitive process we use different procedures on different occasions, and that these procedures have different outcomes. Perhaps, for example, our intuitive notion of justification is tied to a number of prototypical exemplars, and that in deciding new cases we focus in some context sensitive way on one or another of these exemplars, making our decision about justification on the basis of how similar the case at hand is to the exemplar on which we are focusing. This is hardly a fanciful idea, since recent work on the psychological mechanisms underlying categorization suggests that in *lots* of cases our judgment works in just this way.⁸ If it turns out that our judgments about the justification of cognitive processes are prototype- or exemplar-based, then it will be a mistake to look for a property or characteristic that all justified cognitive processes have. It will not be the case that there is any single test passed by all the cognitive processes we judge to be justified. I am partial to a reading of the later Wittgenstein on which this is just what he would urge about our commonsense notion of justification, and I am inclined to suspect that this Wittgensteinian story is right. But I do not pretend to have enough evidence to make a convincing case. For present purposes it will have to suffice to note that this *might* be how our commonsense concept of justification works. If it is, then the neo-Goodmanian program is in for some rough sledding.

A final difficulty with the neo-Goodmanian program is that it assumes, without any evidence, that the test or procedure we use for assessing the justification of cognitive processes exhausts our concept of inferential justification, and thus that we will have characterized the concept when we have described the test. But this is hardly a claim that can be assumed without argument. It might be the case that our procrustean concept of justification is an amalgam composed in part of folk epistemological theory specifying certain properties or characteristics that are essential to justification, and in part of a test or cluster of tests that folk wisdom holds to be indicative of those properties. Moreover, the tests proposed might not always (or ever) be reliable indicators of the properties.⁹ I do not have any compelling reason to believe that our commonsense notion of justification will turn out like this. But I would not be much surprised. Though our understanding of the mechanisms underlying commonsense concepts and judgments is still *very* primitive, as I read the literature it points to two important morals. First, the mental representation of concepts is likely to turn out to be a very messy business. Second, it is no easy job to separate commonsense concepts from the folk theories in which they are enmeshed. All of this bodes ill for the neo-Goodmanian who hopes that the analysis or explication of our concept of justification will yield some relatively straightforward elaboration of the reflective equilibrium test.

Against Analytic Epistemology

The problems posed in the previous section shared a pair of properties. They all turned on empirical assumptions about the nature of our ordinary concept of justification, and they were all targeted fairly specifically at the neo-Goodmanian

project.¹⁰ In the current section I want to set out a very different sort of argument, an argument, which if successful, will undermine not only reflective equilibrium theories but also the whole family of epistemological theories to which they belong.

To give some idea of the range of theories that are in the intended scope of my critique, it will be helpful to sketch a bit of the framework for epistemological theorizing suggested by Alvin Goldman in his book, *Epistemology and Cognition* (Goldman 1986). Goldman notes that one of the major projects of both classical and contemporary epistemology has been to develop a theory of epistemic justification. The ultimate job of such a theory is to say which cognitive states are epistemically justified and which are not. Thus, a fundamental step in constructing a theory of justification will be to articulate a system of rules evaluating the justificatory status of beliefs and other cognitive states. These rules (Goldman calls them *justificational rules* or *J-rules*) will specify permissible ways in which a cognitive agent may go about the business of forming or updating his cognitive states. They "permit or prohibit beliefs, directly or indirectly, as a function of some states, relations, or processes of the cognizer" (Goldman 1986: 60).¹¹

Of course, different theorists may have different views on which beliefs are justified or which cognitive processes yield justified beliefs, and thus, they may urge different and incompatible sets of J-rules. It may be that there is more than one right system of justificational rules, but it is surely not the case that all systems are correct. So in order to decide whether a proposed system of J-rules is right, we must appeal to a higher criterion, which Goldman calls a "criterion of rightness." This criterion will specify a "set of conditions that are necessary and sufficient for a set of J-rules to be right" (Goldman 1986: 64).

But now the theoretical disputes emerge at a higher level, for different theorists have suggested very different criteria of rightness. Indeed, as Goldman notes, an illuminating taxonomy of epistemological theories can be generated by classifying theories or theorists on the basis of the sort of criterion of rightness they endorse. Coherence theories, for example, take the rightness of a system of J-rules to turn on whether conformity with the rules would lead to a coherent set of beliefs. Truth linked or reliability theories take the rightness of a set of J-rules to turn in one way or another on the truth of the set of beliefs that would result from conformity with the rules. Reflective equilibrium theories judge J-rules by how well they do on their favored version of the reflective equilibrium test. And so on. How are we to go about deciding among these various criteria of rightness? Or, to ask an even more basic question, just what does the correctness of a criterion of rightness come to; what makes a criterion right or wrong? On this point Goldman is not as explicit as one might wish. However, much of what he says suggests that, on his view, *conceptual analysis* or *conceptual explication* is the proper way to decide among competing criteria of rightness. The correct criterion of rightness is the one that comports with the conception of justifiedness that is "embraced by everyday thought or language" (Goldman 1986: 58). To test a criterion we explore the judgments it would entail about specific cases, and we test these judgments against our "pretheoretic intuition." "A criterion is supported to the extent that implied judgments accord with such intuitions, and

weakened to the extent that they do not" (Goldman 1986: 66). Goldman is careful to note that there may be a certain amount of vagueness in our commonsense notion of justifiedness, and thus there may be no unique best criterion of rightness. But despite the vagueness, "there seems to be a common core idea of justifiedness" embedded in everyday thought and language, and it is this common core idea that Goldman tells us he is trying to capture in his own epistemological theorizing (Goldman 1986: 58-59).

The view I am attributing to Goldman on what it is for a criterion of rightness to itself be right is hardly an idiosyncratic or unfamiliar one. We saw earlier that a very natural reading of Goodman would have him offering the reflective equilibrium story as an explication or conceptual analysis of the ordinary notion of justification. And many other philosophers have explicitly or implicitly adopted much the same view. I propose to use the term *analytic epistemology* to denote any epistemological project that takes the choice between competing justificational rules or competing criteria of rightness to turn on conceptual or linguistic analysis. There can be little doubt that a very substantial fraction of the epistemological writing published in English in the last quarter of a century has been analytic epistemology.¹² However, it is my contention that if an analytic epistemological theory is taken to be part of the serious normative inquiry whose goal is to tell people which cognitive processes are good ones, or which ones they should use, then for most people it will prove to be an irrelevant failure.

I think the most intuitive way to see this point is to begin by recalling how the specter of culturally based cognitive diversity lends a certain urgency to the question of which cognitive processes we should use. If patterns of inference are acquired from the surrounding culture, much as language or fashions or manners are, and if we can learn to use cognitive processes quite different from the ones we have inherited from our culture, then the question of whether our culturally inherited cognitive processes are good ones is of more than theoretical interest. If we *can* go about the business of cognition differently, and if others actually *do*, it is natural to ask whether there is any reason why we should continue to do it our way. Even if we cannot change our cognitive processes once we have acquired them, it is natural to wonder whether those processes are good ones. Moreover, for many people the absence of a convincing affirmative answer can be seriously disquieting. For if we cannot say why our cognitive processes are any better than those prevailing elsewhere, it suggests that it is ultimately no more than an historical accident that we use the cognitive processes we do, or that we hold the beliefs that those processes generate, just as it is an historical accident that we speak English rather than Spanish and wear trousers rather than togas.

Consider now how the analytic epistemologist would address the problem that cognitive diversity presents. To determine whether our cognitive processes are good ones, he would urge, we must first *analyze* our concept of justification (or perhaps some other commonsense epistemic notion like rationality). If our commonsense epistemic notion is not too vague or ambiguous, the analysis will give us a criterion of rightness for J-rules (or perhaps a cluster of closely related criteria). Our next step is to investigate which sets of J-rules fit the criterion.

Having made some progress there, we can take a look at our own cognitive processes and ask whether they do in fact accord with some right set of J-rules. If they do, we have found a reason to continue using those processes; we have shown that they are good ones because the beliefs they lead to are justified. If it turns out that our cognitive processes do not accord with a right set of J-rules, we can try to discover some alternative processes that do a better job, and set about training ourselves to use them.

It is my contention that something has gone very wrong here. For the analytic epistemologist's effort is designed to determine whether our cognitive states and processes accord with our commonsense notion of justification (or some other commonsense concept of epistemic evaluation). Yet surely the evaluative epistemic concepts embedded in everyday thought and language are every bit as likely as the cognitive processes they evaluate to be culturally acquired and to vary from culture to culture.¹³ Moreover, the analytic epistemologist offers us no reason whatever to think that the notions of evaluation prevailing in our own language and culture are any better than the alternative evaluative notions that might or do prevail in other cultures. But in the absence of any reason to think that the locally prevailing notions of epistemic evaluation are superior to the alternatives, why should we care one whit whether the cognitive processes we use are sanctioned by those evaluative concepts? How can the fact that our cognitive processes are approved by the evaluative notions embraced in our culture alleviate the worry that our cognitive processes are no better than those of exotic folk, if we have no reason to believe that our evaluative notions are any better than alternative evaluative notions?

To put the point a bit more vividly, imagine that we have located some exotic culture that does in fact exploit cognitive processes very different from our own, and that the notions of epistemic evaluation embedded in their language also differ from ours. Suppose further that the cognitive processes prevailing in that culture accord quite well with *their* evaluative notions, while the cognitive processes prevailing in our culture accord quite well with ours. Would any of this be of any help at all in deciding which cognitive processes we should use? Without some reason to think that one set of evaluative notions was preferable to the other, it seems clear that it would be of no help at all.

In the philosophical literature there is a tradition, perhaps traceable to Wittgenstein, that would reject the suggestion that our evaluative notions should themselves be evaluated. Justifications, this tradition insists, must come to an end. And once we have shown that our practice accords with our evaluative concepts, there is nothing more to show. Our language game (or form of life) does not provide us with any way to go about evaluating our evaluative notions. There is no logical space in which questions like "Should we hold justified beliefs?" or "Should we invoke rational cognitive processes?" can be asked seriously. If a person did not recognize that the answers to these questions had to be affirmative, it would simply indicate that he did not understand the logical grammar of words like 'should' and 'justified' and 'rational'.

I am inclined to think that there is at least a kernel of truth in this "Wittgensteinian" stand. Justifications do ultimately come to an end. However, it is, I think, a disastrous mistake to think that they come to an end here. For

there are lots of values that are both widely shared and directly relevant to our cognitive lives, though they are quite distinct from the “epistemic values” that lie behind our ordinary use of terms like ‘justified’ and ‘rational’. It is against the background of these nonepistemic values that our socially shared system of epistemic evaluation can itself be evaluated. Thus, for example, many people attach high value to cognitive states that foster happiness (their own or everyone’s), and many people value cognitive states that afford them the power to predict and control nature. Some people share Mother Nature’s concern that our cognitive lives should foster reproductive success. And, on a rather different dimension, many people care deeply that their beliefs be true.¹⁴ Each of these values, along with many others that might be mentioned, affords a perspective from which epistemic values like justification and rationality can be evaluated. We can ask whether the cognitive states and processes endorsed by our notions of epistemic value foster happiness, or power, or accurate prediction, or reproductive success, or truth. More interestingly, we can ask whether the cognitive states and processes we actually have or use foster happiness, power, or the rest. And if they do not, we can explore alternatives that may do a better job, though there is of course no guarantee that all of these values can be maximized together.¹⁵

At this point, it might be protested that the values I am proposing to use in evaluating our socially shared notions of epistemic evaluation are themselves lacking any deeper justification. If someone can accept *these* as ultimate values, why could not someone do the same for justification or rationality? My reply is that of course someone could, but this is no objection to the view I am urging. There are many things that people might and do find ultimately or intrinsically valuable. Some of these values might be rooted more or less directly in our biological nature, and these we can expect to be widely shared. Other values, including intrinsic, life-shaping values, might be socially transmitted, and vary from society to society. Still others might be quite idiosyncratic. It is entirely possible for someone in our society to attach enormous value to having justified beliefs or to using rational inferential strategies—that is, to having beliefs or inferential processes that fall within the extension of ‘justified’ or ‘rational’ as they are used in our language. Similarly, it is entirely possible for someone in another society to attach enormous value to having cognitive states that fall within the extension of the terms of cognitive evaluation current in that society. In each case the evaluation may be either instrumental or intrinsic. A person in our culture may value the states and processes that fall within the extension of ‘rational’ or ‘justified’ because he thinks they are likely to be true, to lead to happiness, and so forth, or he may value them for no further reason at all. And a person in another culture may have either sort of attitude in valuing what falls within the extension of his language’s terms of cognitive evaluation. Where the value attached is instrumental, there is plenty of room for productive inquiry and dialogue. We can try to find out whether rational or justified cognitive processes do lead to happiness or power or truth, and if they do we can try to understand why. But where the value accorded to one or another epistemic virtue is intrinsic, there is little room for debate. If you value rationality for its own sake, and the native of another culture values some rather different cognitive characteristic (“shmashinality” as Hilary Putnam might put it) for its own sake, there is not

much you can say to each other. Moreover, there is not much I can say to either of you, since on my view the fact that a cognitive process is sanctioned by the venerable standards embedded in our language of epistemic evaluation, or theirs, is of no more interest than the fact that it is sanctioned by the venerable standards of a religious tradition or an ancient text—unless, of course, it can be shown that those standards correlate with something more generally valued.¹⁶ But I do not pretend to have any arguments that will move the true epistemic xenophobe. If a person really does attach deep intrinsic value to the epistemic virtues favored by folk epistemology, then dialogue has come to an end.

Finally, let me say how all of this relates to analytic epistemology. The analytic epistemologist proposes to arbitrate between competing criteria of rightness by seeing which one accords best with the evaluative notions “embraced by everyday thought and language.” However, it is my contention that this project is of no help whatever in confronting the problem of cognitive diversity unless one is an epistemic xenophobe. The program of analytic epistemology views conceptual analysis or explication as a stopping place in disputes about how we should go about the business of cognition. When we know that a certain cognitive process falls within the extension of our ordinary terms of epistemic evaluation—whatever the analysis of those terms may turn out to be—we know all that can be known that is relevant to the questions of how we should go about the business of reasoning. But as I see it, the only people who should take this information to be at all relevant to the question are the profoundly conservative people who find intrinsic value in having their cognitive processes sanctioned by culturally inherited standards, whatever those standards may be. Many of us care very much whether our cognitive processes lead to beliefs that are true, or give us power over nature, or lead to happiness. But only those with a deep and free-floating conservatism in matters epistemic will care whether their cognitive processes are sanctioned by the evaluative standards that happen to be woven into our language.

Notes

A previous version of this chapter appeared in *Synthese* 74 (1988): 391-413. Reprinted with kind permission from Kluwer Academic Publishers and Stephen Stich.

This paper has been evolving for a long time. Earlier versions were presented in my seminars at the University of Sydney, the University of Maryland, and the University of California, San Diego, and in colloquia at the University of Adelaide, La Trobe University, the Australian National University, the University of Illinois at Chicago, the University of Vermont, Tulane University, the University of Southern California and the University of Colorado. Suggestions and criticism from these varied audiences have led to more changes than I can remember or acknowledge. My thanks to all who helped, or tried. Special thanks are due to Philip Kitcher, David Stove, and Joseph Tolliver.

1. Nor are these the only alternatives. There are lots of characteristics that are innate (not part of our cultural heritage) though they differ substantially from one group to another. Sex, hair color, and blood type are three obvious examples.

2. Actually, the issue is not so straightforward if we compare languages at very different stages of development, or languages involving different theoretical assumptions. It is only when the choice is between languages that are more or less inter-translatable with our own that we are inclined to judge that one is as good as another. Thanks to Paul Churchland for reminding me of this point.

3. L. J. Cohen (1981) seems to read Goodman this way since he exploits Goodman's notion of reflective equilibrium in giving an account of good reasoning.

4. Well, I will argue it a little. Note first that according to Goodman the only justification needed for either rules or inferences "lies in" the agreement achieved by the reflective equilibrium process. This talk of justification *lying in* the agreement strongly suggests the constitutive reading. Moreover, on the nonconstitutive reading, Goodman's doctrine would be an oddly incomplete one. It would present us with a test for justification without telling us why it was a test or giving us any account of what it is that is being tested for. On the constitutive reading, by contrast, no such problem arises. We have in one tidy package both an analysis of the notion of justification and an unproblematic explanation of the relation between justification and the process Goodman describes.

5. The writer was Henry Coppee (1874). Here is a brief quote:

Thus, in throwing dice, we cannot be sure that any single face or combination of faces will appear; but if, in very many throws, some particular face has not appeared, the chances of its coming up are stronger and stronger, until they approach very near to certainty. It must come; and as each throw is made and it fails to appear, the certainty of its coming draws nearer and nearer. (162)

6. For an excellent survey of the literature in this area see Nisbett and Ross (1980); a number of important studies are collected in Kahneman, Slovic, and Tversky (1982).

7. As Conee and Feldman (1983) point out, the situation is actually a bit worse for the version of the expert reflective equilibrium analysis that Nisbett and I offered. On that account, different groups may recognize different people as experts. And it is surely at least possible for a group of people to accept as an expert some guru who is as bonkers as he is charismatic. But we certainly do not want to say that the followers of such a guru would be rational to invoke whatever wild inferential principle might be in reflective equilibrium for their leader.

8. For a good review of the literature, see Smith and Medin (1981).

9. For some insightful observations on the potential complexity of common-sense concepts and the ways in which intuitive tests can fail to capture the extension of concepts, see Rey (1983).

10. Actually, the last three of my four objections might, with a bit of reworking, be generalized so as to apply to all of analytic epistemology, as it is defined below. But I do not propose to pursue them since, as we shall see, analytic epistemology has more pressing problems.

11. For the reader who wants a more hands-on feel for Goldman's notion of a J-rule, the quote continues as follows:

For example, J-rules might permit a cognizer to form a given belief because of some appropriate antecedent or current state. Thus, someone being 'appeared to' in a certain way at *t* might be permitted to believe *p* at *t*. But someone else not in such a state would not be so permitted. Alternatively, the rules might focus on mental operations. Thus, if *S*'s believing *p* at *t* is the result of a certain operation, or sequence of operations, then his belief is justified if the system of J-rules permits that operation or sequence of operations.

12. For an extended review of part of this literature see Shope (1983). As Shope notes, relatively few of the philosophers who have tried their hands at constructing an "analysis" of knowledge (or of some other epistemic notion) have been explicit about their objectives (see pp. 34-44). However, absent indications to the contrary, I am inclined to think that if a philosophical project proceeds by offering definitions or "truth conditions," and testing them against our intuitions about real or imaginary cases, then the project should be viewed as an attempt at conceptual analysis or explication. Unless one has some pretty strange views about intuitions, it is hard to see what we could hope to gain from capturing them apart from some insight into the concepts that underlie them.

13. Evidence on this point, like evidence about crosscultural differences in cognitive processes, is hard to come by and hard to interpret. But there are some intriguing hints in the literature. Hallen and Sodipo (1986) studied the terms of epistemic evaluation exploited by the Yoruba, a west African people. It is their contention that the Yoruba do not have a distinction corresponding to our distinction between knowledge and (mere) true belief. They do, however, divide beliefs into two other categories: those for which a person has immediate, eyewitness evidence, and those for which he does not. In the standard Yoruba-English dictionaries, the Yoruba term for the former sort of belief, 'mo', is translated as 'knowledge' while the term for the latter sort, 'gbagbo', is translated as 'belief'. However, Hallen and Sodipo argue that these translations are mistaken, since 'mo' has a much narrower extension than 'knowledge'. Most of what we would classify as scientific knowledge, for example, would not count as 'mo' for the Yoruba, because it is based on inference and second-hand report. Since the Yoruba do not draw the distinction between knowledge and (mere) true belief, they have no use for our notion of epistemic justification, which earns its keep in helping to draw that distinction. Instead, the Yoruba presumably have another notion that they exploit in distinguishing 'mo' from 'gbagbo'. Hallen and Sodipo do not indicate whether the Yoruba have a single word for this notion, but if they do, it would be a mistake to translate the word as "(epistemic) justification". Clearly, if Hallen and Sodipo are right, the Yoruba categories of epistemic evaluation are significantly different from our own.

14. I should note, in passing, that I think it is a mistake to include truth on the list of intrinsically valuable features of one's cognitive life. But that is a topic for another paper (see Stich in preparation), and I will ignore the point here.

15. The point I am making here is really just a generalization of a point made long ago by Salmon (1957), Skyrms (1975), and a number of other authors. Strawson (1952) argued that the rationality or reasonableness of inductive reasoning was easy to demonstrate, since being supported by inductive inference is part of what we *mean* when we say that an empirical belief is *reasonable*. To which Salmon replied that if Strawson is right about the meaning of 'reasonable' it is not at all clear why anyone should want to be reasonable. What most of us do care about, Salmon notes, is that our inferential methods are those that are "best suited to the attainment of our ends" (Salmon 1957: 41). "If we regard beliefs as reasonable simply because they are arrived at inductively and we hold that reasonable beliefs are valuable for their own sake, it appears that we have elevated inductive method to the place of an intrinsic good" (Salmon 1957: 42). The analytic epistemologist elevates being within the extension of our ordinary terms of epistemic evaluation to the place of an intrinsic good. In so doing, the analytic epistemologist embraces a system of value that few of us are willing to share.

16. Let me try to head off a possible misunderstanding. Some analytic epistemologists claim that our ordinary notions of epistemic evaluation are conceptually linked to truth. On Goldman's account, for example, the rightness of a set of J-rules is

a function of how well the processes sanctioned by those rules do at producing truths. If this is right, then a person who attached intrinsic value to having true beliefs would, of course, have reason to be interested in whether his cognitive states and processes were sanctioned by the standards embedded in our language. But here it is the appeal to truth that is doing the work, not the appeal to traditional standards. For if Goldman is wrong in his conceptual analysis and '(epistemic) justification' is not conceptually tied to truth, the person who values truth will stay just as interested in whether his cognitive processes reliably lead to truth, though he may have no interest whatever in how traditional notions of epistemic evaluation judge his cognitive processes. Thanks to Steven Luper-Foy for the query that prompted this note.

Chapter 7

Reflection on Reflective Equilibrium

Robert Cummins

Reflective Equilibrium and Scientific Method

As a procedure, reflective equilibrium (RE) is simply a familiar kind of standard scientific method with a new name. (For descriptions of reflective equilibrium, see Daniels 1979, 1980b, 1984; Goodman 1965; Rawls 1971.) A theory is constructed to account for a set of observations. Recalcitrant data may be rejected as noise or explained away as the effects of interference of some sort. Recalcitrant data that cannot be plausibly dismissed force emendations in theory. What counts as a plausible dismissal depends, among other things, on the going theory, as well as on background theory and on knowledge that may be relevant to understanding the experimental design that is generating the observations, including knowledge of the apparatus and observation conditions. This sort of mutual adjustment between theory and data is a familiar feature of scientific practice. Whatever authority RE seems to have comes, I think, from a tacit or explicit recognition that it has the same form as this familiar sort of scientific inference.

One way to see the rationale underlying this procedure in science is to focus on prediction. Think of prediction as a matter of projecting what is known onto uncharted territory. To do this, you need a vehicle—a theory—that captures some invariant or pattern in what is known so that you can project it onto the unknown. How convincing the projection is depends on two factors: (i) how sure one is of the observational base, and (ii) how sure one is that the theory gets the invariants right. The two factors are not independent, of course. One's confidence in the observational base will be affected by how persuasively the theory identifies and dismisses noise; one's confidence in the theory, on the other hand, will depend on one's confidence in the observations it takes seriously. Prediction is important as a test of theory precisely because verified predictions seem to show that the theory has correctly captured the general in the particular, that it has got the drift of the observational evidence in which our confidence is ultimately grounded. Falsified prediction seems to show that it has not. We are justified in accepting a theory to the extent that we are justified in thinking it properly transfers our confidence concerning observed cases to those that have not been observed. Theory is certainly more than a vehicle for sophisticated inductive inference, but it needs to be at least that if it is to count as more than mere speculation.