Philosophy and Phenomenological Research
Vol. L I, No. 1, March 1991

Evaluating Cognitive Strategies: A Reply to Cohen, Goldman, Harman, and Lycan

STEPHEN P. STICH
Rutgers University

I. Reply to Cohen

There is one point in Cohen's paper that I can cheerfully endorse, and a second that I find completely baffling. In a third passage Cohen seems to be hinting at a deep and difficult question that I have no idea how to answer. I'll take these up in that order.

The point I endorse is Cohen's insistence that it is not *easy* to show that experimental subjects are reasoning fallaciously. We must always be sensitive to the fact that the subjects may not understand the question in the way the experimenter intends, or that they may be importing additional premises. Also, as Cohen has noted elsewhere, it is often difficult to determine which aspects of a subject's performance should be attributed to his underlying inferential competence, and which to various other factors that interact with competence to produce performance. Motivation, attention, short term memory limitations and a host of other factors may produce faulty reasoning without impugning the subject's underlying "psycho-logic." For the better part of two decades, Cohen has helped to keep the field honest by proposing explanations of experimental findings that are at odds with those favored by most psychologists. The debates these challenges have engendered have been enormously useful; they've also been great fun.

So I'm prepared to concede that in any particular case it might turn out that we were wrong to think some strange or fallacious inferential principle is part of someone's underlying cognitive competence. And if passing the reflective equilibrium test is a reliable indicator that an inferential principle is part of someone's underlying competence, then in any particular case we might be wrong in thinking that the principle in question passes the test. What I find deeply puzzling is Cohen's insistence

that we *must* be wrong—that it couldn't *possibly* be the case that some fallacious principle passes the reflective equilibrium test for anyone. According to Cohen,

> [Every] party to the controversy about rationality—indeed, everyone with an intellectual conscience—is convinced that he or she can in principle, if not always in practice, teach anyone else what the correct inferential rules are in relevant contexts.

Cohen doesn't tell us what he means by "in principle," nor does he say how he would distinguish teaching from other ways of changing people's minds—like drugs or psycho-surgery. But on any interpretation of these notions that does not render the claim completely trivial, I think it is plainly false. Some years ago I said to Cohen, jokingly, that he would be less sanguine about human reasoning if he spent as much time as I do talking to undergraduates. But whatever one thinks about the prospects of getting undergraduates to reason better, it seems quite wild to insist that "in principle" we can "teach" correct inferential rules to every mental patient, every severely retarded child, and every senile inhabitant of a geriatric ward. If this is really what Cohen is urging, then I am simply baffled.

On Cohen's view, there may well be "alternative and equally coherent reflective equilibria" when ethical or political principles are involved, but not when inferential or "cognitive" principles are involved. Unfortunately, he does not spell out his argument for this view in any detail. But he does offer an intriguing hint.

> Different people may value different possible worlds. But there is only one world to be *known*—the actual world—and in that sense there is only one subject-matter for cognitive inquiry.

The suggestion seems to be that pluralism about good reasoning is incompatible with the view that there is only one world to be known. I think we would be well advised to be very cautious here, since the history of philosophy is littered with implausible attempts to wring ontological conclusions from epistemological premises. Still, I suspect Cohen is on to something. In my book I said nothing at all about the metaphysical implications of my epistemological relativism. The reason was not that I don't think there are any, but rather that I didn't know what to say.

II. Reply to Goldman

Goldman characterizes "analytic epistemology" as "the attempt to identify epistemic value by analyzing epistemic terms found in everyday language." He claims that my "critique of analytic epistemology is largely

based on the specter of cultural diversity in epistemic concepts." But this is a mistake. Though I exploit the idea of cultural diversity in an attempt to make the problem vivid, the essential point is not the cultural diversity of epistemic concepts but their *idiosyncrasy*. Whether or not anybody actually uses them, there are *lots* of alternatives to the evaluative notions embedded in our everyday language, and there is nothing special about the ones we use except for the fact that we use them.

In a spirit of ecumenical pluralism, Goldman goes on to suggest that if there really are other notions of epistemic evaluation exploited in other languages, then "we might simply *add* the others to our repertoire." This suggestion can be interpreted in two ways. On one reading, Goldman is proposing that we could learn to *apply* alternative evaluative notions, in much the same way that we can learn to apply the exotic kinship terms exploited in certain other cultures. And on this reading, I have no objection to the proposal. But it is one thing to be able to apply other systems of evaluation, and quite another thing to *care* about the evaluations that the exotic system produces. Perhaps some culture values beliefs in a particular domain to the extent that they are consistent with the views contained in a certain venerated text. With appropriate training I might become quite adept at making these evaluations and applying their evaluative language. But I attach no particular value to beliefs that are highly ranked on that evaluative scheme. If Goldman is suggesting that "we might simply add" all exotic evaluative categories to the list of things that we ourselves value, then I must say that I find the suggestion bizarre.

There is, I suspect, a deep difference in perspective lurking here. Goldman is an analytic epistemologist and has spent much of his career trying to analyze the notions of epistemic evaluation exploited in our language. As I see it, the project of analytic epistemology, and indeed of analytic philosophy more generally, is best viewed as a sort of domestic cognitive anthropology. It is of a piece with the anthropological analysis of the kinship notions, religious beliefs or value systems prevailing in a culture. I certainly have no quarrel with this sort of anthropological inquiry. Indeed, I often find its results to be quite fascinating. But the descriptive questions that preoccupy analytic epistemologists are not the ones that I find most pressing in epistemology.

My interest in epistemology is motivated by normative concerns that are both very pragmatic and very traditional. Like Bacon, Descartes, Hume and Popper, I worry that people often do a rather poor job of reasoning, and I think epistemological theory has a very practical role to play in *improving* this situation. However, I don't think the findings of analytic epistemologists are going to be of much help in deciding how we should

educate our children, or restructure our intellectual institutions, or reform our own cognitive habits. In pursuing these projects, (or, more whimsically, in trying to decide which of the Genie's offers to accept) it does no good at all to be told that one strategy of reasoning is ranked most highly by an evaluative notion found in the Yoruba language, another is ranked highest on an evaluative scale embedded in an Australian aboriginal language, and a third is ranked highest by some evaluative term in our own language. In making practical epistemic decisions, I don't much care what traditional Yoruba epistemic norms say; nor do I care what modern American norms say. What I really want to know is which decision is most likely to lead to states of affairs that I find intrinsically valuable.

Ironically, I think that analytic epistemology does have an important contribution to make in these pragmatic normative inquiries. But it is a subversive contribution. Painstaking analyses by Goldman and other analytic epistemologists have underscored the quirky idiosyncrasy of the notions of epistemic evaluation that we have inherited from our culture. And this, in turn, helps to undermind the conviction that it is desirable or valuable to have cognitive states that are highly ranked by those standards. Gilbert Harman has suggested that this sort of challenge to our traditional epistemic norms is similar to the "deconstructionist" critique of other traditional values and practices. If that's what deconstructionism is, I'm all for it.

III. Reply to Harman

Harman distinguishes the "immanent" notion of truth from the "transcendent" notion. He claims that my arguments about the desirability of true beliefs will not go though on the immanent notion. I'm not convinced. One strand of my argument against the intrinsic desirability of having true beliefs turns on the idiosyncrasy of the transcendent notion of truth. This is the strand I sketched earlier. But there is a second line of argument developed in my book, which turns on the fact that the transcendent notion of truth is applicable only in a very *limited* domain. There are many possible systems of mental storage and processing to which our intuitive interpretation function simply does not apply. If the elements out of which someone's mental "representations" are built are not hooked on to the world by the right sort of causal connections, our intuitive psychosemantics provides no interpretation at all for them. And if the constructions used to build complex "representations" are not close relatives of the constructions we currently exploit, then once again our intuitive psycho-semantics falls silent. (The argument for these claims is set out in Section 5.4 of *The Fragmentation of Reason.*) But, of course, if

our intuitive interpretation function does not apply to some set of mental "representations," then those "representations" have no truth conditions (thus the scare quotes around the word "representations"); and if a mental state has no truth conditions, it is neither true nor false. It follows that valuing true beliefs is a very conservative thing to do. Some systems of mental storage and processing that are significantly different from the one we currently employ may have quite extraordinary pragmatic virtues to recommend them. But the "representations" exploited in those systems have no truth conditions, and thus cannot be true.

Thusfar I have been talking about the transcendent notion of truth. But the same conclusion applies to the immanent notion—with a vengeance! For a person's immanent notion of truth simply does not apply to a belief that invokes a concept the person does not now have. Consider the case of the scientist and the schoolboy. The scientist says, "Hydrogen has a valence of +1." But since the schoolboy does not have the concept of valence, he cannot formulate the immanent truth conditions for the scientist's belief. As Harman notes, the immanent notion of truth "applies only to sentences *within* [one's] own language of thought....[It] has no application to sentences in any other language of thought." If we value having beliefs that are immanently true, we should be rather reluctant to acquiring new concepts, since the immanent notion of truth (remember it is *our current* immanent notion of truth) has no application to the beliefs we might build from those concepts.

One way to think about this is to imagine the Genie offering the schoolboy the following deal: The Genie will modify the boy's mind in such a way that the boy will quickly and easily acquire lots of new scientific concepts, and a healthy serving of new scientific beliefs. These new beliefs will be shared by leading scientists. "No thanks," says the boy. "I value immanent truth, and none of those new beliefs could possibly be immanently true." People who think that having transcendently true beliefs is intrinsically valuable are epistemic conservatives, but those who think that having immanently true beliefs is intrinsically valuable are epistemic reactionaries.

Let me turn, briefly, to a second issue raised by Harman. According to Harman,

> It is easy to see how [Stich's] argument against *caring* about whether one's beliefs are true can be converted into an argument against caring that one's desires come true. (This is left as an exercise for the reader.)

Goldman makes a similar suggestion, and several other people having raised the issue in conversation. I confess that when I was writing the book I wor-

ried about this objection for a long time. Still, I very much wish that Harman had not left the details as an exercise for the reader, since every time I try to fill in those details, my efforts are unsuccessful. Remember that in asking whether we really care whether our beliefs are true, my concern is with the evaluation of belief forming processes, and ultimately with the modification of those processes. My thought experiment invoking the Genie is intended as a way of focusing in on this concern. It is important to keep in mind that what the Genie is offering to change is our own cognitive processes. He is not offering to modify the world in a way that will make more of our beliefs true. Rather, he is offering to modify our belief-forming mechanism in a way that will result in it producing more beliefs that the intuitive interpretation function maps on to truth conditions that actually obtain. Some people think that our belief forming mechanisms would be improved by such a change, though I disagree.

Now let's think about the analogy with desire. The Genie is not offering to change the world in such a way that more of our desires are satisfied. (That's another Genie. He lives in a lamp, and on the traditional version of the story he offers you three wishes.) Our Genie is offering to change our desire-forming mechanisms in such a way that the world (which remains largely unchanged) satisfies more of our desires. It takes a lot of work to persuade people that the Genie's offer to change our belief forming mechanisms is not a particularly attractive one. But when what is on offer is a change in our desire forming mechanisms, it generally takes very little work. What the Genie is pushing is something akin to Stoicism (or perhaps it's Buddhism): If the world isn't the way you want it to be, he'll help you change what you want. For most of us, that's not even tempting.

IV. Reply to Lycan

My reply to Lycan will be the briefest of the four, because I find so little to disagree with in his paper. Lycan agrees that "once we see just what 'truth' is and once it has been compared to some of its close cousins, the notion is of dubious philosophical interest." But he thinks that perhaps we can make a case for truth that parallels the case to be made for continuing to drive on the right. It doesn't really matter which side we drive on, though it does matter that we all do the same thing. Since the notion of truth does a variety of interpersonal-communicative jobs for us, we should stick with it, rather than "gratuitously and expensively converting to some other truth-like notion."

I have two concerns with this strategy for defending truth. First, I'm less convinced than Lycan is that we make much use of the notion of truth in interpersonal communication. Second, I think the analogy is a bit

misleading. On the highway, there are only two viable options—driving on the left or driving on the right. And, as the Swedes learned a decade or two ago, switching from one to the other is difficult and expensive in all sorts of ways. In semantics, by contrast, there are *lots* of alternatives. Thus it may be possible to *tinker* with whatever semantic notion we need for interpersonal communication. If there is a useable semantic notion that has some of the virtues of Lycan's notion of gruth, we might hope to locate it, and end up using it, by making a series of small changes in the notion we use now.

Lycan would not be very sanguine about this senario, I suspect. It's not that he doubts there are alternatives to truth that would be more adaptive (or pragmatically preferable for other reasons). His notion of gruth is an example. But gruth "is not a viable design option." Natural selection could not have "fashioned a cognizer to have consistently grue beliefs, since the vagaries and vicissitudes of people's life histories vary unpredictable and far too widely across the species." Here Lycan is surely right. However, he goes on to speculate that Mother Nature couldn't do "much better than to satisfice by making our belief-forming methods truth-conducing." I find Lycan's epistemic pessimism a bit puzzling, particularly since he agrees that the "causal/historical/functional relations of brain states to external things" that truth requires is quirky and idiosyncratic. Indeed, it is such an odd relation that it is hard to imagine Mother Nature could do a very good job at making our belief forming methods truth-conducing. I suspect that She has made our cognitive mechanisms much more plastic than Lycan imagines, and much less likely to produce true beliefs. But ultimately I must agree with Lycan. It is a brutely empirical question, and neither one of us have much by way of evidence.