

HILARY KORNBLITH

## EPISTEMIC NORMATIVITY\*

**ABSTRACT.** This paper examines the source and content of epistemic norms. In virtue of what is it that epistemic norms have their normative force? A semantic approach to this question, due to Alvin Goldman, is examined and found unacceptable. Instead, accounts seeking to ground epistemic norms in our desires are argued to be most promising. All of these accounts make epistemic norms a variety of hypothetical imperative. It is argued that such an account may be offered, grounding our epistemic norms in desire, which nevertheless makes these imperatives universal. The account is contrasted with some recent work of Stephen Stich.

In 1969, Quine advocated an approach to epistemological questions which he called epistemology naturalized. On Quine's view,

[e]pistemology, or something like it, simply falls into place as a chapter of psychology and hence of natural science. It studies a natural phenomenon, viz., a physical human subject. This human subject is accorded a certain experimentally controlled input – certain patterns of irradiation in assorted frequencies, for instance, and in the fullness of time the subject delivers as output a description of the three-dimensional world and its history. The relation between the meager input and the torrential output is a relation that we are prompted to study for somewhat the same reasons that always prompted epistemology; namely, in order to see how evidence relates to theory, and in what ways one's theory of nature transcends any available evidence. (1969, pp. 82–83)

For many, this approach seemed to involve rejecting the normative dimension of epistemological theorizing, and, in so doing, abdicating at least one central role which epistemology has traditionally played. Now it is true that nowhere in 'Epistemology Naturalized' does Quine specifically say that there is no normative role for epistemological theorizing to play; but passages like the one quoted above surely do encourage this reading. If epistemology is to become nothing more than a chapter of psychology, then, on one straightforward account of what psychology is all about, the resulting discipline will become merely descriptive, and thereby lose all normative force.

Quine has since clarified his account of naturalistic epistemology, and he has repudiated the suggestion that there is no place for normativity within epistemological theorizing.

Naturalization of epistemology does not jettison the normative and settle for the indiscriminate description of ongoing procedures. For me normative epistemology is a branch of engineering. It is the technology of truth-seeking, or, in a more cautiously epistemological term, prediction . . . There is no question here of ultimate value, as in morals; it is a matter of efficacy for an ulterior end, truth or prediction. The normative here, as elsewhere in engineering, becomes descriptive when the terminal parameter is expressed.<sup>1</sup> (1986, pp. 664–65)

This passage is admirably clear on the point that naturalistic epistemology does not abandon normative theorizing. It is, however, far less clear than one might like on the source of this normativity. Once we accept truth as our goal, there are, clearly, some empirical questions to be answered in what Quine calls the “technology of truth-seeking”, and questions of this sort have occupied many who think of themselves as naturalistic epistemologists. But how is it that truth acquires this status as our goal and thereby confers normative force on the recommendations to pursue certain strategies of belief acquisition and retention, namely, those which are conducive to achieving it? Does Quine mean to be making a sociological observation here, that many people do in fact have this goal? Or is there some deeper fact about true belief that somehow recommends it to us? It is especially important to address these questions when many in the naturalistic tradition have, on the one hand, suggested that truth is not the only goal of our epistemic activity (see, e.g., Field 1982; Goldman 1986; Elgin 1988), or, on the other, that the goal of truth should be abandoned in favor of other goals, such as fitness (Lycan 1981; 1988, esp. Ch. 7) or the totality of things we value exclusive of truth (Stich 1990). If we are to be in any position rationally to adjudicate among these competing views, we cannot rest content with Quine’s seemingly innocent suggestion that epistemic norms “become[ ] descriptive when the terminal parameter is expressed”, for we need to know what the source of this terminal parameter is. What, ultimately, is the source of epistemic normativity?

This question about the source of epistemic normativity is not just a question for epistemologists of a naturalistic turn of mind. Insofar as epistemologists endorse epistemic norms, it is incumbent upon them to explain the source of this normativity. How is it that epistemic norms come to have their normative force? What we are looking for here is not necessarily a naturalistically acceptable answer to this question, but any acceptable answer. What we need is an account of the source of epistemic normativity which does not make a mystery of it. A divine

command theory of epistemic norms, for example, seems unpromising. But what are the alternatives? That is what I hope to explore here. The project I set myself is to make clear what the live options are for making sense of epistemic normativity.

I begin by considering a semantic approach to this problem implicit in the work of Alvin Goldman. I argue in Section 1 that this kind of approach cannot provide us with a satisfying solution. In Section 2, I explain how norms might be grounded in desire, and in the sections which follow I pursue a number of different ways of working out this suggestion. All such accounts make epistemic norms a variety of hypothetical imperative: they tell us how we should acquire our beliefs if we meet certain conditions. Accounts of this sort run the risk of being parochial. If the conditions are met by few people, then epistemic norms should be of little interest to most of us. Even if the conditions are met by many, norms construed as hypothetical imperatives seem inevitably to lack the universality which we want our epistemic norms to have.

The account I believe to be most promising, while a species of the hypothetical imperatives approach, nevertheless makes epistemic norms universal. I argue that there are certain substantive constraints which we will want our cognitive systems to meet simply on the condition that we value anything at all. Since everyone values something or other, this condition is universally met; the imperatives to have cognitive systems meeting my constraints thus apply to everyone. Such an account provides, I believe, much of what we want in an account of the source of epistemic normativity. I am not at all convinced that it is possible to give an account of epistemic norms which provides more than this.

## 1.

I want to begin by examining the way in which Alvin Goldman makes room for normativity in his naturalistic epistemology. Although Goldman's account is not the sort with which Quine would be sympathetic, it provides, for many, an extremely attractive way of explaining the source of epistemic norms.

Goldman is rightly famous for urging epistemologists to make more room for empirical concerns in general, and psychological concerns in particular, within their epistemological theorizing. But in *Epistemology and Cognition*, empirical concerns play no role at all in explaining the

source of epistemic normativity. In spite of Goldman's urging that "epistemology should be a multidisciplinary affair, not the province of pure, *a priori* philosophy" (1976, p. 1), Goldman makes room for an important *a priori* component in his epistemology, and it is precisely here that epistemic normativity gets its footing.

In what Goldman conceives of as the foundational part of epistemology, we engage in autonomous inquiry into the meaning of various epistemic terms. We do this by way of testing proposed analyses against our intuitions, and we thereby attempt to capture our ordinary concept of, for example, justification or knowledge. Thus, Goldman says,

I suggest that the meaning of the term 'justified' (in its epistemic use) is fixed by certain things that we *presume* about the world, whether we are right or not. (1986, p. 108)

At this foundational level of inquiry, what the world is actually like is irrelevant, for we are attempting to discover the contours of our ordinary concept of, say, justification, and this concept is constructed against a backdrop of assumptions about the world (see 1986, p. 107). We are not attempting to discover whether this concept is, in some sense to be specified, adequate or accurate; we are merely attempting to discover what the concept is. Conceptual analysis does not require empirical information about the world around us. This is, I believe, a very familiar conception of what at least one part of epistemology is all about. Goldman's break with traditional epistemology in *Epistemology and Cognition*, and the features of that account which make it a variety of naturalism, are not to be found here. Rather, they are to be found in Goldman's insistence that this is not all there is to epistemology, and in his important and distinctive ways of illuminating the contributions which psychology makes to epistemological projects of obvious importance.

Nevertheless, it is in the foundational part of epistemology, in the part which is investigated by examining our concepts, that epistemic normativity is located. The terms of epistemic appraisal are evaluative terms: to say that a belief is justified, or that it is a case of knowledge, is to say that it is good in some sense.

[E]pistemology is an evaluative, normative, or critical discipline. Let me now address the possible scope and nature of such evaluation. First, what do we mean by 'evaluation,' or 'norm'? We mean a judgment that pronounces something good or bad, right or wrong, proper or improper, and the like. (1986, p. 20)

Moreover, we discover through conceptual analysis what conditions must be satisfied if a belief is to be justified, or a case of knowledge, or whatever. It is thus a matter of conceptual analysis both that the class of justified beliefs are good, right or proper, and that they are identified by certain empirical conditions (roughly, being produced by processes which are truth-conducive). Suppose then that someone were to ask what makes truth-conduciveness a good thing: What is good, this person says, about justified belief so understood? Goldman seems committed to the following answer. To be justified is simply to be something good; this is a matter of the meaning of the term. Similarly, and to a first approximation, for a belief to be justified just is for it to be the product of a truth-conducive process; this, too, is a matter of the meaning of the term. Thus, it is merely a matter of the meaning of the term 'justified' that truth-conduciveness is a good thing. A similar argument may, of course, be constructed for other epistemic terms of approval. Normative force seems to derive from semantic considerations alone.

Now there is something terribly unsatisfying about this. Imagine someone, such as Stephen Stich, who claims that conduciveness to truth is of no value at all; that a belief which is the product of a truth-conducive process is not, in virtue of that very fact, a good belief to have. On Stich's view, it is not merely that truth by itself does not make a proposition worthy of belief; rather, the truth of a belief, or the truth-conduciveness of the process which produced it, does not even count in its favor. Now if Stich were to confront Goldman and ask why he should favor true beliefs, or beliefs produced by truth-conducive processes, it seems that Goldman should offer him the semantic argument above. It also seems that this should not convince Stich at all.

Indeed, this is exactly what Stich argues. Stich is concerned that even if Goldman has given the proper conceptual analysis of our epistemic terms, the fact that he has uncovered the proper conceptual analysis carries no normative force. In particular, as Stich points out, suppose there were another culture whose epistemic terms embodied a different set of standards. Suppose that in this other culture beliefs meeting different conditions were approved of. What reason is there to approve of beliefs meeting our conditions rather than those meeting the conditions of this other culture? It will hardly do to point out that this is what our terms mean. As Stich puts it,

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 for most of us it would be of no help at all. (1990, pp. 92–93)

What Goldman says on this score only serves to underline the difficulty of the position he is in. Goldman offers a reliability account of justification and knowledge, and there are a number of different ways to explicate the notion of reliability. Here is what Goldman says about his preferred account, the normal worlds account:

My proposal to judge reliability by reference to normal worlds is made in the spirit of trying to elucidate the *ordinary* conception of justifiedness. I am prepared to be persuaded that this ordinary conception can be improved upon. I would lend a receptive ear to proposals to ‘regiment’ the concept of justifiedness so as to judge rightness by reliability in the *actual* world, or by reliability in the possible world of the belief in question. Either of these approaches might seem preferable from a systematic or theoretical point of view. Nonetheless, they do not seem to be what is implied by the ordinary conception as it stands; and that is all I am currently trying to capture. (1986, p. 109)

When Goldman acknowledges that conceptions of justifiedness other than the ordinary one might be preferable to it, he is granting Stich’s point. Semantic consideration alone thus cannot explain the normative force of epistemic terms. In particular, the kind of semantic analysis that Goldman practices cannot explain the force of his own remark that he might find standards different from the ones captured by our ordinary use of epistemic terms “preferable”, for of course only the very standards we have would be recommended by the standards we currently employ. What we really want to know, as Stich rightly urges, is what standards we ought to have, and semantic analysis cannot answer this question. Semantic arguments cannot explain the source of epistemic normativity.

Goldman has recently rejected his suggestion that an a priori investigation of our ordinary concepts can tell us about the nature of justification and knowledge (Goldman 1988), but not because he rejects the enterprise of conceptual analysis. Goldman has suggested instead that conceptual analysis is not to be understood as an a priori discipline; it is to be investigated by straightforwardly empirical means. Goldman’s current psychologized account of the meanings of terms, where lexical

items are said to have psychologically real semantic representations, fits in more easily, I believe, with his overall naturalistic orientation. It is important to see, however, that this kind of move does nothing to solve the problem Stich poses. Whether the meanings of our terms are discovered by a priori investigation or by psychological experimentation is simply irrelevant to the issue under discussion. The fact remains that semantic considerations alone cannot explain the source of the normative force of epistemic terms. Whatever I might mean by the term 'justified', and whatever I might currently approve of, there is a substantive question to be asked about why I should approve of certain sorts of beliefs, and this question is not answered by pointing out what it is I mean by the term 'justified'. Goldman's accounts, both old and new, are in no position to explain the source of epistemic normativity.<sup>2</sup>

## 2.

Any account which does explain the source of epistemic normativity must explain how it is that epistemic claims have normative force. If you tell me that a belief of mine is unjustified, this gives me reason to give up that belief. The epistemic claim is something about which I should care, and an account of the source of epistemic norms must explain why it is that I should care about such things. Since having a desire for something gives one a reason to care about it, it is well worth considering whether desire might serve as the source of epistemic normativity.

There are a number of different ways in which desire might be pressed into service as the source of epistemic norms. I briefly chart the territory here, and in the sections which follow I examine the various possibilities in greater detail.

Stich's account of the relationship between epistemic normativity and desire is certainly the boldest. Stich suggests that epistemic evaluation is grounded in desires for whatever we intrinsically value. Thus:

In evaluating systems of cognitive processes, the system to be preferred is the one that would be most likely to achieve those things that are intrinsically valued by the person whose interests are relevant to the purposes of evaluation. In most cases, the relevant person will be the one who is or might be using the system. So, for example, if the issue at hand is the evaluation of Smith's system of cognitive processes in comparison with some actual or hypothetical alternative, the system that comes out higher on the pragmatist account of cognitive evaluation is the one that is most likely to lead to the things that

Smith finds intrinsically valuable . . . there is no mystery why Smith should care about the outcome of this evaluation. (1990, pp. 131–32)

In light of the diversity of the things which we find intrinsically valuable, this kind of position leads to a form of relativism: the standards which determine what is epistemically right for me are likely to be quite different from the standards which determine what is epistemically right for you. No sense can be made of any attempt to adjudicate among these standards.

Grounding epistemic norms in desire does not, of course, require this radical a position, nor does it require any sort of relativism. Those who seek to avoid relativism in epistemic evaluation, while simultaneously grounding norms in desire, will construe norms as imperatives which apply given that certain conditions are met. Such a view may take any of the following three forms: (1) imperatives may be endorsed which are simply conditional on having certain desires or goals, while acknowledging that these goals are not universally shared; relativity is thereby avoided at the price of a loss of universality; (2) it may be argued that although epistemic imperatives are conditional upon having certain particular goals, these goals are in fact universally held; and (3) it may be argued that epistemic imperatives are conditional upon having any goals at all.

I proceed as follows. In Section 3, I examine the view that epistemic norms are imperatives which apply given that one has a certain particular desire or goal; this encompasses both possibilities (1) and (2) above. I argue that the first of these views, which abandons universality, is not implausible, but that the cost of abandoning universality is not negligible either. The second of these views, however, which would have us regain universality, is extremely implausible. In Section 4, I consider Stich's view that we should instead allow the totality of our concerns to dictate epistemic norms. I argue that Stich's pragmatic view is entirely untenable. In coming to understand why this view is untenable, however, we are provided with the basis for an account of epistemic norms which explain them as imperatives contingent upon having any goals whatever. This view is, I believe, extremely promising, and I examine it further in Section 5.

## 3.

Let us then consider the suggestion that epistemic norms be understood as imperatives contingent upon having certain goals. This is, perhaps, the most natural way of understanding Quine's comment that, "[t]here is no question here of ultimate value, as in morals; it is a matter of efficacy for an ulterior end, truth or prediction". What is the best way to fill out this view?

It might be thought that as long as epistemic norms are being construed in this way, we may simply stipulate any end for epistemic activity we like, and, precisely because the end is merely stipulated, it itself requires no justification.<sup>3</sup> Thus, if Quine is interested in the extent to which belief-producing processes are truth-conducive, he may certainly evaluate them in that way. Others who value the truth as well will find this illuminating; those who do not will find this kind of evaluation of less interest. Similarly, if someone wants to evaluate epistemic activity from the point of view of its conduciveness to truth plus something else, or just something else entirely, that person is free to do so. The norms which issue from such schemes of evaluation are merely directed at select audiences. It should be noted that for those, like Quine, who choose a goal which is very widely held, the interest of their normative scheme is assured. Since many people do clearly care about the truth of their beliefs,<sup>4</sup> Quinean epistemic norms, construed as imperatives contingent upon valuing truth, will carry normative force for a great many people. This would surely explain much of what needs to be explained about the force of epistemic norms.

Although this is certainly one way to fill out the view that epistemic norms are hypothetical imperatives, I do not believe that this way of filling it out does justice to the concerns of many of those involved in epistemological theorizing.

There have, after all, been debates about the ends toward which epistemic norms are directed. Some, like Quine, favor truth and nothing else. Some have urged that the goal of truth must be balanced against other important goals, such as comprehensiveness and speed. Still others have suggested that the goal of truth should be dropped entirely in favor of, for example, biological fitness. How are we to understand these debates?

If the ends which epistemic norms posit are merely stipulated in the

way described above, there is no room for substantive debate among these different schemes of evaluation. Epistemic activity may be evaluated for its conduciveness to various ends, some of which are widely held and some of which are more parochial; but these different schemes of evaluation are not competitors. On this account, it is not as if those who offer different schemes were each trying to characterize a single notion. The problem with this account, of course, is that those who offer different schemes of evaluation typically have seen other schemes of evaluation as competitors. The tenor of this debate suggests that the different parties to it believe themselves to be attempting to characterize a single notion, with some parties to the debate getting it right and others getting it wrong. Now it is certainly possible that the various parties to this debate are simply confused, and this, of course, is what the stipulation account suggests. But it is at least worth exploring whether there might be some account available here which would lend substance to these debates rather than merely dismissing them as conceptual confusions.

Fortunately, there is. The various parties to this debate do not act as if they are merely offering their favorite end as a way of giving substance to epistemic evaluation. Instead, there are activities in which humans are frequently involved, and the favored ends are offered as ways of making sense of and accounting for these activities. For many epistemologists, it is making sense of science that is at issue. Truth is thus frequently offered as at least one of the goals of epistemic activity because science is seen as a paradigm of such activity, and we can understand what science is about when we see it as motivated by a desire for truth. To the extent that this makes sense of scientific activity, it thereby makes sense of a system of evaluation which measures success by its conduciveness to truth. There is room for substantive disagreement here when others agree about the kind of activity they wish to evaluate, and yet disagree about the kinds of motivations which make sense of it. Norms which arise from such evaluations remain merely hypothetical, for they are contingent upon valuing the activity which is made sense of by the indicated concerns. On this view, then, the suggestion that these norms are hypothetical does not make them entirely idiosyncratic, for the activity which gives rise to them may be very widely valued; nor does it make substantive disagreement about them impossible, for there are substantive questions about what the aims or goals of the activity actually are.

This is, I believe, the most plausible version of the view that epistemic norms are imperatives contingent upon valuing certain particular ends.<sup>5</sup> Nevertheless, there can be no denying that it leaves out much that might be desired in an account of epistemic normativity. Why is it, for example, that science is so frequently seen as the paradigm activity from which epistemic norms might be derived? On the account just given, the answer must be merely that many people do value it. Those who wish to say more, for example, that those who do not value science nevertheless ought to, do not see the norms which issue from scientific practice as merely optional. It would certainly be nice if we could have an account of epistemic norms which gave substance to the suspicion that the force of such norms is not so easily avoided.

It is important to note as well that any attempt to gain universal applicability by appeal to goals that all humans in fact have will almost certainly run afoul of the facts. Human beings are a very diverse lot; some of us are quite strange. It is hard to imagine making a plausible case for any particular goal or activity which is genuinely universally valued. There is good reason to believe, at a minimum, that science is not such an activity.<sup>6</sup> I think it is safe to say that anyone who attempts to derive universally applicable norms in this way, namely by combining hypothetical imperatives with particular universally shared goals, has a very substantial burden of proof to meet.

#### 4.

Let us turn then to Stich's view, that epistemic evaluation is just the determination of the extent to which our cognitive states or processes are conducive to the totality of things we value intrinsically.

What is the attraction of a view like this? As Stich indicates, such a view has an important advantage over semantic views like Goldman's. Unlike semantic views, Stich is able to explain straightforwardly why it is that anyone should care about the epistemic status of his cognitive states and processes. Precisely by identifying cognitive evaluation with conduciveness to things one cares about, Stich assures that everyone will care about the outcome of cognitive evaluation. The extent to which one should care about cognitive evaluation on Stich's view, however, may be more a source of difficulty than it is a point in its favor.

If someone tells me that a belief of mine is unjustified, this seems to me to count against my holding it. On Stich's view, however, this is

not merely one consideration against it; it provides a conclusive reason for rejecting it. The judgment that a belief is epistemically unacceptable is, on Stich's view, no different from the judgment that all things considered, it is unacceptable; for on Stich's view epistemic evaluation already takes account of everything an agent values. Now there are two things that are quite strange about this result, and I believe that they are related.

(1) We commonly assume that epistemic evaluation is only one kind of evaluation among many. A candidate belief may fare badly when it comes to epistemic evaluation, but fare well when it comes to various other kinds of evaluation, say, aesthetic or moral. Now it is not that Stich has no room for other kinds of evaluation; he clearly does. Given his strategy of gaining normative force by tying norms to desires, these other norms would presumably derive from desires an agent might have. But if these norms are to be any different from epistemic norms, as surely they must, then they will derive their normative force from some proper subset of the agent's desires for things to which he attaches intrinsic value, rather than the totality of such desires, as is the source of epistemic norms. But if Stich has room for legitimate evaluation which derives only from a subset of an agent's desires, why is it that he insists on tying epistemic norms to the totality of an agent's desires for things to which he attaches intrinsic value? Isn't epistemic evaluation, like many other kinds, more plausibly viewed as directed by only certain concerns and not others? It seems strange that epistemic evaluation should be so all-encompassing. Thus, for example, it seems that I might recognize that having a certain belief would be epistemically ill-advised, and yet have good reason, all things considered, for trying to come to have the belief. If I could assure world peace by committing some epistemic impropriety, surely it would be worth the price. By identifying epistemic propriety with all-things-considered judgments, Stich makes this thought self-contradictory.

(2) This leads me to my second point. What is it, on Stich's account, that makes epistemic evaluation epistemic? I recognize that the objects of evaluation, namely cognitive states and processes, lend some epistemic flavor to this mode of evaluation. But these same objects may be evaluated from, for example, an aesthetic perspective. The mere fact that we are evaluating beliefs does not make our evaluation an epistemic one. And the range of desires people have surely does not do much to lend an epistemic flavor to the evaluation relative to all of one's con-

cerns. Stich's own examples here do not help his case. His chief examples of things one might intrinsically value are health, happiness, and the well-being of one's children (1990, p. 131). It is hard to see how evaluation relative to these concerns is rightly termed epistemic.

Indeed, it seems to me that the natural way to describe Stich's pragmatic view is to say that it is eliminativist about epistemic evaluation: there is nothing distinctively epistemic about the kind of evaluation Stich proposes. If this is correct, then Stich's pragmatic approach falls outside the purview of this essay, for I am interested in finding what room can be made for genuinely epistemic evaluation. If Stich's account should prove correct, it would be because there is no room for such evaluation.

I do want to suggest, however, that Stich's position is not available to those seeking an account of epistemic evaluation, even as a fallback position. More importantly, by seeing how Stich's view fails, we are given a basis for a more substantive and satisfying account of epistemic evaluation. In order to see why this is so, we must imagine that we are actually trying to implement a Stichean evaluation. Here is how Stich tells us to do this.

To assess the comparative merits of a pair of cognitive systems that a person might exploit requires that we compute the expected value of adopting each system. To do that, we must try to determine the probability of each option leading to various possible outcomes and then multiply those probabilities by cardinal number indices of the values we have assigned to the outcomes. The consequences that are important for a pragmatic evaluation will be things that the person in question takes to be intrinsically valuable. (1990, p. 134)

Stich bases this account on standard cost-benefit models of decision. But the presuppositions of cost-benefit calculations undermine Stich's attempt to turn this into an account of cognitive evaluation.

Consider an unproblematic case in which the cost-benefit approach is applied. If I am deciding between two toasters and I wish to use the cost-benefit model, I will begin by determining the consequences of buying each of the candidate toasters. I assign values to each of these consequences, and I do some simple arithmetic. The toaster which has the higher expected value is the toaster I should buy.

In doing all of this, I make use of my cognitive system. I need to figure out the relevant consequences; I need to assign values to each of them; I need to do some arithmetic. The cost-benefit account assumes that these will be done accurately, otherwise the fact that one toaster

is assigned a higher number by this procedure is of no interest. So it is assumed that my cognitive system is generating truths about the toasters,<sup>7</sup> truths about what I value, and accurately computing certain arithmetic functions. These assumptions are perfectly legitimate ones to make when trying to devise a decision procedure for the purchase of toasters. It is how we figure out which toaster better serves our interests, whatever those interests may be.<sup>8</sup>

Now Stich proposes that we evaluate cognitive systems in much the same way. We do not, Stich tells us, value truth, but we do value a great many things: health, happiness, the welfare of our children, and so on. So in evaluating our cognitive systems, we should choose those which favor the things that we value (rather than those that favor truth), just as we do when choosing among toasters.

Now it will not be unfair to Stich if we assume that cognitive systems which are effective in producing happiness and so on are quite different, both in their inferences and in the beliefs they ultimately produce, from cognitive systems which are effective in producing truths. Indeed, it would be miraculous if all cognitive systems produced inferences and beliefs in very much the same way, regardless of the ends they were effective in serving. Moreover, were this the case, it would rob Stich's position of its interest, for Stich means to be endorsing systems very different from those which are favored by truth-based accounts. So we may safely assume, without unfairness to Stich, that those cognitive systems which satisfy his epistemic standards produce beliefs which, by and large, are not true.

Now it seems to me that if we accept a cognitive system of the sort Stich commends, we will have undermined our project of satisfying our desires, whatever those desires may be. For let us imagine now that we are faced, once again, with the prospect of choosing between two toasters, and let us suppose that we do not have a cognitive system which is effective in getting at the truth. Let us instead suppose that we possess a cognitive system which favors happiness, for we are terribly simple folk and care about nothing but happiness. In choosing between the two toasters, once again, we must figure out the consequences of the two purchases; we must assign values to each of them; we must do some arithmetic. If we performed this calculation by using a cognitive system which gave us true beliefs, we would thereby be informed

about the actual consequences of purchasing each toaster, what it is we actually value, and the extent to which these consequences actually produce those things we value. We would thus come to know which toaster better serves our interests, whatever those interests may be.

But Stich does not commend such a cognitive system to us. Instead, he endorses cognitive systems which themselves serve our interests. And we have seen that we may, without unfairness to Stich, assume that such cognitive systems produce very different beliefs than systems which are truth-conducive. But this means that when the happiness-conducive cognitive system is turned to the task of toaster evaluation, it will not tell us what the actual consequences are of purchasing each toaster; instead it will tell us what we would be happiest to believe the consequences to be. Similarly, the happiness-conducive cognitive system will not accurately tell us what it is we value; it will tell us instead what it is we would be happiest to believe that we value. Finally, it will not tell us accurately what will, all things considered, serve our interests, but instead what would make us happiest to believe will, all things considered, serve our interests. As we have seen, it is not unfair to Stich to assume that in each of these cases the happiness-conducive system will generate different results from the truth-conducive system. To put the point only slightly differently, the happiness-conducive system will not tell us which toaster will actually make us happier. Allowing our cognitive systems to be determined by the totality of our interests exclusive of truth thus undermines our ability to make choices, outside the cognitive realm, which are conducive to those very interests.

It is thus safe to say, I believe, that Stich's proposed method of epistemic evaluation does not do the job it was meant to do. Ironically enough, Stich's attempt to devise a pragmatic scheme of cognitive evaluation runs into difficulty precisely where a pragmatic account should be strongest: namely, in allowing us to act so as to serve whatever interests we may care about. It is only by evaluating cognitive systems without regard for the effect such evaluation would have on our actions that Stich is able to endorse the cognitive systems he does.<sup>9</sup> It seems that someone who cares about acting in a way which furthers the things he cares about, and that includes all of us, has pragmatic reasons to favor a cognitive system which is effective in generating truths, whether he otherwise cares about truth or not. We should thus

adopt a method of cognitive evaluation which endorses truth-conducive processes.

## 5.

The argument I just gave against Stich seems to provide the basis for an account of the source of epistemic norms which would allow, on the one hand, that they are derived from our desires in a way which removes any mystery surrounding them, and, on the other, that they are universal in their applicability and not merely contingent upon having certain values. Since this seems to provide us with almost everything one could reasonably want in an account of epistemic norms, we will want to examine this view with some care.

The problem for Stich arose because we need to make evaluations of alternative courses of action and, whatever we care about, we need these evaluations to be done accurately, i.e., by a cognitive system which generates truths. If we have been making revisions in our cognitive system which make it unable to serve this function, then by our own standards we will have done ourselves a disservice. It is thus of the first importance that our cognitive systems remain suitable for the purpose of performing the relevant cost-benefit calculations. And what this requires is that our cognitive systems be accurate, that is, that they reliably get us at the truth.

This suggests that epistemic evaluation takes on a special role. Such evaluation cannot be, as Stich suggests, all-things-considered evaluation; it cannot be so all-encompassing. Precisely because our cognitive systems are required to perform evaluations relative to our many concerns, and to perform these evaluations accurately, the standards by which we evaluate these cognitive systems themselves must remain insulated from most of what we intrinsically value, whatever we may value. This provides a reason to care about the truth whatever we may otherwise care about. It also provides us with a reason to evaluate our cognitive systems by their conduciveness to truth. And this is precisely what epistemic evaluation is all about. Truth plays a pre-eminent role here.

Have I assumed here that epistemic evaluation is measured by conduciveness to truth and nothing else? I have not. I have argued that truth is pre-eminent here; that any account of epistemic evaluation which does not give truth a central role to play is inadequate. There

may still be a good deal of room for other factors to play a role. My argument for the importance of truth turned on its being implicated in certain cost-benefit calculations; the calculations which we need to perform in making choices among alternative courses of action must be done accurately. But various goals other than truth are likely to be implicated in this task as well. A system of evaluation which was perfectly accurate but could not perform its evaluations in real time<sup>10</sup> would be of little value. The task of evaluation thus brings with it certain demands. There is room for substantive discussion about just what these demands are. And it is just this kind of discussion which lends substance to debates about the dimensions of epistemic evaluation.

## 6.

I have argued that epistemic evaluation finds its natural ground in our desires in a way which makes truth something we should care about whatever else we may value. This provides us with a pragmatic account of the source of epistemic normativity, but an account which is universal and also allows truth to play a central role. Pragmatists have typically suggested that epistemic evaluation will have little to do with truth; but if I am right, it is for pragmatic reasons that truth takes on the importance it does in epistemic evaluation.

Some will, I believe, hanker after a stronger grounding for epistemic normativity, an account which would make the injunction to seek the truth not merely hypothetical, even if universal, but categorical instead. Such an account would entail that the value of truth is not merely instrumental, as I have suggested, but intrinsic, and that attaching intrinsic value to truth is not merely optional, but required. I would not be hostile to such an account, but I do not currently see any way of giving substance to it. As things stand, I believe the account of epistemic normativity I offer allows us to make sense of much of what a categorical account would provide, while simultaneously removing the mystery from epistemic norms. As far as the remainder goes, those things that a categorical account would provide which my account does not, it remains to be seen whether any real sense can be made of them.<sup>11</sup>

## NOTES

\* I am indebted to Derk Pereboom, Lynne Baker, George Sher, William Talbott, and an anonymous referee for helpful comments on a previous draft. David Christensen provided me with advice which resulted in major changes in the structure of the argument here; I am especially indebted to him. In addition, I received helpful comments from audiences at Concordia University, SUNY at Albany, and the 1991 Western Division meeting of the A.P.A. After this paper was written, Alvin Goldman sent me a draft of his (1991). There is a good deal of overlap between what Goldman has to say about Stich and Section 4 above.

<sup>1</sup> For similar remarks, see the discussion of this issue in Quine (1990, pp. 19–21).

<sup>2</sup> In Goldman (1992), Goldman now makes room for the revision of our “folk concepts” of justification and knowledge, thereby leaving room for the kinds of concerns urged here.

<sup>3</sup> I myself do not believe that this is a proper account of what is involved in stipulation, but because so many use that term as if it carried with it no requirement of justification, I will follow that practice here.

<sup>4</sup> Or at least believe that they do. Stich denies that many people really do attach intrinsic value to the truth of their beliefs.

<sup>5</sup> Room should be made as well for interaction between views about the goals which make sense of an activity and the activity itself. On discovering that an activity I value only makes sense relative to certain goals, I may modify the activity itself by modifying the goals. This kind of reflective equilibrium account fits well with the view of norms as imperatives contingent upon having certain goals.

<sup>6</sup> A referee suggested that science might be defended here by reference to its contribution to technology and the control of nature; this, it was claimed, is a plausible candidate for something with universal appeal. But technology and the control of nature do not have universal appeal. There are certainly many who sincerely claim to disvalue technological achievement and who see the desire to control nature as somehow perverse or pathological. I would not defend these values, but there are people who hold them.

<sup>7</sup> Or accurate probabilities. This complication will not help Stich.

<sup>8</sup> A referee reminded me that Stich would deny this. It is, indeed, the burden of Chapter Five of Stich (1990) to argue that the notion of truth is an idiosyncratic notion, of no cognitive significance. I cannot possibly do justice to that argument of Stich here. Those who favor cost-benefit analyses, however, have traditionally made the assumptions I attribute to them in the text. It would be interesting to see precisely what a full-blown Stichean reconstruction of cost-benefit analyses would look like.

<sup>9</sup> This may suggest that I believe Stich’s method of cognitive evaluation could be effective in initiating modifications in cognitive systems, and that difficulties only arise when the outputs of those cognitive systems are used in evaluating choices for action. Now this would be cold comfort to Stich or any other pragmatist, but I do not believe that Stich’s system of evaluation would work on even this extremely limited scale. Imagine someone evaluating his own cognitive system to see, as Stich suggests we ought, the extent to which it is conducive to things he cares about. For reasons like those presented in the argument above, such an evaluation will only be accurate if the agent uses a cognitive system which generates truths. If the agent abandons such a cognitive system in favor of

some non-truth-conducive system, then when the agent is prompted to re-evaluate his cognitive system relative to others, he will not accurately determine the comparative merits of the two systems, even relative to his own standards. Having abandoned a truth-conducive cognitive system for one which satisfies the many things he values, he will be no more able to assess accurately the merits of his cognitive system than he will be able to assess accurately the merits of toasters. So Stich's system of cognitive evaluation could not be put to work on successive cognitive evaluations, even if one were wholly unconcerned about the ways in which beliefs bear on actions.

<sup>10</sup> Cherniak (1986) is filled with illuminating discussion of this point.

<sup>11</sup> One kind of account of epistemic norms I have not considered here, as a referee pointed out to me, would derive from a theory of epistemic virtue or epistemic responsibility. I myself tried to motivate such an account (in Kornblith 1983); see also Code (1987) and Sosa (1990). I do not currently see a way to develop such an account without signing on to the instrumental account of epistemic value offered here. Others who favor the epistemic virtue approach, however, may have an independent account of epistemic value in mind.

## REFERENCES

- Cherniak, Christopher: 1986, *Minimal Rationality*, MIT Press, Cambridge, MA.
- Code, Lorraine: 1987, *Epistemic Responsibility*, University Press of New England, Hanover, NH.
- Elgin, Catherine: 1988, 'The Epistemic Efficacy of Stupidity', *Synthese* **74**, 297–311.
- Field, Hartry: 1982, 'Realism and Relativism', *Journal of Philosophy* **79**, 553–67.
- Goldman, Alvin: 1986, *Epistemology and Cognition*, Harvard University Press, Cambridge, MA.
- Goldman, Alvin: 1988, 'Psychology and Philosophical Analysis', *Proceedings of the Aristotelian Society* **LXXXIX**, 195–209.
- Goldman, Alvin: 1991, 'Stephen P. Stich: *The Fragmentation of Reason*', *Philosophy and Phenomenological Research* **LI**, 189–193.
- Goldman, Alvin: 1992, 'Epistemic Folkways and Scientific Epistemology', in his *Liasons: Philosophy Meets the Cognitive and Social Sciences*, Bradford Books/MIT Press, Cambridge, MA, pp. 155–75.
- Kornblith, Hilary: 1983, 'Justified Belief and Epistemically Responsible Action', *Philosophical Review* **XCII**, 33–48.
- Lycan, William: 1981, "'Is" and "Ought" in Cognitive Science', *Behavioral and Brain Sciences* **4**, 344–45.
- Lycan, William: 1988, *Judgment and Justification*, Cambridge University Press, New York.
- Quine, W. V. O.: 1969, 'Epistemology Naturalized', in his *Ontological Relativity and Other Essays*, Columbia University Press, New York, pp. 69–90.
- Quine, W. V. O.: 1986, 'Reply to Morton White', in Lewis Hahn and Paul Schilpp (eds.), *The Philosophy of W. V. Quine*, Open Court, La Salle, IL, pp. 663–65.
- Quine, W. V. O.: 1990, *The Pursuit of Truth*, Harvard University Press, Cambridge, MA.

Sosa, Ernest: 1990, *Knowledge in Perspective: Selected Essays in Epistemology*, Cambridge University Press, New York.

Stich, Stephen: 1990, *The Fragmentation of Reason*, MIT Press, Cambridge, MA.

Department of Philosophy  
University of Vermont  
70 South Williams Street  
Burlington, VT 05401  
U.S.A.