

Epistemology as Engineering

1. Introduction

Wild guessing and wishful thinking do not produce knowledge, even if your guess is lucky or your wish comes true. For a belief to constitute knowledge, it needs something more than mere truth. Let us use ‘justification’ to name that property of beliefs which, together with truth, suffices for a belief to be knowledge. When a belief has this property, we will call it “justified,” and we will call it “unjustified” when it lacks justification.

Philosophers disagree widely about the nature of justification, and even about whether ‘justification’ is an apt name for the property.¹ Cutting across this disagreement, however, is widespread agreement that justification, whatever it is, entails the proper or correct use of one’s cognitive resources.² One way of showing that a proposed account of justification fails is to show that it could be satisfied even in circumstances where a believer has not used her cognitive resources correctly. If this widespread agreement is on the right track -- and I will not dispute it here -- then it would follow that the notions of justification and knowledge are normative; they essentially involve a distinction between good or correct uses of one’s mind and bad or incorrect uses of it. It also seems to follow that an adequate theory of knowledge needs to be a normative theory. As Jaegwon Kim has put it, “Epistemology is a normative discipline as much as, and in the same sense as, normative ethics” (2000/1988: 302).

The normativity of epistemology poses a problem for a certain version of epistemological naturalism. According to such naturalists as W. V. Quine, an adequate theory of knowledge must be an empirical, scientific theory. Instead of doing epistemology a priori and in isolation from our substantive, scientific understanding of the world, Quine thinks we should conceive of

epistemology as continuous with science, and he believes that epistemologists should employ the best methods and results contemporary science makes available (Quine 1969).³ But science is a paradigmatically descriptive enterprise. It is, we can suppose, very good at providing us with information about how things are in the world, but it is hard to see how empirical science could tell us anything about how good things are or about how things should be. Naturalism thus appears to be a necessarily inadequate approach to epistemology. Naturalists might be able to provide descriptive theories about knowledge, but their commitment to studying knowledge empirically seems to prevent them from producing normative theories of knowledge. And since epistemology requires the latter, naturalistic epistemology seems to be no epistemology at all.⁴

I call this the ‘normativity objection’ to naturalism, and my primary concern in this paper is to examine one way in which naturalists have tried to answer it. Quine has suggested that, once we see ‘theoretical’ epistemology as a ‘chapter of science,’ we are in a position to see ‘normative epistemology’ as an engineering discipline, concerned with questions about how best to suit our cognitive means to our cognitive ends (Quine 1998: 664-5). I will refer to this move as the ‘engineering reply’ to the normativity objection. For the most part, naturalists have not worked out the engineering reply in very much detail, and that lack of detail makes the reply hard to evaluate. My aim in this paper, then, is to work out a version of the engineering reply in considerably more detail than has previously been done, and to identify some of its limitations. My ultimate evaluation of the reply is that it is sufficient to undermine the least sophisticated version of the normativity objection, but much more work is needed before it can provide a full-dress defense of normative, naturalistic epistemology.

In Section 2, I summarize the least sophisticated normativity objection and I describe the basic shape of the engineering reply. Since both the objection and the reply turn on claims about

whether or not scientific theories can be normative, my aim in Section 3 is to set out a plausible account of what it could mean to call theories ‘normative’ or ‘descriptive.’ I apply that account to epistemology and engineering in Section 4, and in Section 5 I look at some of the limitations of the engineering reply.

2. Naturalism and the Normativity Objection

In ‘Epistemology Naturalized,’ Quine calls for epistemologists to give up, among other things, the project of justifying scientific theories and methods a priori. Instead, he thinks we should address our philosophical questions about knowledge by studying it as a natural phenomenon, using the same scientific theories and empirical methods we would use to study any other natural phenomenon. Once we have done this, he says, ‘epistemology, or something like it, ... falls into place as a chapter of psychology’ (1969: 82). It studies ‘a physical human subject’ and seeks answers to questions about how such a subject derives beliefs about ‘the three-dimensional external world and its history’ from the ‘meager input’ of sensory stimulation (1969: 82-3).

Such advocates of the normativity objection as Jaegwon Kim and Wilfrid Sellars, however, see Quine-style proposals as calls to give up on epistemology altogether. According to Kim, the scientific study of knowledge can tell us only about how people do arrive at their beliefs, but it cannot answer normative questions about what justifies those beliefs or about how people should use the evidence sensory stimulation provides them. Thus, Kim believes, to ‘settle for psychology,’ as Quine proposes (Quine 1969: 75), is to stop trying to produce theories of justification. But, says Kim, ‘If justification drops out of epistemology, knowledge itself drops out

of epistemology.... For epistemology to go out of the business of justification is for it to go out of business” (2000/1988: 305-306).

Sellars expresses much the same view in ‘Empiricism and the Philosophy of Mind’ (1956: Sect. 36):

The essential point is that in characterizing an episode or state as that of knowing, we are not giving an empirical description of that episode or state; we are placing it in the logical space of reasons, of justifying and being able to justify what one says.

For Sellars, an adequate understanding of knowledge requires an understanding of how the normative, ‘logical space of reasons’ is structured. That is, it requires an understanding of what can count as a good reason for what. The most he thinks we can get from empirical studies of knowledge or belief, however, is information about the place of beliefs in the causal-nomological order of the world. But in Sellars’ view, the space of reasons is irreducible to the causal-nomological order, and so understanding knowledge requires us to go beyond the empirical information science makes available to us.⁵

In its least sophisticated form, then, the normativity objection is as follows:⁶

P1. Naturalistic theories of knowledge are empirical, scientific theories.

P2. An adequate theory of knowledge must be normative.

P3. Empirical, scientific theories are not (or cannot be) normative.

C. Therefore, naturalistic theories of knowledge are not (or cannot be) adequate. They are either false or incomplete.

How should naturalists respond to this argument? Though there are philosophers who would reject them, we can assume that P1 and P2 are unobjectionable.⁷ On that assumption, naturalists

must respond to the objection by providing a principled way of rejecting P3. That is exactly what Quine's suggestions about the engineering reply are meant to do.

In his 'Reply to Morton White,' Quine writes (1998: 664-5):

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 more cautiously epistemological term, prediction. Like any technology, it makes free use of whatever scientific findings may suit its purpose. It draws upon mathematics in computing standard deviation and probable error and in scouting the gambler's fallacy. It draws upon experimental psychology in exposing perceptual illusions, and upon cognitive psychology in scouting wishful thinking. It draws upon neurology and physics, in a general way, in discounting testimony from occult or parapsychological sources. 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.

Quine sees science's capacity to address normative questions in terms of its ability to address questions about which means are most (or least) likely to promote which ends causally. The ends relevant to epistemology are such things as true belief and the prediction of sensory stimulation. Just as engineers can use science to tell the difference between better and worse bridges, rockets, and skyscrapers, epistemologists, in Quine's view, can draw on science to tell the difference between better and worse uses of people's cognitive resources.

The basic idea behind the engineering reply is that naturalistic epistemology can be normative in the same way that engineering is, by studying the causal relationships between our cognitive

means and our cognitive ends. For the reply to be satisfactory, however, it needs to be worked out in much more detail. We need to be clearer about (i) what it is for theories (scientific or otherwise) to be “normative,” (ii) the sense in which engineering is “normative science,” and (iii) why it is legitimate to assimilate normative epistemology into engineering. Only when we are sufficiently clear about these things will we be in a position to determine how successful the engineering reply is.

3. Normative Theories

Let us begin by trying to see what it could mean to call a theory “normative.” To make the engineering reply work, naturalists need an account that does two things. First, it should preserve our intuitive classifications of certain theories as normative or descriptive. Otherwise, the account would be too implausible to ground a useful defense of naturalism. Second, it should allow for the possibility of normative scientific theories. Otherwise, the account would ground no defense of naturalism at all. In this section, I sketch what I take to be a plausible account that meets both these conditions.

Consider the following three sentences:

(D) Roses are red.

(E) Violets are pretty.

(P) One ought to be patient with folks from the city.

Classified intuitively, (D) is descriptive, (E) is evaluative, and (P) is prescriptive. Also intuitively, (E) and (P) are normative sentences, though (D) is not. The present problem is to say something intelligible about how (E) and (P) differ from (D). Once we are clear about what makes sentences normative or descriptive, it will be possible to extend the analysis to apply to theories.

Ordinarily, (D), (E), and (P) would be used to do different things. One would ordinarily use a sentence such as (D) to assert that things in the world stand in a certain way, without committing oneself to any particular attitude of approval or disapproval of things' standing in that way. When one utters (D), one is ordinarily in the business of making a report. In contrast, the ordinary point of uttering (E) is not just to report that violets have a certain property, prettiness. Rather, the point of (E) is express one's approval of the way violets look. This is not to say that a sentence such as (E) expresses no cognitively contentful claim, but only to point out that one who utters (E) would ordinarily be trying to do more than just to assert that claim.⁸ Part of the difference between (D) and (E), then, is that utterances of sentences like (E) ordinarily function as manifestations of approval or disapproval, while utterances of (D) do not.

A person who utters (P) would ordinarily be endorsing a rule or a pattern of action. To be patient with city folk is to conform to the rule; to be impatient with them is to violate it. As with (E), there is an element of approval and disapproval present here that is absent in (D). Uttering (P) is a way to manifest one's approval of patience with city people and one's disapproval of impatience with them.

Of course, neither (D), (E), nor (P) must be used in the ways I have described as "ordinary." Each of them might be uttered by actors in a play or never uttered at all. Such accidents, however, should not affect our judgments of the sentences' "ordinary" uses. In English and other languages, there are familiar grammatical and lexical devices whose function is to indicate the ordinary point of uttering a sentence. For example, the imperative mood and certain uses of 'ought' are markers of prescription. Though both descriptive and evaluative sentences often occur in the indicative mood, evaluative speech acts are often marked by such vocabulary as 'good', 'bad', 'nice', and

'ugly'. I define the ordinary uses of sentences to be those indicated by the features of a sentence that, according to existing linguistic convention, function to mark kinds of speech acts.

In this technical sense, any sentence in the imperative mood has prescription as an ordinary use. This is because the imperative mood is a grammatical device for marking prescriptive speech acts. It is important not to assume that every sentence has exactly one ordinary use. A single sentence with the markers of description, evaluation, and prescription might have three ordinary uses, one derived from each of the kinds of markers present. For example, there is no good reason not to count:

Roses are not red, though violets are pretty, and one ought to be patient with folks from the city

as "ordinarily" used to assert that the world is a certain way, to manifest one's preferences, and to express endorsement of a rule.⁹ It is also important to note that I do not contend that these are the only possible ordinary uses of sentences.

I define the descriptiveness, evaluativeness, and prescriptiveness of sentences as follows:

Descriptiveness. A sentence is descriptive iff one of its ordinary uses is to assert that things in the world are a certain way (i.e., have certain properties or stand in certain relations to one another).

Evaluativeness. A sentence is evaluative iff one of its ordinary uses is to manifest or express one's liking/approval or disliking/disapproval of something in some respect.

Prescriptiveness. A sentence is prescriptive iff one of its ordinary uses is to endorse a rule or to express one's approval/disapproval of actions (types or tokens) or patterns of behavior.

With these definitions in hand, we can characterize as normative exactly those sentences that are either evaluative or prescriptive.¹⁰ Figure 1 summarizes the way in which these definitions carve up the space of sentences.

We can extend the distinction between normative and descriptive sentences to apply to theories. Accepting a theory commits a person to assent to certain sentences. For example, imagine I have adopted Newton's theory of gravitation. My endorsement of the theory commits me to assent to 'My coffee cup exerts an attractive force on the planet Mars'. I cannot consistently deny that sentence while maintaining my endorsement of Newton's theory. Notice that this is so whether we treat theories as sets of sentences or not. The point is that whatever a theory is, adopting one commits a person to assent to some sentences. Also note that the sentences in question need not be "part" of the theory in any particularly robust sense. The sentence about my coffee cup, for example, is not part of Newton's theory; for the sentence, but not the theory, implies that I have a coffee cup. Finally, we should recognize that "assent to a sentence" here applies to more than merely holding true the cognitive contents of declarative sentences. Accepting Kant's theory of morality, for example, commits one to assent to such prescriptive sentences as the various formulations of the Categorical Imperative.

I define the consequences of a theory as follows:

Sentence S is a consequence of theory T for person x iff (i) x would be committed to assent to S were x to accept T, and (ii) there is a subset B of x's background beliefs such that B alone does not commit x to S but B and T together do commit x to S.¹¹

The purpose of condition (ii) is to ensure that x's acceptance of T is responsible for x's commitment to S, rather than x's background beliefs alone. If a theory is a deductively closed set of sentences, the members of that set are consequences of the theory for everyone. Because not

all the consequences of a theory are parts of it, whether a given S is a consequence of a given T turns on what difference adding T to an already existing system of background beliefs would make. To the extent that background beliefs vary from person to person, there will be differences in a theory's consequences from person to person, and to the extent that there is uniformity in background beliefs, there will be uniformity in the consequences of a theory. Whenever a sentence is a consequence of a theory for everyone or for nearly everyone, I say that the sentence is a consequence of the theory, and I drop the relativization to persons.

A theory is descriptive for a person when it has descriptive consequences for her, normative when it has normative consequences for her. It is descriptive or normative full stop accordingly as it has descriptive or normative consequences for everyone or nearly everyone. I thus count many theories, perhaps most, as both descriptive and normative, but that is not a defect in the account. The intuitive distinction between descriptive and normative theories distinguishes the properties of being descriptive and being normative. On the present account, these properties remain distinct, if often coinstantiated. When a theory is both normative and descriptive, it is normative in virtue of having one set of consequences and descriptive in virtue of having another.

The Special Theory of Relativity, on this way of drawing the distinction, is descriptive because it has as a consequence 'Nothing travels faster than light', which is a descriptive sentence. Kantian ethics, on the other hand, is normative because it has 'One ought not treat other human beings merely as means' as a consequence.

Among normative theories, it is useful to distinguish those that are strongly normative from those that are weakly normative. To be a weakly normative theory is just to be a normative one, i.e., to have normative consequences. Strongly normative theories are, in a certain sense, about their normative consequences. It is part of the point in adopting the theory that it provides

guidance about what stances to take towards evaluative or prescriptive sentences. To put it another way, the theory's normative consequences are the ones we are most interested in when we put the theory to use; we use strongly normative theories to decide whether or not to accept normative sentences. Even though merely weakly normative theories have normative consequences, we tend to use them to settle questions about the truth or falsity of descriptive sentences, rather than questions about what stance to take toward normative sentences.¹²

Suppose, for example, that the theory of relativity has normative consequences. It is merely weakly normative because its purpose is not to settle normative issues.¹³ We intend it to tell us how things are, but we do not usually appeal to it in making evaluations or prescriptions. In contrast, the Kantian theory of morality is strongly normative. Not only does it have normative consequences, but its purpose is to provide guidance about what stances to take toward evaluative and prescriptive sentences.

As well as distinguishing strong and weak normativity, we can distinguish between theories that are ‘hypothetically’ normative and theories that are ‘categorically’ normative. This latter distinction pertains to what kinds of evaluations and prescriptions are the consequences of a theory. Some evaluative sentences ascribe value to things for certain purposes. Steel chairs are good for sitting but bad for starting fires, and claw hammers are good for driving nails, for example. Similarly, some prescriptive sentences are meant to provide instruction only to those who have adopted certain ends or goals. These kinds of prescriptive sentences are often called ‘hypothetical imperatives.’ An example of such an imperative would be ‘Those who want to lose weight ought to eat less and exercise more’. I count a theory as hypothetically normative whenever its consequences include hypothetical imperatives or evaluative sentences that ascribe value relative to particular purposes, ends, aims, or goals. When a theory has consequences that

ascribe value independent of anyone's purposes, ends, aims, or goals ('Murder is bad,' for example), or it has prescriptive consequences meant to bind people regardless of their ends, etc., I call the theory "categorically normative."

Just as a single theory can be both descriptive and normative, a single theory can be both hypothetically and categorically normative. For example, there is a version of Utilitarianism according to which happiness is intrinsically valuable, and all other value is instrumental to happiness. The ascription of intrinsic value to happiness is independent of people's particular ends, aims, or goals, but, according to this theory, anything else that is valuable is valuable only for the purpose of promoting happiness (or avoiding the reverse of happiness). Thus the theory has both the sort of consequences that would make it hypothetically normative and the sort of consequences that would make it categorically normative.

A final distinction will also be useful. To some ears, the following sentence is not normative:

(A) To get the most out of one's heroin, one should shoot up in the presence of friends.

though the formally similar sentence below is normative:

(B) To get the most out of one's dinner, one should eat in the presence of friends.

The account I have given so far counts both sentences as normative (because prescriptive), and it would count as normative any theory with either as a consequence. The most important reason why one would balk at calling (A) normative is that, intuitively, normative sentences should, in some sense, make a claim on people's behavior. But, so long as we assume that getting the most out of one's heroin is not a worthwhile goal (i.e., not a goal with sufficient value to make it worth having or pursuing), (A) seems to make no such claim. In contrast, getting the most out of one's dinner probably is a worthwhile goal, and so (B) does seem to make a claim of some sort on people's behavior.

We can incorporate this intuition into the present account. For any strongly, hypothetically normative theory, let us say that it is well-motivated (or that it exhibits well-motivated hypothetical normativity) if and only if it has evaluative or prescriptive consequences relativized to worthwhile ends. When a strongly, hypothetically normative theory is not well-motivated let us call it merely strong.

Figure 2 summarizes the distinctions that have just been drawn among kinds of theories.

4. Engineering and Normative Epistemology

With the above explication in hand, we can return to the normativity objection and the engineering reply. First, the account shows it to be plainly false that science cannot, in principle, be normative. This is because adopting a scientific theory often commits one to normative sentences. For example, suppose I accept Newtonian mechanics. I am thus committed to the following normative sentence:

Those who want to launch rockets into outer space should load them with enough fuel of the right kind to ensure they accelerate fast enough long enough to reach their escape velocities (as calculated in Newtonian mechanics).

I am also committed to counting rockets that do what is described above as better for going into outer space than rockets that do not. Had I accepted Aristotelian rather than Newtonian mechanics, there would be different normative consequences. There is no such thing as ‘escape velocity’ in Aristotelian mechanics, and the theory does not imply that those who want to launch rockets into space should build rockets that achieve it.

These considerations suggest that not only can scientific theories be normative, but they often are. It is sensible to construe engineering and applied science as largely concerned with working

out the normative consequences of scientific theories. They are interested in determining what the best ways to pursue our goals are, given that the world is the way we think it is. In the case of engineering, at least, where the point of using a theory is often to distinguish better from worse ways of getting a job done, scientific theories are not only normative but strongly normative.

Consider the case of a civil engineer whose job is to design a bridge to carry traffic across a river. Once it is clear what kinds of loads the bridge must support to serve its function, the engineer will draw on a sheaf of scientific theories to determine how the bridge ought to be built. She will use general theories from physics as well as theories specific to bridge-building to determine what designs are sufficiently stable for the task. She will use geology to make decisions about the best ways to connect the bridge to the ground, and she will apply materials science to questions about what kinds of concrete are best for the bridge's support piles and driving surface. She will apply meteorology to decide what kinds of winds the bridge needs to withstand, and physics again to determine what designs can withstand them. Even before the engineer begins her work, others will have drawn on economics and other departments of science to decide how much traffic a good bridge in this location should be capable of handling and how long its life span should be. The result of all this applied science will be a specification of the difference between good and bad bridge designs for a given location and purpose, as well (one hopes) as the construction of a good bridge.

One task of engineering is to apply our best scientific theories of the world to questions about how to suit our means to our ends. There is thus no shortage of strong, hypothetical normativity in applied science. When our ends are worthwhile, the normativity is also well-motivated. Let us now see what it could mean to construe normative epistemology as a "branch of engineering."

The main problem normative epistemology sets out to solve is that of distinguishing better from worse cognitive practices or uses of one's cognitive resources. Normative epistemology modeled on engineering, then, sees this as a problem about how best to suit our cognitive means to our cognitive ends, given our best scientific understanding of ourselves and our world. We want to use our minds to produce certain results, such as belief in important truths, disbelief in dangerous falsehoods, and accurate prediction of the future course of our experience. Science provides us with information about both how the world is and about how our minds process the information we gather from it. Normative naturalized epistemology draws on that science to evaluate the effectiveness of our cognitive means in achieving our cognitive ends.

I think Quine has something like this in mind when he compares normative epistemology to engineering. In one place, he describes how accepting scientific theories can commit one to normative, epistemological principles as follows (1992: 19):

The most notable norm of naturalized epistemology actually coincides with that of traditional [empiricist] epistemology: nihil in mente quod non prius in sensu. This is a prime specimen of naturalized epistemology, for it is a finding of natural science itself, however fallible, that our information about the world comes only through impacts on our sensory receptors. And still the point is normative, warning us against telepaths and soothsayers.¹⁴

'Those who want to have true and not false beliefs should ignore the testimony of soothsayers' is a normative epistemological claim. But it is, as Quine points out, also a consequence of our best current scientific understanding of the world.

Some readers might want a more interesting example of normative naturalized epistemology than Quine's remark about soothsayers. I thus offer the following. Scientists studying the brain

are very interested in determining what kinds of activity occur in which brain areas as subjects perform various experimental tasks. One way in which they do this is by using magnetic resonance imaging (MRI) to track blood flow through the brain.¹⁵ This technique has been in use for many years, but until very recently the legitimacy of its results was not as firmly established as one would like in science.

The legitimacy of the results was unclear because MRI measures blood flow, but, according to our best view of the how the brain works, it does not process information by moving blood around. Information processing in the brain is generally believed to result from the transmission of electrical signals between neurons. Though there has been good reason to believe that increases in blood flow to a brain area are correlated with increased electrical activity there, no one had been able to measure blood flow and electrical activity simultaneously.

Then, in July 2001, Nikos Logothetis and some of his colleagues published the results of a study in which they had managed to simultaneously measure electrical activity and the blood flow properties MRI measures in monkey brains. They did find close correlations between some kinds of electrical activity and blood flow. But they also found that the kinds of blood flow properties MRI measures are much better indicators of certain forms of neural activity than of others.

The epistemological upshot of Logothetis et al.'s work is more than just a vindication of MRI as a way of locating brain activity. It does not only indicate that those who want to find out where the brain is processing information can do so accurately by attending to the results of MRI scans. Since their work also identifies which kinds of neural activity go along with what MRI measures, it also allows scientists to make decisions about whether or not to use MRI on the basis of what kinds of neural activity they want to find out about. In particular, Logothetis et al. claim that MRI measures properties that 'reflect the input and intracortical processing of a given area rather than

its spiking output” (2001: 150). This means that those who want to gather information about an area’s input and intracortical processing should consider using MRI to do so, and those who are more interested in gathering information about output should not rely on MRI alone.¹⁶

To the extent that the normativity objection turns on the view that naturalized epistemology cannot be normative at all because, necessarily, science is merely descriptive, I think it is clearly a failure. Not only is it false that science is necessarily merely descriptive, but it is also false that the findings of empirical science provide no normative guidance to the conduct of our mental lives. That fundamental insight of the engineering reply is enough to undermine the least sophisticated version of the normativity objection, but I do not think it is enough to provide a full-dress defense of normative, naturalistic epistemology. A full-dress defense would require naturalists to do some additional work to answer some further questions and to block more sophisticated versions of the normativity objection. I examine the prospects for those projects in the following section.

5. Completing the Engineering Reply

The least sophisticated version of the normativity objection turns on the claim that science is not (or cannot be) normative. The engineering reply works by showing that that is simply not true. A slightly more sophisticated version of the normativity objection would turn on the claim that even though science is (or can be) normative, that is irrelevant to epistemology. Science, according to this version of the objection, cannot tell us how best to suit our cognitive means to our cognitive ends. The fact that it is science that tells us not to trust soothsayers if we prefer believing truths to falsehoods seems to show that this version of the objection also turns on a mistake.

Naturalists should not rest with having answered just these two versions of the normativity objection. There are some additional issues they must address if they are to vindicate normative naturalistic epistemology fully. I conclude this paper by briefly discussing some of them.

Issue #1: Kinds of Normativity. Naturalistic appeals to engineering show only that scientific theories can be hypothetically normative, and that theories in engineering tend to be strongly hypothetically normative. This has two important implications for naturalized epistemology.

First, the engineering reply does not establish that naturalized epistemology can be categorically normative. Some philosophers will insist, however, that an adequate epistemological theory must be not only normative, but categorically normative.¹⁷ It must tell us about noninstrumental, noncontributory epistemic value, or it must describe the epistemic duties that bind everyone regardless of their actual goals or ends. They might add that this was the real point of the normativity objection all along, and so declare that the engineering reply misses the point. Naturalists, then, need to show either that categorical normativity is within the reach of science (which the engineering reply does not do), or that epistemology does not require categorically normative theories after all. Most often, naturalists have taken the latter course,¹⁸ but there is still a great deal of work to be done on all sides in the debate about whether or not there are categorical epistemic norms.

Second, the engineering reply does not even show that scientific theories of knowledge can be strongly hypothetically normative. Even if we suppose categorical normativity is unnecessary, naturalists should not be satisfied with scientific theories of knowledge that are only weakly normative. You could make a case for the weak normativity of almost any theory, and so that kind of normativity is probably too trivial to answer the demands of those who would press the

normativity objection. I do not think it would be very hard, however, to make the case that strongly hypothetically normative, scientific theories about cognition are possible. All that is required is for us to develop a practice of applying scientific findings to questions about how to pursue our cognitive ends, and naturalists are already working on that.

Issue #2: Identification of Ends. This issue stems from the first. If we suppose that naturalists must reject categorical epistemic normativity, then it is incumbent on naturalists to identify the ends to which their hypothetically normative theories are relativized. This is a project many naturalists have pursued, and there are at least two ways to pursue it.

One is the analytical way. Those who pursue the project in this way suppose the question of what our cognitive ends are can be settled a priori, by examining the concepts of knowledge, belief, inquiry, science, or cognition. Among those who have used this approach are Alvin Goldman (1986; 1992), Susan Haack (1995), James Maffie (1990; 1995a; 1995b), and even Quine (1992).¹⁹ Those who take this approach are likely to say, for example, that if a person is not trying to acquire true beliefs and to avoid false ones, what she is doing simply does not count as inquiry, for it is part of the very concept of inquiry that it is aimed at truth.

A second approach treats the question of what our cognitive ends are as one for empirical science. Such philosophers as Larry Laudan (1984) contend that one cannot properly understand human cognitive activities -- especially science -- by appeal to such vague and metaphysically loaded notions as "truth." If we want to know what our cognitive goals are, they think, we should take a scientific look at our cognitive practices -- especially our scientific practices. We can then learn things about, for example, the conditions under which we take an inquiry to have reached a successful conclusion. It may be unlikely that we will find any overarching goals toward which all our cognitive activities are aimed, but that might be because we do not have any such overarching

goals. We can still discover the more particular cognitive goals we pursue, and we can still examine the ways in which we pursue them.

It is not my aim to argue the virtues and vices of these two approaches. The second has a more naturalistic flavor, but it is not clear that the analytical approach should be a nonstarter. (How could it fail to be part of the concept of inquiry that inquiry is aimed at truth?) Rather, I just want to make three points. First, the question of what our cognitive ends are is a very important one for naturalists. Second, it has proven to be far from obvious what our cognitive ends are. And third, it is not entirely clear how to decide what our cognitive ends are while remaining consistent with naturalistic commitments.

Issue #3: Evaluation of Ends. Suppose naturalists were able to come up with a respectable list of ends to which epistemological norms are relativized. They still might face the objection that they have not yet -- or cannot -- show the resulting theories to have well-motivated normativity. To do that, naturalists would have to show that the ends they are concerned with are worthwhile ends, ends worth pursuing. This is a further normative question, about the value of the putative cognitive ends themselves, and it is again far from clear how naturalists will answer it.

There seem to be at least four available options:

1. Contend that the problem is not really all that important because these are ends we actually take to be worthwhile -- subject to later revision -- and we have to do epistemology from within our ongoing view of things.
2. Argue for the instrumental value of these goals relative to other (worthwhile?) goals we happen to have.
3. Argue for the intrinsic or noninstrumental value of these goals.

4. Argue that these goals are instrumentally valuable with regard to whatever goals we might also have.

Quine (1992) seems to have followed the first two courses. He contends that accurate prediction is worthwhile with regard to our technological goals and the goal of understanding, but he does not defend those further goals. They are just goals we happen to have. Alvin Goldman (1999) and Hilary Kornblith (1993) have each pursued the fourth option. They believe that we have an interest in believing the truth no matter what else we happen to want, because believing the truth will help us to get whatever else we happen to want.

As with the problem of identifying our cognitive goals, I do not intend to review all the virtues and vices of all the approaches to evaluating whatever cognitive goals we have. Instead, I will simply assert that (i) it is far from clear what the best way to approach this problem is, and (ii) none of the existing approaches (including my own stab at the second approach in Author [2001]) is an unqualified success. To answer the objection that naturalism cannot establish the well-motivated normativity of its theories, however, this problem must be solved.

Issue #4: Traditional Epistemological Concepts. The complaint that drives the normativity objection is that, when we stop doing normative epistemology, we are giving up on producing theories of knowledge, epistemic justification, and rationality. To stop producing theories of these things, it is claimed, is to stop producing epistemological theories altogether. To answer this complaint fully, naturalists must do more than just show that their theories can be normative. They must show that the hypothetically normative theories they produce, and the ends to which they are relativized, have some connection to the traditional concepts of knowledge, justification, and rationality that epistemology has addressed. Otherwise, they might be able to claim they have

produced normative, scientific theories about thinking, but they would not be able to claim they have produced normative, scientific, epistemological theories.

Naturalists thus seem to have the burden of showing that their theories are relevant not only to normative questions about how to think, but to the traditional epistemological questions about the nature of knowledge, justification, and rationality. Mere normativity -- even well-motivated normativity -- is not enough. As before, it is not at all clear how naturalists will discharge this burden.

One strategy is eliminativist. A naturalist could claim that our traditional epistemological concepts are based on folk psychology and the mistaken view that epistemology is independent of science. When psychology and epistemology become scientific, perhaps we will stop worrying about the traditional epistemological concepts as well.²⁰ If the eliminativist strategy can be successfully defended -- and it is far from clear that it can be -- naturalists will have excused themselves from having to take the traditional epistemological concepts into account.

Another strategy, however, involves extensive analysis or explication of the traditional concepts to show how the findings of science can help us to understand them. I take it that this is a major aspect of the first part of Goldman's Epistemology and Cognition. What is not clear, however, is whether this strategy is consistent with the naturalist's commitment to producing empirical, scientific theories of knowledge (or justification, or ...). It certainly seems that this strategy constitutes a return to the project of producing an a priori theory of knowledge, and then filling in some of the details with scientific results. One's theory of knowledge is still a priori, though it is supplemented by an empirical theory about knowledge.

Again, the relevance of naturalistic epistemology to our traditional epistemic concepts is something naturalists can and have addressed. But again, they have not been clearly successful.

This is, in my view, another open question for naturalized epistemology, and another question that must be answered if naturalists are to adequately address the problems of epistemic normativity.

There are certainly other issues advocates of the engineering reply would need to address. My point in mentioning these four has not been to undermine the reply, however. Instead, my aim has been to show that the engineering reply alone is enough to answer only some limited, relatively unsophisticated forms of the normativity objection. I have done that by elaborating the engineering reply in considerably more detail than naturalists ordinarily do, and by trying to identify some of the further work that needs to be done. I mean for this to be a contribution to the defense of naturalism, for one must often have a clear view of what needs to be done before one can expect to do it successfully.

Notes

- 1 Alvin Plantinga (1993), for example, calls it “warrant.”
- 2 The idea of using one’s cognitive resources “correctly” is also a matter of widespread dispute. Among the contenders are the views (i) that one has used one’s cognitive resources correctly with regard to a belief b iff b is in accord with one’s evidence (Feldman [2000]), (ii) that one has used one’s cognitive resources correctly with regard to b iff the cognitive mechanism that gave rise to b was in no way malfunctioning (and certain other conditions are met) (Plantinga 1993), (iii) that one has used one’s cognitive resources correctly with regard to b iff the processes whereby one has come to hold b are such as tend to produce true beliefs, either in the actual world or in worlds relevantly similar to the way we take the actual world to be (Goldman 1986), and (iv) that one has used one’s cognitive resources correctly with regard to b iff b is held in conformity with one’s

epistemic norms (Pollock 1986). Note that these are not all the contenders, but merely (what I take to be) a fairly representative sample.

3 This is far from the only view that could claim that name ‘epistemological naturalism’. Alvin Goldman’s naturalism, for example, is just the view that a priori epistemology “needs help” from the sciences to complete its task (Goldman 1986). What Jaegwon Kim (2000/1988) calls ‘psychologistic’ naturalism holds that science is relevant to epistemology just insofar as it is needed to identify the bases on which normative epistemic properties (such as justification) supervene. Still weaker versions of “naturalism,” often identified with David Hume (Strawson 1985), consist of only the rejection of Descartes’ method of radical doubt and the insistence that we can only theorize about knowledge -- or anything else -- from within the context of an existing body of belief and existing habits of inference.

4 For a much more detailed presentation of this objection, see Kim (2000/1988).

5 See McDowell (1996) and Brandom (1994) for extended discussions of the alleged independence of the Sellarsian “space of reasons” and the causal-nomological “realm of law.”

6 More sophisticated versions of the objection identify some particular kind or form of normativity that epistemological theories must address, and they contend that either science in general or the sciences of the mind are ill-suited to address that form of normativity. The engineering reply, I argue in Section 5, cannot dispatch all such versions of the objection as readily as it can dispatch the least sophisticated version.

7 According to the version of naturalism Alvin Goldman endorses, epistemology ‘needs help’ from the natural sciences, but it may not count as one of them (Goldman 1986). P1 might therefore be false of it. According to the version of reliabilist naturalism most often associated

with David Armstrong (1968), knowledge arises from the lawlike connections between states of affairs and one's beliefs that those states of affairs obtain. P2 might therefore be false of it.

8 The question whether evaluative sentences express cognitively contentful claims has been one of the most central issues of twentieth century metaethics, and I do not intend to answer it here. Nor is it necessary for my purposes that this question be answered. All parties to the debate would agree that someone who sincerely asserts that violets are pretty, for example, thereby expresses her liking for the way violets look. They disagree about whether she is also doing something else.

9 I intend the notion of sentences' 'ordinary uses'' to be a simplified version of Ruth Millikan's notion of sentences' 'proper functions'' (Millikan 1984). The differences between the two ideas do not make a difference to any of the points I want to make, and readers familiar with Millikan's work are free to substitute 'proper function' for 'ordinary use' in what follows.

10 One might worry that the idea of ordinary use (and that of linguistic proper function) is normative, and thus worry that the definitions of descriptiveness, evaluativeness, prescriptiveness, and normativity are all defective for containing normative vocabulary. Such worries are misguided, however. The point here is not to give a descriptive definition of these terms or to reduce the normative to the descriptive. Rather, the point here is just to give some clear formulation of what it is for a sentence to be descriptive and what it is for a sentence to be normative. If it turns out that "being descriptive" and "being normative" are both normative properties of sentences (whatever "normative properties" are), then so be it. I have never claimed that they are not.

11 As with the definitions of descriptiveness and normativity for sentences, it is not my intention to give a nonnormative characterization of "consequences" or a nonnormative definition of the

normativity and descriptiveness of theories. However, one might claim that commitment to claims and thus the consequence relation are not only normative but epistemic, and thus that I am illicitly applying normative epistemological notions as I try to explain how naturalized epistemology can accommodate epistemological normativity. I have three responses to this worry. First, though I admit I am taking the notion of commitment to claims as primitive, and though I admit that commitment might turn out to be normative, I do not see why this should be considered illicit. Nonnaturalists routinely make use of normative primitives as well. If using such a primitive means that naturalists cannot really accommodate epistemic normativity, then the same goes for nonnaturalists. Second, if the story I tell in Section 4 is on the right track, and if commitment does turn out to be normative, then naturalists will be in a position to explain its normativity in terms of instrumental efficacy with regard to our cognitive ends. They are thus not left with an unexplained normative explainer. Third, it is far from clear that we must consider commitment to be normative. A plausible first draft of an account of commitment would be to say that a person is committed to a claim that p iff p is more likely than Not- p given her other beliefs. Such an account would work for my purposes, and it does not seem to import any normative concepts.

12 Alternatively, we could relativize the definition of strong normativity to times. A theory would be weakly normative (for a person) at a time iff and only if it has normative consequences (for that person). A theory would be strongly normative (for a person) at a time iff it is being used (by that person) at that time to decide what stance to take toward normative sentences.

13 Strictly speaking, this is true of the Special Theory of Relativity only as used in theoretical physics. When an engineer uses the theory to help find answers to questions about how best to build something, the theory's normative consequences may come to the fore and the theory may become strongly normative.

14 In rough translation, the Latin slogan “Nihil in mente ...” means “There is nothing in the mind that is not first in the senses.”

15 There is more than one kind of magnetic resonance imaging used to study the brain. Here I am concerned exclusively with functional MRI, often called fMRI.

16 This is, of course, not to say that MRI is useless for those who are interested in the spiking output of a brain area. Presumably, one can glean some useful information about output from information about input, and thus MRI can indirectly provide information about output.

17 One such philosopher is Harvey Siegel (1990; 1996).

18 Larry Laudan (1984; 1987; 1990) is a naturalist who insists that hypothetical normativity is enough for epistemology or the philosophy of science. Hilary Kornblith is a naturalist who argues that epistemological norms are universal because they are relativized to cognitive goals everyone has or everyone ought to have (1993). I think it is best to construe him as holding that well-motivated hypothetical normativity -- not categorical normativity -- is enough for epistemology.

19 It might seem surprising that Quine would treat anything as a purely conceptual question to be answered a priori. And it is true that he believes on empirical grounds that prediction is “occasionally the purpose” of science though “howadays the overwhelming purposes of the science game are technology and understanding” (1992: 20). However, he does hold that methodological norms are relativized to the end of making accurate predictions, and he thinks of that as definitive of the “language game” of science (1992: 20).

20 I think the view Stephen Stich defends in The Fragmentation of Reason (1990) is easily interpreted as a radical version of this view. Stich thinks that once we understand how the mind really works, and once we get clearer about what epistemology is and has always been, we will

see that our traditional questions about the nature of knowledge, justified belief, and rationality are not worth answering.

References

- Armstrong, D. M. 1968. A materialist theory of the mind. Rev. ed. London: Routledge.
- Brandom, R. 1994. Making it explicit. Cambridge, MA: Harvard U P.
- Feldman, R. 2000. The ethics of belief. Philosophy and phenomenological research, **60** (3): pp. 667-95.
- Goldman, A. 1967. A causal theory of knowing. Journal of Philosophy. **64** (12): 355-372.
- 1986. Epistemology and cognition. Cambridge, MA: Harvard U P.
- 1992. Liaisons: Philosophy meets the cognitive and social sciences. Cambridge, MA: MIT Press.
- 1999. Knowledge in a social world. Oxford: Oxford U P.
- Haack, S. 1995. Evidence and inquiry. Oxford: Blackwell.
- Kim, J. 2000/1988. What is ‘naturalized epistemology’? In: Sosa, E. and Kim, J., eds., Epistemology: An anthology. pp. 301-13. Malden, MA: Blackwell.
- Kornblith, H. 1993. Epistemic normativity. Synthese. **94**: 357-76.
- Laudan, L. 1984. Science and values. Berkeley: U California Press.
- 1987. Progress or rationality?: The prospects for normative naturalism. American philosophical quarterly. **24**: 19-31.
- 1990. Normative naturalism. Philosophy of science. **57** (1): 44-59.
- Logothetis, N., Pauls, J., Augath, M., Trinath, T., Oeltermann, A. 2001. Neurophysiological investigation of the basis of the fMRI signal. Nature. **412** (12 July 2001): 150-157.
- Maffie, J. 1990. Naturalism and the normativity of epistemology. Philosophical studies. **59**: 333-49.

-----, 1995a. Naturalism, scientism, and the independence of epistemology. Erkenntnis. **43**: 1-27.

-----, 1995b. Towards an anthropology of epistemology. The Philosophical Forum. **26** (3):
218-41.

McDowell, J. 1996. Mind and world. Cambridge, MA: Harvard U P.

Millikan, R. 1984. Language, thought, and other biological categories. Cambridge, MA: MIT
Press.

Plantinga, A. 1993. Warrant and proper function. New York: Oxford U P.

Pollock, J. 1986. Contemporary theories of knowledge. Lanham, MD: Rowman and Littlefield.

Quine, W. V. 1969. Epistemology naturalized. In: Ontological relativity and other essays. pp.
69-90. New York: Columbia U P.

-----, 1992. Pursuit of truth. Rev. ed. Cambridge, MA: Harvard U P.

-----, 1998. Reply to Morton White. In: Hahn, L. and Schilpp, P., eds. The philosophy of W. V.
Quine. 2d ed. pp. 663-5. Chicago: Open Court.

Sellars, W. 1956. Empiricism and the philosophy of mind. In: Feigl, H. and Scriven, M., eds.
Minnesota studies in the philosophy of science. Vol. 2. pp. 225-308. Minneapolis: U Minnesota
Press.

Siegel, H. 1990. Laudan's normative naturalism. Studies in history and philosophy of science. **21**
(2): 295-313.

-----, 1996. Naturalism, instrumental rationality, and the normativity of epistemology.
Protosociology. **8/9**: 97-110.

Stich, S. 1990. The fragmentation of reason. Cambridge, MA: MIT Press.

Strawson, P. F. 1985. Skepticism and naturalism: some varieties. New York: Columbia U P.