

## Quine and Naturalized Epistemology

RICHARD FOLEY

**M**ovements to naturalize are dominant in almost every area of contemporary philosophy. In philosophy of mind, the dominant project is to show either that intentional attitudes are scientifically respectable or that they can be made so; in philosophy of language it is one of how to naturalize content; and in moral philosophy it is one of how to naturalize moral concepts.

In each of these areas, the idea is to make sure that our philosophical theories are compatible with science. Put roughly, and it may be that it cannot be put any other way, this means that in our philosophical theories we are to make use only of those properties that are either reducible to or supervene upon properties that science countenances.

Of course, as science countenances stranger and stranger properties, philosophy is correspondingly free to make use of them. But the guiding idea remains the same. Science constrains philosophy. Philosophical theories, on this view, must be naturalistic, and it is science, ultimately, that will tell us what this amounts to.

Not everyone accepts this view of philosophy, needless to say. Moreover, when the view is stated as vaguely as I have stated it, it is far from clear what its implications are. But for the discussion here, which will concern the naturalization movement in epistemology, I am not going to worry about the plausibility of the view, nor am I even going to try to identify its implications. It is not necessary to do so, since naturalized epistemology cannot be understood as just one more instance of the movement in philosophy to naturalize. In particular, it cannot be understood as an attempt to naturalize the concept of justified belief and other concepts that are central to epistemology. Almost all contemporary epistemologists are engaged in this project, and yet many of these epistemologists are not thought of as doing naturalized epistemology. So, this way of thinking about naturalized epistemology does not succeed in capturing what is supposed to be distinctive about it.

For example, coherentists are not ordinarily regarded as doing naturalized epistemology. Yet they understand epistemic justification in terms of such notions as belief, consistency, and simplicity, and these notions, most would agree, are compatible with a scientific picture of world. Or consider Bayesians who understand epistemic rationality in terms of probabilistic coherence among one's degrees of beliefs. Once again, most would agree that these notions either are or can be made scientifically respectable, and yet ordinarily Bayesians are not thought of as doing naturalized epistemology. Or consider Roderick Chisholm, who is usually thought to be as distant from the naturalized tradition as any contemporary epistemologist. The notion of epistemic justification, says Chisholm, is ultimately to be understood in terms of an ethical requirement to prefer. To say you are more justified in believing a proposition than withholding on it is to say, roughly, that you have a *prima facie* obligation to prefer believing it to withholding on it. But Chisholm then goes on to say that these ethical requirements supervene on non-normative states, most importantly your beliefs and your sense experiences.<sup>1</sup> So, if the movement to naturalize epistemology is thought of simply as a movement that aims to show how epistemic properties are scientifically respectable, then even Chisholm's epistemology is a naturalistic one.

Admittedly, there are those who argue that the notions of belief and sense experience cannot be made compatible with a scientific picture of the world and that, thus, neither the coherentist nor the Bayesian nor Chisholm has succeeded in formulating an acceptably naturalistic epistemology. But in order to be a naturalized epistemologist, one does not have to abandon the notions of belief and sense experience. Indeed, very few of the epistemologists who are most closely associated with naturalized epistemology do so. Think, for example, of W. V. Quine, Fred Dretske, D. M. Armstrong, Alvin Goldman, and Donald Campbell, none of whom altogether discard the notion of belief or sense experience.

But the issue of whose epistemology is, or is not, scientifically respectable is a side issue. The main point I am making is that the program to naturalize epistemology aims at something stronger than merely ensuring that epistemology is compatible with science. It is interested in proposing a stronger, more positive role for science in shaping epistemology. The question is, what exactly is this stronger role?

A number of philosophers have attempted to answer this question. Among the most interesting attempts are those of Jaegwon Kim, Alvin Goldman, Philip Kitcher, Hilary Kornblith, and Hilary Putnam.<sup>2</sup> There is much to be learned from these attempts, but it is also worth noting that these philosophers do not agree on which epistemologists are doing naturalized epistemology, much less on what naturalized epistemology is. I think that there is an explanation for this other than the standard one of philosophers having a special talent for disagreement. Namely, those who see themselves as practicing naturalized epistemology are in fact engaged in different kinds of projects. Thus, it is not surprising that we have trouble agreeing on who is doing naturalized

epistemology, and likewise it is not surprising that we have trouble finding any one thesis or set of theses that definitively characterizes it.

In any event, I am not going to try to provide a general characterization of naturalized epistemology or even try to say who is doing it. I am going to look only at Quine's view of epistemology. After all, if anyone is a naturalized epistemologist, it had better be the case that he is.

Unfortunately, in addition to being one of our most important philosophers, Quine is also one of our most exasperating philosophers. He is legitimately regarded as a great stylist, but in spite of this, indeed in part because of it, he is often difficult to decipher. When faced with the choice of expressing his views with great elegance or maximum clarity, he rarely chooses the latter. And nowhere is his penchant for the stunning phrase more on display than in his epistemology. The result, over the years, is a collection of maxims, mottos, and pronouncements that are none too easy to fit into a coherent whole.

Of course, Quine's broad themes are clear enough. It is clear, for example, that he rejects the analytic-synthetic distinction. This in turn has implications for his epistemology, the most important of which is his rejection of the idea that epistemology can be done *a priori*. But problems begin to crop up when we try to go beyond these familiar Quinean themes to see how his more specific recommendations fit with one another. Indeed, I do not think that they do fit. More specifically, I think that some of Quine's most famous pronouncements on epistemology conflict with what seem to be his settled views.

My aim, however, is not so much to reveal conflicts among the various things Quine has said about epistemology. Nor is it even to evaluate the plausibility of what he has said. Primarily, I just want to get as clear as possible about his epistemology and to get clear also about whether he is doing epistemology in some new and distinctively naturalized way.

The best place to begin is "Epistemology Naturalized." There Quine gives his most famous characterization of naturalized epistemology. He says that its defining mark is that it makes epistemology into a chapter of psychology.

Epistemology, 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 external 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 related to theory, and in what ways one's theory of nature transcends any available evidence.<sup>3</sup>

If this passage is taken at face value, Quine is suggesting that we altogether replace epistemology with a psychology. More specifically, we are to replace it with the project of describing the relationship between our sensory

input and our theoretical output. Within this project, questions of justifying the theoretical output do not arise.

The above passage occurs in the context of a discussion in which Quine argues that a certain kind of project in epistemology has failed—the project of validating science on the basis of sense experience, where validation means proof by logic and set theory. Since theory is underdetermined, it is not possible to validate science. Still, in itself this does not force Quine to give up questions of justification unless he thinks that validation is the only kind of justification worth having. But why would he think that? He never explicitly endorses this position, and he certainly does not give an argument for it.

Nevertheless, the passage suggests that Quine does want to give up epistemology. It suggests that he is an eliminativist when it comes to epistemic justification. Apparently, his idea is not to naturalize this notion. Rather, the idea is to jettison it as unuseful. Or perhaps more accurately, the recommendation is to jettison it as part of the scientific picture of world. Justification has no place in the science game, it seems.

If there is a standard interpretation of Quine's view, this is it. For example, Jaegwon Kim, Alvin Goldman, and Hilary Kornblith all insist that what is distinctive about Quine's view, and for that matter what is also most objectionable about it, is that he is asking us to give up the notion of epistemic justification and with it the normative element in epistemology. Here are some representative passages, the first from Kim, the second from Goldman, and the third from Kornblith:

. . . it is normativity that Quine is asking us to repudiate. Although Quine does not explicitly characterize traditional epistemology as "normative" or "prescriptive", his meaning is unmistakable. Epistemology is to be "a chapter of psychology", a law-based predictive-explanatory theory, like any other theory within empirical science; its principal job is to see how human cognizers develop theories (their "picture of the world") from observation ("the stimulation of their sensory receptors"). Epistemology is to go out of the business of justification. . . . Quine is urging us to replace a normative theory of cognition with a descriptive science.<sup>4</sup>

. . . on W. V. Quine's naturalistic conception, the epistemologist would study how the human subject responds to certain input; how, in response to various stimulus patterns, the subject delivers a description of the external world and its history. In studying the relation between this "meager input" and "torrential output," epistemology "simply falls into place as a chapter of psychology and hence of natural science." . . . But this approach, though perfectly tenable, neglects the evaluative strain pervading most of historical epistemology. Epistemologists have traditionally been interested in whether beliefs about the world are justified or warranted; whether we are rationally entitled to these beliefs. Epistemologists seek to discover or invent proper methods of inquiry

and investigation, often dismissing established procedures as irrational. Clearly, 'justified', 'warranted', and 'rational' are evaluative terms; and the advocacy of particular methods is a normative activity. So traditional epistemology has a strong evaluative-normative strain. I aim to preserve that strain.<sup>5</sup>

. . . according to Quine, foundationalists were asking the wrong question. Once we see the sterility of the foundationalist program, we see that the only genuine questions there are to ask about the relation between theory and evidence and about the acquisition of belief are psychological questions. In this view question 2 [How do we arrive at our beliefs?] is relevant to question 1 [How ought we arrive at our beliefs?] because it holds all the content that is left in question 1. The relation between these two questions is much like the relation atheists believe to hold between questions about God's act of creation and questions about the details of, for example, the big bang; the latter questions exhaust all the content there is in the former questions.<sup>6</sup>

I do not think that this standard view of Quine will ultimately do, but undeniably it is a natural interpretation. The above passage from "Epistemology Naturalized" seems not just to suggest this interpretation but to demand it. Moreover, there are numerous other passages in Quine's work that seem to support the standard interpretation. For example, from *Theories and Things*:

Naturalism does not repudiate epistemology, but assimilates it to empirical psychology. Science itself tells us that our information about the world is limited to irritations of our surfaces, and then the epistemological question is in turn a question within science, the question how we human animals can have managed to arrive at science from such limited information. Our scientific epistemologist pursues this inquiry and comes out with an account which has a good deal to do with the learning of language and the neurology of perception. . . . Evolution and natural selection will doubtless figure in this account, and he will be free to apply physics if he sees a way.<sup>7</sup>

And here is another passage from "Epistemology Naturalized":

But why all this creative reconstruction, all this make-believe? The stimulation of his sensory receptors is all the evidence anybody has had to go on, ultimately, in arriving at his picture of the world. Why not just see how this construction really proceeds? Why not settle for psychology? Such a surrender of the epistemological burden to psychology is a move that was disallowed in earlier times as circular reasoning. If the epistemologist's goal is validation of the grounds of empirical science, he defeats his purpose by using psychology or other empirical science in the validation. However, such scruples against circularity have little point once we have stopped dreaming of deducing science from observations. If we are out

simply to understand the link between observations and science, we are all well advised to use any available information, including that provided by the very science whose link with observation we are seeking to understand.<sup>8</sup>

Despite the seemingly unambiguous intent of these passages, there are two devastating problems for the standard interpretation. The first is that Quine has explicitly rejected it. In his most recent defense of naturalized epistemology, he says he has no quarrel with those critics who object to his retaining the word 'epistemology', since he agrees that the "repudiation of the Cartesian dream is no minor deviation." But then he immediately adds that "they are wrong in protesting that the normative element, so characteristic of epistemology, goes by the board."<sup>9</sup> The matter-of-fact tone with which he announces this looks a little disingenuous in light of the above quotations, but nevertheless he does deny that this was ever his view.

The second devastating problem for the standard view is that at approximately the same time Quine was writing "Epistemology Naturalized," he co-authored *The Web of Belief* with Joseph Ullian, and this book is paradigmatically a work of normative epistemology. Here is a passage from the early pages of the first edition:

The story of the origins and intensities of our beliefs, the story of what happens in our heads, is a very different story from the one sought in our quest for evidence. Where we are rational in our beliefs the stories may correspond; elsewhere they may diverge. The former story is for psychology to tell. On the other hand, our present concern is with grounds, with reasons, with the evidential relations that hold among beliefs.<sup>10</sup>

There is nothing in this passage to suggest that Quine is jettisoning the normative element in epistemology, and likewise nothing to suggest that epistemology is merely a chapter in psychology. On the contrary, he is distinguishing the task of the epistemologist from that of the psychologist and moreover distinguishing them in an altogether traditional way. The psychologist is to tell us what happens in our heads that results in our believing what we do, while the epistemologist is to tell us what we have reasons to believe. And it is the latter project that Quine and Ullian are pursuing in *The Web of Belief*.

So, the standard interpretation will not do. Quine's settled view is not one that simply turns epistemology into a chapter of psychology. Indeed, his own practice illustrates that he thinks that there is something for epistemologists to do other than merely describe the relation between experiential input and theoretical output. They can also be concerned, as always, with what it is rational for us to believe as opposed to what we actually do believe. This is what Quine himself is concerned with in *The Web of Belief*.

Why, then, does Quine say that epistemology is to be a chapter of psychology? Part of the answer, no doubt, is Quine's fondness for shocking aphorisms. But it is not just that. He sees epistemology as continuous with science in

general. Indeed, it is not an exaggeration to say that he sees epistemology as a part of science, or if you prefer, a chapter of science. On the other hand, it is an exaggeration to say, as he does, that epistemology is a chapter of one particular science, viz., psychology. This might seem to be a minor distinction, but in fact it is far from minor. It is crucial for understanding how epistemology can be normative in Quine's view.

To explain why, I need first to discuss Quine's view of normativity in epistemology. Then I will return to the distinction between epistemology as a chapter of science in general and epistemology as a chapter of psychology.

Quine thinks that the normative element in epistemology is ultimately a matter of identifying effective means to a valued end, where in epistemology the relevant valued end is truth, or more cautiously, accurate predictions. He never defends the idea that the ability to predict observations has value. He apparently takes it as obvious that it does, and thus he also takes it as obvious that whatever helps us to make more accurate predictions likewise has value—instrumental value. So, put crudely, the view is that we are justified in using a certain method insofar as it helps us to generate theories that accurately predict. In effect, this means that the norms operative within the domain of epistemology are ones with which we can “engineer” our way to accurate theories. But then, insofar as science gives us information about which methods are reliable and which are not, it is providing us with the information we need to solve this engineering problem. In so doing, science itself has normative import.

Here are a couple of relevant passages:

There is no question here [in epistemology] 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 has been expressed.<sup>11</sup>

. . . normative epistemology gets naturalized into a chapter of engineering: the technology of anticipating sensory stimulation.<sup>12</sup>

These passages, which emphasize the normative element of epistemology, contrast with those passages in “Epistemology Naturalized” where Quine seems to be recommending that we replace epistemology with psychology. But it is also important to remember that in that same article, Quine deplores what he takes to be the epistemic nihilism of Kuhn, Polanyi, and Hanson. He seems to be complaining that these philosophers repudiate traditional epistemology and its normative notion of justification without replacing it with anything else. By contrast he sees himself as extracting and thereby salvaging the scientifically respectable part of normative epistemology—what he later calls the ‘engineering’ part.

But exactly where and how does this normative element, understood as the engineering of accurate theories, enter into Quine's epistemology? I do not think there is an easy answer to this question, but Quine's most recent words on epistemology may help provide at least the beginnings of an answer. On the

very first page of *Pursuit of Truth*, Quine says that in looking at the relation between sensory stimulation and scientific theory, we can separate out some things to do without paying much attention to neurology or psychology.

Within this baffling tangle of relations between our sensory stimulation and our scientific theory of the world, there is a segment that we can gratefully separate out and clarify without pursuing neurology, psychology, psycho-linguistics, genetics, or history. It is the part where theory is tested by prediction. It is the relation of evidential support.<sup>13</sup>

This sounds exactly like traditional epistemology. It sounds as if Quine is going to tell us about the relation of evidential support, and he is going to do so from his armchair, without relying on the findings of science. In other words, it sounds as if he is going to be doing epistemology in an *a priori* manner. He would no doubt insist that it is not really *a priori*. On his view, even armchair theorizing draws upon a large background of empirical information. So, it is not truly *a priori* even if it is not truly neurology, psychology, psycho-linguistics, genetics, or history either. But it is the Quinean counterpart of *a priori* theorizing.

In any event, after having announced that he is going to tell us about the relation of evidential support, he never does. He says nothing at all about this relation, although he does spend quite a bit of time on the notion of evidence, or at least its surrogate within the Quinean system, observation.<sup>14</sup>

Quine's struggles with the notions of observation and observation sentences are legendary, but for purposes here it is another question that is more pressing. Namely, why is it, after telling us that we can clarify the relation of evidential support without doing science, that Quine does not tell us anything about how observation evidentially supports theory?

There is a simple answer to this question. There is nothing to tell. There is nothing to tell because in Quine's system observation does not support theory; it only tests theory. Or expressed another way, observation is capable of generating only negative evidence. Quine is a Popperian: strictly speaking, observations can only refute theories; they cannot support them.<sup>15</sup>

The qualification 'strictly speaking' is needed because Quine need not deny that within a theory there are relations of positive evidential support. He can say, for example, that the litmus turning pink is evidence for, or confirms, that the liquid is an acid. But this sense of confirmation is theory-relative for Quine. The pinkness confirms acidity relative to our best overall theory of the world, and that theory is not confirmed by our observations; it is only tested by them.

It is clearly true, moreover, that one continually reasons not only in refutation of hypotheses but in support of them. This, however, is a matter of arguing logically or probabilistically from other beliefs already held. . . . Some of those supporting beliefs may be observational, they contribute support only in company with others that are theoretical. Pure observation

lends only negative evidence, by refuting an observation categorical that a proposed theory implies.<sup>16</sup>

So, according to Quine, there is no theory-independent relation of positive evidential support between observation and theory. Accordingly, one cannot appeal to such a relation in order to introduce a normative element into epistemology. For example, the normative element cannot be a matter of having an obligation to believe only those theories for which we have adequate, positive empirical evidence. Locke thought this was the way in which epistemology became normative. In *An Essay Concerning Human Understanding*, he talks of our duty as rational creatures to conform our beliefs to the evidence.<sup>17</sup> Moreover, the view also has distinguished contemporary adherents, e.g. Laurence Bonjour and Roderick Chisholm.<sup>18</sup>

But if this is not how the normative element enters into Quine's epistemology, how and where does it enter? One natural answer is that it enters with the testing of theories. After all, Quine does think that theories can be tested by observations even if they are not positively supported by them, and antecedently one would think that the testing of theories must involve normative elements.

But surprisingly, this is not Quine's view. He is explicit that it is not in the testing of theories that the normative element enters into epistemology. Here is the crucial passage:

I have been treating the testing of a theory after it has been thought up, this being where the truth conditions and empirical content lie; so I have passed over the thinking up, which is where the normative considerations come in.<sup>19</sup>

And shortly after this passage, Quine adds:

. . . when I cite prediction as the checkpoints of science, I do not see that as normative. I see it as defining a particular language game, in Wittgenstein's phrase: the game of science, in contrast to other good language games such as fiction and poetry.<sup>20</sup>

I say that this view is surprising, because there is a tradition in epistemology and philosophy of science that insists that the distinctively normative element of epistemology is not concerned with the context of discovery but rather the context of justification. In effect, Quine reverses this, although he prefers to talk not of the context of justification and the context of discovery but rather of the process of testing theories and the process of thinking them up. But he does say that the former is not normative whereas the latter is.

More specifically, Quine's view is that insofar as we are doing science, we have to make predictions that are testable. Moreover, if the predictions are not as we say they will be, something has to be revised. There is no choice in this and nothing at all normative about it. It is just part of the science game that when things do not turn out as we have predicted, something has to give. There is, of course, a normative issue as to whether to do science at all, but once we

have agreed to do science, it is the constitutive rule of the game that something has to give when our predictions go wrong. Normative considerations come in with respect to what is to give. We have choices about what to revise when our predictions do not come out as we expect. It is here that there is room for norms. The norms govern what Quine calls 'the thinking up of theories', only it is important to realize that as he uses the phrase, the thinking up of theories can be a matter of either inventing altogether new theories or revising existing ones in the face of unsuccessful predictions.

Moreover, Quine identifies the normative considerations that in his judgment should govern this process of thinking up new theories and revising old ones. In *The Web of Belief*, Quine and Ullian listed five virtues to seek in a new or revised hypothesis: conservatism, generality, simplicity, refutability, and modesty. In his latest work, *Pursuit of Truth*, Quine has pared the list down to two fundamental norms: the maximization of simplicity of our hypotheses and the minimalization of the mutilation of old hypotheses.<sup>21</sup>

Of course, often enough these two overarching norms will be at odds with one another, in which case they must be balanced in some appropriate way. Quine tells us nothing about how to go about doing this, even though his view commits him to regarding this as one of the most fundamental normative issues in epistemology. Perhaps he thinks that there just is not anything specific to say about how an appropriate weighting can be achieved. We just have to muck around and do the best we can.

By way of contrast, Quine does think that at a lower level it is possible to generate specific advice about how to conduct our theorizing, but to do so we need to draw upon the detailed findings of our various sciences. We need these findings, because, as I pointed out before, Quine thinks that the norms that govern our theorizing are engineering norms, and in engineering, we are free, indeed required, to draw upon the findings of science. Here is a relevant passage, a portion of which I cited earlier:

Naturalization of epistemology does not jettison the normative and settle for indiscriminate description of the on-going process. For me normative epistemology is a branch of engineering. It is the technology of truth-seeking, or, in more cautiously epistemic terms, prediction. Like any technology, it makes free use of whatever scientific finding 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 exploring 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 and 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 has been expressed.<sup>22</sup>

So, if we are looking for specific intellectual advice, we must look at mathematics, cognitive psychology, neurology, statistics, and any other relevant science. Notice that the point is not that these are the places to look if what we want are philosophical criteria for justified belief or rational belief. The point, rather, is that these are places to find one of the other things that epistemologists have traditionally wanted—namely, concrete advice about how to get at the truth, or more cautiously accurate predictions.

Notice also that psychology is just one of the sciences that we look to for advice. If naturalized epistemology were simply in the business of explaining how we get from observation to theory, then psychology would be the only, or at least the principally, relevant science. But naturalized epistemology, Quine insists, is normative, and precisely because of this psychology has no privileged position with respect to it. Epistemology is not simply a chapter of psychology. Any science that is relevant to the technology of truth seeking is *ipso facto* relevant for epistemology.<sup>23</sup>

This link between epistemology and science explains why Quine thinks he is not jettisoning the normative in epistemology, but on the other hand it cannot be this that makes his epistemology distinctively naturalistic, since it is not even controversial that science is linked in this way with epistemology. Neither Descartes nor Locke nor any other non-naturalized epistemologist need deny that the particular sciences are capable of providing us with information that is relevant to what Quine calls 'the technology of truth seeking'. They would have no trouble agreeing, for example, that probability theory can warn us about the gambler's fallacy, psychology about perceptual illusions, and physics about testimony from the occult. What they will deny is that science is capable of telling us anything interesting about the most fundamental epistemological norms—for example, the norms that guide us in doing science itself.

This suggests that if we are to find something distinctively naturalized in Quine's epistemology, we must look at the status of the most general epistemological norms within his system. According to Quine, these norms are ones that tell us to maximize simplicity and minimize mutilation in thinking up theories. So, the crucial question to ask is, what on Quine's view makes these norms the correct ones?

Regrettably, Quine never says. In fact, he never even tries to cite considerations in favor of these norms. He simply posits them from his philosophical armchair, much as Descartes posited his method of doubt and Locke his way of ideas from their armchairs. The appearance, in other words, is that at the most basic level, Quine's way of doing epistemology is not different in kind from the traditional way of doing epistemology. He is proposing general norms of inquiry and he seems to be doing so *a priori*.

But of course, this cannot be right. If Quine stands for anything, he stands for the rejection of *a priori* epistemology. So, there is a puzzle here. Quine's official pronouncements on epistemology make it sound as if he sees himself doing epistemology in a radically new way, and yet his actual practice in doing normative epistemology looks altogether familiar.

One way of solving the puzzle would be to insist that Quine is merely saying that when scientists revise their theories and think up new ones, they in fact try to maximize simplicity and minimize mutilation. This reading has the advantage of making Quine's claim about simplicity and mutilation straightforwardly empirical. Of course, it also makes it incumbent upon him to cite some empirical evidence indicating that scientists do act in this way. However, he neither cites such evidence nor even displays much interest in empirical studies of science. Besides, on this interpretation, his epistemology would once again lose its normative force. It would be simply reporting how scientists go about their theorizing, while remaining neutral on whether this is a desirable way to theorize. But as we have seen, Quine sees himself as doing normative epistemology.

How, then, are we to interpret Quine's recommendation that we maximize simplicity and minimize mutilation? The only interpretation consistent with Quine's rejection of *a priori* epistemology is one that has the recommended norms as themselves being engineering norms. They too are part of the technology of truth seeking, and as such they are the products of science, broadly construed. To be sure, Quine never tries to show that this is so. He simply posits the norms without defense. Nonetheless, he has no choice but to regard them as conclusions of science, albeit ones with normative import for the technology of truth seeking.

A good way to see how this might work within the Quinean system is to look at what he says about another norm. In addition to the norms that tells us to maximize simplicity and minimize mutilation, Quine also defends an empiricist norm, which he identifies using the Latin phrase: *nihil in mente quod non primus in sensu*. He also expresses the norm in a less refined way by putting the motto of the Sherwin-Williams paint company to his own use: save the surface and you save all. Saving the surface is here equated with saving observation sentences.

According to Quine, empiricism is a position that tells us where to look for the content and truth conditions of our theories—we are to look at observation sentences. But in addition, Quine thinks of empiricism as a normative position that tells us to do science insofar as we are interested in truth. What this means, most generally expressed, is that we are to formulate hypotheses that are empirically testable, and if these hypotheses together with our other beliefs generate faulty predictions, we are to change something. The overarching norms then tell us how to make the changes. We are to maximize simplicity and minimize mutilation.

So, in addition to the simplicity and non-mutilation norms, Quine also endorses an empiricist norm, which he conceives to be identical with the scientific method broadly conceived. This then is one of the senses in which Quine's epistemology can be regarded as a scientific epistemology: the norms of rational belief, as he sees it, just are the norms of science broadly conceived.

There are many questions that can be raised about this view, chief among them are whether empiricism can be plausibly identified with the scientific

method, even broadly construed, and whether empiricism, so conceived, is really a plausible candidate to be a fundamental norm of rational belief. But I am going to glide over these questions, since my main concern is not so much the content of Quine's recommendations as his way of doing epistemology. The objective is to identify something distinctively naturalistic about his way of doing epistemology, and this objective is not met by pointing out that on Quine's view the norms of rational belief are identical with the norms of science broadly conceived. This only shows that the content of Quine's normative recommendations is different from the content of, say, Descartes's and Chisholm's recommendations, not that his way of doing epistemology is distinct in some interesting way. The interesting conception is one that makes his epistemology a part of science, but this requires that even his most fundamental norms—the empiricist norm as well as the simplicity/non-mutilation norms—be products of science. We have seen that Quine thinks that psychology, statistics, physics, and other sciences provide us with information that we can use to fashion a technology of truth seeking. In so doing, they provide us with lower-level epistemic norms. What is not yet so clear is how psychology, statistics, physics, or any of the other sciences could generate the most fundamental norms. It is not clear, in other words, in what sense these norms are engineering norms.

As a way of trying to get an answer to this question, let's ask another one—how, if at all, are these norms revisable? This question is a pressing one given Quine's rejection of the analytic-synthetic distinction and his endorsement of epistemological holism. These positions may not quite entail that all of our opinions are in principle revisable, but they do make the view hard to resist. And of course, it is a view that Quine accepts. Thus, it had better be the case the most fundamental parts of his epistemology are revisable.

And indeed, Quine says precisely this about the empiricist norm. He says the norm is itself a finding of science and hence revisable. In particular, it is a finding of science that our information about the world comes from our five senses, but this point, he adds, is normative, since it warns us against any purported source of information that is not rooted in observations. It warns against telepathers and soothsayers, for instance.<sup>24</sup>

But can it really be the case that empiricism is a finding of science? Science, in Quine's view, is defined by its empirical methods. It is essentially a matter of constructing theories that can be refuted by observations. But then, isn't empiricism a presupposition of science rather than a finding of it?

Quine does not explicitly address this question, but it is clear what he should say. He should say that empiricism is both a presupposition and a finding of science. It is a presupposition in that science is defined, on Quine's view, by its empirical methods, but it is also a finding of science in that science, according to Quine, tells us that our most reliable information about the world comes to us through our senses. It tells us that we are not telepathic and that we are not clairvoyant and that we do not have any other reliable, non-observational access to the world.

But isn't it inevitably question begging to make use of empirical methods to argue that empirical methods are our most reliable ways of gaining information? Perhaps, but we should expect question begging when the issue concerns our most fundamental methods of inquiry. If a method is not fundamental, it can be defended employing methods that are more fundamental than it. But epistemology is most centrally concerned with our fundamental methods of inquiry, and our fundamental methods, if they are to be defended at all, must be defended using those very methods. We can say that this begs the questions, if we so wish, but some questions need to be begged, if they are to be answered at all, and questions about the reliability of our most fundamental methods of inquiry are just such questions.

Besides, it is no trivial matter for a method to beg the question in its own defense. Not every proposed method in every situation will be able to do so, since it is possible for a method to generate evidence that undermines its own reliability. Indeed, it is just this that accounts for the revisability of the empirical norm, in Quine's eyes. To be sure, he thinks that at the current, relatively advanced stage of inquiry, it would take an extraordinary turn of events to convince us (or at least most of us) that we have a faculty of clairvoyance and that this faculty is a more reliable source of information about the world than our eyes, ears, and other senses. But it is at least thinkable for things to go this way, and if they did, Quine admits, even the empiricist norm could be rejected. This norm is derived from science, and like the rest of science it is fallible. This is the price of Quine's rejection of epistemology as a first philosophy, and it is a price that Quine thinks we have no choice but to pay. There simply is no helpful *a priori* intellectual advice to be given.<sup>25</sup>

Having said this, it is also important to see that Quine is proposing something that is very, very close to an *a priori* epistemology, a kind of Quinean counterpart to it. To be sure, Quine is clear about his rejection of *a priori* epistemology. For him there is no difference in kind between the truths sought in philosophy (and hence epistemology) and those sought in science. But he is also well known for pointing out that some opinions are more central to the web of our beliefs than others and that we will be more reluctant to give up these beliefs. Since they are so central, we will and should do everything we can to protect them. It is this that gives some of these beliefs the appearance of being necessarily true. We shield such beliefs from revision by exercising our freedom to reject other beliefs in the web instead, and this creates the appearance of necessity.<sup>26</sup> We do this in mathematics, and we can do it in epistemology as well. We can shield the opinion that undergirds the empiricist norm—viz., the opinion that insofar as we are interested in truth, it is best for us to look for hypotheses that can be tested empirically—and this can create the appearance that we are treating it as a necessary truth that we know *a priori*.

As a devout empiricist, this is precisely the appearance that Quine creates. He talks about the empiricist norm almost as if it were a necessary truth known *a priori* and that as such it is beyond dispute. In turn, this creates the appearance that he is doing epistemology pretty much as usual. Indeed, for all

practical purposes he *is* doing epistemology as usual. In practice, he is no more concerned with the empirical support for his empiricist norm than Descartes was concerned with the empirical support for his method of doubt. However, unlike Descartes and unlike a good many other traditional epistemologists, Quine is committed to the view that this norm, and the belief upon which it is based, are revisable. But this really is not much more than a bare possibility. The sense in which they are revisable is the same as the sense in which elementary truths of mathematics are revisable, which is just to say that for all practical purposes, they might as well not be revisable at all.

Exactly the same must be said about Quine's other fundamental norms, the ones that tell us to maximize simplicity and minimize mutilation. Although Quine never explicitly says so, they too had better be based on findings of science and hence revisable. In particular, they must correspond to an empirical claim about the technology of truth seeking—viz., that a policy of maximizing simplicity and minimizing mutilation helps us in our engineering task of constructing theories that accurately predict observations. Quine apparently takes it as utterly obvious that this is the case. So, he never tries to marshal empirical evidence in defense of the claim. And no doubt he would be prepared to shield it from revision at almost any cost, thus creating the appearance of necessity. But again, he is committed to its revisability in principle, even if this does not have much effect on his practice of epistemology.

This, then, is the overall structure of Quine's epistemology. The various norms he endorses correspond to various beliefs he has, and he hopes that the rest of us have as well, about the technology of truth seeking. Since these beliefs are revisable, the norms that are derived from them are also revisable. However, we will exhibit different degrees of stubbornness with respect to different norms.

It is relatively easy to imagine circumstances in which we would give up norms that are derived from specific findings in the sciences; we need only imagine circumstances in which we would abandon or revise the findings on which they are based. For example, currently our best theory of dreaming takes REM as the most reliable indicator of dream activity, and our epistemic norms reflect this theory. We tend to discount first-person negative testimony about dreaming; even if in the morning I do not remember dreaming, you are entitled to conclude that I dreamed if I displayed REM during the night. Even so, it is not all that difficult to imagine dream research developing in such a way that our best future theory will imply that REM produces dreams only in conjunction with some other factor F, and, surprisingly, subjects almost always remember dreaming when REM + F are present. If so, our epistemic norms will be correspondingly revised; first-person negative testimony about dreaming will be taken more seriously.

By contrast, it is much harder to imagine our giving up the empiricist norm, since on Quine's view it is constitutive of the scientific method as we know it. Still, he is committed to the position that science could generate evidence that undermines science as we know it. More accurately, it could

generate evidence that make the costs of hanging on to science unacceptably high, where the costs here would be a matter of decreased simplicity and increased mutilation of our web of beliefs.

Similarly, Quine is also committed to saying that the policy of maximizing simplicity and minimizing mutilation might generate evidence that undermined this very policy. Again, it would take a remarkable turn of events for this to happen, but it is at least thinkable that in maximizing simplicity and minimizing mutilation we worsen, or at least do not help, our abilities to predict observations. Hence, it is at least thinkable that we will come to believe that moderately unstable and complex belief systems serve us as well or better in predicting the course of observations as maximally stable and simple ones. Indeed, in the worst case scenario, we might become convinced that every available way of conducting inquiry, whether it be one that emphasizes simplicity and non-mutilation or any other, will be met with widespread predictive failures, in which case we might well resign ourselves to skepticism. If so, then no way of conducting inquiry would be rational for us.

So, where does this leave us in our attempt to identify a sense in which Quine is a naturalized epistemologist, whereas Descartes, Locke, Carnap, and Chisholm are not. It leaves us, I think, with two possible ways of doing so, although one, I have argued, is not especially interesting.

First, the uninteresting way: within Quine's system, the most fundamental norms are identical, as Quine sees it, with the norms of natural science. So, the canons of rational belief just are the canons of science, broadly conceived. But all this means is that the content of Quine's recommendations is different from other epistemologists, not that his way of doing epistemology is different.

The second and more interesting way: Quine rejects the analytic-synthetic distinction and with it any distinction between *a priori* knowledge and *a posteriori* knowledge and any distinction between necessary and contingent truths. So, the fundamental epistemic norms in his system cannot be known *a priori*, nor are they necessary. Rather, they are continuous with science. They are part of our overall theory of the world, and like any other part of that theory, they are revisable in principle. Any appearance to the contrary is to be explained by our freedom to make revisions elsewhere in our web of belief when things go wrong, thus shielding these norms from revision.

Still, what is remarkable is how little difference this second point makes to Quine's practice of epistemology. Admittedly, it does make some difference. The rejection of the analytic-synthetic distinction leads Quine to endorse holism, and one implication of his holism is that the primary objects of epistemic justification are not individual beliefs but rather our overall theory of the world—in effect, our overall web of beliefs. So, this becomes the focus of his epistemology. But once it is settled that the primary epistemological norms are to be ones that govern the revision of our overall theory, Quine proceeds in the usual manner. The specific norms Quine favors are ones that he recommends from his philosophical armchair, with little or no concern for an empirical defense of them.

So, Quine's practice as a naturalized epistemologist is not discernibly different from that of any other famous epistemologist in the history of the subject. His rejection of the analytic-synthetic distinction does make for a significant difference when we are thinking about the metaphysical status of his claims, but it has little relevance when Quine is actually doing epistemology.

An interesting final question is whether something analogous is true of other philosophers who are regarded as doing naturalized epistemology. I think that it is. More specifically, I think that insofar as their concerns are normative and fundamental, their practice, like Quine's, will not be discernibly different from that of other epistemologists in the history of the subject. But to establish this would require a case by case examination of their views. So, for the time being, I will be content with limiting my claim to Quine, who is widely recognized as the founder of naturalized epistemology: he is not doing epistemology in a fundamentally new way.

#### NOTES

1. Roderick Chisholm, *Theory of Knowledge*, 3rd ed. (Englewood Cliffs, N.J., 1989).
2. Jaegwon Kim, "What Is Naturalized Epistemology?" in *Philosophical Perspectives*, edited by J. Tomberlin vol. 2 (Atascadero, Calif., 1988), 381-405; Alvin Goldman, *Epistemology and Cognition* (Cambridge, Mass., 1986); Philip Kitcher, "The Naturalists Return," *Philosophical Review* 101 (January 1992): 53-114; Hilary Kornblith, "Introduction: What Is Naturalized Epistemology?" in *Naturalizing Epistemology*, edited by Hilary Kornblith (Cambridge, Mass., 1985), 1-14; and Hilary Putnam, "Why Reason Can't be Naturalized," *Synthese* 52 (1982): 3-23.
3. Quine, "Epistemology Naturalized," in *Naturalizing Epistemology*, 25.
4. Kim, "What Is Naturalized Epistemology?" 389.
5. Alvin Goldman, *Epistemology and Cognition*, 3. But also see footnote 5, where Goldman admits that it is not always clear whether Quine wants to repudiate altogether the normative element of epistemology, but he does insist that Quine's "actual characterizations of naturalized epistemology do not expressly introduce this dimension."
6. Kornblith, "Introduction: What is Naturalized Epistemology?" 4.
7. Quine, *Theories and Things* (Cambridge, Mass., 1981), 72.
8. Quine, "Epistemology Naturalized," 19.
9. Quine, *Pursuit of Truth* (Cambridge, Mass., 1990), 19.
10. Quine and J. S. Ullian, *The Web of Belief* (New York, 1970), 7. The second and significantly revised edition (1978) does not contain this passage.
11. Quine, "Reply to White," in *The Philosophy of W. V. Quine*, edited by L. A. Hahn and P. A. Schilpp (LaSalle, Ill., 1986), 663-65.
12. Quine, *Pursuit of Truth*, 19.
13. Quine, *Pursuit of Truth*, 1.
14. Quine does not define evidence in terms of observations. He talks rather of observations, or more exactly observation sentences, as being "the vehicle of scientific evidence." See, e.g., *Pursuit of Truth*, 5.
15. See *Pursuit of Truth*, 12-13.
16. *Ibid.*, 13.
17. Locke, *An Essay Concerning Human Understanding*, edited by P. H. Nidditch (Oxford, 1975); see especially IV, xvii, 24.
18. See Chisholm, *Theory of Knowledge*, 3rd ed., and Bonjour, *The Structure of Empirical Knowledge* (Cambridge, Mass., 1985).
19. *Pursuit of Truth*, 20.

20. Ibid.

21. This was also the way he expressed the point in *Word and Object*, which was written ten years before *The Web of Belief*: "The last arbiter is the so-called scientific method, however amorphous . . . a matter of being guided by sensory stimuli, a taste for simplicity in some sense, and a taste for old things" (Quine, *Word and Object* [Cambridge, Mass., 1960], 22).

22. "Reply to White," in *The Philosophy of W. V. Quine*, 663–65.

23. Susan Haack makes a similar point. See Haack, "Two Faces of Quine's Naturalism," *Synthese* 94 (1993): 335.

24. *Pursuit of Truth*, 19.

25. "I admit to naturalism and even glory in it. This means banishing the dream of a first philosophy and pursuing philosophy rather as a part of one's system of the world, continuous with the rest of science" (Quine, "Reply to Putnam," in *The Philosophy of W. V. Quine*, 430–31).

26. See, e.g., Quine, *Pursuit of Truth*, 15.