Epistemology as Methodology

Gilbert Harman

June 27, 2005

What is distinctive about my views in epistemology? One thing is that my concern with epistemology is a concern with methodology. Furthermore, I reject psychologism about logic and reject the idea that deductive rules like *modus ponens* are in any way rules of inference. I accept a kind of methodological conservatism and reject methodological theories that appeal to special foundations, analytic truth, or a priori justification. Although I believe that there are significant practical aspects of theoretical reasoning, I reject the suggestion that theoretical reasoning is a special case of practical reasoning applied to a special epistemic goal. I also believe that a methodological epistemology that is concerned with the reliability of inferential methods can benefit from an appreciation of important relevant concepts and results about reliability in statistical learning theory.

1 Logic and Epistemology

Logic is a theory of implication and inconsistency. It has an abstract subject matter, being concerned with relations among propositions. It does not have a psychological subject matter, nor a normative subject matter. Logical principles are exceptionless necessary principles about what follows from

what, not default principles about what can be inferred from what. From contradictory beliefs, anything and everything follows, but it is not the case that from contradictory beliefs it is rational or reasonable to infer anything and everything.

Sometimes, the fact that S follows from things you believe is part of what makes it reasonable for you to believe that S. But not always. It may be more reasonable to abandon one of the beliefs from which it follows that S. Furthermore, infinitely many things follow from your beliefs hardly any of which are worth adding to your beliefs.

Reasoning normally involves change in view, often giving up things previously accepted as well as coming to accept new things, so reasoning is not in general thinking through a proof or argument in logical order, with premises, intermediate steps, and final conclusion. Although in reasoning you may construct an argument, your construction of the argument may start with the conclusion and work backward to the premises, or you may start in the middle and work in both directions. If the argument is part of an explanation that you infer as the best explanation of some data, the conclusion of the argument may be the beginning of your reasoning and the premise of the argument may be the conclusion of your reasoning.

Deductive rules are rules of implication, not rules of inference or reasoning. They are perhaps rules about what makes something a proof or argument, but they are not rules that you could follow in constructing a proof or argument, except in the sense that they characterize what you are constructing.

Some theorists respond to these obvious points by saying that they show

only that there is a harmless ambiguity in the phrase "rule of inference," one meaning of which is "rule of implication." But the "ambiguity" is not harmless. It has served to hide the difficult issue of saying how recognition of implications does or ought to figure in inference, it makes it much harder for good students to learn logic than it should, and it has been responsible for many confused attempts to develop various logics corresponding to aspects of the way people reason: inductive logic, because there is inductive reasoning; practical logic, because there is practical reasoning; relevance logic, because reasoning is sensitive to relevance; a logic of belief revision because reasoning is change in view; and so on. (Similarly for attempts to state "closure principles" for knowledge.) These failed projects have been motivated entirely by the thought that rules of logic are "rules of inference." This "terminology" has not been "harmless" and should be abandoned.

2 Methodological Conservatism

In reasoning, you start where you are with your current beliefs, plans, and goals, and your current methods or procedures for modifying these plans and methods. It is not reasonable for you to make any changes in your starting points except to resolve tensions within them and to answer questions which you have reasons to answer. Any reasonable methodology has to be conservative in this sense.

In a certain (methodological) sense, your initial beliefs, plans, goals, and methods have an immediate default or prima facie "justification." They are the "foundations" for your reasoning—foundations in the sense of starting

points. Philosophers sometimes advocate more limited foundations, where the strictest theories allow as foundations only things you cannot doubt and where your nonfoundational beliefs, goals, plans and methods are to be retained only to the extent that they have argument like justifications ultimately based on foundational beliefs, goals, plans, and methods. Such strict foundations theories do not seem plausible as methodological theories, since you do not normally associate argument like justifications with the beliefs, goals, plans, and methods you currently accept, which means these views imply you should abandon almost everything you believe.

Even apart from that consideration, it is obvious that the strictest methodological foundations theories imply that you ought to give up almost all your beliefs, goals, plans, and methods, as lacking sufficient justifications given the strictest standards. This result can be avoided only by adding to the foundations, making them less strict, for example, allowing into the foundations certain perceptual beliefs about the environment, and certain memory beliefs about the past, as well as inferential methods like enumerative induction, inference to the best explanation, and reliance on testimony. These less strict foundations theories are closer to general conservatism. In as much as the move to less strict theories is motivated by the initial implausibility of the skeptical implications of the stricter theories, the move appears to be based on a background acceptance of something like general conservatism.

General conservatism is a widely accepted methodology in philosophy outside of epistemology, for example in Rawls' discussion of trying to put the principles of justice you accept into "reflective equilibrium" with your judgments about particular cases.

Rawls and others emphasize that the relevant reflective equilibrium should be "wide" rather than "narrow," also taking into account what other views there are and what can be said to proponents of these other views. More generally, as Thagard has emphasized, you should consider how your current methods look in relation to the best scientific reasoning, what your goals in reasoning are, and how reliable various methods are. Furthermore, it is methodologically useful to maintain a Popperian questioning attitude toward the views you accept. These further methodological principles do not conflict with general methodological conservatism to the extent that they are themselves currently accepted or are principles there are reasons to accept given current beliefs, goals, plans, and methods.

People sometimes object to methodological conservatism by considering a case in which you irrationally come to accept a new belief, perhaps randomly. The objection is that it seems wrong to suppose your belief is justified as soon as it is accepted.

But methodological conservatism does not imply that your belief is justified as soon as it is accepted. It implies that the belief is prima facie of default justified. Whether it is fully justified (in the relevant methodological sense) depends on whether you are aware of how you came to accept that belief. If you are aware that you did not have good reason to accept it then that awareness is in tension with your continuing to accept the belief. Similarly, if your other beliefs and methods give you now reason to conclude that you did not have good reason to accept the belief in question, then there is a tension in your beliefs that might point to getting rid of that belief. On the other hand, if you are unaware of how you came to have the belief, that

by itself is not enough to give you reason to abandon the belief, because almost all your beliefs have the property that you are unaware of how you came to have them.

Remember that general conservatism is a methodological principle, offering methodological advice of a sort a person can take. It is not a principle of justification in some other sense of "justification."

3 Irrelevance of Apriority and Analyticity

Philosophers sometimes suppose that a limited number of beliefs and methods (and maybe some goals and plans) are a priori justified but I am unaware of any satisfactory accounts of what these suppositions might mean.

There are of course those beliefs and methods (and goals and plans) that are *prior* to a given inquiry, meaning those beliefs and methods (and goals and plans) that are taken for granted at the beginning of the inquiry. While I agree that all such beliefs and methods (and goals and plans) are (default) justified a priori, in the sense of prior to a given inquiry, this is not what is meant by the philosophers I have in mind, because they suppose that only a limited number of beliefs, etc. have the relevant sort of a priori justification.

There may also be beliefs and methods (and goals and plans) that are prior in the sense of prior to birth—innate or innately programmed to arise. But the fact that a given belief or method is innate does not make it a priori justified according to the philosophers I have in mind, because it is for them an "empirical question" whether a belief or method is innate and a further empirical question whether it is innate in a way that makes it likely to be

true, where the fact that these are empirical questions means that the beliefs and methods in question are not a priori justified in the relevant sense.

I know only two ways to try to explain the relevant sense of *a priori*. One presupposes a strict foundations theory of justification. The idea is that everything you are justified in believing is associated with your justification for believing it, and similarly for methods, goals, and plans, where there are strict standards for foundational beliefs, methods, goals, and plans, which must be directly justified either by your present sensory or perceptual experiences or by some sort of a priori insight. In this view, something is a priori justified if it is directly justified—not just justified by other beliefs, methods, goals, and plans for which one has justification—and not justified empirically by present sensory or perceptual experiences. One problem with this explanation is that the characterization of a priori justification is completely negative. A related problem is that strict foundations theories of this sort imply extreme skepticism and so, if interpreted methodologically, imply that you should give up almost all your current beliefs, methods, plans, and goals, which is really quite crazy!

The other way I know of to explain a relevant sense of *a priori* supposes that some claims are (at least default) justified simply because of the meanings of the words used to make those claims, perhaps because one intends to use those words in such a way that those claims are justified. But it would seem that accepting any claims is in part accepting that the claims are justified, so it would seem that for any such claim there is a sense in which one intends to be using words in such a way that the claim is justified. Furthermore, given general conservatism, any claim one accepts is at least

default justified, so it is unclear what methodological difference analyticity is supposed to be making.

Those who have been brought up to believe in an epistemically important analytic-synthetic distinction may well point out that these brief remarks do not refute the importance of the distinction. At best they raise a show a need for an account of the distinction and its epistemological significance. However, I have never seen an account that isn't obviously defective. So I have my doubts.

4 Pragmatism

There is an important distinction between epistemic reasons to believe something and nonepistemic reasons to believe something. A study showing that smoking cigarettes is not statistically associated with higher risk of cancer might provide an epistemic reason to believe that cigarette smoking does not cause of cancer. Learning that the company you would like to work for only hires people who believe that cigarette smoking does not cause cancer might provide you with a nonepistemic reason to believe it. Although this nonepistemic reason to have a belief is a practical reason, it is not obvious that a practical reason to believe something is always a nonepistemic reason. For example, considerations of conservatism and simplicity may be both practical and epistemic.

Some theorists suppose that epistemic reasons are actually practical reasons associated with an "epistemic" goal, such as the alleged goal of believing something if and only if it is true. In this view, theoretical reasoning, which

affects beliefs, is properly seen as a special case of practical reasoning. This is actually quite implausible, however, given certain obvious differences between theoretical and practial reasoning, the most serious being that good practical reasoning almost always involves arbitrary decisions in a way that good theoretical reasonng cannot do so. In deciding where to go after work it is not enough to decide to go home, because there may be several equally good possible ways to get there, so an arbitrary choice must be made and it would be irrational not to make such choices. But in trying to deciding which route home someone else took, where there are several possible routes with nothing favoring one rather than another, it would be irrational arbitrarily to believe he took a certain particular route.

Furthermore, as Hempel observed in his classic paper, "Inductive Inconsistencies," the alleged goal of believing something if and only if it is true might be interpreted as implying that you should believe nothing whose probability on your evidence is less than .5 and everything whose probability is more than .5. And, since this is practical reasoning, you can arbitrarily choose whether to believe those propositions with probability of exactly .5.

Various other suggestions have been made about what supposed epistemic goals there might be, none very plausible. So, I have trouble believing that theoretical reasoning can be reduced to practical reasoning. On the other hand, I do think that certain practical considerations are directly relevant to theoretical reasoning and epistemic reasons. I have already mentioned conservatism and simplicity. Furthermore, what it is reasonable to reason about can depend on practical considerations. And whether it is reasonable to end inquiry and accept certain conclusion can depend on practical

considerations concerning how urgent it is to get an answer, what resources are available for further inquiry, and what other questions you want to answer.

5 Reliability

You use inferential methods that in some sense you take to be reliable. The thought that a belief of your is the result of using an unreliable method would be in tension with the belief in question, providing some reason to abandon either the thought or the belief. Furthermore, in trying to reach an appropriately wide reflective equilibrium, you need to take into account how reliable are the various methods you have used in arriving at your various beliefs.

The reliability of certain methods has been studied empirically and compared with human performance, with such results as that simple linear models often do better than experts at predicting the success of medical interventions, predicting criminal recidivism, predicting tomorrow's weather, predicting academic performance, predicting loan and credit risk, or predicting the quality of a French wine vintage.

Statistical learning theory provides a more theoretical and philosophically deeper approach to the study of inferential methods. It provides an account of reliability in terms of a (usually unknown) background probability distribution. It provides a clear account of why inferential methods must have some sort of inductive bias favoring some hypotheses over others. It distinguishes different sorts of inductive inference with importantly different

properties. It provides useful worse case results concerning various inductive methods.

In one method of infering a hypothesis from data, you do not simply accept that hypothesis that best fits the data but you balance that consideration against another consideration sometimes called "simplicity." Popper argued that simplicity is connected with the "falsifiability" of the claim that the true hypothesis is a member of a certain class of hypotheses. Popper claims that the class of linear functions of one variable is simpler than the class of quadratic hypotheses, because the number of data points needed to get a result that does not perfectly fit some linear hypothesis is less than the number needed to get a result that does not perfectly fit some quadratic hypothesis. In statistical learning theory, this idea is captured more precisely by the important notion of "VC dimension."

Popper thinks that what he calls the falsifiability of a class of hypotheses is connected with the number of parameters that have to be determined in order to specify a particular member of the class. Statistical learning theory rejects that part of Popper's theory, noting for example that sine curves can be specified using two or three parameters even though the "falsifiability" or VC dimension of the class is is very high. One thing this means is that, if sine curves are examples of simple curves, simplicity is not what should be balanced against fit to data.

Statistical learning theory has many other aspects, including discussions of learning in feed-forward neural networks, with which philosophers have some familiarity, and the much more effective support vector machines, about which philosophers tend to be ignorant. Statistical learning the-

ory also discusses "transductive methods" that do better than inductive methods that use labeled data to come up with general rules for classifying new unlabeled cases, where the transductive methods use information about what new unlabled cases have come up to be classified. The theory of transduction has implications for discussions of moral particularism and for the question whether all inductive inference involves some sort of inference to the best explanation. But I do not have space here to elaborate.