
Epistemology and the Psychology of Human Judgment

by Michael A. Bishop and J.D. Trout

Chapter 1: Laying Our Cards on the Table

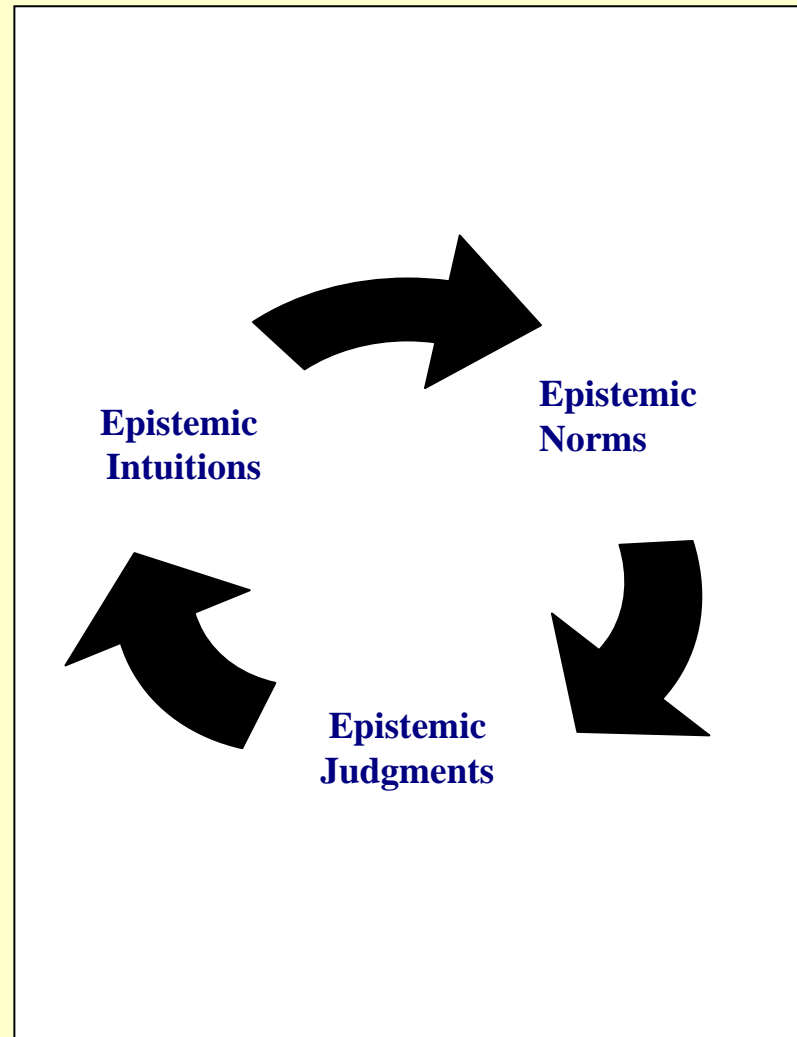
- In this chapter Bishop and Trout make it clear that they believe that SAE (Standard Analytic Epistemology) is misconceived.
 - They argue that epistemology, like ethics, should result in practical advice for improving our cognitive practices, and that the methods use in SAE simply do not achieve that.
-

The Stasis Requirement

- Bishop and Trout begin by pointing out an interesting methodological assumption in SAE, viz., that our theory of epistemic justification should leave our epistemic situation largely unchanged.
 - Quoting Jaegwon Kim
 - “it is expected to turn out that according to the criteria of justified belief we come to accept, we know, or are justified in believing, pretty much what we reflectively think we know or are entitled to believe”
 - You will recognize here the appeal to intuition which we’ve already had occasion to criticize.
-

Reflective Equilibrium

- The basic method of analytic philosophical inquiry is what Nelson Goodman called termed “reflective equilibrium”, which basically looks like this.
- The idea is essentially that captured by Neurath’s image of the sailor who must always repair his ship while at sea. We must presume that that the ship is seaworthy. Only minor tinkering at any one time can be done without putting the vessel at risk.
- The idea of reflective equilibrium is that we make small adjustments to our intuitions, norms, and judgments to maintain a functioning belief system.
- The fundamental working principle here is conservatism.



Gettier redux

- B&T use the Gettier example in much the way Kornblith did, to show that SAE philosophers tend to assume that their intuitions are sufficient to determine whether or not something is knowledge.
 - The authors are not actually assaulting conservatism so much as pointing out that sometimes conservatism impedes progress. They think SAE has reached a critical point where progress simply is not being made because it is not capable of recognizing the possibility that our cognitive practices are sometimes deeply misguided.
-

The philosophy of science approach

- B&T claim that epistemology is part of the philosophy of science. This is interesting for two reasons.
- First, though as a result of being in this class you are not in a position to appreciate it, most SAE texts make almost no reference to the philosophy of science. It is downright eerie, that in an age where just about everyone stipulates that we gain knowledge through scientific inquiry, a text about knowledge would take almost no interest in questions about scientific method.
- Second, even if you think of philosophy of science as part of science, thinking of epistemology as a branch of philosophy of science is a little different than thinking of it as a branch of empirical psychology. If philosophy of science is dedicated to the normative study of scientific methods, then the idea here is that we would expect improvements in everyday thinking to be informed by a proper understanding of scientific method.
- But for B&T, it turns out that epistemology really is a branch of empirical psychology after all. They call it “ameliorative psychology” which they define as the use of scientific and statistical methods to improve upon our reasoning.

Predictive modeling

- B&T introduce the concept of a statistical prediction rule (SPR). The basic idea here is that it has now been well established that very simple statistically informed rules are much more reliable than highly trained human experts.
 - On P. 13 they provide a list of citations to give you an idea what they are talking about. Here are just a few:
 - A model that used past criminal and prison records was more reliable than expert criminologists in predicting criminal recidivism.
 - SPR's predict academic performance better than admissions officers.
 - SPR's are better than forensic psychologists at predicting future violence.
-

Predictive modeling 2

- It's worth noting here that the things Bishop and Trout are talking about here are starting to make an enormous impact on practice. Consider two examples not discussed here:
 - Baseball scouting has recently been completely revolutionized by the realization that SPR's vastly outperform professional scout's abilities to identify star talent. (This is the thesis of the book *Moneyball*, by Michael Lewis.)
 - Though this fact is vociferously denied by investment agencies, private investors are far better off following a basic SPR than listening to the advice of a broker. (The SPR, btw, is simple, invest in index mutual funds and leave it there until you need it.)
 - A very simple SPR has been shown to be dramatically more effective at producing correct treatment for cardiac patients than the informed judgment of highly trained cardiologists. (This is SPR is called the Goldman rule, and is discussed in Malcolm Gladwell's book: *Blink*)
 - The authors also discuss Goldberg rule, which is a simple SPR for deciding whether a patient should be diagnosed as neurotic or psychotic. It vastly outperforms the independent judgment of clinical psychologists.
-

SAE vs. Ameliorative Psychology

- Bishop and Trout claim that SAE and AP are based on very different ideas concerning the subject matter of epistemology.
 - SAE: What is the nature of knowledge and epistemic justification?
 - AP: How do we achieve reasoning excellence?
- Interestingly, these aren't really the radically different questions that B&T make them out to be. After all, epistemic justification subsumes the idea of reasoning excellence.
- The point really is one of methodology. AP counsels paying attention to empirical results not just intuitions. The fact is that in the past it never occurred to the typical SAE philosopher that the results of SPR's would be relevant to anything in epistemology, and that is just because epistemology was not regarded as an empirical discipline.
- So here is the shocking pronouncement: How good your reasoning processes are is an empirical question. The question is: How reliable are these processes at producing correct results?

Expert attitudes

- Unsurprisingly, experts whose livelihoods are threatened by SPR's tend to be quite dismissive of them. In other words, we find the same kind of disinterest in empirical matters in experts who call themselves scientists as in experts who call themselves philosophers.
 - Experts in the relevant empirical disciplines are going to be tempted to explain away the success of SPR's by special pleading. Maybe most "experts" can be outperformed by simple statistical rules, but not the best ones, not them.
 - SAE philosophers may be tempted to a similar gambit. Maybe SPR's are just better for ordinary people who aren't competent to understand the a priori methods they prescribe.
-

The structure of a healthy epistemology

- B&T identify three components of a healthy epistemological tradition:
 - Practical: It must issue in useful advice on how to think better.
 - Theory: Epistemological theory should *explain* why the practical advice is successful.
 - Social: There must actually be an effective mechanism for communicating practice and theory to the people who would benefit from them. (This is a shot at the insularity and basic irrelevance of SAE to scientists and ordinary people.)
-

Naturalized normativity redux

- In section 4 B&T touch on the objection we have encountered in various forms, that empirical methods can't yield normative results unless they are grounded in some a priori epistemology.
 - On p. 19 they formulate this point as follows:
 - Suppose our epistemological theory begins with empirical claims about Ameliorative Psychology. Presumably, we have to make some decisions about which empirical claims to trust. So we have to decide which views are the epistemically good ones. But such decisions require a prior epistemological theory. So one can not begin ones epistemological speculations with empirical claims.
 - Of course, everyone should now be seasoned enough at naturalistic thinking to be able to anticipate the answer: this is based on the assumption that we can discover epistemologically good views *a priori*. Since this has never been done, the burden of proof is on those who claim it is possible.
 - B&T argue that there is really no serious disagreement about how to proceed. Even though we can not justify it in non circularly, everyone agrees that the best reasoning strategies to use are ones that have worked in the past, and the best reasoning strategies to avoid are the ones that have failed.
-

Example: the shoe size rule

- B&T ask us to consider a rule for deciding whether a criminal is likely to be a repeat offender. Consider the shoe size rule: If the shoe size of the offender is a whole number (8, 9, 10, etc.) then he is unlikely to be a recidivist. If not (8.5, 9.5, 10.5) then he is.
 - This is a sensationally stupid rule, but it's stupidity can not be determined a priori. The reason is simply that the rule make worse predictions than a monkey flipping a coin would make.
-

Reliabilism and jumping the shark

- There's a sweet quotation on p.22
 - Somewhere around the third epicycle on a counterexample involving reliable clairvoyants, back-up electrical generators, or an environment full of objects that are phenomenologically identical but ontologically distinct, SAE jumped the shark.
 - B&T also show that they belong to a somewhat more radical naturalistic camp than Kornblith. They say on the same page:
 - We should admit that reliabilism has achieved some of epistemology's reason-guiding potential. But as long as reliabilism remains wedded to the goals and methods of SAE it is doomed.
-