

The Pathologies of Standard Analytic Epistemology *

Michael Bishop J.D. Trout

Standard Analytic Epistemology (SAE) names a contingently clustered class of methods and theses that have dominated English-speaking epistemology for about the past half-century. The major contemporary theories of SAE include versions of foundationalism (Chisholm 1981, Pollock 1974), coherentism (Bonjour 1985, Lehrer 1974), reliabilism (Dretske 1981, Goldman 1986) and contextualism (DeRose 1995, Lewis 1996). While proponents of SAE don't agree about how to define naturalized epistemology, most agree that a thoroughgoing naturalism in epistemology can't work. For the purposes of this paper, we will suppose that a naturalistic theory of epistemology takes as its core, as its starting-point, an empirical theory. The standard argument against naturalistic approaches to epistemology is that empirical theories are essentially descriptive, while epistemology is essentially prescriptive, and a descriptive theory cannot yield normative, evaluative prescriptions. In short, naturalistic theories cannot overcome the is-ought divide.

Our main goal in this paper is to show that the standard argument against naturalized epistemology has it almost exactly backwards. On the one hand, it is the theories of Standard Analytic Epistemology that have at their core essentially descriptive, empirical theories that face the serious challenge of a yawning is-ought divide. On the other hand, some of the best psychological science of the past half-century is normative in a way that can serve as a natural starting point for a naturalistic epistemology with prescriptive force. We will argue for the following five theses:

1. The dominant theories of Standard Analytic Epistemology (foundationalism, coherentism, reliabilism, contextualism) have at their core a descriptive theory.
2. This descriptive theory aims to capture the considered epistemic judgments of a small group of idiosyncratic people.
3. The standard charge leveled against naturalistic epistemology can also be leveled against the dominant theories of Standard Analytic Epistemology: They attempt to extract prescriptions from descriptions.
4. Some of the best psychological science of the past half-century is deeply normative and makes specific recommendations about how to improve our reasoning about matters of great practical significance.
5. An approach to epistemology that takes seriously these psychological findings is better suited to overcoming the is-ought gap than are the theories of SAE.

This paper is organized as follows. In section 1, we document the attractions of some of the empirical science that is devoted to passing normative judgments and prescribing new and better ways to reason. Because this literature is so wide-ranging, it will be useful to give it a name. We call it Ameliorative Psychology. In section 2, we argue that the theories of SAE have at their core a descriptive theory. As a result, the theories of SAE face the same is-ought challenge faced by naturalistic theories. Further, we offer some reasons for thinking that the prospects for SAE overcoming the is-ought challenge are not good. In section 3, we argue that the approach to epistemology that takes seriously Ameliorative Psychology is superior to that of SAE because it is much more likely to provide a motivated way of overcoming the is-ought divide. The normative recommendations and evaluative theses of Ameliorative Psychology can receive confirmation by the best science of the day. And some of these recommendations have been impressively confirmed, in the form of documented results and a proven method for securing them. Standard Analytic Epistemology, on the other hand, has a long tradition and the loyalty of its enthusiasts.

1. A new approach to naturalistic epistemology

What makes our approach to epistemology naturalistic is that we begin our epistemological investigations with a descriptive core and work out from there.¹ (We take this to be sufficient for an approach to be naturalistic; we don't know whether it is also necessary.) Ameliorative Psychology is the descriptive core of our epistemological theory. It is an empirical inquiry (or, rather, a set of empirical inquiries) that aims to give positive advice about how we can reason better. It includes work in psychology, statistics, machine learning and Artificial Intelligence. A significant part of Ameliorative Psychology is known as “predictive modeling” and it includes discussion of models such as linear models, multiple regression formulas, neural networks, naïve Bayes classifiers, Markov Chain Monte Carlo algorithms, decision tree models and support vector machines. It also includes significant parts of more traditional psychology, including the well-known heuristics and biases program launched by the groundbreaking research of Kahneman and Tversky (Kahneman, Slovic, and Tversky 1982).

The core of our naturalistic epistemological theory consists of the descriptive findings of Ameliorative Psychology. But Ameliorative Psychology is also deeply normative in the sense that it makes (implicitly or explicitly) evaluative “ought” claims that are intended to guide people’s reasoning. Epistemic prescriptions are an essential feature of Ameliorative Psychology. To see this, it will be useful to consider three examples of reason-guiding prescriptions derived from Ameliorative Psychology.

The Goldberg Rule is perhaps the most well-documented success of Ameliorative Psychology. It predicts whether a psychiatric patient is neurotic or psychotic on the basis of a MMPI profile. Lewis Goldberg (1965) found that the following rule outperformed 29 clinical judges (where L is a validity scale and Pa, Sc, Hy and Pt are clinical scales of the MMPI):

$$x = (L + Pa + Sc) - (Hy + Pt)$$

If $x < 45$, diagnose patient as neurotic.
If $x \geq 45$, diagnose patient as psychotic.

When tested on a set of 861 patients, the Goldberg Rule had a 70% hit rate; clinicians’ hit rates varied from a low of 55% to a high of 67%. (13 of the 29 clinical judges in the above study were experienced Ph.D’s, while the other 16 were Ph.D students. The Ph.D’s were no more accurate than the students. This is consistent with the findings reported in Dawes 1994.) So here we have a prediction rule that could literally turn a smart second-grader into a better psychiatric diagnostician than highly credentialed, highly experienced psychologists – at least for this diagnostic task. Almost four decades after the appearance of Goldberg’s results, making an initial diagnosis on the basis of a MMPI profile by using subjective judgment rather than the Goldberg Rule would bespeak either willful irresponsibility or deep ignorance.

Another example of Ameliorative Psychology is a 1995 paper by Gigerenzer and Hoffrage entitled, “How to Improve Bayesian Reasoning Without Instruction: Frequency Formats” (emphasis added). As the title of the paper suggests, Gigerenzer and Hoffrage show how people charged with making high-stakes diagnoses (e.g., about cancer or HIV) can improve their reasoning. They suggest a reasoning strategy that enhances reasoners’ ability to identify, on the basis of medical tests, the likelihood that an individual will have cancer or HIV. A third particularly successful example of Ameliorative Psychology is credit scoring. Many financial institutions no longer rely primarily on financial officers to make credit decisions – they now make credit decisions on the basis of simple SPRs developed as the result of research by psychologists and statisticians (Lovie and Lovie 1986).

Goldberg’s Rule, frequency formats and credit scoring are only a small sample of the successes of Ameliorative Psychology. SPRs have been shown to be superior to experts in making a number of very high-

stakes predictions, including predicting the success of electroshock therapy, criminal recidivism, academic performance, the presence of progressive brain dysfunction, the presence, location and cause of brain damage, proneness to violence, and psychiatric diagnoses on the basis of MMPI profiles. Ameliorative Psychology has 50 years' worth of success in the investigation and discovery of new and better ways of reasoning about issues of great practical significance. (For an overview of some of these findings, see Grove and Meehl 1996.)

Ameliorative Psychology tends to focus on the assessment of reasoning strategies in terms of their reliability or their relative reliability (i.e., their reliability relative to other potential reasoning strategies) and their tractability (i.e., how easy it is for people to implement them). These sorts of findings can naturally yield straightforward normative, epistemological prescriptions. People ought to use the Goldberg Rule in making preliminary diagnoses of psychiatric patients. People ought to use frequency formats when diagnosing rare conditions on the basis of well-understood diagnostic tests. And people ought to use a credit scoring model when making predictions about someone's creditworthiness. Ameliorative Psychology is normative in the straightforward sense that it makes (implicitly or explicitly) evaluative "ought" claims that are intended to guide people's reasoning.²

Since Ameliorative Psychology offers reason-guiding prescriptions, it must presuppose epistemological principles that guide those recommendations. On our view, the central goal of naturalistic epistemology is to articulate the normative framework of Ameliorative Psychology. This makes epistemology a branch of the philosophy of science. Just as a philosopher of biology might uncover and articulate the metaphysical assumptions of evolutionary theory, the epistemologist uncovers and articulates the normative, epistemic assumptions of the long and distinguished tradition of Ameliorative Psychology.

One might reasonably wonder whether the recommendations of Ameliorative Psychology are really normative in the same way as the recommendations of SAE are normative. Admittedly, there is one telling difference. People outside academia have on occasion actually changed the way they reason about significant matters as a result of the normative recommendations of Ameliorative Psychology.

2. The naturalistic theories of Standard Analytic Epistemology

We contend that the central theories of SAE (foundationalism, coherentism, reliabilism, contextualism) are structurally analogous to the naturalistic approach described above. They have at their core a descriptive theory, and from that descriptive theory, proponents of SAE draw normative, epistemological prescriptions. In 2.1, we will examine the curiously conservative method of SAE. In 2.2, we argue that this method is geared to give us descriptive knowledge about the epistemic judgments of a relatively small group of idiosyncratic people. In 2.3, we note with some irony the fact that proponents of SAE have offered anti-naturalistic arguments that apply at least as well to their own views as to any naturalistic epistemology. In 2.4, we consider an obvious way the proponent of SAE might try to bridge the is-ought gap. We argue that the prospects for this argument are not very promising.

2.1. Standard Analytic Epistemology's Commitment to Stasis

For most proponents of SAE, the primary goal of epistemology is to provide an account of knowledge and epistemic justification. In taking on any goal-oriented project, it is often useful to ask what would count as success in that project. In the case of SAE, what would it be for an account of justification to be successful? What property or properties distinguish a successful account from the unsuccessful accounts? In a typically clear and careful article, Jaegwon Kim identifies a number of criteria that any account of justification must meet in order to succeed. For our purposes, the most important of these success conditions is what we will call the stasis requirement:

...Although some philosophers have been willing to swallow skepticism just because what we regard as correct criteria of justified belief are seen to lead inexorably to the conclusion that none, or very few, of our beliefs are justified, the usual presumption is that our answer to the first question should leave our epistemic situation largely unchanged. That is to say, 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. (Kim 1988, 382)

It is worth noting that this requirement – that the right account of justification “leave our epistemic situation largely unchanged” – is profoundly conservative. In particular, it is extraordinary that SAE should have built right into it a requirement that makes it virtually impossible that a successful epistemological theory would force us to radically alter our epistemic judgments.

Of course, proponents of SAE will not typically suggest that they are trying to provide an account of their naïve epistemic judgments, but of their considered epistemic judgments. There may be many ways to spell out what it is for our judgments to be considered, but a standard approach would appeal to reflective equilibrium. Nelson Goodman introduced reflective equilibrium as a process that involves aligning our judgments about particular instances with our judgments about general principles. “The process of justification is the delicate one of making mutual adjustments between rules and accepted inferences; and in the agreement achieved lies the only justification needed for either” (1965, 64). Narrow reflective equilibrium is the process of bringing our normative judgments about particular cases into line with our general normative prescriptions and vice versa. Wide reflective equilibrium differs from narrow reflective equilibrium by including our best theories in the mix. So wide reflective equilibrium is the process of bringing into alignment our best theories as well as our normative judgments about particular cases and our general normative prescriptions (Rawls 1971, Daniels 1979). From our perspective, it doesn’t matter whether the theories of Standard Analytic Epistemology are supposed to capture our judgments in wide or narrow reflective equilibrium – or in some other sort of state (Bealer 1996, 28 [fn 3]). We have described the notion of reflective equilibrium because we think it is a reasonably popular view among analytic philosophers and it is plausible to suppose that the methods of SAE are geared toward capturing our judgments in reflective equilibrium.

Now, back to the stasis requirement: If an epistemic theory forced us to radically alter our considered epistemic judgments (e.g., our epistemic judgments in reflective equilibrium), then ipso facto that theory is unacceptable. While perhaps not all proponents of SAE embrace the stasis requirement (e.g., see Unger 1984), we think that Kim is right in identifying it as a success condition that most proponents of SAE place on epistemological theories. But it is not a requirement that is often explicitly stated. So where do we find it? We suggest that the commitment to stasis is embodied in the method of SAE. Philosophers accept or reject an epistemological theory on the basis of whether it accords with their considered judgments. Consider the Gettier Problem, the archetype of method and substance in SAE for more than a generation. Gettier’s (1963) paper is a classic because it describes clear and compelling examples in which the justified true belief (JTB) account of knowledge is at odds with our considered judgments about knowledge. One of Gettier’s famous cases involves a man named Smith who has overwhelming evidence, and so justification, for believing that Jones will get a job and that Jones has 10 coins in his pocket. On the basis of these beliefs, Smith infers that the man who will get the job has 10 coins in his pocket. It turns out that unbeknownst to Smith, he will get the job and he has 10 coins in his pocket. His belief that the man who will get the job has 10 coins in his pocket is true and justified. But Gettier insists that it is “clear” that Smith’s belief is not knowledge (Gettier 1963, 122). For proponents of SAE, the Gettier examples are important because they show that the JTB account can’t be right because it does not “leave our epistemic situation largely unchanged.” And herein lies our primary objection to Standard Analytic Epistemology. The theories of SAE are often rejected solely on the grounds that they violate our considered epistemic judgments. We suspect that most readers familiar with the literature will grant that this is how SAE

works. So rather than explore any more of the countless and wonderfully rococo counterexamples prevalent in the SAE literature, let's look at how some of these counterexamples end.

However, it is perfectly apparent that I know nothing of the sort. (Lehrer and Paxson, 1969, 235).

Even if S correctly predicts that he is going to lose, we would deny that he knew he was going to lose if the only basis he had for this belief was the fact that his chances of winning were so slight. (Dretske 1971, 3).

The situation is a peculiar one, and my intuitions, and I would suppose other people's, are not completely clear on the matter. But it seems, on the whole, that we ought not to speak of knowledge here... (Armstrong 1973, 181).

But, to make such an assumption is counterintuitive. In everyday situations we do not regard deception as precluding rationality. Likewise, we do not regard the fact that we have been deceived, or will be deceived, or would be deceived, as precluding rationality. (Foley 1985, 192).

...And, surely, we do not want to say that the fact that his friend has a generator in his basement prevents S from having knowledge that the company's generators are causing the lights to be on. (Pappas and Swain 1973, 66)

In the above passages (and we could have chosen literally hundreds of others) we are urged to share the philosopher's considered judgments about the case. And we usually do. The noteworthy feature of these passages is that the acceptance or rejection of a theory turns entirely on whether or not it violates our considered epistemic judgments.

Rejecting theories solely because they do violence to our considered judgments is a shockingly conservative principle of theory choice. This may only become clear if we compare it to methods in other fields of inquiry. The special theory of relativity does extreme violence to our considered judgments about simultaneity. But that is hardly a reason to reject it. If physics had been burdened with such a conservative method, we wouldn't have relativity, quantum mechanics or even Copernicanism! If biologists had embraced the stasis requirement, we certainly would not have Darwinism. If cultural studies had been taken over by such a conservative method, we wouldn't have postmodernism.

Ok, so sometimes conservatism is a good thing.

Behind this joke is an important point. Conservative methods work very well when applied to theories or propositions for which have overwhelming evidence. It is perfectly reasonable to be conservative about the commitments of theoretical chemistry reflected in the periodic table, or about the core attachments of contemporary physics or biology. That doesn't mean we rule out the possibility that new developments will force us to abandon them. Conservatism isn't mulishness. Conservatism is appropriate in the case of the core commitments of these theories because we have so much evidence in their favor that in absence of extraordinary counterevidence, they deserve our allegiance. But while conservatism is fine for excellent theories, it is poison in domains where progress awaits deep and durable changes in method and outlook. The alchemist's attachment to conservatism was ill-advised; it only protracted the alchemist's crippling (and it turns out, thanks to mercury and lead, fatal) ignorance.

This raises an obvious concern. The stasis requirement is an extraordinarily conservative principle. Do the theories of SAE have a track record (akin to our best theories of physics or chemistry) that makes it

reasonable for us to apply such a conservative principle? Surely not. Why then does philosophical practice embody the stasis requirement? We believe that there is a reasonable answer to this question. The stasis requirement is meant to tell us what counts as evidence for the theories of SAE. It is not a requirement that protects theories of SAE (e.g., reliabilism, coherentism, foundationalism) from negative evidence. The stasis requirement tells us something about the nature of the theories of SAE: They are meant to account for “our” considered epistemic judgments. If this is correct, the theories of SAE have at their core an essentially descriptive theory. This core is a theory that is meant to yield epistemic judgments that capture the judgments of a certain group of people.

2.2. The Descriptive Core of Standard Analytic Epistemology

Given the conservative nature of SAE, it is reasonable to ask: What kind of knowledge do we get when we do SAE? We will argue that in the first instance, we get knowledge of descriptive matters from SAE. Of course, it is perfectly possible that we might also get knowledge of normative matters from SAE. But at the center of all the major perspectives of SAE is a descriptive theory. One way to see this is to reverse engineer the conservative methods of SAE. By looking at how SAE works, we can illuminate the kind of knowledge it’s supposed to give its practitioners. The argument pattern so prevalent in the SAE literature (theory – counterexample – revision) is somewhat reminiscent of thought experiments in science. They are written primarily for and by people who have received idiosyncratic educations and who have a highly specialized set of skills. This education significantly affects the concepts, categories, and inferential patterns one uses in thinking about the world (certainly in physics, see 2.4, and we suspect in philosophy as well). So far, so good. In science, however, the understanding gained by thought experiments is typically tested against the world. But how is the understanding generated by counterexamples in SAE tested? Against the well-considered judgments of other (similarly trained) philosophers.

One needn’t be a sociologist to recognize that philosophers as a group are a relatively small and idiosyncratic sample of folks. Philosophers’ median education and intelligence are surely well above average. We speculate that philosophers’ median scores on various MMPI scales (e.g., social alienation, hypersensitivity, social introversion) might be above average as well. So the methods of SAE are not geared to make us experts about how regular folk think about justification. Nor are the methods of SAE geared to delivering generalizations about how a wide variety of people who lead flourishing lives – people in a wide range of stations, in different cultures, in different times – reason about important matters. So the conservative method of SAE does not seem especially well-designed to giving us expertise about what (if anything) distinguishes the reasoning of flourishing people from the reasoning of others. On the other hand, if epistemologists themselves tend to lead particularly successful lives, then perhaps providing people with their epistemic autobiographies would be useful. It is not obvious, however, that when socio-economic factors are controlled for, epistemologists as a group lead significantly more (or less) meaningful or flourishing lives than other folk.

A close examination of the primary tools of SAE suggests that the information SAE is geared to give us is information about the reflective epistemic judgments of a group of idiosyncratic, non-representative people who have been trained to use highly specialized epistemic concepts and patterns of thought. (By ‘highly specialized’ we mean that people who have not received the relevant training would find at least some of those concepts and patterns of thought strange, foreign or unfamiliar.) The conservative goals and methods of SAE are suited to the descriptive task of providing an account of the considered epistemic judgments of (mostly) well-off Westerners with Ph.Ds in Philosophy.

As an attempt to divine the nature and structure of the philosophers’ epistemic concepts, we have reasons to believe that the methods of psychologists are superior to those of SAE. This point is not essential to our argument. But it is worth noting that psychologists develop models of our concepts all the time. These

models mimic our categorization judgments. (These models can mimic concepts with “fuzzy boundaries” and indeterminate instances.) If we want an account of justification that mimics our considered judgements about what beliefs are justified, one reasonable way to do this is to find a psychologist who is willing to model our epistemic concepts. That model would be the account (e.g., Smith & Medin 1981, Keil 1989). Indeed, if what a philosopher really wants is a descriptive account of his concepts, he could save a lot of time, energy and expense by employing a few psychology graduate students.

A particularly dramatic way to see that the core of SAE is a descriptive theory of analytic epistemologists’ own epistemic judgments is to consider how SAE might be different if it were conducted by a very different group of people. In a fascinating study, Weinberg, Nichols, and Stich (2001) found that people in different cultural and socioeconomic groups make significantly different epistemic judgments. A group of Western subjects and non-Western subjects were given the following Gettier-style example:

Bob has a friend, Jill, who has driven a Buick for many years. Bob therefore thinks that Jill drives an American car. He is not aware, however, that her Buick has recently been stolen, and he is also not aware that Jill has replaced it with a Pontiac, which is a different kind of American car. Does Bob really know that Jill drives an American car, or does he only believe it?

REALLY KNOWS ONLY BELIEVES

A large majority of Western subjects gave the answer sanctioned by SAE (“only believes”) but a majority of East Asians and a majority of subjects from India gave the opposite answer (“really knows”) (2001, 443). Weinberg, Nichols and Stich also found cases in which there were significant differences between the epistemic judgments of people of high socioeconomic status (SES) and of low SES (2001, 447-448).

The possibility that there is considerable variation across cultures and within cultures in people’s epistemic judgments makes it plausible to believe that we learn about the epistemic judgments of an idiosyncratic group of people when we do SAE. This is as descriptive a fact as there could possibly be. Indeed, it suggests that SAE is actually an odd kind of cultural anthropology: Building theories that describe how privileged Westerners with PhDs in Philosophy engage in epistemic assessment. Weinberg, Nichols and Stich call this endeavor “ethno-epistemology” (2001, 454). If SAE is but anthropology, it is unclear on what grounds its proponents can reasonably make universal normative claims about the nature, origin, and limits, of human knowledge. To make universal claims – to claim SAE is more lofty than anthropology – has the uncomfortable feel of brute cultural imperialism.

2.3. Standard Analytic Epistemology: Throwing Stones in Glass Houses

We have argued that the descriptive core of SAE is a theory that captures the considered epistemic judgments of philosophers. Some proponents of SAE believe that those judgments are best captured by a coherentist theory, others believe that they are best captured by a foundationalist theory, others believe that they are best captured by a reliabilist theory, etc. We opened this paper by suggesting that what is idiosyncratic about a full-fledged naturalistic theory of epistemology is that its core, its starting-point, is an empirical theory. If this characterization has merit, then we have argued that the major theories of SAE are examples of naturalized epistemology. And so the obvious challenge for the theories of SAE is the same as for any naturalized epistemology: How are they to extract normative consequences from a descriptive theory? How are they to overcome the is-ought gap? In the next section (2.4), we will argue that the prospects for the theories of SAE bridging the is-ought gap are not good. Our goal in this section is to turn the tables on Standard Analytic Epistemology. For decades, proponents of SAE have wielded arguments against naturalized epistemology; we will show that these arguments can be readily adapted to show that their own preferred theories are doomed to

fail because they cannot be normative. We emphasize that we do not endorse these arguments. Our point here is that proponents of SAE have for too long been throwing stones at naturalistic epistemology from glass houses.

Michael Williams argues that the normative nature of epistemology makes it impossible to fully “naturalize” it.

[Epistemic claims] depend on meeting certain norms or standards which define, not what you do do, but what you must or ought to do. To characterize someone’s claim as expressing or not expressing knowledge is to pass judgment on it. Epistemic judgments are thus a particular kind of value-judgment. It is far from obvious that investigations with such a strongly normative component can be fully ‘naturalized’ (Williams 2001, 11).

Williams’s argument applies equally to the theories of SAE: When it comes to epistemic judgments, the theories of SAE define (or at least attempt to define) what we “do do” not what we “must or ought to do.” They try to tell us how we do make epistemic judgments, but they don’t tell us how we must or ought to make such judgments. (And by ‘we’, we mean the tiny fraction of the world’s population who has studied SAE.)

Richard Feldman maintains that psychology and philosophy must at best co-exist, because psychology can’t ask or answer the distinctively normative questions epistemology sets for itself.

The original epistemological questions seem to be perfectly good questions, well worthy of our attention. It is difficult to see, then, why the availability of this other field of study [psychology], concerning how we reason, is a suitable replacement for the evaluative questions that are at the heart of epistemology. (Feldman 2003, 168)

Feldman’s point can be made against the theories of SAE: “It is difficult to see, then, why the availability of this other field of study” (namely, SAE, which aims to describe how some people make evaluative judgments) “is a suitable replacement for the evaluative questions that are at the heart of epistemology.” The proponent of SAE is replacing normative questions about how to evaluate reason and belief with descriptive questions about how proponents of SAE evaluate reason and belief.

In a similar vein, Lawrence BonJour argues that any epistemology subsumed by psychology does not have the resources to evaluate, positively or negatively,

beliefs about alleged occult phenomena of various sorts, such as astrological or phrenological beliefs. For just as naturalized epistemology can say nothing positive about the justification of science or common sense, and is thus impotent in the face of skepticism, so also it can say nothing distinctively negative about the justification of these less reputable sorts of belief (BonJour 2002, 244).

We can once again turn the tables on the proponent of SAE. Any theory of SAE that accurately describes how a certain group of people make certain evaluative judgments “can say nothing positive about the justification of science or common sense, and is thus impotent in the face of skepticism, so also it can say nothing distinctively negative about the justification of” disreputable beliefs. Of course, such a theory might tell us how some people evaluate those beliefs. But that’s not the same as actually evaluating those beliefs.

Here is one of many proponents of SAE who take on Quinean naturalism in epistemology.

[T]he most extreme version of naturalism in epistemology eschews normativity altogether, seeking to replace traditional epistemology (with its concern with justification, rationality, reasonability, and their

normative colleagues) by descriptive psychology; this seems to be Quine's suggestion. [fn deleted]
(Plantinga 1993, 45)

By now, we're confident our argumentative strategy is wearing thin, but here it is anyway: SAE "eschews normativity altogether, seeking to replace traditional epistemology (with its concern with justification, rationality, reasonability, and their normative colleagues) by descriptive psychology" – a psychology that describes how certain people make certain normative judgments.

Proponents of SAE have argued that radically naturalistic theories of epistemology cannot succeed because they cannot be normative. Our aim has been to show that these standards and arguments can be used against the theories of SAE themselves. We distance ourselves from such arguments. We do not believe that merely pointing out that a theory faces the problem of bridging the is-ought divide damns that theory. And this is a good thing for proponents of SAE. We all start our normative musings with psychology. Proponents of SAE start by describing a certain group's epistemological judgments, and we start with Ameliorative Psychology. When it comes to bridging the is-ought gap, everybody has work to do.

2.4. How SAE Might Try to Get Normative Prescriptions from Its Descriptive Core

Let's suppose that a breakthrough in SAE results in wide agreement that a certain kind of foundationalism captures perfectly well our considered epistemic judgments (e.g., our judgments in narrow or wide reflective equilibrium). Does anything follow about how we ought to reason or about what beliefs we ought to adopt? The proponent of SAE might argue that in the given scenario, it follows that our cognitive efforts should be aimed at adopting empirical beliefs that are basic or that are appropriately related to basic beliefs – related in the way described by the account that accords with our considered epistemic judgments. But why? What's so great about philosophers' considered epistemic judgments?

Proponents of SAE might respond to this challenge as follows: "We can connect the descriptive results of SAE with normative prescriptions by noting that normative, epistemic claims are a priori. It is natural, therefore, to suppose that figuring out the truth about epistemology will involve the close analysis of our epistemic concepts. To characterize SAE as a descriptive endeavor (as you have done) might be correct, but it is misleading. The theories of SAE aim to describe an essentially normative concept, and that's why SAE is normative. To put it crudely, discovering conceptual truths involves the accurate description of (the "content" or the extension of) concepts. So discovering conceptual truths about the epistemological involves the accurate description of epistemological concepts (their content or their extension). And this is precisely what SAE does. And so even though this endeavor is descriptive (it involves describing our concepts), it nonetheless yields normative, a priori prescriptions. It tells us what it really is for a belief to be justified, and so what we ought to believe."

Let's grant for the sake of argument that epistemic claims are a priori (Bonjour 2002). It doesn't follow that SAE is the proper way to discover such a priori truths. Given that proponents of SAE disagree with each other about the nature of justification and that not all of these views can be true, we can distinguish between a priori beliefs (that are true or false) and a priori knowledge. We are willing to grant for the sake of argument that the theories of SAE give us a priori beliefs. But why suppose that they give us a priori knowledge? The history of mathematics shows that a priori truths cannot always be easily read off of our deeply considered judgments. The diversity findings discussed above (Weinberg, Nichols and Stich 2001) bring home the possibility that proponents of SAE are busy analyzing the wrong epistemic concepts. If there are significant inter-cultural and intra-cultural differences in people's epistemic concepts, then not everyone will be able to read off a priori truths from an accurate account of their use of epistemological expressions (unless, of course, one defends a crude kind of epistemological relativism). If proponents of SAE have not properly grasped the

concepts of epistemic evaluation, then no amount of aligning our general judgments with our particular judgments and no amount of armchair conceptual analysis is going to succeed in uncovering the a priori epistemological truths we seek. The specter we're raising is that proponents of SAE are like the armchair biologist who conscientiously explored his concepts and decided that he knew a priori that whales are great fish.

Consider another analogy. We share a common, folk understanding of physics that has a powerful hold on us (Carey 1985, Spelke 1994). For example, when novices explain what forces are acting on a flipped coin, they typically identify a diminishing upward force as the coin ascends, no forces (or balanced forces) when the coin reaches its peak, and then increasing downward forces as the coin drops (Clement 1982). These subjects employ a "folk" theory that is reminiscent of Aristotelian physics. The flipped coin acquires "impetus" which explains its upward motion; but the impetus is soon sapped by gravity, at which point the coin falls. This way of understanding the problem is so natural that even many who have taken Newtonian mechanics will describe the flipped coin in these terms. But from the perspective of Newtonian mechanics, ignoring air resistance, the only force acting on the coin is gravity, even when the coin is moving upwards. And that force, for all practical purposes, is constant.

Suppose the naïve physicist sits down and carefully analyzes his concept of impetus. He refines, codifies and harmonizes his impetus judgments with great care and Austinian attention to linguistic detail. Now he constructs an account that captures, with total accuracy, his application of his concept of impetus. As an attempt to tell us the truth about something other than his own linguistic predilections, lovingly detailing his naïve concept of impetus is a waste of resources. By comparison, the attempt to clarify the central concepts of highly successful scientific theories – gene, function, superposition, mass – can be of great value. What's interesting and important about these notions, and what draws our attention to them, is that they do (or presume to do) real explanatory, predictive and practical work in a successful theory about how some aspect of the world works. Providing a careful account of a concept can yield worthwhile results – but only when the concept is embedded in a high quality theory. And it is here that the challenge to SAE can be put in sharp relief by comparing it to Ameliorative Psychology. Unlike the normative judgments of Ameliorative Psychology, philosophers' considered epistemic judgments have been incubated in happy isolation from what we have learned about how best to reason about significant matters.

The move to the a priori, by itself, is not going to help the proponent of SAE. Why should we believe that the theories of SAE are capable of yielding correct epistemic judgments? Why should we believe that proponents of SAE are telling us what justification really is rather than merely telling us what they think it is? As far as we can tell, there is only one line of argument for this proposition. The proponent of SAE must make a case for a special kind of expertise for himself in matters of reasoning and belief. This claim to expertise would presumably depend on the claim that the methods of SAE have allowed philosophers to home in on what knowledge or justification really is. Just as medical doctors spend years studying what disease really is and so end up with an expertise in matters of disease that others lack, philosophers with the appropriate training have an insight into knowledge and justification that others lack. Without this sort of claim to epistemological expertise, it is hard to see why one should take the edicts of SAE seriously. (That is, unless philosophers have expertise in epistemic matters, it is hard to see why anyone should care about a description of philosophers' epistemic concepts.)

Pronouncements of expertise require some kind of documentation. If proponents of SAE are experts about justification, then it is reasonable to suppose that they have some kind of acknowledged success in epistemic matters. As far as we know, however, SAE does not have a track record of successful outcomes in epistemic matters. Further, philosophers are not the only ones who study what's involved in good reasoning. Ameliorative Psychologists can also reasonably claim to have some expertise in what's involved in good and

bad reasoning. Further, as we have already noted, these scientists have many documented successes in helping people and social institutions reason better about matters of great importance.

Proponents of SAE have many of the social trappings of expertise. Their work appears in excellent philosophy journals; and within the philosophical community, many are famous and are taken very seriously. Further, as a rule, proponents of SAE are very smart. And some SAE, like some medieval theology, exhibits real intellectual virtuosity. But none of these characteristics give us reason to believe that proponents of SAE have a special kind of expertise in epistemological matters. If proponents of SAE are experts about anything, they are experts about a purely descriptive domain: their own epistemological views (and the views of others who have been similarly trained). Unless proponents of SAE can offer some evidence for thinking that they have some kind of expertise that makes their judgments about epistemic matters more worthy of trust than the judgments of East Asians, Ameliorative Psychologists, or (for that matter) plumbers, it is not clear how the proponent of SAE is going to bridge the is-ought divide.

3. The promise of Ameliorative Psychology

We have considered two different approaches to epistemology. The standard approach taken by most analytic philosophers begins with a theory that aims to describe their considered epistemic judgments. From that descriptive core, they infer (with rather too little support) that this theory yields correct epistemic prescriptions. While there is considerable debate among proponents of SAE about what theory accurately captures their considered epistemic judgments, there is widespread agreement (at least implicit agreement) that their considered judgments are the correct judgments. The alternative approach we have offered begins with Ameliorative Psychology, which investigates and often discovers new ways to reason better about issues of practical significance. Our claim that certain parts of psychology are essentially normative will strike many as surprising (or perhaps outrageous). But we are not the first to have noticed this. Quine's initial failure to discuss the normative, reason-guiding potential of psychology made it easy for proponents of SAE to reject Quine's naturalism. But in some of his later writings, Quine noted the normative impetus of psychology:

Naturalization of epistemology does not jettison the normative and settle for the indiscriminate description of ongoing procedures. For me normative epistemology is a branch of engineering. It is the technology of truth-seeking, or, in a more cautiously epistemological term, prediction... (1986, pp.664-665)

Of psychology, Quine knew only behaviorism, and so the normative role he assigned to a naturalized epistemology could be no more than an instrumental, calculating device. We can do better now because psychology is better. We have Ameliorative Psychology.

In this paper, we have considered two approaches to epistemology that have the same structure – a descriptive core from which normative prescriptions are derived. They also both face the challenge endemic to naturalistic approaches to normative domains: how to overcome the is-ought divide. It is not obvious how the traditional approach of SAE will overcome the naturalist challenge. But what about the approach that takes Ameliorative Psychology seriously? This is a very large issue that we try to tackle elsewhere (see Bishop & Trout, in press). But there is some reason to be optimistic about the prospects for Ameliorative Psychology. That's because we have some reason to think that at least some of the normative recommendations of Ameliorative Psychology are correct. Good reasoning carries with it the reasonable expectation of good outcomes. It doesn't guarantee good outcomes, of course. But we have the reasonable expectation that people who reason better about the world will, generally and in the long run, have more success in the world. Good outcomes are a reliable (but not a perfectly reliable) sign of good reasoning. This empirical hook connecting good outcomes with good reasoning gives us some reason to believe that the recommendations of Ameliorative

Psychology are largely correct. Ameliorative Psychology has many documented cases of recommendations that have led to success. And by ‘success’, we don’t mean anything particularly subtle. For example, Ameliorative Psychology has recommended reasoning strategies that have led to better medical and psychiatric diagnoses. These strategies have ultimately led to better treatment and better outcomes for patients. While there is certainly a lot of work to do in bridging the is-ought gap, there is reason to be hopeful about an approach to epistemology that takes Ameliorative Psychology seriously.

Philosophy may aspire to an eternal subject matter, and its aims are doubtless noble, but it also has a sober social history. This history shapes its methods and the contours of its specializations. At least some of these social forces are nonepistemic, among them, the social isolation of philosophical work (and philosophers) from most other disciplines, a pattern of social backscratching not tied to practical successes, and a tiny repertoire of methods cultivated in graduate training. Once this repertoire has colonized a discipline, it can produce a pathology of method. Our favorite examples are now part of philosophy’s history: The Positivists’ blunt use of a verification principle that condemns the semantic and epistemic status of value claims, the ordinary language philosopher’s exclusive reliance on semantic analysis, and the imitation (both personal and professional) of Wittgensteinian exchange that substitutes arguments with stories.

We suggest that Standard Analytic Epistemology, too, suffers a pathology of method. Ameliorative Psychology now gives us the resources to evaluate the recommendations of SAE by standards that are independent of philosophers’ perhaps idiosyncratic epistemic judgments. We expect there to be positive aversion to the idea that we should judge SAE by applied, empirical standards. But there is no reason to shrink from healthy empirical exposure. Either it will turn out that SAE is as smugly provincial as an Elizabethan explorer’s diary, or it will emerge that psychological science vindicates the standards and methods of SAE. Either way, Ameliorative Psychology should – and in the long run will – guide these judgments.

Notes

* We would like to thank Joe Mendola, Michael Strevens and Mark Wunderlich for very helpful comments on earlier drafts of this paper. We are also grateful to the National Science Foundation for grants SES#0354536 (to MB) and SES#0327104 (to JDT) that have supported this research.

1. There are a number of serious objections to our naturalistic approach to epistemology. One objection has it that our approach is viciously circular; another contends that it inevitably falls victim to a virulent form of skepticism. Our primary goal in this article is to raise serious worries about the theories of Standard Analytic Epistemology and to suggest an alternative approach. We do not have the space here to defend our approach against serious objections. We attempt to tackle this challenge in our book, Epistemology and the Psychology of Human Judgment (in press).

2. Is Ameliorative Psychology really any more normative than any successful science? After all, physics gives us new and improved reasoning strategies all the time. That doesn’t make physics normative in any special sense. The difference is that when there is a theoretical improvement in physics, it improves our thinking only by improving our knowledge of the world. While Ameliorative Psychology does this too (e.g., Goldberg’s Rule improves our thinking about diagnosing psychiatric patients), it also improves our knowledge of ourselves as human reasoners. At its best, Ameliorative Psychology identifies how people reason about a problem and offers better ways to reason about it. And from these findings, we can draw generalizations about how we ought to reason. These generalizations can (in principle at least) put pressure on our deepest epistemological judgments about how we ought to reason.

References

- Armstrong, David. (1973) Belief Truth and Knowledge, Cambridge: Cambridge University Press, 1973.
- Bealer, George. (1996) "On the Possibility of Philosophical Knowledge," Philosophical Perspectives, 10: 1-34.
- Bishop, M., and Trout, J.D. (2005) Epistemology and the Psychology of Human Judgment, New York: Oxford University Press.
- Bonjour, Laurence. (1985) The Structure of Empirical Knowledge, Cambridge: Harvard University Press.
- BonJour, Laurence. (2002) Epistemology: Classic Problems and Contemporary Responses, Lanham, MD: Rowman and Littlefield, p.244
- Chi, M., Glaser, R., and Rees, E. (1982) "Expertise in problem solving", in R. Sternberg, ed., Advances in the Psychology of Human Intelligence. Vol. 1. Hillsdale, NJ: Erlbaum.
- Chisholm, Roderick. (1981) "A Version of Foundationalism," Midwest Studies in Philosophy, vol. 5, Minneapolis: University of Minnesota Press.
- Clement, J. (1982) "Students preconceptions in introductory mechanics," American Journal of Physics, 50 (1): 66-71.
- Daniels, Norman. (1979) "Wide reflective equilibrium and theory acceptance in ethics," Journal of Philosophy, 76:256-82.
- Dawes, Robyn. (1994) House of Cards, New York: The Free Press.
- DeRose, Keith. (1995) "Solving the Skeptical Problem," The Philosophical Review 104(1): 1-52.
- Dretske, Fred I. (1971) "Conclusive reasons," Australasian Journal of Philosophy, 49: 1-22.
- Dretske, Fred I. (1981) Knowledge & the Flow of Information, Cambridge: The MIT Press.
- Feldman, Richard. (2003) Epistemology, Upper Saddle River, NJ: Prentice-Hall.
- Foley, Richard. (1985) "What's wrong with reliabilism?" The Monist, 68: 188-202.
- Gettier, Edmund. (1963) "Is knowldge justified true belief?" Analysis, 23: 121-123.
- Gigerenzer, G., and Hoffrage, U. (1995) "How to improve Bayesian reasoning without instruction: Frequency formats," Psychological Review, 102: 684-704.
- Goldberg, L. (1968) "Simple models of simple processes? Some research on clinical judgments," American Psychologist, 23: 483-496.
- Goldman, Alvin. (1986) Epistemology and Cognition, Cambridge, MA: MIT Press.
- Goodman, Nelson. (1965) Fact, Fiction, and Forecast, Indianapolis, IN: The Bobbs-Merrill Co., Inc.
- Grove, W., and Meehl, P. (1996) "Comparative efficiency of informal (subjective, impressionistic) and formal (mechanical, algorithmic) prediction procedures: The clinical-statistical controversy," Psychology, Public Policy, and Law, 2: 293-323.
- Kahneman, D., P. Slovic, and A. Tversky, eds. (1982) Judgment Under Uncertainty: Heuristics and Biases, Cambridge: Cambridge University Press.
- Keil, F. (1989) Concepts, Kinds, and Cognitive Development, Cambridge, MA: MIT Press.
- Kim, Jaegwon. (1988) "What is Naturalized Epistemology?" in J.E. Tomberlin (Ed.), Philosophical Perspectives 2: Epistemology, Tuscadero, CA: Ridgeview Publishing.
- Lehrer, Keith. (1974) Knowledge, Oxford: Oxford University Press.
- Lehrer, K., and Paxon T. (1969) "Knowledge: Undefeated justified true belief," Journal of Philosophy, 66: 225-237.
- Lewis, David. (1996) "Elusive Knowledge," Australasian Journal of Philosophy, 74: 549-57.
- Lovie, A. D., and Lovie, P. (1986) "The flat maximum effect and linear scoring models for prediction," Journal of Forecasting, 5: 159-168.
- Pappas, G., and Swain. M. (1973) "Some Conclusive Reasons Against 'Conclusive Reasons'," in G. Pappas and M. Swain (eds.) Essays on Knowledge and Justification, Ithaca, NY: Cornell University Press, 1978, pp.61-66.
- Plantinga, Alvin. (1993) Warrant and Proper Function, New York: Oxford University Press.
- Pollock, John. (1974) Knowledge and Justification, Princeton: Princeton University Press.
- Quine, W.V.O. (1986) "Reply to Morton White," in Lewis Hahn and P. Schilpp (eds.) The Philosophy of W. V. Quine, La Salle, IL: Open Court, pp.664-665.
- Rawls, John. (1971) A Theory of Justice, Harvard University Press.
- Smith, E., and Medin, D. (1981) Categories and Concepts, Cambridge, MA: Harvard University Press.
- Unger, Peter. (1984) Philosophical Relativity, Minneapolis, MN: University of Minnesota Press.
- Weinberg, J., Nichols, S. and Stich, S. (2001) "Normativity and epistemic intuitions," Philosophical Topics, 29: 429-460.
- Williams, Michael. (2001) Problems of Knowledge: A Critical Introduction to Epistemology, New York: Oxford University Press.