

Social Epistemology: Theory and Applications

Alvin I. Goldman

Rutgers University

[To appear in Royal Institute of Philosophy, supplement]

1. Mainstream Epistemology and Social Epistemology

Epistemology has had a strongly individualist orientation, at least since Descartes. Knowledge, for Descartes, starts with the fact of one's own thinking and with oneself as subject of that thinking. Whatever else can be known, it must be known by inference from one's own mental contents. Achieving such knowledge is an individual, rather than a collective, enterprise. Descartes's successors largely followed this lead, so the history of epistemology, down to our own time, has been a predominantly individualist affair.

There are scattered exceptions. A handful of historical epistemologists gave brief space to the question of knowing, or believing justifiably, based on the testimony of others. Testimony-based knowledge would be one step into a more social epistemology. Hume took it for granted that we regularly rely on the factual statements of others, and argued that it is reasonable to do so if we have adequate reasons for trusting the veracity of these sources. However, reasons for such trust, according to Hume, must rest on personal observations of people's veracity or reliability.¹ Thomas Reid took a different view. He claimed that our natural attitude of trusting others is reasonable even if we know little if anything about others' reliability. Testimony, at least sincere testimony, is always prima facie credible (Reid, 1970: 240-241). Here we have two philosophers of the 18th century both endorsing at least one element of what nowadays is called "social epistemology." But these points did not much occupy either Hume's or Reid's corpus of philosophical writing; nor were these passages much studied or cited by their contemporaries and immediate successors.

Fast forward now to the second half of the 20th century. Here we find intellectual currents pointing toward the socializing of epistemology. Several of these movements, however, were centered outside of philosophy and never adopted the label of "social epistemology," or adopted it only belatedly. I have in mind movements in the social studies of science and cultural studies. In science studies, the most influential figure was

Thomas Kuhn, whose Structure of Scientific Revolutions (Kuhn, 1962) was itself a revolution in the interpretation of science. Kuhn influenced other investigators in the history and sociology of science to view science as just another social institution, not as a paragon of objective rationality, the stance that had been standard among positivist philosophers of science. According to Kuhn, the replacement of one scientific paradigm by a rival does not occur because the old paradigm is rationally overwhelmed by new evidence. Rival paradigms, according to Kuhn, aren't even (evidentially) "commensurable." It appeared, in other words, that Kuhn denied any rational basis for scientific revolution. Instead, it has a social basis. Although Kuhn's precise commitments were ambiguous, many thought that, according to Kuhn, paradigm change is just a matter of "mob psychology" (Lakatos, 1970: 178). In the wake of Kuhn, many historians and sociologists of science abandoned a rational perspective on science in favor of a sociological one. An earlier champion of this idea was Ludwik Fleck (1979).

Kuhn also expressed doubt about truth as the goal of science, or at least as an achievable goal, and other sociologists of science followed suit. Thus, in the 1970s, the "strong programme" in the sociology of science emerged, centered in Edinburgh, that sought to study science without any assumption that science uses methods that are either rational or superior avenues to truth. Indeed, along with postmodern thinkers like Michel Foucault, many of these authors contended that so-called scientific "facts" or "truths" are mere social constructions. What happens in scientific laboratories isn't the discovery of scientific truths, but their creation or "fabrication." Latour and Woolgar wrote:

[W]e do not conceive of scientists ... as pulling back the curtain on pre-given, but hitherto concealed truths. Rather, objects are constituted through the artful creativity of scientists. (1986: 128-129)

So-called truth, facts, rationality – they are all a matter, not of mind-independent, or society-independent nature, but of social negotiation or politics.

Most writers in sociology of science have shied away from calling their enterprise "social epistemology," but this is the label chosen by Steve Fuller, both for the title of his first book (Fuller, 1988) and for a journal he founded. Is this an apt label for this group of ideas? Well, what is epistemology? It is typically defined as the theory of knowledge. This invites the question: What is knowledge? Mainstream epistemologists universally

agree that knowledge implies truth, that knowledge is factive. If your belief isn't true, it isn't a piece of knowledge. Social constructivists, though they talk about knowledge, are characteristically dismissive or disparaging of truth. There are no facts, they maintain, only what is believed by this or that individual or community. For social constructivists, then, knowledge is simply what is believed, or at least what is communally believed.

In my book on social epistemology, Knowledge in a Social World (Goldman, 1999) I introduced a term to describe postmodernists and social constructivists of the foregoing persuasion. I called them veriphobes, because they display an aversion or abhorrence of truth. (The prefix 'veri' in 'veriphobe' is derived from the Latin 'veritas'.) The affliction from which they suffer is called veriphobia. Let me now introduce an obvious antonym of 'veriphobia,' viz., veriphilia. Mainstream epistemologists are lovers of truth; at least they are comfortable doing epistemology with the truth concept in hand. No doubt there are many philosophical problems concerning truth, both logical and metaphysical. Still, traditional epistemologists (of the last 50-60 years) help themselves to the assumption that some propositions are true, others are false, and what makes them true or false are (generally) mind-independent and community-independent facts, which we may call "truth-makers." The exact nature of truth-makers is controversial, but their characteristic independence of human construction or fabrication is taken as given.

2. Veriphobic Social Epistemology

What would social epistemology consist in under the aegis of veriphobia and veriphilia respectively? Veriphobes, at least those within the social studies of science camp, are usually interested in explanation. They want to provide social explanations, whether historical or sociological, of various knowledge-related episodes in social life, especially in science. These are knowledge-related episodes under their construal of 'knowledge,' which, as explained above, comes to little more than belief, or collective belief. The proposed explanations would deliberately make no reference to any truth-values of the beliefs in question. This was a core tenet of the strong programme, as championed by Barry Barnes and David Bloor (1982). They enunciated the "symmetry" principle, under which the same types of causes would explain both true and false beliefs, both rational and irrational beliefs.

Is a sociology of science that totally rejects truth really feasible? Do veriphobes consistently adhere to such a project? No. An explanation of any phenomenon – a genuine explanation, rather than a merely putative one – must invoke events and processes that truly occurred, and if it's a causal explanation, it must be true that those events and processes had a causal influence on the phenomenon to be explained. This is tacitly appreciated by the relatively careful scientific work that social constructivists conduct. They do minute observations of laboratory life in all (or many) of its details. They do careful historical work to unearth the causal factors that contributed to this or that episode in the history of science. What makes sense of all this diligent effort if not the goal of getting the story right, of telling it as it was, i.e., of getting the truth? So, pervasive rejection of truth cannot co-exist with their own scientific projects.

Here is a second familiar problem with the veriphobic rejection of truth or factuality. Fact-constructivism runs into the obvious problem that the world did not begin with us humans. The Big Bang and the formation of the Earth occurred before we were around. So how could *we* have constructed them? Bruno Latour was ready to bite the bullet on this question, at least on one occasion. When French sociologists working on the mummy of Ramses II, who died in 1213 BC, concluded that Ramses probably died of tuberculosis, Latour denied that this was possible. “How could he pass away due to a bacillus discovered by Robert Koch in 1881?” As Latour boldly put it, “Before Koch, the bacillus had no real existence.”²

If, as social constructivists say, a bacillus comes into existence when the scientific community comes to believe in its existence, how is a case to be handled in which scientists come to believe in a bacillus (or other scientific posit) and later abandon this belief? Did the bacillus exist for a while and then cease to exist? Or was its existence permanently assured because at one time the scientific community believed in it? What about multiple scientific communities taking different stances on the question? Does the consensus of a single community trump the counter-consensus of a second? Or do we have to count numbers of scientists? Under the latter scenario, if 273 biochemists accept the existence of a certain bacillus and 271 are skeptical, then it exists. And what if five biochemists lost in the woods and presumed dead are now found to have survived, and all reject the bacillus? Does this imply that it never existed in the first place? These are

among the conundrums that descend upon us if we adopt the crazy position of fact-constructivism.

Rather less bizarre is a somewhat analogous view about rationality or justification. Social constructivists are equally concerned to dispute the objectivity of these notions. Suppose that two people – say Galileo and Cardinal Bellarmine, his Vatican prosecutor – disagree about whether the earth circles the sun. Each claims to be justified in his belief and denies that the other is so justified. Presumably, if they were pressed to spell out the detailed basis of their justification, each might spell out a system of epistemic principles, or rules of belief justification, that countenance the present belief as justified. Although the two systems might overlap in some of their principles, there would undoubtedly be differences among them. Galileo’s system would license belief in astronomical matters based on what one sees when one looks through a telescope. It might deny the evidential relevance to astronomy of what is found in Holy Scripture. Cardinal Bellarmine’s system would feature opposing principles, principles that endorse the evidential relevance (indeed, decisiveness) of Holy Scripture and dispute the evidential relevance of telescopic observation to the properties of heavenly bodies. So, which system of epistemic principles, or rules of justification, is correct? More fundamentally, is there a matter of correctness, or more or less correctness, in the matter of epistemic systems? Are there facts of justification independent of what individuals or communities say or think? This is a serious question for epistemology.

This problem of objective rationality or justifiedness is pressed by Richard Rorty (1979). Rorty holds that there is no “objective” basis for adopting one system, or set of epistemic standards, rather than another (1979: 331). Galileo *created* the notion of “scientific values,” or standards, but the question of whether he was (objectively) “rational” in doing so is out of place (1979: 331). Paul Boghossian (2006) spells out Rorty’s position in order to criticize it. Construing Rorty as an epistemic relativist, he spells out the relativist thesis, offers an argument on behalf of the relativist, and then critiques the position. The linchpin of the argument is that there is no way for either Galileo or Cardinal Bellarmine to justify the acceptance of their respective epistemic systems. If either of those epistemic systems were objectively right, it should be possible to justify a belief in its rightness. Why can’t a belief in its rightness be justified?

Because, argues the relativist (as Boghossian presents him), any justification would have to rest on an epistemic system, presumably the same system that the protagonist begins with. But to defend a system by appeal to itself is circular, and hence illegitimate. According to the relativist, then, no such justification can be provided, and no such objective (or “absolute”) fact of justification obtains.

How successful is this argument for justificational relativism? Boghossian offers several lines of criticism, which are too complex to pursue here. I’ll offer a different line of criticism. The relativist’s argument against the possibility of justifying one’s own epistemic system rests on the charge that it involves a kind of circularity -- not “premise circularity,” in the language of epistemologists, but “epistemic circularity” -- plus the assumption that this kind of circularity is a bad or illegitimate thing. Epistemic circularity can be illustrated in connection with inductive inference. Suppose someone asks you to justify induction by showing that it is generally reliable, i.e., it generally leads to true conclusions. You reply, “Well, I used induction on occasion O_1 and it led to a true conclusion; I used induction on occasion O_2 and it again led to a true conclusion; and so forth. Therefore [conclusion], induction generally leads to true conclusions [to be interpreted as referring to past, present and future].” This attempted justification uses induction to justify induction. That’s an instance of epistemic circularity, which is a bad thing according to this relativist argument. What kind of “bad thing” is it, allegedly? Presumably, it’s bad as a tool of justification. If this is right, the relativist is assuming that there is some fact of the matter about the illegitimacy of certain patterns of inference. The relativist isn’t entitled, however, to appeal to any such justificational fact. Objective facts in matters of justification are precisely what the relativist is denying! So relativism about justificational facts is difficult to sustain and hasn’t yet been done successfully. I won’t pursue further forays into this territory.

3. Veriphiliac Social Epistemology

By my lights, justification and rationality are trickier topics than truth. So the style of objectivist social epistemology I have tried to develop (especially in Knowledge in a Social World) emphasizes truth over justification and rationality. I’ll continue this

emphasis here, without for a moment denying that justification and rationality (including group rationality) are important and legitimate topics for the field.

It is commonplace among (mainstream) epistemologists to say that our twin intellectual goals are to obtain true belief and avoid false belief (error). Instead of talking of “goals,” one can talk of “values.” One might say that it is more valuable, from an intellectual point of view, to have a true belief on some selected question than to have a false belief. If the question is whether P or not-P is the case, and the truth of the matter is P, then it’s better to believe P than to believe not-P. Another stance one might take on the question of P versus not-P is agnosticism, indecision, or withholding of judgment. On the standard line, such an attitude would be of intermediate value between believing a truth (P, in this case) and believing a falsehood (not-P, in this case). If there are such attitudes as “degrees of belief” (or degrees of confidence), we could extend these ideas and say that believing P to degree .90 has more value than believing P to degree .70. Thus, we have a ranking of possible attitudes toward a truth, such that the highest degree of belief toward the truth (1.0, full belief) has the greatest value, and every weaker degree of belief toward the same truth has a lesser value (perhaps negative, beyond a certain point).

What kind of value is this? In Knowledge in a Social World I called it veritistic value. ‘Veritistic’ connotes truth-centeredness, as opposed to a concern with justification or rationality. The latter epistemic notions express one or more different kinds of epistemic value, where the precise connections to veritistic value are controversial. In KSW, and in the remainder of my remarks here, I focus on veritistic value, or notions closely affiliated with it.

How does veritistic value link up with social epistemology? There is nothing social about a single agent having a true or false belief. Fair enough; but a wide variety of social practices and institution can have causal impacts, often immense causal influences, on the attitudes of individuals, tilting them either toward true beliefs or toward false ones. Large sectors of social interchange involve the transmission of communications – often embodying information, misinformation, partial information. The practices of communication that take place in these social networks can be studied from the vantage point of their impact on the veritistic-value states of multiple

individuals. This is how I conceive of social epistemology, at least veritistic social epistemology.³

“Sectors” of society might be divided into the (relatively) private and the (relatively) public. Perhaps these can be arranged on a continuum. At one end is the purely private sector, featuring conversational practices in which individuals convey their beliefs to other individuals (‘testimony’) or engage in argumentation to persuade others of their views. At the other end of the continuum are highly regulated public practices, for example, the carefully structured proceedings of a courtroom, where a judge oversees the speech of attorneys and witnesses, and controls the items of purported evidence that are admitted into court. Somewhere between the ends of the continuum are the communications that occur in various electronic platforms. Some facilitate individual-to-individual communication that differs little from face-to-face communication. Other electronic platforms feature more in the way of “supervisors” or “gatekeepers” of communication. Other media, ranging from conventional newspapers to weblogs, fall somewhere along the continuum. All of these sectors involve “social” practices, in an inclusive sense of the term.

In the rest of this paper I’ll concentrate on what can reasonably be considered “institutions” involving communication. In each such institution, there are indefinitely many possible ways to structure them, indefinitely many rules or procedures that might govern communicative exchange. Veritistic social epistemology is interested in how to design rules or procedures that improve veritistic outcomes. Like traditional epistemology generally, it is a normative enterprise, not a purely descriptive or explanatory one, although it may require layers of descriptive materials on which to base its normative recommendations.

4. Laws of Speech and Legal Adjudication

Ordinary statutes, constitutional provisions, common-law practices, and judicial interpretations are obvious examples of institutions that can produce better or worse veritistic outcomes. Some statutes, judicial interpretations, etc. either constitute or have definite bearings on government policies of speech and the press. One such example is a ruling by Britain’s highest court, the Law Lords, concerning British libel law (New York

Times, October 12, 2006). Under British libel law, newspapers being sued are required to prove the truth of the allegations they print – the opposite of the situation in the United States, where the burden of proof falls heavily on plaintiffs. According to many authorities, until now the odds of journalists’ winning libel cases have been stacked against them. In the recent case that prompted the high court’s ruling, the European edition of the Wall Street Journal reported that Saudi Arabia was monitoring bank accounts of prominent Saudi businesses and individuals to trace whether they were being used, possibly unwittingly, to siphon money to terrorist groups. One of the businesses sued the newspaper. The newspaper could not prove the truth of their allegations because, in the nature of things, the existence of surveillance by highly secretive Saudi authorities would have been impossible to prove by evidence in open court. Still, the paper argued that the article was in the public interest. The Law Lords agreed with this contention. One member of the panel wrote: “It is no part of the duty of the press to cooperate with any government ... in order to keep from the public information of public interest ...”. Several commentators agreed that this decision should make it easier for newspapers in the U.K. to publish serious stories where they cannot prove that allegations are true, as long as articles are responsibly reported, including the use of confidential sources.

What will be the veritistic outcomes of this change in judicial policy? Before the policy change, so it is argued, stories were not being printed – presumably true stories – because of constant fear of lawsuits. Even people from abroad sued in English courts because English judges were so sympathetic to libel plaintiffs. The judges were presumably motivated to prevent false and defamatory stories from being printed, thereby generating false beliefs. But the result of favoring libel plaintiffs was to impede the publication of true stories (in the public interest). The change in policy, therefore, will arguably have positive veritistic consequences on balance.

The veritistic analysis of legal adjudication systems focuses on a particular division of a legal system, the division responsible for determining guilt or innocence, liability or non-liability, of defendants. I shall assume that, when the law is sufficiently precise, and the true facts of the case fall determinately on one side or other of the law, then each charge brought against a defendant is either true or false. Finally, I assume (for

a defense, see Knowledge in a Social World) that the principal aim of the adjudication arm of the law is to reach accurate verdicts on the charges, given the law and the genuine facts of the case. For any given adjudication system, then, we can ask how well it succeeds in this veritistic task. How frequently does it generate truths rather than falsehoods with respect to guilt or innocence, liability or non-liability? We can also ask comparative questions of the same sort. How reliable is one style of system as compared to a different style of system, e.g., the adversary system of the Anglo-American tradition as compared with the so-called “inquisitorial” system (a very bad label, of course) of the Continental tradition? Getting more specific, we can ask how well some present version of the Anglo-American system works as compared to a version that would result if we tweaked its rules in various ways, for example, by changing the jury-selection procedure, or the instructions that judges give to jurors, or by changing some rule of evidence. All this could be asked in the spirit of contemplating actual institutional changes.

5. Problems with Forensic Laboratories: A Model Case of Veritistic Social

Epistemology

Another institution whose proper function is to (help) obtain the truth is forensic science. Unfortunately, several academic treatments indicate that this function is not being well served by current practice. Saks et al. (2001) report that erroneous and fraudulent expert evidence from forensic scientists is one of the major causes, perhaps the leading cause, of erroneous convictions of innocent persons. One rogue scientist engaged in rampant falsification for 15 years, and another faked more than 100 autopsies and falsified dozens of toxicology and blood reports (Kelly and Wearne 1998; Koppl 2006). Shocking cases are found in more than one country.

Can the error rate from forensic laboratory reports be reduced? This is a question of institutional (re-)design discussed by an economist, Roger Koppl, who offers a theoretical analysis and an experimental finding that supports this analysis. Finally, he offers a particular suggestion for improving the veritistic properties of the current system.

Koppl (2006) pinpoints the problem as the monopoly position enjoyed by most forensic laboratories vis-à-vis the legal jurisdictions that hire them. Each jurisdiction is served by one lab, and only that lab delivers reports about crime scene evidence. A

typical report says whether or not there's a match between an evidentiary item from the crime scene and a trait of the defendant, e.g., a match between a DNA sample found at the crime scene and the DNA profile of the defendant. Knowing that prosecutors prefer messages reporting a match, forensic workers have a bias toward reporting matches. Koppl analyzes the situation by means of game-theoretic models of epistemic systems. Each model contains one or more senders who search a message space and deliver a message to one or more receivers. In forensic science the receivers are jurors who hear the message delivered via testimony in open court. The jury then decides whether a fingerprint or some DNA sample left at the crime scene belongs to the defendant. This is one input into the jury's deliberation that culminates in a judgment of guilt or innocence.

On the basis of a game-theoretic analysis, Koppl argues that in the absence of competition with any other forensic lab, the bias toward reporting matches will produce a high incidence of false information. If competition were introduced into the institutional arrangement, however, e.g., by having three forensic labs produce reports, this competition would create new incentives, more unfavorable to the transmission of false information. Koppl and colleagues performed a gaming experiment designed to mimic the scenarios for forensic laboratories. This experiment confirmed a change in behavior in the predicted direction. The three-sender situation reduced the systemic error rate by two-thirds (as compared with the one-sender situation). This is a fine example of what Koppl calls "epistemic systems design," where we study the impact of system re-design on matters of veracity. It contrasts with the standard question in economics that focuses on the efficiency of institutional systems.

6. When Ignorance is Desirable: A Broadened Conception of Veriphiliac Social Epistemology

I have discussed veritistic value in terms of a specific ordering of doxastic attitudes directed at a true proposition. According to this order, higher strengths of belief in a true proposition always confer greater veritistic value with respect to that proposition (or the question that it answers). In different terminology, a state of being informed that P is veritistically preferable to being uninformed that P (e.g., withholding judgment on

P), which is veritistically preferable to being misinformed that P (believing P where it's false).

Our illustrations make it clear that many social institutions have as part of their goal or function to promote veritistically good states among occupants of certain institutional roles (with respect to selected questions). For example, legal proceedings have the goal of promoting veritistically good states in the fact-finder with respect to questions of guilt versus innocence. But not all institutions have such a goal. In fact, there are cases in which an institution ought to promote veritistically bad states in certain individuals or role-players. Being informed is not always better than being uninformed or misinformed, at least for some people in some social settings (and sometimes in purely individual settings).

What are some examples? One type of case arises from the desirability of privacy. It is generally conceded that people have rights or legitimate interests in keeping certain facts about themselves private, which means keeping other people ignorant of those facts. The relevant facts include their social security number, their cash machine PIN, their medical records, what they do in their bedroom, and so forth. If society ought to protect person's X's privacy with respect to fact F, then society should take steps to ensure that quite a few people other than X -- most people, in fact -- are ignorant of F.

A more novel example involves elections and democratic institutions. The United States Supreme Court, in the name of the First Amendment, has struck down efforts to restrict overall spending on election campaigns. The idea is that voters have a right to vote for their favored candidates not only by casting a ballot in the voting booth but also by supporting those candidates' electoral campaigns with dollars. The result, of course, is the corruption that ensues when elected officials "pay off" those interest groups who donated lots of money. Government is bought by the highest bidders. Of course, politicians cannot deliver the goods to their campaign contributors in so obvious a fashion. But there are plenty of opportunities to deliver in more subtle, or deniable, ways.

What is to be done? One solution on which both liberal and conservative reformers have converged is the "full information" idea. Candidates are required to reveal who is bankrolling their campaigns, and how much they are giving. If knowledge

of the bankrollers is shared with the public, the latter will theoretically be in a position to be watchdogs on the winning candidates' conduct in office.

There is also a much less well-known idea, but (by my lights) more promising. Why not require campaign contributions to be anonymous? That way, with candidates not knowing who gave them a lot of money, they won't be in a position to reward the contributors. This has been proposed by Bruce Ackerman and Ian Ayres (2002). Historically, Ackerman and Ayres point out, the secret ballot came to America only during the late nineteenth century. Previously voters cast their ballots in full view of the contesting parties, who carefully monitored each decision. Within this framework, corrupt vote buying was commonplace. The situation was transformed by the secret ballot. Once a voter could promise to vote one way and actually vote another, it wasn't easy for him to sell his vote, because vote-buyers could no longer verify the credibility of a voter's commitment. Suddenly, a voter's promise to sell his vote for money became worthless.

Ackerman and Ayres use the same logic in dealing with campaign contributions. They propose the "secret donation booth." Contributors will be barred from giving money directly to candidates. Instead they must pass their checks through a blind trust. Candidates would get access to the money deposited in their account with the blind trust, but won't be able to identify who provided the funds. Many people will, of course, claim to have contributed vast sums, but none of them will be able to prove it. Just as the secret voting booth disrupts vote buying because candidates are uncertain how a citizen actually voted, anonymous donations would disrupt influence peddling because candidates would be uncertain whether givers actually gave what they say they gave.

There are many details that would have to be handled to make the anonymity process work. Ackerman and Ayres haven't sold their proposal widely as of yet; to my knowledge, it hasn't reached the threshold of public discussion. For the sake of argument, however, suppose it's a good idea that would really work. It is then a case in which ignorance by certain people, viz., political candidates, of certain facts about others, viz., who has contributed to their campaigns and who hasn't, is an institutional desideratum. If Ackerman and Ayres are right, it is preferable from the perspective of

democratic institutions that certain crucial role players, viz., candidates for office, have veritistically inferior, not superior, positions vis-à-vis certain propositions.

How is this relevant to social epistemology? A veritistic social epistemologist might reply as follows: “Such cases should be of no interest to us, because these are cases where veritistic desiderata don’t kick in. So we should simply ignore such cases.” But there’s another possible response, involving a non-trivial re-design of the foundations of social epistemology.

Let us abandon the assumption that social epistemology (SE) should evaluate all outcomes of interest in terms of the outcome ranking discussed earlier. Under that old style of outcome ranking, true belief is always superior to withholding of judgment, for any given proposition. This ranking style is firmly tied to the notion of “veritistic value.” We now propose, however, that SE not be wedded to veritistic value. V-value could remain central to SE, just not essential to it. In some institutional contexts, we might allow, it is desirable to advocate a different informational policy, one that cuts against true belief as the best condition for all parties, one that views ignorance as preferable to knowledge for some individuals. At least this would be the preferable ranking from a social or institutional perspective. This wouldn’t imply that the individuals themselves would prefer ignorance to being informed (on the matters in question).⁴

Is this a tenable proposal for a conception of SE? What would then distinguish SE from other parts of social philosophy or social theory? Wouldn’t the contemplated change divest the enterprise of its distinctively epistemological dimension? What would it have in common with epistemology as usually conceived? Isn’t a preference for truth over error or ignorance just built into the conception of the epistemic?

What is still distinctive to SE is the focus on what I’ll call “veridoxic” states as the states of interest. A veridoxic state is a state with two components. The first component is a doxastic attitude, like belief, disbelief, and withholding of judgment. The second component is a truth value: either truth or falsity. So, each of the states described earlier in our V-value scheme are veridoxic states. Under the new proposal, SE would continue to focus on this class of states. The difference is that the new proposal would no longer have us restrict attention to the “canonical” ranking of veridoxic states associated with veritism. It would not cling to the treatment of true belief as being superior to false

belief or withholding, from a social point of view (or even from an individual point of view). We could distinguish this canonical ranking from alternative rankings, where the latter rankings also concern veridoxic states. By contrast with social constructivists and other fact-relativists, we would insist that the propositional contents of doxastic states are (typically) either true or false. We would not insist, however, that SE take a purely truth preferring (i.e., true-belief preferring) stance for all agents and all societal topics. The desirability of privacy is a sign that no such stance is warranted. Numerous other examples are readily produced. In time of war (just war, at any rate), it isn't incumbent on a society to deliver military secrets to the enemy. It's entirely legitimate to retain its secrets despite the fact that successful secrecy entails ignorance on the part of others.

I can already hear the predictable complaints of my epistemologist colleagues: "OK, it isn't good from the society's viewpoint to deliver its military secrets to the enemy. And if that society's cause is just, delivering military secrets to the enemy isn't good from the standpoint of justice. Nonetheless, it's good from an epistemic point of view to do so. Transmission of truths is always epistemically good, at least truths of interest to the hearers or recipients. That's just the distinctive nature of epistemology and the epistemic."

What shall we say, then, about the following two cases, where veritistic ends demand "anti-veritistic" means? Take Koppl's example of forensic laboratories and their relationships to courts. If Koppl's proposal were adopted of hiring multiple forensic laboratories to report on the same items of evidence, it could be that it would deter biased reporting for each laboratory to be ignorant of what the other laboratories report. Only such ignorance can guarantee that the laboratories not be complicit with one another. Notice, however, that this ignorance is a means to achieve an ultimate state of accurate judgment on the part of the fact-finder (the jury). So we cannot say that our interest in the laboratories' being ignorant of one another's reports is not of social epistemological interest, because we certainly want to regard the forensic laboratory case as a specimen problem for social epistemology.

Similarly, consider the desirability of journalists maintaining the confidentiality of their sources. To maintain confidentiality is keep the public ignorant of who these sources are. Could that possibly be a socially good informational state? Certainly, it

could be (and probably is, in many cases). Moreover, it is socially good because of the larger informational payoffs. If a source would decline to disclose publicly important information to a reporter unless his identity is kept confidential, then the public wouldn't receive the information in question. Surely the whole field of public information policy deserves to be treated under the heading of social epistemology. If keeping sources' identities secret is an epistemically illicit act – from the “get-go”, as it were – the social epistemology of this subject will be unacceptably constrained or circumscribed.

Again, the response of (some) other epistemologists is predictable. “We should distinguish intrinsic versus instrumental epistemic value. A state of affairs is intrinsically valuable from an epistemic point of view only if it has (positive) veritistic value, where true belief is superior to ignorance, for example. But this doesn't preclude the possibility that a state of affairs have instrumentally (positive) epistemic value even by being a state of ignorance rather than true belief. That's still of interest to social epistemology, so long as the final end which the ignorance promotes is a veritistically good state such as knowledge or true belief.”

This is one route that social epistemology might take. But it doesn't strike me as the best route, certainly not a required route. Notice that the sought-after states of knowledge (true belief) in the two cases just cited – the forensic laboratories case and the confidential sources case – are not sought after purely for their own sake. In each case, there is a plausible further end beyond the sought-after veritistic states. In the forensic laboratories case, it is delivering justice with respect to the criminal matters before the court. Justice is the final end, and accurate judgment by the jury is a means to that end. Similarly, the reason one wants vital information reported to the public is so they can act in the public's interest. The sought-after knowledge states are themselves not “final,” intrinsic ends. So it doesn't seem reasonable to admit these cases into the sphere of social epistemology while excluding the campaign-donation anonymity proposal.

Let me try to clarify this proposal for a modified conception of social epistemology by drawing an analogy with engineering. The science (or art) of engineering isn't responsible for the aims that various users might wish to achieve for a sought-after object or system to be engineered. Most people who want a bridge to be built would want the bridge to be very strong and capable of withstanding as much

weight as possible (relative to cost constraints). But there might be exceptions. A small country surrounded by aggressive and highly armed neighbors might prefer to have weak rather than strong bridges built over the rivers that constitute their borders. This might be seen as a means to keep invading tanks from getting across the bridges. Weak bridges would conveniently collapse under the weight of tanks. It would be a good engineering feat to have bridges designed to withstand the weight of ordinary commercial traffic but not tank traffic. In general, engineering deals with the design and production of artifacts that meet specifications independently arrived at. Engineering per se doesn't fix the desired specifications. Similarly, SE would not try to fix the specifications for desirable veridoxic states. For some purposes, ignorance (on the part of some) might be better than knowledge. SE is prepared to work with all sorts of ranking specifications. But it aims to figure out the social practices and institutional arrangements that promote higher attainments on whatever veridoxic rankings are appropriate, using normative considerations independent of SE per se.⁵

References

- Ackerman, Bruce and Ayres, Ian (2002). Voting with Dollars: A New Paradigm for Campaign Finance. New Haven: Yale University Press.
- Boghossian, Paul (2006). Fear of Knowledge. New York: Oxford University Press.
- Barnes, Barry and Bloor, David (1982). "Relativism, Rationalism and the Sociology of Knowledge," in M. Hollis and S. Lukes, eds., Rationality and Relativism. Cambridge, Mass: MIT Press.
- Fleck, Ludwik (1979). The Genesis and Development of a Scientific Fact, eds. T. J. Trenn and R. K. Merton. Chicago, Ill: University of Chicago Press.
- Fuller, Steve (1988). Social Epistemology. Bloomington, Ind: Indiana University Press.
- Goldman, Alvin (1979). "What Is Justified Belief," in G. Pappas, ed., Justification and Knowledge. Dordrecht: Reidel. Reprinted in Goldman, Liaisons: Philosophy Meets the Cognitive and Social Sciences. Cambridge, Mass: MIT Press.
- Goldman, Alvin (1986). Epistemology and Cognition. Cambridge, Mass: Harvard University Press.
- Goldman, Alvin (1999). Knowledge in a Social World. Oxford: Oxford University Press.
- Hume, David (1972). Enquiries Concerning Human Understanding and Concerning the Principles of Morals (1977), ed. L. A. Selby-Bigge, 2nd edn. Oxford: Oxford University Press.
- Kelly, J. F. and Wearne, P. (1998). Tainting Evidence: Inside the Scandals at the FBI Crime Lab. New York: Free Press.
- Koppl, Roger (2006). "Democratic Epistemics: An Experiment on How to Improve Forensic Science." <http://papers.econ.mpg.de/evo/discussionpapers/2006-09.pdf>
- Kuhn, Thomas (1962). The Structure of Scientific Revolutions. Chicago, Ill: University of Chicago Press.
- Lakatos, Imre (1970). "Falsification and the Methodology of Scientific Research Programmes," in I. Lakatos and A. Musgrave, eds., Criticism and the Growth of Knowledge (pp. 91-196). Cambridge: Cambridge University Press.
- Latour, Bruno and Woolgar, Steve (1986). Laboratory Life: The Construction of Scientific Facts. Princeton, NJ: Princeton University Press.

Reid, Thomas (1970). An Inquiry into the Human Mind, ed. Timothy Duggan. Chicago, Ill: University of Chicago Press.

Rorty, Richard (1979). Philosophy and the Mirror of Nature. Princeton, NJ: Princeton University Press.

Saks, Michael et al. (2001). "Model Prevention and Remedy of Erroneous Convictions Act," Arizona State Law Journal 33: 665-718.

Sokal, Alan and Bricmont, Jean (1998). Fashionable Nonsense: Postmodern Intellectuals' Abuse of Science. New York: Picador USA.

¹ Hume wrote: “[T]here is no species of reasoning more common, more useful, and even necessary to human life, than that which is derived from the testimony of men, and the reports of eye-witnesses and spectators... [O]ur assurance in any argument of this kind is derived from no other principle than our observation of the veracity of human testimony, and of the usual conformity of facts to the reports of witnesses.” (Hume 1972: 11)

² See Boghossian (2006: 26), who cites a quotation of Latour by Alan Sokal and Jean Bricmont (1998: 96-97).

³ There is a clear parallel between the social “practices” of veritistic social epistemology and the cognitive “processes” that play a pivotal role in the reliabilist form of individual epistemology I have advocated. (See especially Goldman, 1979, 1986.)

⁴ Sometimes even individuals have reasons to prefer ignorance to knowledge. See the case described in note 4 below.

⁵ Notice that similar considerations apply to purely individual choice situations. There are cases in which an individual might prefer being ignorant to being knowledgeable, even when no “social” desiderata are in play. For example, one might prefer to be ignorant of any intended messages that a potential blackmailer might send him. If the agent doesn’t receive or learn of the blackmailer’s message, he can’t really be blackmailed (at least if the potential blackmailer knows that the agent is ignorant). Thanks to Holly Smith for this point and (Thomas Schelling’s) example.