

Systems-Oriented Social Epistemology*

Alvin I. Goldman

Rutgers University

1. Laying the Groundwork for Social Epistemology

Social epistemology is an expanding sector of epistemology. There are many directions of expansion, however, and the rationales for them may vary. To illustrate the scope of social epistemology, consider the following topics that have occupied either whole issues or single articles in *Episteme, A Journal of Social Epistemology*:

- (1) testimony,
- (2) peer disagreement,
- (3) epistemic relativism,
- (4) epistemic approaches to democracy,
- (5) evidence in the law,
- (6) the epistemology of mass collaboration (e.g., *Wikipedia*), and
- (7) judgment aggregation.

How can social epistemology (SE) be characterized so that all of these topics fit under its umbrella? Why does each topic qualify as epistemology and in what respects is it social? This paper begins by proposing a tripartite division of SE. Under this classification scheme the first variety of SE is highly continuous with traditional epistemology, whereas the second and third varieties diverge from the tradition to a lesser or greater extent. The divergences are not so great, however, as to disqualify their inclusion under the social epistemology heading. After explaining the proposed classification, the paper examines in greater depth the third variety of SE, *systems-oriented* SE, which is the least familiar and most adventurous form of SE.¹

I shall not formulate any unique characterization of either ‘epistemology’ or ‘the social’ with which to delineate the three types of SE. The basic idea, however, is that epistemology involves the evaluation, from an epistemic perspective, of various “decisions” or “choices” by epistemic agents, or, in the case of the systems approach, the evaluation of alternative social systems from an epistemic standpoint. There are variations in the identities of the agents and systems, as well as in the precise terms and grounds of evaluation. All modes of evaluation are *epistemic*, but several kinds and structures of epistemic evaluation are admissible. For greater systematization, I introduce four parameters and possible settings of these parameters.

The four parameters are the following: (I) the options from which agents or systems make choices or selections; (II) the type of agent or system that makes the choices or selections; (III) the sources of evidence used in making doxastic choices; and (IV) the kinds of epistemic outcomes, desiderata, or norms used in the evaluations.

Types of options. In moral philosophy the objects of evaluation are typically overt actions or conduct, such as making a charitable gift or committing a theft. In epistemology, by contrast, the usual “acts” that comprise the target of evaluation are certain mental “choices,” namely, adopting some doxastic attitude toward a proposition. For example, agents choose to believe, reject, or withhold judgment on the question of whether there is extraterrestrial intelligence. (In speaking here of “choices” or “acts,” I do not mean voluntary or deliberate choices. Belief and rejection are not, for the most part, voluntary affairs.) There is more than one way to delineate doxastic options. There is the tripartite classification listed above -- belief, rejection, and withholding – and there are richer ranges of options, for example, graded beliefs or subjective probabilities which can be represented as points on the interval [0, 1]. Mainstream epistemology seeks principles for selecting doxastic attitudes under varying evidential conditions. Thus, both mainstream epistemology and significant parts of SE are interested in epistemic norms for doxastic choice. In addition to doxastic options, however, epistemology may be concerned with (A) choices of whether or what to assert, (B) choices of whether and how to search for evidence, and (C) choices among alternative institutions, arrangements, or characteristics of social systems that influence epistemic outcomes. None of these types of choices or options is purely mental. Type-(C) options are crucial to the third variety of SE, systems-oriented SE, on which this paper concentrates.

Although epistemology acquires its original impetus from questions about particular beliefs, it usually proceeds by focusing on broad categories of belief, for example, belief based on induction, or belief based on perception, memory, or testimony. Similarly, although SE may examine the epistemic properties of specific social systems, for example, the American criminal trial system, SE can be expected to ascend to more theoretical levels as well, by studying the epistemic consequences of more abstract features of social systems, such as how they employ (or decline to deploy) expertise or how they encourage or discourage various forms of communication or divisions of cognitive labor.

Types of Agents or Systems. The epistemic agents in mainstream epistemology are always individuals. This holds too in one branch of social epistemology. But if the epistemic agents are individuals, it may be asked, how does this variety of epistemology qualify as social? Something beyond the nature of the agent must qualify it as social. We return to this “something” below. A second type of SE takes group agents as its subject-matter, collective entities that make doxastic choices or decisions. The third species of SE takes epistemic systems as its subject matter. An epistemic system is a social system that houses a variety of procedures, institutions, and patterns of interpersonal influence that affect the epistemic outcomes of its members. Epistemic systems and their properties can arise and evolve in many ways. Some might be deliberately designed; others might emerge through ill-understood forms of historical evolution. Systems of legal adjudication, for example, are sometimes devised at a constitutional stage. Such systems can be designed with an explicit concern for truth-promoting or error-minimizing properties. Other epistemic systems and their properties are the products of social processes that are difficult to pinpoint. Whatever the historical

process of establishment, theorists and practitioners can engage in the epistemic appraisal of such systems. This is what interests us here.

Types of Evidential Sources. One way of presenting traditional epistemology uses the terminology of epistemic “sources.” Standard examples of such sources are perception, memory, reasoning, and introspection. These sources can be sources of knowledge, of justification, or of evidence. Here we are primarily interested in sources of evidence. What a person perceives or seems to perceive provides evidence for the truth or falsity of external-world propositions, such as whether there is a persimmon on the table. A long-neglected evidential source has become prominent in recent decades, namely testimony, or the statements one hears (or reads) from other persons. If another person testifies to the truth of P, a hearer acquires a new source of prima facie evidence for P. The precise circumstances in which testimony provides prima facie evidence is a central question in the field, but not something we seek to resolve here. To the extent that mainstream epistemology largely ignored testimony for a long period, it was wholly non-social. Contemporary mainstream epistemology, however, regards testimony as an important source of evidence. So a certain portion of contemporary epistemology is already squarely social. This segment of current mainstream epistemology is the first variety of SE distinguished here (see section 2 for details).

Epistemically Valuable States. Epistemology assesses doxastic and other choices as comparatively good or bad, proper or improper, from an epistemic point of view. So there must be schemes of epistemic valuation for making these assessments or judgments. A scheme of epistemic valuation may appeal to a set of fundamental epistemic values, which might include (i) having true beliefs, (ii) avoiding errors, (iii) having justified beliefs, (iv) having rational beliefs (or partial beliefs), and (v) having knowledge. This article adopts an ecumenical approach to SE, in which any of these states of affairs can be taken as fundamentally valuable from an epistemic standpoint.² We also admit both consequentialist and non-consequentialist approaches to epistemic norms and valuations. Systems-oriented SE adopts a predominantly consequentialist framework, in which states of affairs such as (i)-(v) are treated as valued outcomes that epistemic systems can promote or impede to a greater or lesser extent. Systems that generate better sets of epistemic outcomes merit higher epistemic ratings than alternatives. Non-consequentialist approaches are also used, however, including ones that appeal to formal theories or models (e.g., the probability calculus) to articulate norms of rationality.

Some comments are in order about the methodologies available to SE (although methodology is not one of our parameters). Our general approach is again ecumenical. SE employs both informal and formal methodologies. It examines epistemic concepts and phenomena by traditional analytical techniques, as well as by using formal approaches like the axiomatic method or mathematical modeling and computer simulation. Finally, empirical methodologies are eminently appropriate in assessing epistemic outcomes of alternative epistemic institutions. The admissibility of empirical methods indicates that SE is not confined to armchair or a priori methodologies. Thus, disciplines outside of philosophy can make important contributions to SE.

2. First Variety of SE: Individual Doxastic Agents (IDAs) with Social Evidence

With our four parameters in place, we can turn to characterizations of the three types of SE. Agents in the first variety of SE are individuals, and the options from which they choose are doxastic attitudes. Call these agents individual doxastic agents (IDAs). Doxastic choice by IDAs is, of course, the primary topic of traditional epistemology. What is initially puzzling is how part of social epistemology can take individuals as its targets of analysis. Isn't social epistemology's mission to go "beyond" individuals?

We have already responded to this question. In evaluating an IDA's doxastic choices, one commonly considers her evidence. Most evidential sources have little or nothing to do with other people, but some evidential sources do involve them. What qualifies the first sector of SE as social is that it addresses doxastic choices made in the light of social evidence. What is social evidence? For present purposes, evidence possessed by an agent is social evidence if it concerns acts of communication by others, or traces of such acts such as pages of print or messages on computer screens. Additionally, social evidence can consist in other people's doxastic states that become known to the agent.

The terms of evaluation for IDA social epistemology are quite inclusive. IDA SE studies any of the following questions: Under what conditions are social-evidence-based beliefs justified (or warranted)? Under what conditions are they rational? And under what conditions do they qualify as knowledge?

One issue in the justification category is whether testimony is a basic or derived source of evidence. According to David Hume (1977), the evidential worth of testimony arises from personally verifying earlier testimonial claims, remembering those verifications, and drawing inductive inferences from the earlier cases to the present instance of testimony. This is a reductionist view of the evidential power of testimony; it reduces such power to the combined power of observation, memory, and induction. An early anti-reductionist position was defended by another 18th-century Scottish philosopher, Thomas Reid (1983). Anti-reductionism holds that there is a separate and independent principle of testimonial justification, a principle to the effect that one is *prima facie* justified in trusting someone's testimony even without prior knowledge or justified belief about the testifier's competence and sincerity, and without prior knowledge of the competence and sincerity of people in general.³ Whether one adopts a reductionist or anti-reductionist position, an epistemological question of justification is on the table. And since the justification is based on the testimony of another person, a social source of evidence, the topic belongs to the province of IDA SE.⁴

The generic notion of testimony may be subdivided into a variety of special categories. For example, one might consider the special case of a layperson hearing testimony from an expert. How much deference should the layperson accord to the expert? Another serious problem arises when one hears testimony from many different people, including many experts, some of whom might disagree with one another. How can a layperson, L, justifiably determine which speaker has superior expertise and

therefore deserves greater credence? To determine comparative expertise, L might try to assess the speakers' past track records. But can she justifiably establish past track-records, since verifying someone's track-record in a specialized domain requires intellectual training in that domain, which is precisely what a layperson lacks (see Goldman, 2001)?

Another much-discussed problem of social evidence is peer disagreement (see Feldman and Warfield, forthcoming). This is usually discussed under the heading of rationality or reasonability. Is it ever reasonable for two intellectual peers, who take themselves to be such peers, to disagree with one another on a given proposition? For present purposes, two people count as peers (with respect to a given question or proposition) if they share the same relevant evidence and have comparable intellectual skills pertinent to that proposition. If each recognizes the other as a peer, how can they knowingly maintain different attitudes toward the target proposition? Shouldn't each recognize that her peer is as likely to be right as she is, and shouldn't she therefore continually adjust her credence in the peer's direction until they converge on the same attitude? Feldman (2007), Elga (2007) and Christensen (2007) roughly agree with this approach, whereas others -- Kelly (in press), Sosa (in press), and Lackey (in press) -- reject it. Whatever one's theoretical answer, the problem remains in the province of IDA SE, because it concerns doxastic choices for each agent given her social evidence about the peer's opinion (as well as her own).

The problem acquires wider scope by going beyond the special case of peers and considering how weights should be assigned to other people's viewpoints in general. Upon learning that someone else -- peer or non-peer -- has a different opinion than yours, how should you revise or update your credence in the target proposition? This is a special case of the general question of what doxastic attitude to adopt in light of one's total evidence. What others think is part of one's total evidence, a social part of that evidence.⁵

3. Second Variety of SE: Collective Doxastic Agents (CDAs)

The next variety of SE departs more from the mainstream by positing collective doxastic agents (CDAs) and investigating their distinctive properties. For example, it explores the prospects of collective doxastic agents having rational doxastic attitudes toward sets of related propositions. A collective epistemic agent has members, or constituents, who are themselves epistemic agents. Like their members, CDAs make judgments -- collective, or aggregate, judgments -- about the truth-values of propositions. Such collective judgments are presumably determined by its member judgments. CDAs can accept propositions, reject them, or suspend judgment about them.⁶ What makes CDA epistemology social epistemology is the collective nature of the agents under study. In this case, the type of evidence used plays no role in the classification. CDA doxastic decision-making qualifies as social whether the evidence used is social or non-social.

Many philosophers in recent years have defended the plausibility of treating collective entities as subjects of propositional attitudes (Gilbert, 1989; Tuomela, 1990; Bratman, 1993; Pettit, 2003; Schmitt, 1994; Searle, 1995).⁷ Here we are only interested in collective factual judgments, i.e., doxastic attitudes. In everyday life and public affairs something like doxastic states are often ascribed to collective entities. We speak of governments, courts, juries, commissions, corporations, and even political campaigns as “thinking,” “endorsing” or “denying” the truth of specified propositions. A football team might be described as being “confident” of winning an upcoming game; the Council of Economic Advisors might be described as “expecting” the recession to be short-lived. In 2007 the Intergovernmental Panel on Climate Change issued a public statement rating the likelihood levels of various propositions, including propositions about the impact of humans on climate change. Some propositions were rated as “very likely,” others as “likely,” and still others as “more likely than not.” Such statements seemingly express levels of credence or confidence by the panel, or commission, as a whole, qua collective doxastic agent.

Christian List and Philip Pettit (2002, forthcoming) have spearheaded a new research paradigm that focuses on the epistemological properties of collective agents. They dub this general area of research “judgment aggregation” (see List, 2005). Several kinds of epistemological questions can be asked about collective agents. One question is how collective judgments or attitudes are related to their members’ judgments: must they be responsive, in specified ways, to members’ judgments? Another question is how aggregate judgmental rationality is related to member judgmental rationality.

Special problems of rationality emerge when we reflect on the rationality of aggregate judgments. Even reasonable-looking aggregation functions like majority rule can generate an inconsistent set of aggregate judgments over a set of related propositions despite the fact that each individual’s judgment set over the same propositions is consistent. To illustrate this, consider the interconnected set of three propositions -- P, “If P then Q,” and Q -- and consider a group with three members who make judgments about each proposition. Assume that the collective judgment is determined by majority vote of the members on each proposition. It can easily happen that each member (A, B, and C) has a consistent set of attitudes toward the three propositions, yet two members accept P, two accept “If P then Q”, and two reject Q. In these circumstances, the group’s aggregate judgments across the three propositions will be inconsistent (see Figure 1).⁸

	A	B	C	Group
P	Yes	Yes	No	Yes
If P then Q	Yes	No	Yes	Yes
Q	Yes	No	No	No

Figure 1

Thus, inconsistency – and hence irrationality -- arises more easily for collective attitudes than individual attitudes, even under an attractive judgment aggregation function (in this case, proposition-wise majority rule).

What are the prospects for finding a judgment aggregation function that avoids this kind of scenario, a mapping from profiles of individual judgments into collective judgments that always outputs a rational set of collective judgments when the input is a profile of individual judgment sets that are all rational? A number of impossibility theorems in this territory have been proved. Here is one such theorem due to Dietrich and List (2007).

First, here are some standard conditions (constraints) on a plausible aggregation function:

- (1) Universal domain. The aggregation function accepts as admissible input any possible profile of fully rational individual judgment sets.
- (2) Collective rationality. The aggregation function generates as outputs fully rational collective judgment sets.
- (3) Consensus preservation. If all individuals submit the same judgment set, this is also the collective one.
- (4) Independence[systematicity]. The collective judgment on ‘p’ depends only on the individual judgments on ‘p’ [and the pattern of dependence is the same across propositions].

Theorem: Any function satisfying the conditions of universal domain, collective rationality, consensus preservation, and independence [systematicity] is a dictatorship.

Assuming that dictatorship is an unacceptable aggregation function for collective agents, it follows that there is no acceptable judgment aggregation function. This is a surprising and unsettling result, though it is not clear what its epistemically normative consequences are, or should be. Just as epistemologists may disagree about how to resolve skeptical paradoxes, social epistemologists may disagree in deciding how to react to this CDA “paradox.” At a minimum, however, this is interesting fodder for SE. Moreover, since the methodology used in studying this type of phenomenon is a formal methodology, i.e., the axiomatic method, it is an example of formal methods currently being utilized in SE.

4. Third Variety of Social Epistemology: Systems-Oriented SE

As previously indicated, a third type of SE would study a class of entities I shall call epistemic systems. These social systems are to be studied in terms of their effects on epistemic outcomes. Thus, the third variety of SE is a systems-oriented variety, which I shall call the SYSOR conception of SE. This form of SE departs fairly substantially from the tradition. As we said earlier, “epistemic system” designates a social system that houses social practices, procedures, institutions and/or patterns of interpersonal influence that affect the epistemic outcomes of its members. The outcomes typically involve IDAs

as doxastic agents, but in special cases could involve CDAs. Epistemic systems themselves are not usually CDAs, although it is not precluded that some entities might qualify as both. Even if an epistemic system coincides with a CDA, however, analyzing and appraising it as an epistemic system would belong to a different sector of SE than treating it as a CDA.

Paradigm cases of epistemic systems are formal institutions with publicly specified aims, rules, and procedures. Not all epistemic systems, though, are formal in this sense. Among social systems, formal or informal, some have a fairly explicit aim of promoting positive epistemic outcomes in their members. These systems include science, education, and journalism. The core mission of each of these systems is to elevate its community's level of truth possession, information possession, knowledge possession, or possession of justified or rational belief. The legal trial is another institution for which truth determination is a core mission.⁹ In each case social epistemology would examine the systems in question to see whether its mode of operation is genuinely conducive to the specified epistemic ends. It would also identify alternative organizational structures that might be epistemically superior to the existing systems. Systems-oriented SE would proceed in a similar fashion even with systems that do not regard epistemic improvement as their primary mission.

For a concrete illustration of how SYSOR SE might proceed, consider legal adjudication. Historically, many different traditions have evolved by which societies assign responsibility for criminal acts or property damage. In the European countries and their one-time colonies, two major traditions can be identified: the common-law (English) system and the civil-law (Continental) system, also called, respectively, the adversarial system and the "inquisitorial" system. One thing that SYSOR SE might do is compare these general systems in terms of epistemic outcomes. In the common-law tradition, the course of a trial proceeding is substantially managed by attorneys (including the prosecutor) who represent the contending parties, with a judge who serves as arbiter and a lay jury that serves as "fact-finder." In the Continental, or Inquisitorial, tradition, criminal trials are primarily conducted by a panel of professional judges (sometimes a combination of professional and lay judges), which jointly play several roles: investigator, interrogator, judge, and jury, all rolled into one.

Another difference between these types of systems is that the common-law system has highly detailed rules of evidence, many of which bar the admission of certain types of evidence despite their acknowledged relevance and reliability. Some of these exclusionary rules are rationalized on epistemic grounds, as promoting the aim of truth determination. Evidence might be excluded on the ground that it would excessively bias jurors, or because excluding this kind of evidence (e.g., hearsay evidence) might lead to the production of "better" (more probative) evidence. Whether such rules achieve their intended epistemic ends – e.g., fewer false verdicts -- is open to question. Laudan (2006) makes a lively assault on the American system's large-scale adoption of exclusionary rules. This critical assessment from an epistemic perspective is an instance of SYSOR social epistemology, whether pursued by philosophers (like Laudan) or legal scholars.

5. Investigation, Communication, and Trust

In the remainder of the paper we explore additional topics that SYSOR SE might tackle and methods it might employ.¹⁰ In some cases, we shall find, SYSOR SE overlaps to some extent with other disciplines or subdisciplines.

Consider a community with a shared interest in moving to a new location. Perhaps their present ecology has deteriorated through drought, fire, or other natural calamities. The community must first identify a desirable new location. How should it proceed? Like honeybee colonies, they might send out scouts in many directions to survey resettlement prospects and report on their findings. The community (the system) has various options on how to divide its cognitive labor in the search process. It might send single scouts, or messengers, in each of numerous directions. Or it might send teams of scouts, either because a team is needed for defense against enemies or because single individuals cannot provide accurate enough appraisals of a potential site. Perhaps no individual has the expertise to provide accurate (enough) reports on more than one or two dimensions. Or individuals might have biases that render their reports untrustworthy. For example, someone might have relatives who own land in a certain direction. His family would profit if the community locates there, and his report might be biased accordingly. The community must choose the structure and composition of the search team in order to maximize (or satisfice) the quality of information received.

Analogous choices with significant epistemic consequences are made in a variety of social systems. Edwin Hutchins (1995) analyzes the communication network used in ship navigation. He offers a vivid illustration of distributed cognition, in which visual bearings of landmarks are made by certain members of the crew, whose reports are forwarded to higher levels of decision making within the vessel. The overall operation is called the fix cycle, consisting of two major epistemic tasks: determining the present position of the ship and projecting its future position. The organization of ship navigation is a carefully designed social epistemic system, in which assigned roles and carefully honed expertise are relied upon to achieve designated epistemic ends. The general theory of how to distribute search or investigation operations, and how communication networks can optimally be built upon them, is a prime topic for SYSOR social epistemology. At the theoretical level this topic intersects with distributed artificial intelligence.

Ship navigation is not a traditional topic for epistemology. A more common topic, science, can also be approached as a social system. The institutional features of science -- its reward structure, for example -- provide one window on its epistemic characteristics. The sociologist Robert Merton (1973) noted that institutional science uses a priority rule to reward its members, so that honors and prizes (e.g., Nobel prizes) are awarded to the first individual(s) to discover or establish a major scientific fact. This reward system can influence investigational choices made by scientists, as Philip Kitcher (1993) observes (see also Strevens, 2003). Does this incentivizing feature of science have good or bad epistemic consequences? Kitcher argues that scientists with “sullied”

motives -- driven by a quest for priority rather than a “disinterested” goal of helping the scientific community -- will actually do better in terms of the community achieving its epistemic ends. If so, this is a case in which (in Adam Smith’s famous metaphor) an “invisible hand” brings about good outcomes for society as a whole even when actors – in this case scientists – pursue their private ends.¹¹

The pooling of informational resources is a pervasive practice throughout society. It is obvious, however, that people do not invariably convey accurate or sincere information to their peers. Epistemic incompetence and private interest often lead to inaccurate, insincere, deceptive, or incomplete information. To assist people in deciding whom to trust as an informant, some types of people are often designated as “reliable informants.” People who possess certain “indicator properties,” in William Craig’s (1990) terminology, are said to be worthy of credence and trust. Others do not merit such trust, at least not to the same degree. According to Steven Shapin (1994), being a gentleman in seventeenth century England was a positive marker of epistemic trustworthiness or credibility. Gentlemen were regarded as distinctively reliable informants because they had no need, in virtue of their social position, to lie or dissemble.

However, assigning indicator properties is a fallible process. Socially selected indicator properties may or may not correlate with genuine credibility, i.e., truthfulness. Those who possess these properties may not really be so credible, and those who lack them may nonetheless merit high credibility. Critics can properly challenge prevailing indicator properties of a given social system. This is one possible application of social epistemology. Several feminist epistemologists have weighed in on this issue, pointing to failed systems that wrongly deny epistemic credibility to large groups of people, for example, females in general. Miranda Fricker (1998, 2007) characterizes this phenomenon as “epistemic injustice,” because social denial of a due credibility commonly results in many disadvantages. Elizabeth Anderson (2006) argues that an important aspect of democracy is to make optimal use of all of society’s epistemic resources, without ignoring some voices for prejudicial reasons. She adopts an epistemic approach to democracy, in which democracy is fundamentally an epistemic engine.

Is it always inappropriate, then, for society to assign markers of reliability or credibility? This goes too far, I would argue. It would seem to exclude entirely proper activities such as making public the “scorecards” of financial advisors who predict the rise and fall of markets or equities. And it might exclude the journalistic practice of “fact-checking” the political campaign statements of candidates for office. Surely these are not objectionable activities. It may be replied that these examples concern the track records of particular individuals, not the legitimacy of general markers of credibility to be used in the absence of individualized information. But are general markers of credibility to be rejected as a blanket rule? Credentials like education and professional training commonly serve as indicators of competence and credibility. Such indicators are fallible, to be sure, but won’t social systems have better epistemic prospects if their members have clues to others’ reliability? Aren’t some clues better than none at all?

6. Expert Testimony in the Law

Similar issues are encountered in connection with experts and expertise. Reliance on expertise is a pervasive feature of epistemic systems, but distinguishing between genuine and faux expertise is fraught with difficulty. This section looks at problems associated with expertise in the legal sphere. Later we shall see how the cyber-era has ushered in critiques of reliance on experts and proposed replacements for such reliance.

In systems of legal adjudication, several types of actors play important roles in the identification and deployment of putative expertise, especially scientific expertise. When prosecutors seek to introduce forensic testimony into court, it is up to a judge's discretion whether to admit such testimony. Consider forensic evidence, such as fingerprint evidence and breath-analysis evidence, to which forensic "scientists" are prepared to testify. Many species of forensic evidence are entrenched in the legal system and the public's mind, so jurors give high credence to what forensic witnesses say. How reliable is this evidence? And how good are judges in deciding which kinds of forensic evidence and whose forensic testimony is of sufficiently high quality? Do judges make well-informed decisions to admit or exclude such evidence? What are their guidelines for deciding which forensic witnesses and methods should be admitted?

Latent¹² fingerprint evidence was long considered forensic science's gold standard, but it has lately become rather tarnished, as Jennifer Mnookin (2008) explains. It has been subject to increasing scrutiny, including numerous challenges to its reliability. The method of fingerprint examiners, explains Mnookin, offers less than meets the eye; its empirical validation is shockingly limited. Latent fingerprint examiners employ the ACE-V methodology, which stands for "analysis," "comparison," "evaluation," and "verification." This sounds methodical, but what does it really come to? Given a pair of prints for comparison, one found at a crime scene and one taken from a suspect, an examiner looks at them closely ("analysis"), notes both similarities and potential differences ("comparison"), and then evaluates these similarities and possible differences to reach a conclusion about whether they came from the same source ("evaluation"). A second examiner re-analyzes the same pair of prints ("verification"), though in many jurisdictions the second examiner has full knowledge of the original examiner's conclusion, so that verification is hardly independent!

Fingerprint examiners insist that ACE-V is a scientific method and offers a reliable methodology; and many courts have agreed. As Mnookin notes, however, merely labeling a process of careful looking a "methodology" does not make it one, nor does labeling it "scientific" tell us anything about its validity or error rate. In fact, fingerprint examination lacks any formalized specifications of what is required to declare a match: no minimum number of points of resemblance, etc. Moreover, fingerprint examiners employ no statistical information, and have no statistically validated standard to justify how many characteristics must be the same on two prints to warrant a conclusion of a "match." There is a shocking lack of empirical research to substantiate the claim of reliability for their so-called "method."¹³

Why do judges allow such testimony into court? What guidelines are they using? They are supposed to be governed by the Supreme Court decision of 1993, Daubert v. Merrell Dow, in which the Court delineated judges' responsibilities vis-à-vis scientific evidence in federal courts. The Court said that trial court judges must serve as gatekeepers to assure that proffered evidence is genuinely reliable and based on scientific validity. There are two problems here. First, judges may not be well prepared by virtue of their training to make such assessments of scientific validity and expertise. Second, the guidelines offered in this Supreme Court ruling (and others that followed in its wake, attempting to provide clarification) are of questionable adequacy. A substantial body of legal literature criticizes the clarity and adequacy of Daubert. The Daubert criteria and their progeny are a theoretical hodgepodge, drawing on a wide assortment of philosophers and theorists of science of divergent opinion and debatable questionable credentials.¹⁴ In short, many commentators argue that the theory is a morass (see Brewer, 1998; Haack, 2003). Thus, there is reason for grave doubts about the soundness of the relevant part of the legal epistemic system: at the level of fingerprint examiners, at the level of judges, and at the level of the reigning criteria for admissibility. It cannot be said that a sound system for handling science-based testimony in the law is in place. The kind of critique that lawyers like Mnookin perform belongs in the SYSOR category of social epistemology. Although her critique is not spelled out specifically in terms of epistemic outcomes, she does highlight the centrality of a method's reliability as a standard for its admissibility as scientific evidence.)

There are other problems afflicting expert forensic testimony in the courts. According to Roger Koppl and colleagues (2008), crime laboratories are part of an institutional structure that probably fosters bias. Forensic laboratories have a monopoly position on the analysis of any evidence sent to them: once a given laboratory receives and analyzes such a body of evidence, it is unlikely that any other laboratory will examine it. A further serious problem is that crime laboratories are dependent for their business on the police, and therefore have a powerful incentive to "give" the police what they want, namely, testimony of "matches" rather than "non-matches." This does not bode well for the truth-seeking rationale of the criminal trial system. Koppl et al. therefore propose to break the monopoly structure of crime laboratories' relationship with police by periodically and randomly sending evidence to more than one laboratory. This would change the institutional setting in which crime lab reports are produced, and a game-theoretic analysis suggests that it could reduce the existing bias toward matches.

The foregoing discussion registers three worries about self-proclaimed experts. First, their expertise may be far more modest than they claim. Second, systems that utilize proffered experts may have poor methods for discriminating better ones from worse. Third, systems may hide experts' liabilities from the very decision-making agents who rely on their testimony.

Some writers register even more radical doubts about experts and expertise. They regard expertise as a myth or masquerade, behind which ideology rules. Critiques of this sort can be found in the literatures of social theory and cultural studies.¹⁵ These more radical doubts are not pursued here -- since we have our own doubts about the

epistemological bases of these doubts. Nevertheless, worries raised in the bulk of this section provide reasons to explore the prospects for epistemic systems that reduce reliance on experts. As it happens, dispensing with experts is currently advocated in many arenas, especially in systems using digital technologies. These are the subjects of the next section.

7. Pooling Information via the Internet

In this section we do not challenge the reality of expertise or deny the possibility of identifying the experts (see Goldman, 2001). Instead we consider the claim that what experts know (often) pales by comparison to the knowledge dispersed in society at large. By harvesting this dispersed knowledge, a social epistemic engine can foster better epistemic consequences than it can by relying on a small group of experts. Mass collaboration implemented on the Internet enables democratic epistemic systems to reap significant epistemic bounty. This theme is especially prominent among web utopians.

The basic idea behind the contemporary version of this approach is often credited to Friedrich Hayek (1945), an economist and political theorist. A key economic question for Hayek was how to incorporate the unorganized and dispersed knowledge that exists in society, which is far greater than that held even by well-chosen experts. Hayek argued that free economic markets offer the best hope of surpassing the latter's knowledge. A model for the best solution, he suggested, is the price system. Prices in a well-functioning economic market act as an astonishingly concise and accurate signaling device. They incorporate the dispersed knowledge and also publicize it, because the price itself operates as a signal to all.

Using the Internet, such ideas have been realized via prediction markets. Just as horserace odds are set by wagers on an upcoming race, prices in a prediction market are set by bets about the occurrence of a selected future event. These bets reflect widely dispersed information and perspectives concerning the target event. As George Bragues (2009) reports, prediction markets now exist for elections, weekend movie box office receipts, snowfall amounts, scientific discoveries, disease outbreaks, and earthquakes. In a recent empirical test, two websites based on major polls as well as two prediction markets predicted the outcome of the 2008 U.S. Presidential election, on the night before the election. Averaging the polls yielded quite accurate predictions of Obama's victory, but the two prediction markets did even better (Bragues, 2009). InTrade gave Obama a 364-174 margin in the Electoral College, missing the actual margin by just 1. Iowa Electronic Markets priced in a prediction of 7 percentage points for the popular vote margin, coming within 0.2 points of the actual spread. Such successes suggest that prediction markets are a remarkably accurate way of extracting the best available information from a large number of people, and probably exceeds the capacities of single experts or teams of experts.¹⁶

A more familiar instance of Internet-based mass collaboration is the online encyclopedia Wikipedia. Wikipedia is a species of wiki, a Web site that allows any user to add material and to edit and delete what previous users have done. Bol Leuf and Ward

Cunningham (the originator of the wiki concept) explain the rationale behind wikis in explicitly democratic terms: “Wiki is inherently democratic—every user has exactly the same capabilities as any other user” (Leuf and Cunningham, 2001, p. 15). Evidently, this is intended to be the antithesis of an expertise-based mechanism or institution, which is more “elitist.”

How should the Wikipedia system be evaluated based on current evidence? How does it compare with the traditional encyclopedia-construction system, of which Britannica is the best current product? A preeminent strength of Wikipedia is its speed in constructing entries, exploiting an enormous army of volunteers. What about accuracy? The journal *Nature* compared four pairs of articles in the respective works on various scientific topics (Giles, 2005). Here are some data from the study (Magnus, 2009):

- (1) Britannica had a mean error per article of 3.0, with a standard deviation of 2.4.
- (2) Wikipedia had a mean error per article of 3.9, with a standard deviation of 3.5.
- (3) Wikipedia contained more entries than Britannica with zero errors, but two Wikipedia articles were worse than the worst of Britannica.

In addition to having more errors overall, Wikipedia’s entries varied in accuracy more than Britannica’s entries.

What general characteristics of Wikipedia might make its entries generally accurate, and hence productive of epistemically good outcomes (e.g., true beliefs)? The key idea is that errors can be quickly found and corrected in Wikipedia because such a large number of people are working to remove them. As Don Fallis (2008) points out, however, this story is not completely satisfying. Just as errors can be easily corrected, they can also be easily introduced (either intentionally or unintentionally). Another popular story of why Wikipedia should be very reliable is that it is an example of the “wisdom of crowds” (Surowiecki, 2004). Surowiecki presents examples of large groups whose average guess about various quantities -- e.g., the weight of a fat ox in a livestock exhibition -- was extremely accurate. Similarly, when contestants on *Who Wants to Be a Millionaire* call for the assistance of the studio audience, the audience gets the right answer approximately 91 percent of the time (Surowiecki 2004: 4). Surowiecki argues that groups will be reliable when they are large, independent and diverse. This is in line with the Condorcet Jury Theorem. Does this apply to Wikipedia? True, a large number of people contribute to Wikipedia, but typically only a few of these people work on any given entry (Sunstein 2006: 152). Second, it is not clear how diverse or independent are the contributors to any specific entry. Third, the examples of the wisdom of crowds involve aggregation, i.e., either averaging or taking a majority vote of the independent viewpoints. Wikipedia entries, by contrast, are rarely determined in this fashion. Entries are usually edited by single individuals and the form of an entry at each moment is a function of whoever was the last person to edit it before you looked at it. The last editor can therefore be a self-appointed dictator (Sunstein 2006: 158). So the claim that Wikipedia is especially “democratic” is open to debate.

Another social arena in which the “democracy” of the Web might vie with expert-based systems for superior epistemic outcomes is reporting the news. Traditional news media make extensive use of experts, in particular, journalists professionally trained to ferret out the news and commissioned by their news organization to report what they learn. The chief competitor of the traditional media is the blogosphere, a set of Web-based platforms that invite all comers to contribute their thoughts on the affairs of the day. With the eroding economic climate for (print) newspapers, many of which have already closed their operations, it is widely agreed that we are witnessing a transition from one kind of epistemic system for news dissemination to a very different kind of system. What are the consequences for the quality of epistemic outcomes?

Richard Posner (2005) argues that the takeover of the journalism function by the blogosphere is not inimical to the prospects of public knowledge. He puts the point primarily in terms of error detection:

[T]he blogosphere as a whole has a better error-correction machinery than the conventional media do. The rapidity with which vast masses of information are pooled and sifted leaves the conventional media in the dust. Not only are there millions of blogs, and thousands of bloggers who specialize, but, what is more, readers post comments that augment the blogs, and the information in those comments, as in the blogs themselves, zips around blogland at the speed of electronic transmission.

This means that corrections in blogs are also disseminated virtually instantaneously, whereas when a member of the mainstream media catches a mistake, it may take weeks to communicate a retraction to the public.

The charge by mainstream journalists that blogging lacks checks and balances is obtuse. The blogosphere has more checks and balances than the conventional media, only they are different. The model is Friedrich Hayek’s classic analysis of how the economic market pools enormous quantities of information efficiently despite its decentralized character, its lack of a master coordinator or regulator, and the very limited knowledge possessed by each of its participants. (2005: 10-11)

However, Posner ignores (or underplays) a crucial ingredient: investigative reporting. When there are no longer conventional journalistic enterprises, which hire reporters to investigate matters that require months of research, who will undertake this investigation? Where corruption and other public harms are underway, whether in government, business, or you name it, who will unearth these facts and disseminate them? The matter might be formulated in terms of the epistemological metaphor of “foundations” of knowledge. There cannot be “corrections” that defeat or undermine an initial journalistic story unless such a story is first reported by somebody. Unless we are content to let bloggers fabricate whatever comes into their heads, we need initial stories to be based on first-hand observation, or searching interviews with people who have observed the relevant incidents (or ongoing practices) first-hand. People involved in

corruption or other practices inimical to the public good have powerful incentives to remain silent. They will also try to silence anybody who works with them who might otherwise be willing to disclose relevant information. Traditionally, investigative reporters are the people paid and trained to unearth such facts. Abuses of the political system were uncovered by such reporters in the United States in many of the crucial annals of political history of the last 50 years (Watergate being one of the most famous). How would bloggers serve this function? So it is doubtful that the blogosphere, qua social system, can adequately replace the traditional media in terms of epistemic outcomes. The blogosphere “free-rides” on the conventional media by picking up their reportage and commenting on it. But if all of the conventional media disappear, including news-gathering agencies of all sorts (newspapers, wire services, and so on), how will the blogosphere supplant them with unpaid amateurs (Goldman, 2008)?

The Web is a platform that enormously enhances speech opportunities, a feature of cyberspace widely extolled by theorists and enthusiasts. But a closer look may reveal some problems. Freedom of speech, while rightly associated with democratic values, does not automatically solve all problems of public knowledge. Much depends on how such speech is consumed by the listening (or reading) public. Using the terminology of “social evidence,” the ready availability of the Internet implies that there is a vast array of social evidence on offer. But who will encounter which sectors of that evidence, and what use will they make of it? That is an important determinant of the distributed epistemic outcomes.

Cass Sunstein (2008) points to the problem of people increasingly making use of personally designed communication packages, a type of package that Nicholas Negroponte refers to as “the Daily Me.” The components of such a package are fully chosen in advance. When a consumer exercises this sort of control over content, with a corresponding decrease in the power of general interest intermediaries to select the content, the consumer’s prior tastes and points of view greatly narrow the evidence encountered. When reading a city or international newspaper, one comes across stories on topics one didn’t set out to read. This is educational. It is like walking down a public street where one might encounter not only like-minded friends engaged in similar activities as one’s own but a heterogeneous variety of people engaged in a wide array of activities. A system of perfect individual control of the news reduces exposure to the “public sphere,” which is important to epistemic outcomes. Reduced exposure to the public sphere may be a worrisome side-effect of the communicational ascendance of the Web, says Sunstein.

8. Computer Simulations of Social Epistemic Systems

At the end of section 1, we commented on the multiplicity of methodologies that might help ply the enterprise of SE. Among these methodologies is mathematical modelling, possibly accompanied by computer simulations that study the consequences of certain assumptions about successive interactions among individual epistemic agents. This section provides a brief and superficial overview of the work being done using this methodology.

Computer simulation is a widely used technique in social science, where investigators seek to model what transpires in systems of interacting agents. In the present case, as befits SE, the interest focuses on what transpires in a search for truth by the various agents when they receive independent evidence concerning the truth they are after and also revise their beliefs upon learning of the beliefs of selected others. It is the dynamics of these belief changes over time – and their changing relation to the truth -- that interests SE investigators. In particular, under what varieties of assumptions do the several agents converge on the truth? What proportion of the agents approach the truth, and how quickly? For example, in one model -- the “bounded-confidence” model -- it is assumed that agents assign positive weights only to other people whose opinions are “not too far away” from their own opinions.

The best-known models of opinion dynamics are due to Rainer Hegselmann and Ulrich Krause (2006, 2009). In the Hegselmann-Krause models, the truth the agents are after is a numerical value of some parameter. So the belief states of the agents consist of single numerical beliefs. This suggests straightforward ways in which an agent may revise its belief in light of other agents’ beliefs and the evidence it receives. For example, it could take the arithmetic mean of those beliefs and the evidence, which also comes in the form of a number.

Riegler and Douven (2009) study types of epistemic interaction between agents capable of having richer belief states. These agents tend to have numerous beliefs, many of them of a non-numerical nature, which are typically interconnected. Thus, they study truth-seeking agents where the truth is a theory rather than a numerical value, and where the agents receive evidence in varying degrees of informativeness about the truth. Computer simulations are then used to determine how fast and accurately such populations are able to approach the truth under differing combination of settings of the key parameters of the model, for example, the degree of informativeness of the evidence. Here is one conclusion Riegler and Douven draw from their results. Being open to interaction with other agents and giving some weight to their beliefs helps them, on average, track the truth more accurately, it also slows them down in getting within a moderately close distance of the truth as compared to when they go purely by the evidence. Of course, this kind of effect depends on the settings of the parameters. But this illustrates how exploring interesting assumptions about social epistemic systems can illuminate the likely upshots in terms of epistemic outcomes.

9. Relativism vs. Objectivism about Justificational Outcomes

Our illustrations of systems-oriented SE have focused heavily on how social systems can influence epistemic outcomes by influencing the body of social evidence that an individual encounters. Some critics might object, however, that we have ignored a very different way in which epistemic systems can influence epistemic outcomes. In addition to the causal impact epistemic systems can have on people’s epistemic outcomes

(via social evidence), they can also have a constitutive impact on such outcomes. This needs some explaining.

Focus on the justificational dimension of epistemic outcomes (justified beliefs, unjustified beliefs, and so forth). Assume that a belief's justificational status is a function of whether it conforms to the governing set of epistemic norms, norms that permit belief in light of the agent's evidential situation (see Feldman and Conee, 1985; Conee and Feldman, 2008; Goldman, forthcoming b). Throughout the paper, up to this juncture, I have assumed that there is an objectively correct set of norms that does not vary across cultures or communities. This objectively correct set of norms, in conjunction with the agent's situation, determines whether her beliefs are justified or unjustified. However, this assumption may be disputed. Critics may contend that there are no universally correct norms, only "local" norms, tacitly endorsed by this or that society or culture. (Of course, norms may overlap with one another across localities to any degree.) Because (i) a belief's justifiedness is a function of epistemic norms, (ii) local norms are the only norms there are, and (iii) local norms are a function of epistemic systems (which construct or create such norms), therefore whether and when an agent's belief is justified depends constitutively on a local epistemic system. It is not merely causally influenced by such a system or systems. So say the critics.¹⁷

There is no conflict, of course, between the two different ways in which epistemic systems can influence epistemic outcomes. There can be both causal and constitutive influence. Epistemic outcomes of the justificational type may partly be constituted by justificational norms in the sense that whether a given belief is justified is partly a matter of justificational norms (perhaps local norms of a social epistemic system). But whether that belief occurs at all may be a causal upshot of other features of a social epistemic system. There is no inconsistency or incompatibility here.

Furthermore, some types of epistemic outcomes we have delineated – specifically, the veritistic type – are not even partly constituted by epistemic norms. That a belief is true or false is not a function of whether the agent's choice of that belief in her situation conforms to local norms or to objective norms. Thus, true and false beliefs are positive and negative epistemic outcomes (respectively) that can be employed in epistemic consequentialism without worrying about the status of the governing epistemic norms.

Still, this does not help with the justification-related outcomes (or the rationality-related ones). We still face choices, it would seem, between an objectivist approach or a relativist approach, the former favoring appeal to objectively correct norms and the latter favoring appeal to local norms (norms of the agent's own community or culture.) However, why not be ecumenical? Assuming there are both objectively correct norms and local norms, why not offer theorists who wish to evaluate epistemic systems the option of using either objectivist norms or local (relativist) norms? This would not trivialize the evaluations to be made. Judging an agent's doxastic choices by the norms of her own culture or community does not guarantee her the highest grades. She may not always conform to her own culture's norms. Moreover, if we assume that justified belief is a superior epistemic outcome as compared with justified suspension of judgement (or

justified degree of belief 0.60), it is not a trivial task to achieve high-grade justificational outcomes (even judged by one's own culture's norms). This will depend on obtaining the kind of evidence that entitles one to full belief under the governing norms.

Having raised the possibility of local norms, however, readers may feel that it would set standards too high to evaluate epistemic systems by outcomes defined by different norms, viz., objective ones. If a person is acculturated in one set of norms, is it reasonable or fair to judge her doxastic choices by entirely different norms? This is particularly problematic because objective epistemic norms may not be accessible to her, and may require habits of mind that don't come "naturally."

Here is a third possible approach to justificational outcomes. We can mark the first two approaches by speaking of L-justifiedness (local justifiedness) and O-justifiedness (objective justifiedness). We can formulate a sort of "compromise" between the two by introducing a concept that uses only the terminology of O-justifiedness, but in a way that makes indirect reference to L-justifiedness.

Suppose an agent absorbs her culture's epistemic norms by heeding remarks of her parents and teachers. Assume, plausibly enough, that objectively correct norms allow a person's beliefs to be based on the testimony of such elders (who have usually proved trustworthy in the past). Then the agent may well be O-justified in believing that her local norms are correct, even if this is false.¹⁸ What shall we say, then, about the agent's beliefs that conform to local norms but violate objective norms? Well, we have to say that they are L-justified but O-unjustified. At the same time, we can add that she is O-justified in believing that she is O-justified in believing P ($J_o J_o(P)$). This is a matter of iterative O-justifiedness. Notice that being O-justified in believing that one is O-justified does not entail being O-justified (just as being O-justified in believing P does not entail P, for any random proposition P). Iterative O-justifiedness is a third type of justifiedness that is plausibly regarded as a species of positive epistemic value (see Goldman, forthcoming c). We can add it to the list of types of justifiedness that can figure in epistemic outcomes.

A concrete illustration may help. The astronomer-physicist Galileo believed that heavenly bodies move, because he could see such motion through his telescope. But the movement of heavenly bodies contradicted Scripture. So Cardinal Bellarmine, voicing the viewpoint of the prevailing culture, denied the thesis of movement. In the historical story, the Cardinal declined even to look through Galileo's telescope, because he had a better source of evidence about the makeup of the heavens, namely Holy Scripture itself. For illustrative purposes, however, let's amend the story so that the Cardinal looks through the telescope. This is an instructive revision, because now the Cardinal's evidence is very similar to Galileo's. Yet Bellarmine does not believe what he sees. How shall we appraise the justificational statuses of the two opponents? Bellarmine's belief that the heavenly bodies do not move is L-justified, assuming that local norms instruct one to believe in the dicta of Scripture. This belief is also iteratively O-justified, assuming that the objective norms instruct one to believe what one's elders say and that the Cardinal's elders said that Scripture is an objectively correct guide to belief. What

about Galileo? Assuming that objectively correct norms imply that science should be the guide to belief and that observation is a core part of science, Galileo is objectively justified (given the evidence of his eyes, when looking through the telescope) in believing that heavenly bodies do move. But he is not L-justified in believing this; and it is doubtful whether he is iteratively O-justified in believing this. (This depends on further details about Galileo that we won't pursue.) So each agent's belief has some kind of justifiedness property, but the properties differ from one another. Some lend themselves to one or another brand of relativism; others do not (Goldman, forthcoming c).

Which is the most important justifiedness property? That's a question we don't have to settle here. Indeed, this entire discussion is a bit of a digression in terms of the tripartite conception of SE. Some philosophers, however, regard it as critically important to social epistemology.¹⁹ So it should not be ignored or swept under the rug.

11. Conclusion

This article began by delineating three distinct but equally legitimate varieties of social epistemology. It then highlighted the third of these – systems-oriented social epistemology – because it is the least developed or widely understood variety. Systems-oriented SE is a flexible form of epistemological consequentialism that evaluates social epistemic systems in terms of their impact on epistemic outcomes. This variety of SE raises many issues of theoretical interest. But it also demonstrates that ethics is not the only sector of philosophy that can make helpful contributions toward solving real-world problems and therefore invites an applied subfield, viz., applied ethics. There is room in philosophical space for a substantial field of applied epistemology as well.

References

- Anderson, E. (2006). "The Epistemology of Democracy." Episteme, A Journal of Social Epistemology 3(1-2): 8-22.
- Bragues, G. (2009). "Prediction Markets: The Practical and Normative Possibilities for the Social Production of Knowledge." Episteme, A Journal of Social Epistemology 6 (1): 91-106.
- Bratman, M. (1993). "Shared Intention." Ethics 104: 97-113.
- Brewer, S. (1999). "Scientific Expert Testimony and Intellectual Due Process." The Yale Law Journal 107: 1535-1681.
- Burge, T. (1993). "Content Preservation." Philosophical Review 102: 457-488.
- Christensen, D. (2007). "Epistemology of Disagreement: the Good News." Philosophical Review 166: 187-217.
- Coady, C. A. J. (1973). "Testimony and Observation," American Philosophical Quarterly 10: 149-155.
- Coady, C. A. J. (1992). Testimony: A Philosophical Study. Oxford: Oxford University Press.
- Conee, E. and Feldman, R. (2008). "Evidence," in Q. Smith, ed., Epistemology: New Essays (pp. 83-104). Oxford: Oxford University Press.
- Conradt, L. and List, C. (2008). "Introduction: Group Decisions in Humans and Animals." Philosophical Transactions of the Royal Society, B.
- Craig, E. (1990). Knowledge and the State of Nature. Oxford: Oxford University Press.
- Daubert vs. Merrell Dow Pharmaceuticals (1993). 509 U.S. 579.
- Dietrich, F. and List, C. (2009). "The Aggregation of Propositional Attitudes: Towards a General Theory," in T. Gendler and J. Hawthorne, eds., Oxford Studies in Epistemology. (THIS VOLUME). Oxford: Oxford University Press.
- Elga, A. (2007). "Reflection and Disagreement." Nous 41(3): 478-502.
- Fallis, D. (2008). "Toward an Epistemology of Wikipedia." Journal of the American Society for Information Science and Technology 59: 1662-1674.

- Faulkner, P. (2007). "What Is Wrong with Lying?" Philosophy and Phenomenological Research 75(3): 535-557.
- Feldman, R. (200?). "Reasonable Religious Disagreement." In L. Antony, ed., Philosophers without God. Oxford: Oxford University Press.
- Feldman, R. and Conee, E. (1985). "Evidentialism," Philosophical Studies 48: 15-34.
- Feldman, R. and Warfield, T., eds. (in press). Disagreement. Oxford University Press.
- Feyerabend, P. (1978). Science in a Free Society. London: New Left Books.
- Fricker, M. (1998). "Rational Authority and Social Power: Towards a Truly Social Epistemology." Proceedings of the Aristotelian Society 98(2): 159-177.
- Fricker, M. (2007). Epistemic Injustice: Power and the Ethics of Knowing. Oxford: Oxford University Press.
- Gilbert, M. (1989). On Social Facts. London: Routledge.
- Giles, Jim (2005). "Internet Encyclopaedias Go Head to Head." Nature 438: 900-1.
- Goldman, A. I. (1999). Knowledge in a Social World. Oxford: Oxford University Press.
- Goldman, A. I. (2001). "Experts: Which Ones Should You Trust?" Philosophy and Phenomenological Research 63: 85-110.
- Goldman, A. I. (2008). "The Social Epistemology of Blogging," in J. van den Hoven and J. Weckert, eds., Information Technology and Moral Philosophy, pp. 111-122. Cambridge: Cambridge University Press.
- Goldman, A. I. (forthcoming a). "Why Social Epistemology is Real Epistemology," in D. Pritchard et al., eds., Oxford: Oxford University Press.
- Goldman, A. I. (forthcoming b). "Internalism, Externalism, and the Architecture of Justification," Journal of Philosophy.
- Goldman, A. I. (forthcoming c). "Relativism and Reasonable Disagreement," in R. Feldman and T. Warfield, eds., Disagreement. Oxford: Oxford University Press.
- Goldman, A. I. and Shaked, M. (1991). "An Economic Model of Scientific Activity and Truth Acquisition," Philosophical Studies 63: 31-55.
- Haack, S. (2003). "Disentangling Daubert: An Epistemological Study in Theory and Practice." APA Newsletter on Philosophy and Law 31: 118- ?

Hayek, F. (1945). "The Use of Knowledge in Society." American Economic Review 35: 519-530.

Hegselmann, R. and Krause, U. (2006). "Truth and Cognitive Division of Labour: First Steps towards a Computer Aided Social Epistemology," Journal of Artificial Societies and Social Simulation 9(3). <http://jasss.soc.surrey.ac.uk/9/3//10.html>

Hegselmann, R. and Krause, U. (2009). "Deliberative Exchange, Truth, and Cognitive Division of Labour: A Low-Resolution Modeling Approach," Episteme, A Journal of Social Epistemology 6(2):

Hume, D. (1977). An Enquiry Concerning Human Understanding, E. Steinberg, ed. Indianapolis: Hackett.

Hutchins, E. (1995). Cognition in the Wild. Cambridge, MA: MIT Press.

Kelly, T. (in press). "Peer Disagreement and Higher Order Evidence," in R. Feldman and T. Warfield, eds., Disagreement. Oxford University Press.

Kennedy, J. and Eberhart, R. (2001). Swarm Intelligence. San Diego: Academic Press.

Kitcher, P. (1993). The Advancement of Science. New York: Oxford University Press.

Koppl, R., Kurzban, R., and Kobilinsky, L. (2008). "Epistemics for Forensics," Episteme, A Journal of Social Epistemology 5(2): 141-159.

Kusch, M. (2002). Knowledge by Agreement. Oxford: Oxford University Press.

Lackey, J. (in press). "A Justificationist View of Disagreement's Epistemic Significance." In A. Haddock, A. Millar, and D. Pritchard, eds., Social Epistemology. Oxford: Oxford University Press.

Laudan, L. (2006). Truth, Error, and Criminal Law: An Essay in Legal Epistemology. Cambridge: Cambridge University Press.

Lehrer, K. and Wagner, C. (1981). Rational Consensus in Science and Society. Dordrecht: D. Reidel.

Leuf, B. and Cunningham, W. (2001). The Wiki Way: Quick Collaboration on the Web. Boston: Addison-Wesley.

List, C. (2005). "Group Knowledge and Group Rationality." Episteme, A Journal of Social Epistemology 2(1): 25-38.

List, C. and Pettit, P. (2002). "The Aggregation of Sets of Judgments: An Impossibility Result." Economics and Philosophy 18: 89-110.

List, C. and Pettit, P. (forthcoming). Group Agents: The Possibility, Design, and Status of Corporate Agents. Oxford: Oxford University Press.

Magnus, P. D. (2009). "On Trusting Wikipedia," Episteme, A Journal of Social Epistemology 6(1): 74-90.

Merton, R. K. (1973). The Sociology of Science. Chicago: University of Chicago Press.

Mnookin, J. (2008). "Of Black Boxes, Instruments, and Experts: Testing the Validity of Forensic Science". Episteme, A Journal of Social Epistemology 5(3): 344-358.

Moran, R. (2006). "Getting Told and Being Believed," in J. Lackey and E. Sosa, eds., The Epistemology of Testimony. Oxford: Oxford University Press.

National Academy of Sciences (2009). Free Executive Summary of Strengthening Forensic Science in the United States: A Path Forward. Washington, DC: National Academies Press.

Pettit, P. (2003). "Groups with Minds of Their Own." In F. Schmitt, ed., Socializing Metaphysics. Lanham, MD: Rowman and Littlefield.

Posner, R. (2005). "Bad News," New York Times, July 31, book review, pp. 1-11.

Reid, T. (1983). Essay on the Intellectual Powers of Man, in R. Beanblossom and K. Lehrer, eds., Thomas Reid's Inquiry and Essays. Indianapolis: Hackett.

Riegler, A. and Douven, I. (2009). "Extending the Hegselmann-Krause Model III: From Single Beliefs to Complex Belief States," Episteme, A Journal of Social Epistemology 6(2):

Schmitt, F. (1994). "The Justification of Group Beliefs." In F. Schmitt, ed., Socializing Epistemology. Lanham, MD: Roman and Littlefield.

Searle, J. (1995). The Construction of Social Reality. New York: Free Press.

Shapin, S. (1994). A Social History of Truth. Chicago, IL: University of Chicago Press.

Sosa, E. (in press). "The Epistemology of Disagreement," in A. Haddock, A. Millar, and D. Pritchard, eds., Social Epistemology. Oxford: Oxford University Press.

Strevens, M. (2003). "The Role of the Priority Rule in Science," Journal of Philosophy 100(2): 55-79.

Sunstein, C. R. (2006). Infotopia: How Many Minds Produce Knowledge. Oxford: Oxford University Press.

Sunstein, C. R. (2008). "Democracy and the Internet," in J. van den Hoven and J. Weckert, eds., Information Technology and Moral Philosophy, pp. 93-110. Cambridge: Cambridge University Press.

Surowiecki, J. (2004). The Wisdom of Crowds. New York: Doubleday.

Tuomela, T. (1990). "Can Collectives Have Beliefs?" In Acta Philosophica Fennica. Special issue on Language, Knowledge, and Intentionality: Perspectives on the Philosophy of Jaakko Hintikka, 49. L. Haaparanta, M. Kusch, and I. Niiluoto, eds.

Turner, S. (2006). "What Is the Problem with Experts?" in E. Selinger and R. P. Crease, eds., The Philosophy of Expertise (pp. 159-186). New York: Columbia University Press.

Whitcomb, D. (2007). An Epistemic Value Theory. Ph.D. dissertation. Department of Philosophy, Rutgers University, New Brunswick, NJ.

* I am grateful to the following people for many valuable comments and suggestions: Christian List, Jennifer Lackey, Dennis Whitcomb, Frank Jackson, Don Fallis, and Holly Smith. Talks based on earlier versions of this paper were presented at a conference on “The Epistemology of Liberal Democracy,” at the University of Copenhagen, Denmark, and in a symposium on social epistemology at the Pacific Division meeting of the American Philosophical Association, 2009, in Vancouver, Canada.

¹ In a companion paper (Goldman, forthcoming a), I present a slightly different tripartite taxonomy of approaches to SE. The three conceptions of SE presented there are called “Revisionist,” “Preservationist,” and “Expansionist.”

² Knowledge in a Social World (Goldman, 1999) concentrated on a “veritistic” approach to social epistemology, highlighting true belief and error avoidance as the fundamental epistemic values. Neither justifiedness, nor rationality, nor knowledge in the strong sense played important roles in that treatment. The present article, by contrast, assigns no priority to the veritistic scheme of valuation. It is often invoked in illustrating the systems-oriented approach, but in principle any plausible form of epistemic valuation is admissible.

³ In the modern era, Coady (1973, 1992) was perhaps the first to argue for anti-reductionism. Burge (1993) defends a particularly strong form of anti-reductionism.

⁴ Not all theorists of testimony agree that a hearer’s reason for believing a speaker’s testimony is that the testimony constitutes evidence. On the “assurance” view, advocated by Moran (2006) and Faulkner (2007), the reason for belief arises from the speaker’s “assuming responsibility” for the hearer’s belief. I disagree with this alternative approach, and endorse the social evidence interpretation. This is not a weighty matter here, however. Sympathizers with the assurance view might simply be asked to allow a terminological stipulation that extends the notion of “evidence” to the social act of assuming responsibility.

⁵ Lehrer and Wagner (1981) offer a specific approach to this subject. Each agent should start with a weight assignment (a measure of respect) to himself and every other agent. Then the agent should form a new opinion by taking a weighted average of all agents’ opinions from the initial stage. Finally, each agent should iterate this procedure for higher-order weights, where weights are held constant. These proposals of iteration and weight constancy are controversial. The spirit of the approach, however, fits the present variety of SE insofar as it presents principles for a single agent’s rational treatment of social evidence (the other agents’ opinions).

⁶ Perhaps CDAs can also have graded beliefs, or subjective probabilities. In a paper in THIS VOLUME, Dietrich and List (2009) present a general theory of propositional-attitude aggregation that encompasses graded beliefs of collective agents as well as binary beliefs.

⁷ Related publications outside of philosophy treat the epistemic properties of collections of infra-human animals, for example, the “intelligence” of swarms, flocks, herds, schools, and colonies of social insects and other species (Kennedy and Eberhart, 2001; Conradt and List, 2008).

⁸ Specifically, the inconsistency arises under the aggregation function of proposition-wise majority voting.

⁹ See (Goldman 1999) for a defense of the thesis that each of these enterprises (science, education, journalism, and legal adjudication) have true belief promotion (and error avoidance) as their core aim(s). Others may prefer to emphasize justified or rational belief as the primary aim. Yet others may prefer the core aim of knowledge promotion. These distinctions are unimportant under the present ecumenical approach, which does not privilege any one of these choices. Whitcomb (2007) offers a unified account of epistemic values in terms of the preeminent good of knowing. On his approach, other epistemic values like

true belief and justified belief attain their respective statuses in virtue of their “closeness” to knowledge. He calls this approach epistemism.

¹⁰ Some readers might be uncomfortable with our tripartite structure of SE, feeling that it is a heterogeneous and disunified assemblage, with little claim to being a “natural” discipline or sub-discipline. Our structure, however, closely parallels a similar assemblage of inquiries that also spans philosophy and the theoretically oriented social sciences. I have in mind here – following the suggestion of Christian List – the field of rational choice theory. The three branches of rational choice theory are decision theory, social choice theory, and game theory, and these branches bear striking parallels with our three branches of SE. This is especially true if we consider the new form of game theory called “mechanism-design theory.” Like the IDA branch of social epistemology, decision theory studies decision-making by individuals. Like the CDA branch of social epistemology, social choice theory studies the relationship between collective preference orderings and individual preference orderings. Mechanism-design theory earned Economics Nobel prizes for three economists in 2007. It aims to design social systems, institutions, or arrangements that yield some antecedently specified social desideratum, as judged by efficiency or another goal-maximizing standard. Mechanism designers try to devise a “game form” that would best achieve a specified goal under realistic assumptions about the players’ preferences and rationality. In all of these areas of rational choice theory, there is the common theme of evaluating decisions, preferences or institutions (game forms) in terms of rationality – in this case practical rather than epistemic rationality. This assemblage of inquiries strikes one as a reasonably homogeneous and unified assemblage. So I see little reason to despair of the notion that social epistemology, as presently described, is similarly well-motivated and well-integrated. Of course, the proof of the pudding is in the eating. As SE continues to grow, practitioners will see how well it works as a unified (sub-)discipline. There is already a fair amount of consensus -- among practitioners -- that it does so work.

¹¹ A somewhat similar invisible-hand theme is struck by Goldman and Shaked (1991), who show that under certain assumptions about the motives of scientific investigators and the credit-assignment practices of their peers, investigators driven by a credit-earning motive will do about as well as investigators driven by a pure truth-revelation motive in advancing scientists’ truth acquisition.

¹² Latent fingerprints are finger marks not immediately visible to the naked eye. Technicians use fingerprint powder, fuming, and other techniques to expose them and then analyze them.

¹³ Shortly after Mnookin’s article was published, a report was issued by the National Academy of Sciences that contained very similar conclusions (about many parts of forensic science). The Executive Summary of the full report stated: “there is a notable dearth of peer-reviewed, published studies establishing the scientific basis and validity of many forensic methods.” (National Academy of Sciences, 2009: S-6)

¹⁴ The National Academy of Sciences report on forensic science, addressing the impact of the Daubert ruling on forensic science testimony in the courts, remarks: “Daubert and its progeny have engendered confusion and controversy.” (National Academy of Sciences, 2009: S-8.)

¹⁵ For an overview, see Turner (2000). Some theorists are suspicious of expertise because of an apparent tension between it and liberal or democratic principles. Feyerabend (1978) argued that public science education is merely a form of state propaganda for a faction of so-called experts. Postmodernists and cultural studies writers, following Foucault, argue that expert claims about reality are the products of discursive structures that were originally expressions of patriarchy, racism, and the like.

¹⁶ For discussion, see Sunstein (2006: 129-145).

¹⁷ It may be objected that I do not correctly characterize these critics in saying that they deny that there are any universally correct epistemic norms. Don’t they accept as universally correct the following norm: “X’s belief that P is justified if and only if it meets the norms of X’s own society”? Well, perhaps they would accept this statement. But this is not a good example of what I mean by an epistemic norm. For better examples, see Goldman (forthcoming b).

¹⁸ Suppose that she never conceptualizes a distinction between cultural norms and objective norms. Everyone in her culture just assumes that the epistemic norms they accept are correct norms. The prospect of there being other norms that are “really” correct is never contemplated and never raised.

¹⁹ This is because a number of philosophers have thought that a properly socialized epistemology should be a relativistic epistemology. For a recent statement of this kind of position, see Kusch (2002).