The Social Epistemology of Blogging

Alvin I. Goldman

Rutgers University

Department of Philosophy

To appear in: Information Technology and Moral Philosophy

Jeroen van den Hoven and John Weckert, Editors

Cambridge University Press

1. Democracy and the Epistemic Properties of Internet-based Communication

The impact of the Internet on democracy is a widely discussed subject. Many writers view the Internet, potentially at least, as a boon to democracy and democratic practices. According to one popular theme, both e-mail and web pages give ordinary people powers of communication that have hitherto been the preserve of the relatively wealthy (Graham 1999, p. 79). So the Internet can be expected to close the influence gap between wealthy citizens and ordinary citizens, a weakness of many procedural democracies.

I want to focus here on another factor important to democracy, a factor that is emphasized by so-called epistemic approaches to democracy. According to epistemic democrats, democracy exceeds other systems in its ability to "track the truth". According to Rousseau, for example, the people under a democracy can track truths about the "general will" and the "common good" (Rousseau 1962, book 4). Recent proponents of epistemic democracy include Estlund (1990, 1993), Grofman and Feld (1988), and List and Goodin (2001). Their idea is that, assuming certain political outcomes are "right" or "correct", democracy is better than competing systems at choosing these outcomes.

Elsewhere I have proposed a variant on the epistemic approach to democracy (Goldman 1999, chap. 10). Epistemic theorists of democracy usually assume that, on a given issue or option, the same option or candidate is right, or correct, for all voters. A system's competence with respect to that issue is its probability of selecting the correct option. Under the variant I propose, different citizen-specific options may be right for different citizens.[1] In a given election, for example, candidate X may be the best, or right, choice for you (i.e., relative to your desiderata) and candidate Y may be the best, or right, choice for me (relative to my desiderata). Even if we make this assumption, however, we can still say what an overall good result would be from a democratic point of view. Democratically speaking, it would be good if as many citizens as possible get an outcome that is right for them. Now in the electoral situation, it might seem as if this is precisely what majority voting automatically brings about, at least in two-candidate races. If every voter votes for the candidate who is best for them, then the candidate who is best for a majority of voters will be elected, and a majority of voters will get the outcome that is right for them.[2]

But what guarantees that a voter will vote for the candidate who really is best for them? This isn't guaranteed by a procedure like majority rule. Even if candidate X is in fact best for you -- in terms of the results the two candidates would respectively produce, if elected -- you may not know that X is best for you. On the contrary, you may be mistakenly persuaded that Y is best for you. The difficulty of knowing, or truly believing,[3] which one would be best derives in part from the fact that each candidate for office tries to convince as many voters as possible that he or she is the best candidate for those voters, whether or not this is actually so. With each candidate's campaign aimed at persuading you of his or her superiority, it may not be trivial to determine (truly) who would be best according to your desiderata. Hence, it is a crucial part of a democratic framework, or system, that there be institutions, structures or mechanisms that assist

citizens in acquiring and possessing politically relevant information, where by "information possession" I mean true belief and by "politically relevant" information I mean information that is relevant to their political choices.

Which factors determine how well citizens acquit themselves in getting relevant information or knowledge on political matters (where "knowledge", like "information", entails truth)? This partly depends on citizens themselves, in ways we shall explore below. But it also depends partly on the institutional structures used in the communication or transmission of information and misinformation. This is why the media play a critical role in a democracy. It is a commonplace that democracy requires a free press. Why? Because only a free press can ferret out crucial political truths and communicate them to the public. It is the responsibility of reporters and editors to seek and publish the truth about matters of state because, as argued above, citizens' knowing the truth is crucial to their making correct decisions (correct as judged by their own desiderata). The foregoing theme expresses traditional thinking on this topic.

For the acquisition of knowledge to occur, it isn't sufficient that there be a free press that publishes or broadcasts the relevant truths. It is equally critical that members of the public receive and believe those truths. If truths are published but not read, or published and read but not believed, the public won't possess the information (or knowledge) that is important for making correct decisions. In recent years, however, there has been a waning of influence by the conventional news media in the United States, i.e., newspapers and network television news. The daily readership of newspapers dropped from 52.6 percent of adults in 1990 to 37.5 percent in 2000, and the drop was steeper in the 20-to-49-year-old cohort. This cohort is and will probably remain, as it ages, more comfortable with electronic media in general and the Web in particular (Posner 1995). Is the waning impact of the conventional media, staffed by professional journalists, a bad sign for the epistemic prospects of the voting public? Will the public's propensity to form accurate political beliefs be impaired as compared with the past, or compared with what would hold if the conventional media retained the public's trust and allegiance? This raises the question of whether the Web, or the Internet, is better or worse in epistemic terms than the conventional media, in terms of public political knowledge generated by the respective communication structures. This is the central question of this paper.

2. Epistemic Comparisons of the Conventional Media and the Blogosphere

There are many ways in which the Web, or the Internet, is used in communicating information. The Internet is a platform with multiple applications. We are not concerned here with all applications of the Internet, only with one of the more recent and influential ones, viz., blogging and its associated realm, the blogosphere. Richard Posner (2005) argues that blogging is gradually displacing conventional journalism as a source of news and the dissection of news. Moreover, Posner argues – though with some qualifications and murkiness in his message -- that this is not inimical to the public's epistemic good. The argument seems to be that blogging, as a medium of political communication and

deliberation, is no worse from the standpoint of public knowledge than conventional journalism. Posner highlights this point in the matter 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.
>
> In effect, the blogosphere is a collective enterprise – not 12 million separate enterprises, but one enterprise with 12 million reporters, feature writers and editorialists, yet almost no costs. It's as if The Associated Press or Reuters had millions of reporters, many of them experts, all working with no salary for free newspapers that carried no advertising. (Posner 2005, pp. 10-11)

In these passages Posner seems to be saying that the blogosphere is more accurate, and hence a better instrument of knowledge, than the conventional media. But elsewhere he introduces an important qualification, viz., that the bloggers are parasitical on the conventional media.

> They [the bloggers] copy the news and opinion generated by the conventional media, without picking up any of the tab. The degree of parasitism is striking in the case of those blogs that provide their readers with links to newspaper articles. The links enable the audience to read the articles without buying the newspaper. The legitimate gripes of the conventional media is not that bloggers undermine the overall accuracy of news reporting, but that they are free riders who may in the long run undermine the ability of the conventional media to finance the very reporting on which bloggers depend. (Posner 2005, p. 11)

As I would express it, the point to be learned is that we cannot compare the blogosphere and the conventional news outlets as two wholly independent and alternative communication media, because the blogosphere (in its current incarnation, at least) isn't

independent of the conventional media; it piggy-backs, or free-rides, on them. Whatever credit is due to the blogs for error correction shouldn't go to them alone, because their error-checking ability is derivative from the conventional media.

It would also be a mistake to confuse the aforementioned theme of Posner's article with the whole of his message, or perhaps even its principal point. Posner's principal point is to explain the decline of the conventional media in economic terms. Increase in competition in the news market, he says, has brought about more polarization, more sensationalism, more healthy skepticism, and, in sum, "a better matching of supply to demand" (2005, p. 11). Most people do not demand, i.e., do not seek, better quality news coverage; they seek entertainment, confirmation (of their prior views), reinforcement, and emotional satisfaction. Providers of news have been forced to give consumers what they want. This is a familiar theme from economics-minded theorists.

What this implies, however, is that Posner's analysis is only tangentially addressed to our distinctively epistemic question: Is the public better off or worse off, in terms of knowledge or true belief (on political subjects), with the current news market? Granted that the public at large isn't interested – at least not exclusively interested – in accurate political knowledge, that doesn't mean that we shouldn't take an interest in this subject. It is perfectly appropriate for theorists of democracy and public ethics to take an interest in this question, especially in light of the connection presented in section 1 between successful democracy and the citizenry's political knowledge. So let us set aside Posner's larger message and focus on the two mass communication mechanisms he identifies to see how they fare in social epistemological terms, i.e., in terms of their respective contributions to true vs. false beliefs.[4]

3. To Filter or Not to Filter?

Stay a moment longer, however, with Posner. Posner points to the familiar criticism that "unfiltered" media like blogs have bad consequences. Critics complain that blogs exacerbate social tensions by handing a powerful communication platform to extremists. Bad people find one another in cyberspace and gain confidence in their crazy ideas. The virtue of the conventional media is that they filter out extreme views. Expressing a similar idea in terms of truth-relatedness, the conventional media may be said to filter out false views, and thereby do not tempt their readership into accepting these false views, as blogs are liable to do.

Posner rejects this argument for filtering. First, he says that the argument for filtering is an argument for censorship, a first count against it. Moreover, there is little harm and some good in unfiltered media. The goods he proceeds to discuss, however, aren't linked to true belief. One good is that 12 million people write rather than stare passively at a screen. Another good is that people are allowed to blow off steam. Still another good is that it enables the authorities to keep tabs on potential troublemakers. Conceding that these may be goods, they obviously have little or nothing to do with the kind of epistemic good that interests us. The question remains open whether

communication systems that filter or those that don't have superior epistemic properties, specifically, are better at promoting true belief and/or avoiding error.

What exactly is meant by filtering? Perhaps the standard conception of filtering involves a designated channel of communication and a system of people with three kinds of roles. First, there are prospective senders, people who would like to send a message. Second, there are prospective receivers, people who might receive messages that are sent. Third, there is a filterer, or gatekeeper, an individual or group with the power to select which of the proffered messages are sent via the designated channel and which are not. When a gatekeeper disallows a proffered message, this is filtering. Although some might call any form of filtering "censorship", this term is not generally applied to all forms of filtering. Nor is filtering universally regarded as an "infringement" of speech, as censorship perhaps is.

Let me provide some examples to support these claims. Consider conventional scientific journals as examples of communication channels. Scientific journals obviously engage in filtering. Not all articles submitted for publication in a given journal are published. The function of the peer review process is precisely to select those submissions that will be published and those that won't. Nobody considers the process of peer review to be "censorship". Nor does anyone, to my knowledge, consider it an "infringement" of speech.

Another example is the (common-law) system of trial procedure. In this system, the prospective speakers, or communicators, are the parties to the dispute, or their legal counsel, and witnesses called before the court. The principal receiver is the "trier of fact," often a set of jurors. The gatekeeper is the judge, who oversees the communications that occur in court. The judge applies rules of procedure and rules of evidence to decide which speakers may speak and which messages will be allowed during the trial. Only witnesses that pass certain tests are allowed to testify; only items of evidence meeting certain criteria are admitted into court; and only certain lines of argument and rebuttal, only certain lines of questioning of witnesses, are permitted. This is a massive amount of filtering. but nobody describes such filtering as "censorship", nor is it generally called an "infringement" of speech.

Furthermore, these filtering practices are commonly rationalized in terms of (something like) helping the relevant audience determine the truth. Of course, philosophers of science debate the ultimate aims of science. At a minimum, however, geological studies are undertaken to determine the truth about the geological past, and experimental studies of various sorts are aimed at ascertaining truths about causal relationships among variables. Similarly, the overarching (though not exclusive) aim of trial procedures is to enable triers of fact to judge the truth about substantive matters of fact before the court.[5] To the extent that filtering is part and parcel of those trial procedures, filtering is evidently thought to be conducive to the aim of promoting knowledge and avoiding error. Even if the current filtering system for legal evidence isn't ideal (some evidentiary exclusions aimed at truth enhancement don't really help),

most theorists would probably agree that some kind of filtering has salutary effects in terms of truth determination.

The conventional news media also employ filtering techniques. Newspapers employ fact checkers to vet a reporter's article before it is published. They often require more than a single source before publishing an article, and limit reporters' reliance on anonymous sources. These practices seem likely to raise the veritistic quality of the reports newspapers publish and hence the veritistic quality of their readers' resultant beliefs. At a minimum, they reduce the number of errors that might otherwise be reported and believed. Thus, from a veritistic point of view, filtering looks promising indeed. Isn't that an argument for the superiority of the conventional news media over blogging, so long as knowledge and error-avoidance are the ends being considered?

Let us reflect on this argument by reflecting on the nature of filtering. In order for people to believe truths and avoid believing falsehoods, some selections must be made at one or more stages in the processes that ultimately produce belief (or doxastic states generally). But at what stage of a process is selection, i.e., filtering, necessary and helpful? If we are dealing with a process that includes reporting (in philosophy, usually referred to as "testimony"), three different stages of the process may be distinguished: the reporting stage, the reception stage, and the acceptance (believing) stage. "Filtering" normally refers to the reporting stage. Some messages that prospective speakers would like to send are winnowed out by a gatekeeper, so they don't actually get transmitted over the designated channel. But we can also think of filtering as occurring at either the reception or the acceptance stage. Consider a system in which every message anybody wants to send over a designated channel is actually sent. This doesn't mean that no filtering occurs in the process as a whole. On the contrary, potential receivers can choose which messages they wish to receive, i.e., read and digest. They do this by first selecting which channels to tune in to and then selecting which messages aired or displayed on those channels to "consume" (read or listen to). This too is a kind of filtering. Finally, having read a certain number of messages on a given topic – messages with possibly inconsistent contents – readers must decide which of these messages to believe. The ones they reject can be called "filtered out" messages.

In earlier technological eras, before the Internet, public speech was usually limited, at least over the channels with large capacities. A person could stand on his soap-box and deliver his chosen message, but it wouldn't reach many hearers. A person could send a letter to anyone he wished, but only one receiver would get it. Channels reaching larger audiences – e.g., newspapers, radio and television -- were typically limited in the quantity of messages they could convey, so some filtering had to occur at the reporting stage. The Internet has changed all this, so the question arises whether filtering really needs to be done at the reporting stage. Why not eliminate filtering at this stage, as blogging and other Internet applications easily allow? As we have seen, this doesn't mean eliminating all filtering. But why not let the necessary filtering occur at the reception and acceptance stages rather than the reporting stage?

One problem lies with the reliability of the filtering. If receivers are poor at estimating the reliability of the channels over which messages are sent, they won't do a very good filtering job at the reception stage. They may regularly tune in channels with low reliability. If receivers are poor at estimating the accuracy of the specific messages they receive and read, they also won't do a very good filtering job at the acceptance stage. For receivers of this kind, it might be desirable to have filtering performed at the reporting stage – as long as this filtering would be sufficiently reliable. Presumably, the advantage of having news delivered by dedicated, well-trained professionals embedded in a rigorous journalistic enterprise is that the filtering performed before dissemination generates a high level of reliability among stories actually reported. Of course, "high" reliability doesn't mean perfect reliability. If the American public has recently become disenchanted with the press and network news because of well-publicized errors, that disenchantment may itself be an unfortunate mistake. Receivers might not have better, more reliable, sources to which to turn.

Posner is not so worried about people being excessively credulous about communications found in blogs. He is optimistic that they make accurate assessments of the reliability of such unfiltered media:

> [M]ost people are sensible enough to distrust communications in an unfiltered medium. They know that anyone can create a blog at essentially zero cost, that most bloggers are uncredentialed amateurs, that bloggers don't employ fact checkers and don't have editors and that a blogger can hide behind a pseudonym. They know, in short, that until a blogger's assertions are validated (as when the mainstream media acknowledge an error discovered by a blogger), there is no reason to repose confidence in what he says. (Posner 2005, p. 11)

This is unrealistically sanguine. People may vaguely know these things about blogs in general, but they may not be good as applying these precepts to the specific blogs that most appeal to them. Precisely because what these blogs assert often confirms their own prior views or prejudices, they may repose excessive trust in them. Moreover, it is noteworthy that Posner concedes the need to "validate" a blogger's assertions. But how is an assertion to be "validated" except by recourse to a different, and more reliable, source? Posner's example of such a validation is an error concession by a mainstream medium. But this points precisely to the necessity of using a mainstream medium, a filtered medium! If we are trying to compare the veritistic credentials of a pure blogosphere with a pure set of mainstream media, this hardly vindicates the pure blogosphere, because without the mainstream media to appeal to for validation, Posner implicitly concedes that the reader can't know whom to trust.

4. Blogging as an Adversarial Process

Of course, the reliability of the blogosphere shouldn't be identified with the reliability of a single blog. The presumptive virtue of the blogosphere is that it's a system of blogs with checks and balances that are collectively stronger than the filtering mechanisms of the conventional media. Posner draws an analogy to the way the

economic market pools enormous quantities of information without a master regulator. Another analogy worth examining is the adversarial system of the common-law trial procedure. Many blogs are aptly characterized as forums for the zealous advocacy of a particular political position. News is interpreted through the lens of their advocacy position, which involves lots of bias. But this doesn't mean that the blogosphere as a whole is similarly biased. There are blogs for every point of view. Maybe it's a good global system that allows these different advocates to argue for their respective points of view and lets the reader decide. Maybe this is good even in terms of truth-determination. Isn't that, after all, a primary rationale for the adversary process in the British-American trial system? Each contending party in a legal dispute is represented by an attorney who is expected to be a zealous advocate for the party he/she represents. This means arguing factual issues in a way favorable to his/her party. This sort of structure is thought by many to be a very good method of truth-determination. Many essays and quips from historical theorists (e.g., John Milton, John Stuart Mill, Oliver Wendell Holmes[6]) have bolstered the idea of a "free market for ideas" in which contending parties engage in argumentative battle from which truth is supposed to emerge.

However, a little reflection on the adversarial system in the law reveals some non-trivial differences between the system as instantiated there and in the blogosphere. First, the adversarial system in legal proceedings involves oversight by a judge who requires the advocates to abide by rules of evidence and other procedural canons. Nothing like this holds in the blogosphere. Second, the adversaries in a trial address their arguments to a neutral trier of fact, which is chosen (at least in theory) by its ability to be neutral. Advocates are permitted to disqualify potential jurors ("for cause") if they have characteristics likely to tilt them in favor of one party rather than the other. In the case of blog readers, however, there is no presumption of neutrality. Readers may be as partial or "interested" as the most extreme of bloggers. Under this scenario, is it reasonable to expect readers to be led to the truth by an adversarial proceeding? Third, a crucial difference between jurors and blog readers is that jurors are required to listen to the entire legal proceeding, including all arguments from each side. Since the litigants are systematically offered opportunities to rebut their opponents' arguments, jurors will at least be exposed to a roughly equal quantity of arguments on each side. The analogue is dramatically untrue in the case of blog users. Quite the opposite. For one thing, the number of blogs in the blogosphere is so large that readers couldn't possibly read them all even if they wanted to. Moreover, as many commentators indicate, there is a strong tendency for people with partisan positions to seek out blogs that confirm their own prior prejudices and ignore the rest. Nothing constrains them to give equal time to opposing positions. This is an important disanalogy with the trial format, and renders very dubious the truth-determining properties of the adversarial system exemplified by the blogosphere.

4. Social Mechanisms and Users' Psychological States

A major ambition of social epistemology (in the guise I have presented it) is to compare the knowledge consequences of alternative social practices, institutions, or mechanisms. In the theory of legal adjudication, for example, it might compare the

knowledge consequences of having trial procedures accord with the adversarial (common-law) system or the so-called "inquisitorial" (civil-law) system. In the latter system, typical on the Continent, the entire inquiry is conducted by judges, who gather evidence, call witnesses, interrogate the witnesses, and make final judgments. Attorneys are assigned a very secondary role. There is no battle between legal teams, as there frequently is in the common-law tradition. Social epistemology would consider each system and inquire into the accuracy rate of its verdicts. Accuracy rates, of course, are not easy to ascertain, for obvious reasons. But if accuracy is the preeminent aim of an adjudication system, we should do the best we can to gauge the accuracy propensity of each system, so as to adopt the better one (or, if large-scale institutional transformation isn't feasible, at least to make changes to improve the one we've got). This is the kind of paradigm I have been inviting us to use when comparing the conventional news system with the blogging system.

Unfortunately, as hinted earlier, matters are somewhat more complicated. One cannot associate with a system, institution, or mechanism a uniform propensity to generate a determinate level of knowledge. Much depends on the inputs to the system. What I have in mind by "inputs", in the first instance, are patterns of psychological states of its users. The users' motivations, for example, are an important subset of relevant inputs. If a system's users are highly motivated in certain ways, their use of the system may produce a high level of knowledge consequences. If they are less motivated, the resultant level of knowledge consequences may be lower (or perhaps higher).

How would this work in the case of blogging? Citizens who are highly polarized on the political spectrum will tend to want to make the opposition look bad. This seems to be true in today's America, where there is an unusually high level of polarization. One consequence of this polarization is that many people are highly motivated to gather evidence and report that evidence in a public forum, such as blogs. Assuming this evidence is genuine (true), the unearthing and publication of it over the Internet presumably increases the general level of knowledge on that particular subject. Less polarized citizens will be more passive; they won't devote as much energy to the collection of evidence, or they won't bother to transmit it over the Internet. So the epistemic power of blogging may depend in non-trivial ways on motivations that vary with the political climate. This is not equally true, arguably, with the conventional news system. In such a system, journalists and editors are motivated by their jobs and careers to perform well, and this doesn't change with the political wind. Blogging isn't a career, so the volume and intensity of blogging activity is more dependent on political drive, which is, plausibly, a more variable matter.

Another issue of motivation is people's precise reasons for reading news and news analysis in the first place. Posner (along with other commentators) claims that most people today aren't interested in knowing the truth whatever it may be, at least in political matters. In particular, they don't seek to be exposed to information that might force them to revise or overthrow their earlier opinions. They only want to have their prior opinions confirmed, or articulated more authoritatively, perhaps in order to help defend those views to others. This motivation would explain their propensity to consult only those

channels or sites where they expect to find corroborating material. This isn't said to be everybody's motivation; it isn't a universal human trait. So we are talking about a kind of motivation that is variable across times and individuals.

If this is correct, it has a theoretical bearing on the kinds of statements that can or should be made by social epistemologists (of a veritistic stripe). Statements of the following simple sort should not (commonly) be made: "System S is veritistically better (better in terms of knowledge-production) than system S*". This may be taken to imply that S is better than S* across all sets of system inputs. Since this will rarely be the case, we shall usually want to confine ourselves to a more circumspect formula: "System S is veritistically better than S* for sets of inputs of types I_1, I_2, …, I_k". Once this is clarified, the program of veritistic social epistemology can proceed as before, just more cautiously, or with greater qualification. This implies that it may be unwise to offer a categorical comparison, in veritistic terms, of conventional news versus blogging. Relativization to input sets is probably required. But this doesn't undermine the program of veritistic social epistemology; it just makes it more complex. That is hardly surprising.

References

Estlund, David (1990). Democracy without preference. Philosophical Review 49: 397-424.

Estlund, David (1993). Making truth safe for democracy. In D. Copp, J. Hampton, and J. Roemer, eds., The Idea of Democracy (pp. 71-100). New York: Cambridge University Press.

Federal Rules of Evidence. St. Paul, MN: West Group.

Goldman, Alvin (1999). Knowledge in a Social World. Oxford: Oxford University Press.

Goldman, Alvin (2001). Social epistemology. E. Zalta, ed., Stanford Encyclopedia of Philosophy. URL = http://plato.stanford.edu/archives/spr2001/entries/epistemology-social/.

Goldman, Alvin (2002). What is social epistemology? A smorgasbord of projects. In A. Goldman, Pathways to Knowledge, Private and Public (pp. 182-204). New York: Oxford University Press.

Goldman, Alvin (2004). The need for social epistemology. In B. Leiter, ed., The Future for Philosophy (pp. 182-207). Oxford: Oxford University Press.

Goldman, Alvin (2005). Legal evidence. In M. P. Golding and W. A. Edmundson, eds., The Blackwell Guide to the Philosophy of Law and Legal Theory (pp. 163-175). Malden, MA: Blackwell.

Graham, Gordon (1999). The Internet, A Philosophical Inquiry. London: Routledge.

Grofman, Bernard and Scott Feld (1983). Rousseau's general will: a Condorcetian perspective. American Political Science Review 82: 567-576.

Holmes, Oliver Wendell (1999). Abrams vs. United States 250 U.S., 616 (dissenting).

List, Christian and Robert Goodin (2001). Epistemic democracy: generalizing the Condorcet jury theorem. The Journal of Political Philosophy 9: 277-306.

Milton, John (1959). Areopagitica, a speech for the liberty of unlicensed printing (1644). In E. Sirluck, ed., Complete Prose Works of John Milton.

Posner, Richard (2005). Bad news. The New York Times, July 31, 2005. Book Review, pp. 1-11.

Rousseau, Jean-Jacques (1762). The Social Contract. In G. D. H. Cole, trans., The Social Contract and Discourses. London: Everyman/Dent.

[1] Thanks to Christian List (personal communication) for suggesting this formulation of how my approach differs from standard approaches to epistemic democracy.

[2] See Goldman 1999, chap. 10 for a detailed defense of this claim.

[3] Here, and in Goldman 1999, I understand "knowledge" in the sense of "true belief", which I consider a weak sense of "knowledge". The notion that there is such a weak sense of knowledge (in addition to a strong sense of knowledge, more commonly explored by epistemologists) is briefly defended in Goldman 1999. I expect to offer additional defense of this thesis in future writing.

[4] More precisely, this is the conception of social epistemology that I commend in Goldman 1999. I call this conception veritistic social epistemology. For discussions of alternative approaches to social epistemology, see Goldman 2001, 2002, 2004.

[5] For extended defenses of this truth-oriented, or veritistic, account of the rationale for trial procedures, see Goldman 1999, chap. 9 and Goldman 2005. In partial support of this interpretation, consider the following statement of the Federal Rules of Evidence, as a fundamental basis for the rules that follow: "These rules [of evidence] shall be construed to secure … promotion of growth and development of the law of evidence to the end that the truth may be ascertained and proceedings justly determined." (Rule 102).

[6] Milton wrote: "Let [Truth] and Falsehood grapple; who ever knew Truth put to the wors, in a free and open encounter" (1959, p. 561). Holmes wrote: "[T]he best test of truth is the power of the thought to get itself accepted in the competition of the market." (1919, p. 630)