

SOCIAL EPISTEMOLOGY

ALVIN I. GOLDMAN
University of Arizona

1. *Individual and Social Epistemology*

Epistemology has historically focused on individual inquirers conducting their private intellectual affairs independently of one another. As a descriptive matter, however, what people believe and know is largely a function of their community and culture, narrowly or broadly construed. Most of what we believe is influenced, directly or indirectly, by the utterances and writings of others. So social epistemology deserves at least equal standing alongside the individual sector of epistemology. I do not challenge the integrity or propriety of individual epistemology. I am prepared to concede that much of our perceptual knowledge, memorial knowledge, and introspective knowledge is achieved on a purely individualistic basis.¹ But given the weight and significance of social causes for a very large sector of our beliefs, these social causes should receive a much larger proportion of epistemological attention than

¹ What fixes belief *contents*, or a proper characterization of those contents, is not here at issue. In speaking of an “individualistic” basis for perceptual, memorial, and introspective beliefs I mean to identify only the causal factors that immediately produce these token *believings*.

they have traditionally received. Social factors play an increasingly important role in current theories of semantical concepts, types of theories that lie outside the scope of the current essay. A rising interest in social factors is also visible in the recent epistemological literature,² but as yet there is no consensus on how the field of social epistemology should be constructed or conceived.

A skeptic of social epistemology might reply as follows: “You base your appeal for a social epistemology on the fact that many beliefs have social causes or determinants. This conclusion might be plausible if epistemology were merely a description of the causal antecedents of our beliefs. But do you really intend this? Is epistemology merely a descriptive enterprise, which identifies the sources of our beliefs? What then happens to the traditional normative mission of epistemology, which tries to identify conditions in which credal agents are *normatively entitled* to believe, or *warranted* in believing, a proposition?” My answer is the following. I do not reject a normative conception of social epistemology; indeed, I insist upon it. Nonetheless, my conception of social epistemology is not primarily concerned with warrant or justification. It is chiefly concerned with the pursuit of truth; so I call it a *veritistic* conception of social epistemology. Truth, of course, is not wholly divorced from justification. Justified belief is often a *means* to true belief. But normative social epistemology, under my conception, is not exhausted by the theory of justification. It examines all kinds of social practices —e.g., practices

² E.g., C.A.J. Coady, *Testimony*, Oxford University Press, Oxford, 1992; Alvin Goldman, *Liaisons: Philosophy Meets the Cognitive and Social Sciences*, MIT Press, Cambridge, MA, 1992, part iii; Philip Kitcher, *The Advancement of Science*, Oxford University Press, New York, 1993; Helen Longino, *Science as Social Knowledge*, Princeton University Press, Princeton, 1990; and Frederick Schmitt (ed.), *Socializing Epistemology*, Rowman & Littlefield, Lanham, MD, 1994.

of speakers and practices of speech regulators— that have a causal influence on true or false beliefs, whether or not these practices affect the justificational status of hearers' beliefs.³

In what way is this approach normative? It is normative, evaluative, or critical because it does not merely ask what social practices are actually at work in this or that community or culture, or in this or that discipline. Instead it tries to inquire into the truth-obtaining (“veritistic”) *quality* of various practices. It asks: Is a given social practice *good* from a truth-acquisitional point of view? Does the operation of that practice *improve* or *impair* the level of knowledge of those who use it or who are affected by it? The question might be posed in comparative terms: Which of several possible alternative practices (only one of which, perhaps, is actual) would be *best* from a truth-acquisitional point of view? This conception is both evaluative, in a fairly traditional sense, and yet social insofar as it focuses on social practices.

In saying that my approach is evaluative, I do not imply, of course, that it invokes the most important, encompassing, or dispositive of values. Epistemology is appropriately concerned with only a particular species of value, intellectual value. Intellectual value is not necessarily the weightiest in the pantheon of values; it can probably be trumped, in certain contexts, by other types of values. Nonetheless, it is a value that deserves sustained attention, and is the proper focus of epistemic theorizing.⁴

³ My use of the term ‘practice’ is not derived from Wittgenstein’s conception of a practice. For present purposes, think of a social practice as any type of event sequence involving two or more agents that is involved in the production of beliefs (or other doxastic states).

⁴ For further details on this point, and on most other points in this paper, see my treatment in *Knowledge in a Social World*, Oxford

2. *Social Constructionism*

Essential to this approach is the assumption that the truth-value of a proposition —its being true or false— is (normally) independent of whether it is believed. So I assume a realist conception of truth, something quite close, indeed, to a correspondence conception of truth.⁵ This view of truth, of course, is very much out of joint with our times. Setting aside epistemologists and philosophers of language, I have in mind postmodernists and social constructionists, whose ideas have enormous popularity and influence throughout the humanities and many of the social sciences. They regard the notion of (realist) truth as a relic of a by-gone era. They agree, of course, that “the social” influences everything, indeed, that all sorts of things are mere “social constructions”. Working from a library catalogue, Ian Hacking has produced an alphabetical list of things alleged by recent writers to be socially constructed, all found in the titles of their books: Authorship, Brotherhood, the Child viewer of television, Danger, Emotions, Facts, Gender, Homosexual culture, Illness, and so forth.⁶ The air is filled with social-constructionist tales.

For epistemological purposes, especially my own “veritistic” perspective, the most significant strand of social constructionism is its characteristic denial of the extra-social existence of *facts*, or *truths*. Facts are mere “fabrications”, say social constructionists. Bruno Latour and Steve Woolgar put it this way: A scientific fact is created or constituted when a statement is no longer contested or argued

University Press, Oxford, 1999. The bulk of the material in this paper is drawn from that work.

⁵ See *Knowledge in a Social World*, chap. 2.

⁶ Hacking, *The Social Construction of What?*, Harvard University Press, Cambridge, MA, 1999.

about by the scientific community in question.⁷ “We do not conceive of scientists [...] as pulling back the curtain on pre-given but hitherto concealed truths. Rather, objects [...] are constituted through the artful creativity of scientists.”⁸

How do Latour and Woolgar arrive at their view? I would present their thinking as follows. Latour and Woolgar studied scientists in a neuroendocrinology laboratory. These scientists found traces on their recording instruments, e.g., myographs or amino acid analyzers. They created documents based on these traces, and debated and exchanged documents with scientists at other laboratories. All of this activity was social activity, at the end of which they stopped contesting certain sentences appearing in their documents. So certain statements became matters of *consensus*. These unchallenged statements are what Latour and Woolgar call “facts”. But does this prove that there are no “out-there” facts or objects, which existed prior to the scientists’ consensus and render these agreed-upon statements true or false? Not at all. There may well be facts of a belief-independent sort that make the no-longer-contested statements true; or facts of a realistic sort that make the no-longer-contested statements false. Latour and Woolgar have no way of establishing that there are no such “out-there” facts (though they claim that there are none). Latour and Woolgar rightly feel that the social domain provides a *partial* explanation of why scientists believe what they do. Undoubtedly, debate within the scientific community was the *proximate* cause of the scientists’ ultimate agreement. But these proximate causes do not exclude the existence

⁷ Bruno Latour and Steve Woolgar, *Laboratory Life: The Construction of Scientific Facts*, 2nd edition, Princeton University Press, Princeton, 1986, p. 87.

⁸ *Ibid.*, p. 129.

of more *remote* causes, viz., underlying objects or facts responsible for the traces on the scientists' instruments. These unexcluded objects or facts presumably make the scientists' hypotheses either true or false. So there is no good argument here against truth and falsity in a realistic sense.

Have I done justice to social constructionist arguments? In the case of science, some would say, there are truths that exist independently of social creation. But in many other domains, the very categories in question are socially, i.e., *conventionally*, constructed. Here, there is nothing in "nature" that makes statements true or false. It is all just social construction. Here too I would differ. Consider legal states of affairs, which seem to be paradigm examples of social fabrications or constructions. The natural world does not come subdivided into acts of manslaughter and acts of non-manslaughter; what qualifies as manslaughter is obviously socially constructed. So how can realist truth gain a foothold in this arena?

I reply as follows. What is socially or conventionally constructed are legal *categories* and their associated *criteria of application*. What qualifies as manslaughter is not ordained by nature, but by humanly made criteria or definitions. Still, this does not conflict with the fact that certain actions in the actual world *realize* or *exemplify* the criteria of manslaughter and other actions do not. So it may be true to say that Rodriguez engaged in an act of manslaughter and false to say that Martinez did. To take another example, days of the week are social artifacts unknown to nature. Nature does not divide time into Sundays, Mondays, Tuesdays, and so forth. But this does not prevent its being *true* that, for example, Flores delivered his talk at the Congress on Tuesday and *false* that he delivered the talk on Monday. Even though days of the week are conventional, there are true and false statements about them.

3. *A Framework for Social Epistemology*

How exactly should a normative framework for social epistemology be constructed? If one started from the idea that the primary goal in intellectual matters —at least in *social* affairs of the intellect— is consensus, then social practices could be evaluated by their comparative success in producing agreement. But what is so good, from an intellectual perspective, about consensus? No doubt, it maintains social peace and harmony, values not to be scoffed at. But are these *intellectual* values? I think not. A consensus on the flatness of the earth is not a notable intellectual achievement. This is why I prefer a truth-oriented, or veritistic, conception of social epistemology. Intellectually good social practices are ones that, on average, increase the truth-possession of their users, while intellectually bad social practices are ones that, on average, decrease the truth-possession of their users.

What do I mean by “truth-possession”, and is it the sort of thing that can be increased or decreased? By “possessing” a truth, I mean (roughly) *believing* a truth. On the subject of belief states, epistemologists work in two different traditions. Some assume that there are three types of doxastic states or attitudes one can adopt toward a proposition: believe it (outright), disbelieve it (i.e., believe its negation), or suspend judgment (have no opinion). Another tradition views belief as coming in degrees, perhaps any degree between 0 and 1.0. Such degrees of belief are often described as subjective probabilities, or degrees of confidence. Although my idea will work for either approach, it is clearer when applied to degrees of belief. In this framework, an increase in truth-possession occurs when some proposition p is true and a cognitive agent increases his degree of belief in p . For example, he might start with a degree of belief of .40 in p , and as a result of certain

evidence increase his degree of belief to .80. Given the truth of p , this would count as a gain in truth-possession (or degree of knowledge) vis-a-vis p . If p is false, the same increase in degree of belief in p would comprise a loss in truth-possession.

Can any practices be identified as making a positive contribution, on average, to improvements in truth-possession? Yes. Consider a practice that consists of Bayesian reasoning in accordance with *objective* likelihoods. That is, suppose that an agent observes some evidence or experimental outcome E , and considers how to revise his degrees of belief in some hypothesis H in light of E . He starts with some prior degree of belief attached to H and some conditional probabilities or likelihoods: the likelihood of getting E if H is true and the likelihood of getting E if H is false. We further assume that in addition to these *subjective* likelihoods, there are *objective* likelihoods as well. For example, an emergency room patient is given a certain diagnostic test and the result is positive. There may be an objective likelihood that the result would be positive if the patient is suffering from disease D , and another objective likelihood that the result would be positive if the patient is not suffering from disease D . Moshe Shaked and I have proved that when a credal agent's subjective likelihoods *match* the objective ones, and he uses Bayesian conditionalization from observed evidence, then there is always an objectively expected *positive change* (increase) in truth-possession.⁹

The practice thus far described has no claim to being considered "social". Consider, however, a sub-category of this practice in which the evidence in question is the tes-

⁹ This assumes that his prior is neither 0 nor 1, and the two likelihoods are not identical. See Alvin Goldman and Moshe Shaked, "An Economic Model of Scientific Activity and Truth Acquisition", *Philosophical Studies*, 63, pp. 31–55, 1991, and Goldman, *Knowledge in a Social World*, chap. 4.

timony of other people, or their written records. Bayesian updating based on the evidence of other people's testimony is plausibly considered a social practice. When the reasoner's subjective likelihoods match the objective ones, it is a practice that has an expected positive effect on truth-possession. This is not an easy practice to employ, because it isn't easy to make one's subjective likelihoods match objective ones. Nonetheless, it is a possible practice, however easy or difficult to perform, that would (on average) increase truth-possession, or "veritistic value". This is the sort of practice that veritistic social epistemology should seek to identify.¹⁰

4. *A Global Practice: The Free Marketplace of Ideas*

Now that we have some notion of what a veritistically desirable social practice might be, let us turn to some social domains in which one might try to find such practices. One interesting place to look is at a very inclusive or global level, viz., at social policies that aim to govern all speech and communication activities. This topic has been prominently addressed by Jürgen Habermas. However, Habermas's theory does not seem entirely appropriate or convincing, at least for our purposes. This is so for three reasons. First, Habermas's theory is aimed at "practical" discourse, i.e., discourse in the ethical or political arena, sometimes referred to as "normative dialogue". This restriction is unsuitable for a general social epistemology, which encompasses discourses of a purely factual, nonnormative nature.

¹⁰ Of course, practices that are comparatively easy to instantiate or execute correctly are more worth identifying than practices comparatively hard to execute correctly, like the one discussed in the text. It would be a mistake, however, to refuse to count something as a *bona fide* practice unless one can be assured of executing it correctly whenever one tries. Few things would qualify as practices under that stringent requirement.

Second, the rationale for Habermas's proposal is different from the truth-oriented, or veritistic, one being advocated here. Habermas writes: "Thus all arguments [. . .] require the same basic form of organization, which subordinates the eristic means to the end of developing *intersubjective conviction* by the force of the better argument."¹¹ In other words, Habermas is interested in a speech or dialogue policy that will maximize intersubjective agreement, or consensus. But that is not the standard that we are pursuing, given our previous arguments. Suppose, however, that we neglect these features of Habermas's view. Suppose we look at his specific speech-governing proposal from the vantage point of truth promotion, and suppose it governs all factual as well as normative subject-matter. Is it promising?

Habermas proposes the idea of an "ideal speech situation". This is succinctly described by Seyla Benhabib: "The procedural constraints of the ideal speech situation are that each participant must have an equal chance to initiate and to continue communication; each must have an equal chance to make assertions, recommendations, and explanations; all must have equal chances to express their wishes, desires, and feelings [. . .]."¹² Although this may formulate a very nice ideal, we certainly should not demand that all speech situations conform to these principles. For example, is it wrong for a newspaper to give special publication opportunities to hired editorial writers or columnists rather than ordinary citizens? In many newspapers, of course, ordinary citizens have opportunities to express their views via letters to the editor or "op ed" pieces. But ordinary citizens do not have an *equal* opportunity to express their

¹¹ *The Theory of Communicative Action*, vol. 1, Beacon Press, Boston, 1984, p. 36; italics added.

¹² "Models of Public Space: Hannah Arendt, the Liberal Tradition, and Jürgen Habermas", in C. Calhoun (ed.), *Habermas and the Public Sphere*, MIT Press, Cambridge, MA, 1992, p. 89.

thoughts and feelings as compared with hired columnists. Nor do many philosophy journals offer equal opportunities to all citizens to publish in their pages. Publication opportunities are restricted to people whose work meets certain professional standards. Does this make this speech policy, or communication policy, objectionable? That would be an extreme response, I think. So what principles can be stated for overall speech (or communication) policies? In particular, which principles, if adopted, would maximize truth-possession?¹³

A proposal whose roots go back to John Milton and John Stuart Mill is the familiar idea that *freedom of speech* is the overall policy that has the best prospects for truth acquisition.¹⁴ In the 20th century, this idea has been given an economic interpretation, especially in the United States under the influence of Supreme Court justices such as Oliver Wendell Holmes. In a frequently quoted dictum, Justice Holmes wrote: “the best test of truth is the power of the thought to get itself accepted in the competition of the Market”.¹⁵ This theme is echoed in many Supreme Court opinions, for example, in a 1969 ruling: “It is the purpose

¹³ Of course, Habermas’s own aim in advancing the image of an ideal speech situation is to articulate institutions that would promote democratization. I do not at all take issue with this ideal as a means to, or partial realization of, democratic dialogue. The question *here*, however, is whether an ideal speech situation uniquely promotes *veritistic* value, the type of value with which epistemology is distinctively concerned. To say this is not to imply that *veritistic* value is more important than that of democracy. On the contrary, the opposite *may* be true. The point is that an ideal speech situation is not obviously an optimal scheme for obtaining knowledge, whatever the ultimate importance of that goal may be.

¹⁴ Milton, “Areopagitica: A Speech for the Liberty of Unlicensed Printing”, in E. Sirluck (ed.), *Complete Prose Works of John Milton*, 1959; Mill, “On Liberty” in *On Liberty, Representative Government, The Subjection of Women*, Oxford University Press, New York, 1960.

¹⁵ *Abrams v. United States* (1919).

of the First Amendment [of the U.S. Constitution] to preserve an uninhibited marketplace of ideas in which truth will ultimately prevail.”¹⁶ This argument for free speech construed as a free “marketplace” of ideas is obviously inspired by economic theory. It is so formulated, for example, by Frederick Schauer (who does not himself endorse it): “Just as Adam Smith’s ‘invisible hand’ will ensure that the best products emerge from free competition, so too will an invisible hand ensure that the best ideas emerge when all opinions are permitted freely to compete.”¹⁷ The context makes it clear that by “best ideas” Schauer means *true* ideas. So here, perhaps, we have a global social practice that will maximize truth possession: simply guarantee that there will be a free economic market for speech. Make sure that government does not interfere with the market, and this practice will maximize truth-possession.

There is a second possible interpretation of the marketplace image. The second interpretation construes the term “market” or “marketplace” metaphorically or figuratively, not as an economic market in the sense specified by economic theorists, but as a market-like arena in which debate is wide open and robust. This kind of debate arena may or may not result from an *economic* market mechanism free from government control. In fact, there are defenders of this second version of the marketplace approach who contend that *more* government interference rather than less will contribute toward a more open and robust debate. This is because government might enable those agents with fewer economic resources to speak more frequently and more prominently than would happen under a purely laissez-faire arrangement. (For example, gov-

¹⁶ *Red Lion Broadcasting Co. v. FCC* (1969).

¹⁷ Schauer, *Free Speech: A Philosophical Enquiry*, Cambridge University Press, New York, 1982, p. 16.

ernment might require television channels to give free air time to *all* positions.)

I am not going to disclose my final assessment of either of these proposals. Both interpretations of the marketplace idea are discussed in *Knowledge in a Social World*, and I want to keep you in a bit of suspense in order to encourage you to read the book itself. However, the free marketplace idea, under either interpretation, is unquestionably a good illustration of the *sort* of social practice—a very high-level social practice, of course—that seeks to fulfill the aim of veritistic social epistemology, however adequate or inadequate it proves to be.

5. *Experts and Selection Among Experts*

Another central topic for social epistemology concerns experts and authorities. Obviously, a social route to knowledge is often a more effective means to social knowledge as compared with every individual seeking knowledge through personal observation. If there is a certain truth that all citizens of a community should know—say, that there is a chemical hazard or poison at a certain place in town—it is much easier for a single person to learn this truth and disseminate it via some public communication channels (newspapers, radio, television, etc.) rather than expect each member of the community to learn this fact entirely on their own. Truth-acquisition by communication is cheaper and quicker than each individual duplicating the same direct discovery process. However, since reports are not always reliable or credible, the interest of truth will only be well served if the speakers who possess and transmit truths are usually believed, whereas speakers who report falsehoods either through incompetence or deceit are not usually believed by their hearers. This will only happen, on

a regular basis, if hearers can distinguish the truth-sayers from the falsehood-sayers.

Some cases of reliance on the say-so of others involves eye-witness testimony of the sort that anybody could accomplish, as long as they were in the right place and had adequate observational powers. Reporting an auto accident on the highway, for example, requires no more than this. But in other cases, reports may require some sort of special expertise or authority, which is not shared among the general public. Identifying a chemical hazard might be such a case. The question then arises: How are hearers supposed to detect which speakers have the expertise or intellectual skill to have accurately determined what they claim to know? Whom should we trust and whom should we doubt? How can we tell who merits our credence?

The toughest cases are ones in which supposed experts, or self-proclaimed experts, disagree with one another. Which one (if any) should be believed? When hearers are *novices*—that is, people with little or no expertise of their own—do they have any ways of telling which of the self-proclaimed experts have greater expertise and are therefore more credible (assuming equal sincerity or honesty)? This might be called the *expert-novice puzzle*.

John Hardwig has argued that novices, precisely by virtue of being novices, cannot tell who has greater authority or expertise among those who claim to have it.¹⁸ People must simply rely on *trust*, even in science! This is not very satisfactory, especially when novices are forced to choose among competing authorities, who say contrary or contradictory things. In such a situation, not everyone's claims can be accepted. So what methods should novices use to allocate trust? It might be said: Let them consult *meta-*

¹⁸ Hardwig, "Epistemic Dependence", *Journal of Philosophy*, 82, pp. 335–349, 1985.

experts, people who are experts *about* the experts. But this just pushes the problem one step back without solving it. How are novices supposed to tell who is a reliable *meta-expert*? The very same problem arises once again, and we seem to be caught in a vicious regress.

This problem is not purely abstract or theoretical. There are many real-life arenas where it is encountered. In contemporary American society, the problem surfaces in two highly controversial areas: the legal arena and the educational arena. In the legal arena, personal injury law suits, against large corporations in particular, often involve technical issues about causes of accidents or illnesses that require scientific expertise. Both parties in such disputes introduce competing experts who testify in support of contrary propositions. How is a jury composed of novices supposed to decide whom to trust? In the educational arena, there are disputes between so-called authorities who have different views on, say, the status of evolutionary theory, or what happened in history. How is an ordinary citizen to decide whom to trust? Is evolution a dubious theory that deserves little credence? Or is it as well-established as other reputable scientific theories and therefore deserving of being taught in public schools? Although I think I know the straightforward answer to this question, the theoretical underpinning of my confidence is none too clear, since I rely heavily on the authority of certain experts rather than others, and a defence of my choice of experts would be an intricate matter.

There is some discussion of this theoretical problem in *Knowledge in a Social World*,¹⁹ though I certainly would not claim to have “solved” it. But once again, both in the interest of space and to further pique the reader’s curiosity, I won’t present everything I have to say on this subject. For

¹⁹ See especially pp. 267–271.

purposes of this paper, I again set the topic on the agenda for social epistemology, with the hope that others, like myself, will find this a challenging philosophical terrain that is well worthy of our attention.

Recibido: 23 de septiembre de 1999

RESUMEN

Este artículo describe una concepción “veritística” de la epistemología social, de acuerdo con la cual las prácticas sociales serán evaluadas en términos de su propensión a generar creencias verdaderas. Esto da como resultado una concepción normativa, o crítica, de la epistemología social. Esta propuesta se basa en una concepción realista de la verdad, y discrepa con las versiones populares del construccionismo social. Muchas prácticas sociales admiten una evaluación relacionada con la verdad. Entre éstas están, primero, la práctica global de acordar oportunidades lingüísticas mediante un “mercado libre” y, segundo, los métodos mediante los cuales los principiantes pueden elegir en qué expertos confiar de entre los muchos autodenominados expertos que ofrecen sus puntos de vista sobre el tema.

[Traducción: Claudia Chávez A.]