

Expertise

Alvin I. Goldman¹ 

Published online: 14 July 2016
© Springer Science+Business Media Dordrecht 2016

Abstract This paper offers a sizeable menu of approaches to what it means to be an expert. Is it a matter of reputation within a community, or a matter of what one knows independently of reputation? An initial proposal characterizes expertise in dispositional terms—an ability to help other people (especially laypersons) get answers to difficult questions or execute difficult tasks. What cognitive states, however, ground these abilities? Do the grounds consist in “veritistic” (truth-linked) states or in terms of evidence or justifiedness? To what extent is expertise a matter of superior knowledge or other factors? Some authors seek to debunk the notion of expertise entirely. The present approach resists this stance, but doesn’t dispute the variability and fluidness of the concept. Even more challenging is the problem of how laypersons can determine who is the superior expert, especially when experts disagree.

Keywords Expertise · Social epistemology · Veritism · Identifying experts

1 Who Qualifies as an Expert?

It is widely (if not universally) agreed that there are experts in many fields. What makes someone (deservedly) count as an expert in a given field, D? In other words, what conditions must a person satisfy to qualify as expert? At least two types of conditions might be central to expertise. One type of condition would pertain to the person’s knowledge

or information, and the other to his/her skill or performance ability. I shall make room for both kinds of conditions in due course. But I begin with a different concern that arises from many discussions of expertise in the literature.

Many writers insist that there is something essentially *social* in the phenomenon of expertise. To understand expertise, it is said, we must invoke the special position experts occupy in their social environment. On this approach, or family of approaches, one doesn’t qualify as an expert simply in virtue of what’s in one’s head. The general idea is expressed in a number of (non-equivalent) ways. One representative formulation is found in an article by Matt Stichter, who writes: “If someone loudly declares, ‘I’m an expert,’ then we can always reply, ‘Only if we say you are’” (2015: 126). Essentially the same idea is expressed in saying that being an expert is a *reputational* phenomenon. A person *is* an expert only if s/he has a *reputation* for being one. As will be explained shortly, I regard this proposal as erroneous. Nonetheless, there seems to be something right in the idea that some kind of social factor figures in the concept of expertise.

Christian Quast (2016) approaches the matter by discussing the *function* of expertise and linking it to goals possessed by the expert’s clients. He writes:

“The conceptual function (point, role) of expertise is to substantially improve the social deployment of available agential resources apt for client[s]’ relevant ends ...” (Quast 2016: Sect. 7).

Without exploring the details of Quast’s proposal, or endorsing it in full, I advance a definition of expertise that (by my lights) is in roughly the same territory. Specifically, this definition highlights what expertise *is* by reference to what experts can *do* for laypersons by means of their special knowledge or skill. Here is my proposal:

✉ Alvin I. Goldman
goldman@philosophy.rutgers.edu

¹ Rutgers The State University of New Jersey, New Brunswick, NJ, USA

[CAP] S is an expert in domain D if and only if S has the capacity to help others (especially laypersons) solve a variety of problems in D or execute an assortment of tasks in D which the latter would not be able to solve or execute on their own. S can provide such help by imparting to the layperson (or other client) his/her distinctive knowledge or skills.

Although CAP makes explicit reference to certain (potential) *relations* between experts and non-experts—and is therefore a “social” definition of sorts—it does not endorse a *reputational* approach to expertise. Nor does CAP endorse or imply reputationalism. To *be* an expert in D, according to CAP, it is neither necessary nor sufficient to have a reputation for being able to help others solve problems or execute tasks in D. On the non-necessity side, you may actually possess a *capacity* to serve or assist others in the indicated ways without having earned a reputation for having such an ability. Perhaps you just arrived in town, or in the country; and people know nothing about you. It doesn’t mean you aren’t an expert in D. On the non-sufficiency side, you might have acquired such a reputation illegitimately, without actually possessing the indicated capacities. Thus, mere reputation has little or nothing to do with actual expertise.

Perhaps not all capacities for serving or assisting others in solving problems or executing tasks should qualify as expertise. At least two types of such capacities, however, certainly do qualify. First, an expert’s informational and intellectual capacities can lead an expert to answer *truly* one or more questions that other people—especially laypeople—cannot answer on their own, at least not with their current body of information and problem-solving techniques. Second, the skill capacities of an expert—e.g., in golf or tennis—would enable an expert to teach or show run-of-the-mill players how to execute certain athletic maneuvers that they are initially unable to execute. These are among the “serving” or “assisting” abilities that make someone qualify as an expert in selected domains, under CAP.¹

Let us put the performance domain to the side and focus exclusively on informational or cognitive domains. CAP identifies conditions for expertise in dispositional or functional terms. But mustn’t there be *categorical* states that ground, or underpin, the dispositional or functional requirements of expertise? What must obtain or transpire in an expert’s head that enables him/her to have question-answering abilities—specifically, abilities to find *correct*

answers (with sufficient frequency)? Surely, that is what lay-persons, or anyone else seeking help from an expert, will be seeking, at least in the informational domain. Mustn’t there be some distinctive cognitive states of affairs—e.g., true beliefs and sound inferential practices—that pertain to the domain in question and thereby serve as the *ground* or *basis* of the expert’s relevant capacities?

Although we continue to focus on the informational domains, we do not restrict expertise to “technical” knowledge. The definition does not confine expertise to scientific domains, nor require experts to have specialized training (although such training is a common pre-condition of many types of expertise). Someone interested in the history of opera, for example, might become an expert on opera under her own scholarly steam rather than through a series of academic classes.

Can we be more specific as to what kinds of cognitive states serve as categorical grounds of an expert’s capacities? I recommend a “truth-linked” approach to this problem. Here is an initial sample of this approach.

[TL1] S is an expert about domain D if and only if S has more true beliefs and fewer false beliefs in propositions pertaining to D than do most people.²

According to this definition, expertise is a *comparative* state of affairs. It’s a matter of how the candidate expert compares to others in terms of truth-possession and error-avoidance (or similar notions formulated in terms of credences rather than categorical beliefs). It also presents expertise as an *objective* matter rather than a *reputational* matter. Under this definition, it is possible for S to *be* an expert even if he or she hides this expertise from friends, acquaintances, and the public at large. This comports with our previous rejection of the reputational approach to expertise.

A possible modification of TL1 would add an absolute, i.e., non-comparative, element to the account. Here is the rationale for such an addition. Suppose that almost every sentient being in the universe is ignorant of a certain planet’s existence in a distant galaxy. You, however, are an advanced astronomer, who recently discovered this planet’s existence, although thus far you know *very little* about it. Everything you do know about it is true; but you have only three such true beliefs. Is it reasonable to consider you an *expert* about this planet? This is dubious. Simply having more true beliefs than others doesn’t suffice

¹ As an anonymous referee for this paper points out, expertise alone does not guarantee the ability to teach others. The latter is, arguably, a separate skill. Nonetheless, granted that two separate types of skill are required for the teaching, it is the difficulty of executing the difficult or rare task itself that standardly earns one the label of “expert”.

² An improvement on this formula would include a proviso that the propositional contents in each comparison class be comparatively *important* for understanding the domain in question. Although importance is indeed important, there is no easy way I know of to spell it out, so I won’t explore it in any detail here. Thanks to Christian Quast, however, for a reminder of the importance of importance.

for being an expert. A handful of truths just isn't enough. Thus, some (vague) *threshold* condition ought to be added, an *absolute* condition rather than a comparative one. Thus, we might want to give TL1 a little tweak, to yield TL2:

- [TL2] S is an expert about domain D if and only if
 (A) S has more true beliefs (or high credences) in propositions concerning D than most people do, and fewer false beliefs³; and
 (B) the absolute number of true beliefs S has about propositions in D is very substantial.

How to precisify the absolute condition stated in clause (B) is unclear. However, whatever might be said to tighten this condition needn't affect the general type of approach proposed here. In particular, it would remain a truth-linked, or veritistic, type of approach.

If one is comfortable with vagueness, a third definition might be proposed that would retain condition (B) above but jettison (A). For a high enough absolute threshold of true belief (and a comparably high threshold of error avoidance) one might say that everyone who meets this condition is automatically an expert, even if they don't *exceed* most other people in the relevant domain. For example, perhaps the vast majority of (native) English speakers are experts about English even if they don't markedly exceed most other (native) English speakers in this domain. On this approach to "expertise," the comparative element in (A) drops out, and the criterion of expertise becomes fully absolute. This alternative approach might be co-tenable with [TL1]. Indeed, it is plausible that "expert" is such a fluid term that different criteria for it are used in different contexts.⁴

Return now to the original, comparative sense of "expertise," TL1. One might ask about the scope or extent of the interpersonal comparisons. Who are the "other people" with whom S should be compared? Everyone in the entire world, throughout its history? Or only people in S's local society or social network? Does the comparison class consist exclusively of people living in S's life-time? Or

might it include people from previous eras as well? Plausibly, the best eighth-century physician qualified as a medical expert (for his period) despite having a large number of false medical beliefs, many more than today's merely *average* physician. Again, these questions have no clear-cut answers, given the fluidity of the term "expert." Speakers may, however, generate comparison classes "on the fly." (In the present example, a speaker may restrict the comparison class for the target eighth-century physician to other eighth-century physicians.)

Some readers might prefer a rather different approach to the definition of expertise, one featuring epistemological notions such as *justification* or *evidence* in place of truth and falsity; or perhaps these notions *plus* truth and falsity. Here is an example of the first type, using evidence exclusively.

- [E1] S is an expert with respect to domain D if and only if S possesses substantially more and/or better evidence concerning propositions in D than most people in the relevant comparison class.

Another possible definition would be a "mixed" one:

- [E2] S is an expert with respect to domain D if and only if
 (A) S possesses more and/or better evidence pertaining to propositions in D than most people in the relevant comparison class; and
 (B) S regularly aligns her degrees of belief in such propositions in accord with her total evidence.

What are the advantages of an evidential (or justificational) approach in contrast to a veritistic approach? One advantage is the absence of any appeal to objective truth or falsity concerning what the candidate expert believes. Appealing to objective truth, many would argue, has two liabilities.

One liability is that we shall often be unable to *tell* (with confidence) whether someone is or isn't an expert until decades, or even centuries, after he is dead—because we may not be able to tell how many of the propositions he believed are true or false. I would respond as follows. This might indeed be a common problem; but such is life. That we are challenged when trying to *know* the truth is hardly news. But this doesn't suggest that there are no facts of the matter. Facts may simply be elusive; especially when they pertain to non-observable matters. But in itself this does not undercut the appropriateness of veritistic criteria, especially when viewed from a "God's eye" perspective.

A second objection is that speakers often credit people with expertise on a certain topic even when they don't genuinely consider their beliefs to be correct. A speaker might call Sam an expert because Sam is a member of an elite class, caste, or traditional position in society. Doesn't

³ Admittedly, counting of the number of beliefs a person has is a problematic matter. But we can't solve all problems here. In many cases, it is perfectly clear that X has more true beliefs about a certain matter than Y, even if the precise number of such beliefs is problematic. On a separate topic, perhaps it should not be required of the leading-edge astronomer that she has *fewer false beliefs* about the newly-found planet than her colleagues. After all, the colleagues who remain totally ignorant of the planet's existence may have zero false beliefs about it. We might still want to consider her an expert about the planet. Surely having only one false belief about it (rather than none) shouldn't disqualify her from being an expert. (Perhaps experts should only be required to have a comparatively high *ratio* of true to false beliefs to qualify as experts.)

⁴ Thanks to Holly M. Smith for championing this approach (in conversation).

this neglect of genuine correctness refute a truth-linked approach to expertise?

I respond by pointing out that speech is not always fully revelatory of a speaker's mind-set. For example, people may engage in what is called "protagonist projection." As explained by Richard Holton (1997, 626), for example, protagonist projection occurs when one uses a sentence to express, not one's own opinion, but the opinion of other people whose "perspective" one adopts for the moment and the conversation at hand. The speaker projects into the minds of the people in question the assumption that the target individual is a genuine truth-holder, even though the speaker doesn't honestly share this assessment. The speaker just plays the socially approved "game" by showing deference to the rank or social standing of the target. That this phenomenon occurs doesn't really undermine the objectivist account of expertise presented here.

Another advantage of a veritistic approach is that it makes smooth and straightforward sense of the familiar, longstanding activity of consulting experts when we have questions to which we do not know the answers ourselves (with anything approaching confidence). In the digital age, most of us make a habit of consulting Google. What is Google if not an expert, or at least a "meta-expert," that can readily direct you to authoritative sources? The standard goal of people seeking informational experts is precisely to learn the *true* answers to their questions. Consulting such experts (and trusting what they say) makes perfect sense if such expertise carries a reasonably high probability of obtaining truths from them. Of course, this definitional matter *per se* offers no guarantee that your average expert-seeker will correctly identify the genuine experts. But this is the *expert identification problem*, not a *definitional* problem. (We turn to this identification problem in Sect. 3 below.)

What about the earlier proposal to incorporate (the requirements of) *justification* and/or *evidence* into the definition of an expert? Combining this approach with a veritistic one could circumvent possible counterexamples to a purely veritistic approach, for example, where a putative expert forms true beliefs by lucky guesswork. A downside to incorporating justification into the definition is the risk of opening the door to epistemic relativism, which would constitute a roadblock to objectivism. Many epistemologists argue that different epistemic *systems*—i.e., normative systems for justification and evidence evaluation—can readily yield conflicting appraisals of one and the same epistemic situation. A candidate expert's belief would be justified under one system and unjustified under another. If this is right, and if we incorporate justification or evidence possession into a criterion of expertise, there is potential trouble on this front.

A variety of responses to this threat are possible. First, one may defend the "uniqueness thesis" (Feldman 2007), which holds that there is only one correct epistemic system. Second, one might allow that conflicts sometimes occur between (reasonable) epistemic systems, but not very frequently. If one doesn't insist upon precise appraisals of evidential support, there may be sufficient overlap across epistemic systems to yield plenty of clear cases of experts and clear cases of non-experts. This would suffice for present purposes.⁵

To avoid getting fixated on these issues, we proceed in what follows by concentrating on purely veritistic approaches, using the sorts of criteria sketched earlier featuring true belief and error-avoidance, with a combination of comparative and absolute standards. We shall not attempt to fix any precise *margins* or absolute *thresholds* that must be met for someone to qualify (categorically) as an expert. Nonetheless, we shall continue to use categorical language ("expert" vs. "non-expert") for the sake of convenience. Again, the term "expert" is vague and fluid, so it is fully appropriate for an analysis of it to have a comparable amount of vagueness.⁶

2 Debunking Expertise

The debunking of expertise has become a fairly popular enterprise within the academy. Such debunking enterprises emerge in different quarters and come in different flavors.⁷ To my mind, however, none of them poses serious grounds for abandoning the expertise-related projects that are sketched here.

One variety of debunking is associated with a spectrum of problems pertaining to expertise, specifically, its use in political contexts.⁸ As championed by Michel Foucault, Jurgen Habermas, and many of their followers, the thesis is that the knowledge that so-called experts claim to have merely masquerades as neutral fact, but is actually just another ideology: a false theory. Portrayals of reality by putative experts produce discursive structures—i.e., ideologies—that are actually expressions of patriarchy, racism, and the like. What is presented as fact is merely what people *agree* to *count* as a fact.

Of course, if ideologies are sets of propositions merely "agreed" to be facts—and by implication aren't facts at

⁵ For further discussion in depth, see Seidel (2014).

⁶ Cf. Goldman (1979), in which a similar rationale is offered for endorsing a vague analysis of justifiedness.

⁷ "Debunking" may not be the most appropriate term to characterize the critical stances discussed in this section. But it should serve well enough for present purposes.

⁸ This paragraph follows the exegesis of Stephen Turner (2001).

all—then the perpetrators of the referenced deceptions are not genuine experts by our present lights, but merely pseudo-experts. They don't genuinely possess any greater body of truths in the relevant domain that they claim to have. In this case it seems wrong to indict expertise per se for the pretensions of these pseudo-experts. Moreover, it doesn't seem feasible or plausible to debunk all forms and domains of expertise on the basis of certain aspirational experts who don't really merit the classification. This is not for a moment to deny the massively pernicious effects of views that have been advanced by self-proclaimed and publicly accepted experts. A large part of the problem, however, may lie with society's failure to generate and empower more genuinely expert—i.e., accurate—specialists to contradict the first set. Furthermore, when public decision-making ventures are in question, non-specialists in addition to specialists are obviously needed to participate in the relevant decision-making body. A global rejection of all experts and expertise is unnecessary, and indeed over the top.

Skepticism about expertise has pervaded even scientific circles, although this source of skepticism may be short-lived. In 2005 Philip Tetlock published a book on political forecasting (Tetlock 2005). Whatever his intent at the time, readers interpreted its take-home message as saying that the average expert is no more accurate than a dart-throwing chimp. This became the buzz-word about expertise in some academic circles. In a more recent book, however, Tetlock and Gardner (2015) largely reject the debunking stance that readers of the earlier book imputed to it. Even in the area of forecasting expertise (expertise about the future), new experiments offer grounds for optimism. New data appear to show that there are techniques by which some people become legitimate “superforecasters” (although their superiority over others is relatively modest). The authors also construct a list of “ten commandments” for ways to improve one's expertise, ways that are touted as “experimentally demonstrated to boost accuracy in real-world forecasting contests” (2015: 277). Plainly, Tetlock is no longer a thorough-going skeptic about forecasting expertise (if he ever was one).

Notice, moreover, that forecasting is only one domain out of many in which expertise might be displayed. It is, admittedly, an unusually challenging domain. Depending on what is to be predicted, how far into the future the prediction is aimed, and under what circumstances, its outcomes are highly variable. In the domain of weather prediction, 1- and 2-day forecasts are quite accurate, whereas 8-day forecasts are not. From the perspective of a general theory of expertise, it is doubtful that forecasting deserves pride of place. People can be experts about the past, about the present, or about non-datable subject matters such as the laws of physics, and so forth. Economists

are less than super-successful at predicting future economic developments, because so many independent variables influence such developments. But this weakness has little carry-over into all genres of economic expertise. To be brief, there seems to be no solid scientific basis for mounting a debunking crusade against expertise in general.

Another source of skepticism or puzzlement about expertise emerges in the literature on moral expertise. Many writers express doubts that deference to moral expertise is a plausible, sensible, or legitimate maneuver. Moral cognition should be autonomous rather than borrowed from other people, i.e., experts. However, as Sarah McGrath (2011) argues convincingly, to the extent that this reaction to moral deference seems right, it seems to be peculiar to the proper basis for moral judgment rather than to the absence of expertise in the moral terrain. In any case, there are no (good) grounds for generalizing from skepticism about moral expertise to skepticism about expertise across the domains.

3 Identifying Experts

We turn now from the nature of expertise to the epistemology of expertise, specifically, the *social* epistemology of expertise. I count the epistemology of expertise as part of social epistemology because it naturally falls within that sphere according to a taxonomy I have previously offered (Goldman 2010, 2011). The taxonomy is comprised of three branches. The first branch studies individual doxastic agents who make doxastic choices with the help of social evidence. By “social evidence” I mean evidence associated with statements, beliefs, thoughts, etc. of other people. This includes social evidence that bears on the question of who is or isn't an expert.

Consider two prominent examples of issues for social epistemology: testimony and peer disagreement. In the problem of testimony, a hearer receives testimony from another person and has to decide whether or not to believe it. Since the speaker's assertion of P is a piece of social evidence for the hearer, the latter's task is to decide how this evidence should be exploited. Similarly, in the problem of peer disagreement, the hearer's task is to decide whether and how to revise his belief vis-à-vis P based on a distinctive type of social evidence, i.e., evidence that although the speaker possesses the same prior evidence as he (the hearer) does, she holds a contrary belief. Should the hearer remain steadfast in his prior belief in the face of this new evidence? Or should he somehow revise his doxastic attitude in light of it?

Utilizing the testimony of experts is a sub-problem in the theory of testimony. Consider this following case involving two putative experts. Serena recently arrived in

town. Winter is approaching and she needs a good winter coat. She seeks advice from two acquaintances about where to shop. Each acquaintance professes to be an apparel expert, especially with respect to the two main stores in town. One acquaintance says that Ozzie's is the best store for outerwear; the other says that Izzy's is the best store for outerwear. Which expert should Serena trust? She has no prior opinion of her own; she is prepared to defer to her acquaintances' expertise. But which acquaintance is the superior expert? How can Serena (justifiedly) determine which one merits greater trust?⁹

Our new problem is definitely an epistemological one. It doesn't concern that *nature* of expertise, but rather the methods for identifying the superior experts. In a previous paper (Goldman 2001) I pursued this question in some depth, and I won't re-visit all the nuances here. But I will briefly consider some of the main factors, and probe a couple of them that concern us here.

One way for Serena to try to determine which acquaintance has greater expertise is to have the pair debate one another. Listening to a debate, Serena might be able to assess her acquaintances' respective reliability. This seems feasible because there is not much that is arcane about outerwear. But in other arenas the situation may be different. Experts in specialized fields commonly use technical language to which laypersons aren't privy. Even if the layperson hears (or reads) a debate between them, assessing the relative quality of the debaters' arguments will be a challenging task. The layperson might have to attend to factors such as enthusiasm or glibness of speech, which are far from perfect indicators. There may also be a problem of insufficient background knowledge by which to adjudicate the truthfulness of the consultants' competing claims.

These are some sample cases of the epistemic challenges that confront a layperson when trying to assess two experts' levels of expertise. Of course, professionals also hang out shingles, displaying their academic degrees and the institutions where they were earned. But can a layperson reliably assess the significance of such credentials? How much reliable information is she likely to get about the training offered at a given institution or the specific program in which a putative expert's degree was earned?

There is another potentially promising way for laypersons to assess target experts' veritistic statuses. This is to

⁹ One might object to this example because what is better or worse outerwear is partly a matter of taste rather than fact, and therefore not a matter on which expertise can be assessed. If this is correct, the reader might prefer to substitute a different example entirely, or interpret the present example so that only objective facts about the two stores' products are under discussion (e.g., the warmth they provide, their durability, and their comparative prices).

study each putative expert's *track record* in the relevant domain. Checking a weather forecaster's track record can be relatively straightforward. The layperson need only "observe" a forecaster's prediction each evening for the following day, and then "observe" the actual weather on the day in question. Using pairs of observations for multiple days, it isn't difficult to compute a meteorologist's track record. This can then be compared with the track records of other meteorologists covering the same locale on the same days.

Not all experts, however, make predictions about observable phenomena. An expert might predict what will happen in a given type of experiment, for example, one that employs a particle accelerator. But laypeople won't be able (on their own) to detect whether the read-out from the experiment does or does not confirm the prediction. Nonetheless, for a certain sub-class of cases, laypeople can make relevant observations. Also they can consult other (undisputed) experts to determine whether the target physicist's prediction was verified. In general, *sometimes* there are ways to get reliable evidence about a putative expert's expertise.

4 Expertise, Norms of Speech, and Social Epistemology

The previous section explored a question about experts that arises in the *first* branch of social epistemology. It raises the question of how, and with what prospect for success, laypersons can advance an epistemic project with an expert's help. What sub-tasks do the inquirer and the expert need to perform (respectively) in order to complete a successful social-epistemological transaction, in which the first party learns a true answer to his question by enlisting the second party's assistance? First, the inquirer must select a (candidate) expert from an available pool, and pose his question to this (putative) expert. Second, the expert must apply her knowledge to the question and generate an answer. Third, the inquirer must decide whether to accept the answer. Obviously, the transaction can go awry at any of these junctures. The inquirer may get misleading information about the pool of experts, and therefore select a sub-optimal one. Second, even if a high-quality expert is selected, this doesn't seal the deal. Even the best available expert may have erroneous beliefs pertaining to a particular question. Third, the expert may be too busy with other tasks to pay close attention to the question, or may invest less effort than the answer requires. Thus, the proffered answer might be wrong. Fourth, although the expert may deliver the correct answer, the inquirer may find it unpersuasive. Finally, the expert may have self-interested reasons to withhold the correct answer, and present it in a

misleading way. Social epistemologists may be interested in any or all of these potential misadventures.

Let us focus on the choice of answer that the expert—or any speaker, for that matter—may provide. This will enable us to explore some possible relationships between social epistemology and political philosophy (one of the less-established links in the traditional conduct of philosophy).

Epistemology is not confined to the study of belief-formation, inference, or other mental activities of an individual epistemic subject. It may also address speech or discourse among individuals. This especially holds of social epistemology. In particular, when it comes to the making of statements, claims, or assertions, epistemology advances or explores “norms of assertion.” A norm of assertion specifies when it is permissible for an epistemic agent to assert a particular proposition. While some candidate norms may be relativized to hearers or audiences, I confine the discussion (at least at the outset) to norms that abstract from any such relativization.

For present purposes we can consider two approaches. Some epistemologists (e.g., DeRose 2002; Hawthorne 2004; Stanley 2005; Williamson 2000) contend that *knowledge* is the norm of assertion. It is epistemically wrong (impermissible) to assert P unless one *knows* that P. Other writers (e.g., Lewis 1996; Douven 2006; Lackey 2007) endorse a weaker, *reasonable belief*, rule of assertion.

A crucial component of liberal democratic thought holds that freedom or liberty of speech is a basic human right. But what is freedom of speech if not a norm that grants people permission to assert and to publish whatever they wish (subject to a variety of exceptions). Does this traditional stance have implications for the norm of assertion? It would seem that it does. Doesn't it imply that the norm of assertion must be extremely permissive? “Assert whatever you wish to assert.” On its face, then, there is an apparent conflict between any of the more restrictive norms of assertion that contemporary epistemologists advance and the more permissive norms in the legal sphere. Is this a problem?

If there is a problem, it may not be very severe. After all, the free speech norm is readily construed as a norm directed at governments, namely, “Exceptions aside, do not criminalize people's assertions (or other forms of speech).” Stated this way, it seems to be a political norm rather than an epistemic one. So where is the problem? There are just two types of norms, derived from different sources, and addressed to different parties. Unfortunately, this simple solution runs into problems. The trouble arises when one reflects on the most influential rationale for freedom of speech, the so-called “argument from truth.” This rationale dates back to John Milton in the seventeenth century

(Milton 1968), John Stuart Mill in the nineteenth century (Mill 1975), and a number of American judges in the early twenty century, who used it to interpret and defend the First Amendment to the United States Constitution. The Supreme Court justices in question were Holmes, Brandeis, Frankfurter, and Hand (for references, see Schauer 1981). Milton argued that the absence of governmental restrictions on publishing will enable society to locate truth and reject error. Mill argued that the elimination of speech suppression would increase the likelihood of exchanging error for truth. Holmes argued that the “best test of truth is the power of the thought to get itself accepted in the competition of the market” (Holmes 1919).

The anomaly here is that the argument from truth is clearly an *epistemic* species of argument. It fits especially well with what some social epistemologists have advanced as a template for social epistemology, i.e., appraising social practices in terms of their “veritistic” properties (Goldman 1999). As we have seen, the argument from truth purports to provide an epistemic rationale for extensive freedom of speech. But other epistemic norms of assertion, e.g., the knowledge norm, render conflicting judgments. When such conflicts exist—ostensibly *within* epistemology—which approach is right? Which norm trumps the other? The conflict cannot be resolved by saying that one norm is epistemic while the other is political; because the so-called *political* norm itself turns out to have an *epistemic* basis.

Furthermore, the argument from truth is itself subject to serious doubts, including worries of a broadly epistemological nature. The defensibility of the argument from truth hinges on the assumption that truth will prevail when placed side-by-side with falsity. But is this correct? Frederick Schauer (1981) argues persuasively that this is a shaky thesis. It is hardly a matter of logic, he rightly insists, “that truth will be accepted and falsehood rejected when both are heard.” (1981: 25) “It is one thing to say that truth is likely to prevail in a select group of individuals trained to think rationally and chosen for that ability. It is quite another to say that the same process works for the public at large. Only if the process is effective throughout society can the argument from truth support a Free Speech Principle to limit government power.” (1981: 26)

Some writers try to strengthen the argument from truth by appeal to the “free marketplace of ideas” theme, with explicit reference to economic theory. In a 1969 ruling the Supreme Court wrote: “It is the purpose of the First Amendment to preserve an uninhibited marketplace of ideas in which truth will ultimately prevail” (*Red Lion Broadcasting Co. v. FCC* 1969: 390). This appeal to economic theory, however, is unavailing. The thesis is that according to economic theory competitive markets are efficient in the sense of being maximizing institutions. Free markets are thought to be optimal at producing and

distributing goods, so why shouldn't they also be optimal at producing and distributing knowledge? However, a careful analysis of this claim demonstrates that economic theory entails no such thesis (Goldman and Cox 1996; Goldman 1999: chap. 7).

Of course, there are other possible rationales for the free speech principle. We shall not pursue this complex issue further here. Suffice it to say that the primacy of the Free Speech Principle over other norms of assertion cannot be treated as a given. These are complex matters.

There are plenty of other reasons, in addition, to think that principles for speech and communication that may have been adequate for earlier centuries are inadequate for the digital age. A governmental commitment to refrain from inhibiting the speech of individuals hardly guarantees that the truth will be known when government simultaneously permits unlimited spending by large corporations whose overwhelming resources ensure that corporate opinions will be distributed loudly and massively to the public. Milton and his successors had reason to fear state power, and to hold that protection from government suppression would suffice to get the truth to emerge. However, in an age when corporate power may rival or exceed that of the state, mere freedom from *state* interference in personal speech may comprise only a fraction of the problem. Expert opinion has not managed to convince the public (nor a large number of their representatives) that climate change is the most serious current threat to humanity. No doubt this is the product of many factors. But it shouldn't be blithely assumed that free speech by itself is the solution.

These reflections touch only lightly on a set of concerns that requires another occasion for deeper exploration. But at least this is a beginning of an attempt to put social epistemology and political philosophy into essential contact with one another on topics of mutual interest.

References

- DeRose K (2002) Assertion, knowledge, and context. *Philos Rev* 111:167–203
- Douven I (2006) Assertion, knowledge, and rational credibility. *Philos Rev* 115:449–485
- Feldman R (2007) Reasonable religious disagreement. In: Antony L (ed) *Philosophers without gods*. Oxford University Press, New York
- Goldman A (1979) What is justified belief? In: G Pappas (ed) *Justification and knowledge*. Dordrecht: Reidel. Reprinted in Goldman, Reliabilism and Contemporary Epistemology (2012). Oxford University Press, New York
- Goldman A (1999) *Knowledge in a social world*. Oxford University Press, Oxford
- Goldman A (2001) Experts: Which ones should you trust? *Philos Phenomenol Res* 63(1):85–110
- Goldman A (2010) Systems oriented social epistemology. In: Gendler T, Hawthorne J (eds) *Oxford studies in epistemology*, vol 3. Oxford University Press, Oxford
- Goldman A (2011) A guide to social epistemology. In: Goldman A, Whitcomb D (eds) *Social epistemology: essential readings*. Oxford University Press, New York
- Goldman A, Cox J (1996) Speech, truth, and the free market for ideas. *Leg Theory* 2:1–32
- Hawthorne J (2004) *Knowledge and lotteries*. Oxford University Press, Oxford
- Holmes OW (1919) Dissent in *Abrams v. United States*, 250 U.S. 616
- Holton R (1997) Some telling examples: a reply to Tsohatzidis. *J Pragmat* 28:624–628
- Lackey J (2007) Norms of assertion. *Nous* 41:594–626
- Lewis D (1996) Elusive knowledge. *Aust J Philos* 74(4):749–767
- McGrath S (2011) Skepticism about moral expertise as a puzzle for moral realism. *J Philos* 108(3):1–47
- Mill JS (1975) *On liberty* (D. Spitz, ed). W.W. Norton, New York
- Milton J (1968) *Areopagitica* (J. C. Suffolk, ed). University Tutorial Press, London
- Quast C (2016) Expertise: A Practical Explication. *Topoi*. doi:10.1007/s11245-016-9411-2
- Schauer F (1981) *Free speech: a philosophical enquiry*. Cambridge University Press, Cambridge
- Seidel M (2014) *Epistemic relativism: a constructive critique*. Palgrave-Macmillan, Basingstoke
- Stanley J (2005) *Knowledge and practical interests*. Oxford University Press, Oxford
- Stichter M (2015) Philosophical and psychological accounts of expertise and experts. *Hum Mentis J Philos Stud* 28:105–128
- Tetlock P (2005) *Expert political judgment: How good is it? How can we know?* Princeton University Press, Princeton
- Tetlock P, Gardner D (2015) *Superforecasting: the art and science of prediction*. Crown Publishing, New York
- Turner S (2001) What is the problem with experts? *Soc Stud Sci* 31(1):123–149
- Williamson T (2000) *Knowledge and its limits*. Oxford University Press, Oxford