

Experts: Which Ones Should You Trust?

ALVIN I. GOLDMAN
University of Arizona

1. *Expertise and Testimony*

Mainstream epistemology is a highly theoretical and abstract enterprise. Traditional epistemologists rarely present their deliberations as critical to the practical problems of life, unless one supposes—as Hume, for example, did not—that skeptical worries should trouble us in our everyday affairs. But some issues in epistemology are both theoretically interesting and practically quite pressing. That holds of the problem to be discussed here: how laypersons should evaluate the testimony of experts and decide which of two or more rival experts is most credible. It is of practical importance because in a complex, highly specialized world people are constantly confronted with situations in which, as comparative novices (or even ignoramuses), they must turn to putative experts for intellectual guidance or assistance. It is of theoretical interest because the appropriate epistemic considerations are far from transparent; and it is not clear how far the problems lead to insurmountable skeptical quandaries. This paper does not argue for flat-out skepticism in this domain; nor, on the other hand, does it purport to resolve all pressures in the direction of skepticism. It is an exploratory paper, which tries to identify problems and examine some possible solutions, not to establish those solutions definitively.

The present topic departs from traditional epistemology and philosophy of science in another respect as well. These fields typically consider the prospects for knowledge acquisition in “ideal” situations. For example, epistemic agents are often examined who have unlimited logical competence and no significant limits on their investigational resources. In the present problem, by contrast, we focus on agents with stipulated epistemic constraints and ask what they might attain while subject to those constraints.

Although the problem of assessing experts is non-traditional in some respects, it is by no means a new problem. It was squarely formulated and addressed by Plato in some of his early dialogues, especially the *Charmides*. In this dialogue Socrates asks whether a man is able to examine another man

who claims to know something to see whether he does or not; Socrates wonders whether a man can distinguish someone who pretends to be a doctor from someone who really and truly is one (*Charmides* 170d-e). Plato's term for posing the problem is *techné*, often translated as "knowledge" but perhaps better translated as "expertise" (see Gentzler 1995, LaBarge 1997).¹

In the recent literature the novice/expert problem is formulated in stark terms by John Hardwig (1985, 1991). When a layperson relies on an expert, that reliance, says Hardwig, is necessarily *blind*.² Hardwig is intent on denying full-fledged skepticism; he holds that the receiver of testimony can acquire "knowledge" from a source. But by characterizing the receiver's knowledge as "blind", Hardwig seems to give us a skepticism of sorts. The term "blind" seems to imply that a layperson (or a scientist in a different field) cannot be *rationaly justified* in trusting an expert. So his approach would leave us with testimonial skepticism concerning rational justification, if not knowledge.

There are other approaches to the epistemology of testimony that lurk in Hardwig's neighborhood. The authors I have in mind do not explicitly urge any form of skepticism about testimonial belief; like Hardwig, they wish to expel the specter of skepticism from the domain of testimony. Nonetheless, their solution to the problem of testimonial justification appeals to a minimum of *reasons* that a hearer might have in trusting the assertions of a source. Let me explain who and what I mean.

The view in question is represented by Tyler Burge (1993) and Richard Foley (1994), who hold that the bare assertion of a claim by a speaker gives a hearer *prima facie* reason to accept it, quite independently of anything the hearer might know or justifiably believe about the speaker's abilities, circumstances, or opportunities to have acquired the claimed piece of knowledge. Nor does it depend on empirically acquired evidence by the hearer, for example, evidence that speakers generally make claims only when they are in a position to know whereof they speak. Burge, for example, endorses the following Acceptance Principle: "A person is entitled to accept as true something that is presented as true and that is intelligible to him, unless there are stronger reasons not to do so" (1993: 467). He insists that this principle is not an empirical one; the "justificational force of the entitlement described by this justification is not constituted or enhanced by sense experiences or perceptual beliefs" (1993: 469). Similarly, although Foley does not stress the *a priori* status of such principles, he agrees that it is reasonable of people to grant *fundamental* authority to the opinions of others, where this means that it is "reasonable for us to be influenced by others even when we have no special information indicating that they are reliable" (1994: 55). Fundamental

¹ Thanks to Scott LaBarge for calling Plato's treatment of this subject to my attention.

² In his 1991 paper, Hardwig at first says that trust must be "at least partially blind" (p. 693). He then proceeds to talk about knowledge resting on trust and therefore being blind (pp. 693, 699) without using the qualifier "partially".

authority is contrasted with *derivative* authority, where the latter is generated from the hearer's *reasons for thinking* that the source's "information, abilities, or circumstances put [him] in an especially good position" to make an accurate claim (1994: 55). So, on Foley's view, a hearer need not have such reasons about a source to get *prima facie* grounds for trusting that source. Moreover, a person does not need to acquire empirical reasons for thinking that people generally make claims about a subject only when they are in a position to know about that subject. Foley grants people a fundamental (though *prima facie*) epistemic right to trust others even in the absence of any such empirical evidence.³ It is in this sense that Burge's and Foley's views seem to license "blind" trust.

I think that Burge, Foley, and others are driven to these sorts of views in part by the apparent hopelessness of reductionist or inductivist alternatives. Neither adults nor children, it appears, have enough evidence from their personal perceptions and memories to make cogent inductive inferences to the reliability of testimony (cf. Coady 1992). So Burge, Foley, Coady and others propose their "fundamental" principles of testimonial trustworthiness to stem the potential tide of testimonial skepticism. I am not altogether convinced that this move is necessary. A case might be made that children are in a position to get good inductive evidence that people usually make claims about things they are in a position to know about.

A young child's earliest evidence of factual reports is from face-to-face speech. The child usually sees what the speaker is talking about and sees that the speaker also sees what she is talking about, e.g., the furry cat, the toy under the piano, and so forth. Indeed, according to one account of cognitive development (Baron-Cohen 1995), there is a special module or mechanism, the "eye-direction detector", that attends to other people's eyes, detects their direction of gaze, and interprets them as "seeing" whatever is in the line of sight.⁴ Since seeing commonly gives rise to knowing, the young child can determine a certain range of phenomena within the ken of speakers. Since the earliest utterances the child encounters are presumably about these *speaker-known* objects or events, the child might easily conclude that speakers usually make assertions about things within their ken. Of course, the child later encounters many utterances where it is unclear to the child whether the

³ However, there is some question whether Foley can consistently call the epistemic right he posits a "fundamental" one, since he also says that it rests on (A) my justified *self-trust*, and (B) the *similarity* of others to me—presumably the *evidence* I have of their similarity to me (see pp. 63–64). Another question for Foley is how the fundamentality thesis fits with his view that in cases of conflict I have more reason (*prima facie*) to trust myself than to trust someone else (see p. 66). If my justified trust in others is really fundamental, why does it take a back-seat to self-trust?

⁴ Moreover, according to Baron-Cohen, there is a separate module called the "shared attention mechanism", which seeks to determine when another person is attending to the same object as the self is attending to.

matters reported are, or ever were, within the speaker's ken. Nonetheless, a child's early experience is of speakers who talk about what they apparently know about, and this may well be a decisive body of empirical evidence available to the child.

I don't want to press this suggestion very hard.⁵ I shall not myself be offering a full-scale theory about the justification of testimonial belief. In particular, I do *not* mean to be advancing a sustained defense of the reductionist or inductivist position. Of greater concern to me is the recognition that a hearer's evidence about a source's reliability or unreliability can often *bolster* or *defeat* the hearer's justifiedness in accepting testimony from that source. This can be illustrated with two examples.

As you pass someone on the street, he assertively utters a sophisticated mathematical proposition, which you understand but have never previously assessed for plausibility. Are you justified in accepting it from this stranger? Surely it depends partly on whether the speaker turns out to be a mathematics professor of your acquaintance or, say, a nine-year-old child. You have prior evidence for thinking that the former is in a position to know such a proposition, whereas the latter is not. Whether or not there is an a priori principle of default entitlement of the sort endorsed by Burge and Foley, your empirical evidence about the identity of the speaker is clearly relevant. I do not claim that Burge and Foley (etc.) cannot handle these cases. They might say that your recognition that the speaker is a math professor *bolsters* your *overall* entitlement to accept the proposition (though not your *prima facie* entitlement); recognizing that it is a child *defeats* your *prima facie* entitlement to accept the proposition. My point is, however, that your evidence about the properties of the speaker is crucial evidence for your overall entitlement to accept the speaker's assertion. A similar point holds in the following example. As you relax behind the wheel of your parked car, with your eyes closed, you hear someone nearby describing the make and color of the passing cars. Plausibly, you have *prima facie* justification in accepting those descriptions as true, whether this *prima facie* entitlement has an a priori or inductivist basis. But if you then open your eyes and discover that the speaker is himself blindfolded and not even looking in the direction of the passing traffic, this *prima facie* justification is certainly defeated. So what you empirically determine about a speaker can make a massive difference to your overall justifiedness in accepting his utterances.

The same obviously holds about two putative experts, who make conflicting claims about a given subject-matter. Which claim you should accept (if either) can certainly be massively affected by your empirical discoveries about

⁵ For one thing, it may be argued that babies' interpretations of what people say is, in the first instance, constrained by the assumption that the contents concern matters within the speakers' perceptual ken. This is not an empirical finding, it might be argued, but an a priori posit that is used to fix speakers' meanings.

their respective abilities and opportunities to know the truth of the matter (and to speak sincerely about it). Indeed, in this kind of case, default principles of the sort advanced by Burge and Foley are of no help whatever. Although a hearer may be *prima facie* entitled to believe each of the speakers, he cannot be entitled *all things considered* to believe both of them; for the propositions they assert, we are supposing, are incompatible (and transparently incompatible to the hearer). So the hearer's all-things-considered justifiedness vis-à-vis their claims will depend on what he empirically learns about each speaker, or about the opinions of other speakers. In the rest of this paper I shall investigate the kinds of empirical evidence that a novice hearer might have or be able to obtain for believing one putative expert rather than her rival. I do not believe that we need to settle the "foundational" issues in the general theory of testimony before addressing this issue. This is the working assumption, at any rate, on which I shall proceed.⁶

2. *The Novice/Expert Problem vs. the Expert/Expert Problem*

There are, of course, degrees of both expertise and novicehood. Some novices might not be so much less knowledgeable than some experts. Moreover, a novice might in principle be able to turn himself into an expert, by improving his epistemic position vis-à-vis the target subject-matter, e.g., by acquiring more formal training in the field. This is not a scenario to be considered in this paper, however. I assume that some sorts of limiting factors—whether they be time, cost, ability, or what have you—will keep our novices from becoming experts, at least prior to the time by which they need to make their judgment. So the question is: Can novices, while remaining novices, make justified judgments about the relative credibility of rival experts? When and how is this possible?

There is a significant difference between the novice/expert problem and another type of problem, the expert/expert problem. The latter problem is one

⁶ Some theorists of testimony, Burge included, maintain that a hearer's justificational status vis-à-vis a claim received from a source depends partly on the justificational status of the source's own belief in that claim. This is a *transpersonal, preservationist, or transmissional* conception of justifiedness, under which a recipient is not justified in believing *p* unless the speaker has a justification and entitlement that he *transmits* to the hearer. For purposes of this paper, however, I shall not consider this transmissional conception of justification. First, Burge himself recognizes that there is such a thing as the recipient's "proprietary" justification for believing an interlocutor's claim, justification localized "in" the recipient, which isn't affected by the source's justification (1993: 485-486). I think it is appropriate to concentrate on this "proprietary" justification (of the recipient) for present purposes. When a hearer is trying to "choose" between the conflicting claims of rival speakers, he cannot appeal to any inaccessible justification lodged in the heads of the speakers. He can only appeal to his *own* justificational resources. (Of course, these might include things *said* by the two speakers by way of defense of their contentions, things which also are relevant to *their own* justifications.) For other types of (plausible) objections to Burge's preservationism about testimony, see Bezuidenhout (1998).

in which experts seek to appraise the authority or credibility of other experts. Philip Kitcher (1993) addresses this problem in analyzing how scientists ascribe authority to their peers. A crucial segment of such authority ascription involves what Kitcher calls “calibration” (1993: 314–22). In *direct* calibration a scientist uses his own opinions about the subject-matter in question to evaluate a target scientist’s degree of authority. In *indirect* calibration, he uses the opinions of still other scientists, whose opinions he has previously evaluated by direct calibration, to evaluate the target’s authority. So here too he starts from his own opinions about the subject-matter in question.

By contrast, in what I am calling the novice/expert problem (more specifically, the novice/2-expert problem), the novice is not in a position to evaluate the target experts by using his own opinion; at least he does not think he is in such a position. The novice either has no opinions in the target domain, or does not have enough confidence in his opinions in this domain to use them in adjudicating or evaluating the disagreement between the rival experts. He thinks of the domain as properly requiring a certain expertise, and he does not view himself as possessing this expertise. Thus, he cannot use opinions of his own in the domain of expertise—call it the *E-domain*—to choose between conflicting experts’ judgments or reports.

We can clarify the nature of the novice/expert problem by comparing it to the analogous listener/eyewitness problem. (Indeed, if we use the term “expert” loosely, the latter problem may just be a species of the novice/expert problem.) Two putative eyewitnesses claim to have witnessed a certain crime. A listener—for example, a juror—did not himself witness the crime, and has no prior beliefs about who committed it or how it was committed. In other words, he has no personal knowledge of the event. He wants to learn what transpired by listening to the testimonies of the eyewitnesses. The question is how he should adjudicate between their testimonies if and when they conflict. In this case, the E-domain is the domain of propositions concerning the actions and circumstances involved in the crime. This E-domain is what the listener (the “novice”) has no prior opinions about, or no opinions to which he feels he can legitimately appeal. (He regards his opinions, if any, as mere speculation, hunch, or what have you.)

It may be possible, at least in principle, for a listener to make a reasonable assessment of which eyewitness is more credible, even without having or appealing to prior opinions of his own concerning the E-domain. For example, he might obtain evidence from others as to whether each putative witness was really present at the crime scene, or, alternatively, known to be elsewhere at the time of the crime. Second, the listener could learn of tests of each witness’s visual acuity, which would bear on the accuracy or reliability of their reports. So in this kind of case, the credibility of a putative “expert’s” report can be checked by such methods as independent verification of whether

he had the opportunity and ability to see what he claims to have seen. Are analogous methods available to someone who seeks to assess the credibility of a “cognitive” expert as opposed to an eyewitness expert?

Before addressing this question, we should say more about the nature of expertise and the sorts of experts we are concerned with here. Some kinds of experts are unusually accomplished at certain skills, including violinists, billiards players, textile designers, and so forth. These are not the kinds of experts with which epistemology is most naturally concerned. For epistemological purposes we shall mainly focus on cognitive or intellectual experts: people who have (or claim to have) a superior quantity or level of knowledge in some domain and an ability to generate new knowledge in answer to questions within the domain. Admittedly, there are elements of skill or know-how in intellectual matters too, so the boundary between skill expertise and cognitive expertise is not a sharp one. Nonetheless, I shall try to work on only one side of this rough divide, the intellectual side.

How shall we define expertise in the cognitive sense? What distinguishes an expert from a layperson, in a given cognitive domain? I’ll begin by specifying an objective sense of expertise, what it is to *be* an expert, not what it is to have a reputation for expertise. Once the objective sense is specified, the reputational sense readily follows: a reputational expert is someone widely believed to be an expert (in the objective sense), whether or not he really is one.

Turning to objective expertise, then, I first propose that cognitive expertise be defined in “veritistic” (truth-linked) terms. As a first pass, experts in a given domain (the E-domain) have more beliefs (or high degrees of belief) in true propositions and/or fewer beliefs in false propositions within that domain than most people do (or better: than the vast majority of people do). According to this proposal, expertise is largely a comparative matter. However, I do not think it is wholly comparative. If the vast majority of people are full of false beliefs in a domain and Jones exceeds them slightly by not succumbing to a few falsehoods that are widely shared, that still does not make him an “expert” (from a God’s-eye point of view). To qualify as a cognitive expert, a person must possess a substantial body of truths in the target domain. Being an expert is not simply a matter of veritistic superiority to most of the community. Some non-comparative threshold of veritistic attainment must be reached, though there is great vagueness in setting this threshold.

Expertise is not all a matter of possessing accurate information. It includes a capacity or disposition to deploy or exploit this fund of information to form beliefs in true answers to new questions that may be posed in the domain. This arises from some set of skills or techniques that constitute part of what it is to be an expert. An expert has the (cognitive) know-how, when presented with a new question in the domain, to go to the right sectors

of his information-bank and perform appropriate operations on this information; or to deploy some external apparatus or data-banks to disclose relevant material. So expertise features a propensity element as well as an element of actual attainment.

A third possible feature of expertise may require a little modification in what we said earlier. To discuss this feature, let us distinguish the *primary* and *secondary* questions in a domain. Primary questions are the principal questions of interest to the researchers or students of the subject-matter. Secondary questions concern the existing evidence or arguments that bear on the primary questions, and the assessments of the evidence made by prominent researchers. In general, an expert in a field is someone who has (comparatively) extensive knowledge (in the weak sense of knowledge, i.e., true belief) of the state of the evidence, and knowledge of the opinions and reactions to that evidence by prominent workers in the field. In the central sense of “expert” (a strong sense), an expert is someone with an unusually extensive body of knowledge on both primary and secondary questions in the domain. However, there may also be a weak sense of “expert”, in which it includes someone who merely has extensive knowledge on the secondary questions in the domain. Consider two people with strongly divergent views on the primary questions in the domain, so that one of them is largely right and the other is largely wrong. By the original, strong criterion, the one who is largely wrong would not qualify as an expert. People might disagree with this as the final word on the matter. They might hold that anyone with a thorough knowledge of the existing evidence and the differing views held by the workers in the field deserves to be called an expert. I concede this by acknowledging the weak sense of “expert”.

Applying what has been said above, we can say that an expert (in the strong sense) in domain D is someone who possesses an extensive fund of knowledge (true belief) and a set of skills or methods for apt and successful deployment of this knowledge to new questions in the domain. Anyone purporting to be a (cognitive) expert in a given domain will claim to have such a fund and set of methods, and will claim to have true answers to the question(s) under dispute because he has applied his fund and his methods to the question(s). The task for the layperson who is consulting putative experts, and who hopes thereby to learn a true answer to the target question, is to decide who has superior expertise, or who has better deployed his expertise to the question at hand. The novice/2-experts problem is whether a layperson can *justifiably* choose one putative expert as more credible or trustworthy than the other with respect to the question at hand, and what might be the epistemic basis for such a choice?⁷

⁷ In posing the question of justifiedness, I mean to stay as neutral as possible between different approaches to the concept of justifiedness, e.g., between internalist versus

3. *Argument-Based Evidence*

To address these issues, I shall begin by listing five possible sources of evidence that a novice might have, in a novice/2-experts situation, for trusting one putative expert more than another. I'll then explore the prospects for utilizing such sources, depending on their availability and the novice's exact circumstance. The five sources I shall discuss are:

(A) Arguments presented by the contending experts to support their own views and critique their rivals' views.

(B) Agreement from additional putative experts on one side or other of the subject in question.

(C) Appraisals by "meta-experts" of the experts' expertise (including appraisals reflected in formal credentials earned by the experts).

(D) Evidence of the experts' interests and biases vis-à-vis the question at issue.

(E) Evidence of the experts' past "track-records".

In the remainder of the paper, I shall examine these five possible sources, beginning, in this section, with source (A).⁸

There are two types of communications that a novice, N, might receive from his two experts, E₁ and E₂.⁹ First, each expert might baldly state her view (conclusion), without supporting it with any evidence or argument whatever. More commonly, an expert may give detailed support to her view in some public or professional context, but this detailed defense might only appear in a restricted venue (e.g., a professional conference or journal) that does not reach N's attention. So N might not encounter the two experts' defenses, or might encounter only very truncated versions of them. For example, N might hear about the experts' views and their support from a second-hand account in the popular press that does not go into many details. At the opposite end of the communicational spectrum, the two experts might engage in a full-scale debate that N witnesses (or reads a detailed reconstruction of). Each expert might there present fairly developed arguments in

externalist approaches to justifiedness. Notice, moreover, that I am not merely asking whether and how the novice can justifiably decide to accept one (candidate) expert's view *outright*, but whether and how he can justifiably decide to give *greater* credence to one than to the other.

⁸ I do not mean to be committed to the exhaustiveness of this list. The list just includes some salient categories.

⁹ In what follows I shall for brevity speak about two experts, but I shall normally mean two *putative* experts, because from the novice's epistemic perspective it is problematic whether each, or either, of the self-proclaimed experts really is one.

support of her view and against that of her opponent. Clearly, only when N somehow encounters the experts' evidence or arguments can he have evidence of type (A). So let us consider this scenario.

We may initially suppose that if N can gain (greater) justification for believing one expert's view as compared with the other by means of their arguments, the novice must at least understand the evidence cited in the experts' arguments. For some domains of expertise and some novices, however, even a mere grasp of the evidence may be out of reach. These are cases where N is an "ignoramus" vis-à-vis the E-domain. This is not the universal plight of novices. Sometimes they can understand the evidence (in some measure) but aren't in a position, from personal knowledge, to give it any credence. Assessing an expert's evidence may be especially difficult when it is disputed by an opposing expert.

Not every statement that appears in an expert's argument need be epistemically inaccessible to the novice. Let us distinguish here between *esoteric* and *exoteric* statements within an expert's discourse. Esoteric statements belong to the relevant sphere of expertise, and their truth-values are inaccessible to N—in terms of his personal knowledge, at any rate. Exoteric statements are outside the domain of expertise; their truth-values may be accessible to N—either at the time of their assertion or later.¹⁰ I presume that esoteric statements comprise a hefty portion of the premises and "lemmas" in an expert's argument. That's what makes it difficult for a novice to become justified in believing any expert's view on the basis of arguments per se. Not only are novices commonly unable to assess the truth-values of the esoteric propositions, but they also are ill-placed to assess the support relations between the cited evidence and the proffered conclusion. Of course, the proponent expert will claim that the support relation is strong between her evidence and the conclusion she defends; but her opponent will commonly dispute this. The novice will be ill-placed to assess which expert is in the right.

At this point I wish to distinguish *direct* and *indirect argumentative justification*. In direct argumentative justification, a hearer becomes justified in believing an argument's conclusion by becoming justified in believing the argument's premises and their (strong) support relation to the conclusion. If a speaker's endorsement of an argument helps bring it about that the hearer has such justificational status vis-à-vis its premises and support relation, then the hearer may acquire "direct" justification for the conclusion via that speaker's

¹⁰ It might be helpful to distinguish *semantically* esoteric statements and *epistemically* esoteric statements. (Thanks to Carol Caraway for this suggestion.) Semantically esoteric statements are ones that a novice cannot assess because he does not even *understand* them; typically, they utilize a technical vocabulary he has not mastered. Epistemically esoteric statements are statements the novice understands but still cannot assess for truth-value.

argument.¹¹ As we have said, however, it is difficult for an expert's argument to produce direct justification in the hearer in the novice/2-expert situation. Precisely because many of these matters are esoteric, N will have a hard time adjudicating between E_1 's and E_2 's claims, and will therefore have a hard time becoming justified vis-à-vis either of their conclusions. He will even have a hard time becoming justified in trusting one conclusion *more* than the other.

The idea of indirect argumentative justification arises from the idea that one speaker in a debate may demonstrate dialectical superiority over the other, and this dialectical superiority might be a plausible *indicator*¹² for N of greater expertise, even if it doesn't render N directly justified in believing the superior speaker's conclusion. By dialectical superiority, I do not mean merely greater debating skill. Here is an example of what I do mean.

Whenever expert E_2 offers evidence for her conclusion, expert E_1 presents an ostensible rebuttal or defeater of that evidence. On the other hand, when E_1 offers evidence for her conclusion, E_2 never manages to offer a rebuttal or defeater to E_1 's evidence. Now N is not in a position to assess the truth-value of E_1 's defeaters against E_2 , nor to evaluate the truth-value or strength of support that E_1 's (undefeated) evidence gives to E_1 's conclusion. For these reasons, E_1 's evidence (or arguments) are not directly justificatory for N. Nonetheless, in "formal" dialectical terms, E_1 seems to be doing better in the dispute. Furthermore, I suggest, this dialectical superiority may reasonably be taken as an indicator of E_1 's having superior expertise on the question at issue. It is a (non-conclusive) indicator that E_1 has a superior fund of information in the domain, or a superior method for manipulating her information, or both.

Additional signs of superior expertise may come from other aspects of the debate, though these are far more tenuous. For example, the comparative quickness and smoothness with which E_1 responds to E_2 's evidence may suggest that E_1 is already well familiar with E_2 's "points" and has already thought out counterarguments. If E_2 's responsiveness to E_1 's arguments displays less quickness and smoothness, that may suggest that E_1 's prior mastery of the relevant information and support considerations exceeds that of E_2 . Of course, quickness and smoothness are problematic indicators of informational mastery. Skilled debaters and well-coached witnesses can appear better-informed because of their stylistic polish, which is not a true indicator

¹¹ By "direct" justification I do not, of course, mean anything having to do with the basicness of the conclusion in question, in the foundationalist sense of basicness. The distinction I am after is entirely different, as will shortly emerge.

¹² Edward Craig (1990: 135) similarly speaks of "indicator properties" as what an inquirer seeks to identify in an informant as a guide to his/her truth-telling ability.

of superior expertise. This makes the proper use of indirect argumentative justification a very delicate matter.¹³

To clarify the direct/indirect distinction being drawn here, consider two different things a hearer might say to articulate these different bases of justification. In the case of direct argumentative justifiedness, he might say: "In light of this expert's argument, that is, in light of the truth of its premises and the support they confer on the conclusion (both of which are epistemically accessible to me), I am now justified in believing the conclusion." In indirect argumentative justifiedness, the hearer might say: "In light of the way this expert has argued—her argumentative *performance*, as it were—I can infer that she has more expertise than her opponent; so I am justified in inferring that her conclusion is probably the correct one."

Here is another way to explain the direct/indirect distinction. Indirect argumentative justification essentially involves an *inference to the best explanation*, an inference that N might make from the performances of the two speakers to their respective levels of expertise. From their performances, N makes an inference as to which expert has superior expertise in the target domain. Then he makes an inference from greater expertise to a higher probability of endorsing a true conclusion. Whereas *indirect* argumentative justification essentially involves inference to the best explanation, *direct* argumentative justification need involve no such inference. Of course, it *might* involve such inference; but if so, the topic of the explanatory inference will only concern the objects, systems, or states of affairs under dispute, not the relative expertise of the contending experts. By contrast, in indirect argumentative justifiedness, it is precisely the experts' relative expertise that constitutes the target of the inference to the best explanation.

Hardwig (1985) makes much of the fact that in the novice/expert situation, the novice lacks the expert's reasons for believing her conclusion. This is correct. Usually, a novice (1) lacks all or some of the premises from which an expert reasons to her conclusion, (2) is in an inferior position to assess the support relation between the expert's premises and conclusions, and (3) is ignorant of many or most of the defeaters (and "defeater-defeaters") that might bear on an expert's arguments. However, although novice N may lack (all or some of) an expert's reasons R for believing a conclusion p, N *might* have reasons R* for believing *that* the expert has good reasons for believing p; and N might have reasons R* for believing that one expert has *better* reasons for believing her conclusion than her opponent has for hers. Indirect argumenta-

¹³ Scott Brewer (1998) discusses many of the same issues about novices and experts canvassed here. He treats the present topic under the heading of novices' using experts' "demeanor" to assess their expertise. Demeanor is an especially untrustworthy guide, he points out, where there is a lucrative "market" for demeanor itself—where demeanor is "traded" at high prices (1998: 1622). This practice was prominent in the days of the sophists and is a robust business in adversarial legal systems.

tive justification is one means by which N might acquire reasons R* without sharing (all or any) of either experts' reasons R.¹⁴ It is this possibility to which Hardwig gives short shrift. I don't say that a novice in a novice/2-expert situation invariably has such reasons R*; nor do I say that it is easy for a novice to acquire such reasons. But it does seem to be possible.

4. *Agreement from Other Experts: The Question of Numbers*

An additional possible strategy for the novice is to appeal to further experts. This brings us to categories (B) and (C) on our list. Category (B) invites N to consider whether other experts agree with E₁ or with E₂. What proportion of these experts agree with E₁ and what proportion with E₂? In other words, to the extent that it is feasible, N should consult the numbers, or degree of consensus, among all relevant (putative) experts. Won't N be fully justified in trusting E₁ over E₂ if almost all other experts on the subject agree with E₁, or if even a preponderance of the other experts agree with E₁?

Another possible source of evidence, cited under category (C), also appeals to other experts but in a slightly different vein. Under category (C), N should seek evidence about the two rival experts' relative degrees of expertise by consulting third parties' assessments of their expertise. If "meta-experts" give E₁ higher "ratings" or "scores" than E₂, shouldn't N rely more on E₁ than E₂? Credentials can be viewed as a special case of this same process. Academic degrees, professional accreditations, work experience, and so forth (all from specific institutions with distinct reputations) reflect certifications by other experts of E₁'s and E₂'s demonstrated training or competence. The relative strengths or weights of these indicators might be utilized by N to distill appropriate levels of trust for E₁ and E₂ respectively.¹⁵

I treat ratings and credentials as signaling "agreement" by other experts because I assume that established authorities certify trainees as competent when they are satisfied that the latter demonstrate (1) a mastery of the same methods that the certifiers deem fundamental to the field, and (2) knowledge of (or belief in) propositions that certifiers deem to be fundamental facts or laws of the discipline. In this fashion, ratings and conferred credentials ultimately rest on basic agreement with the meta-experts and certifying authorities.

When it comes to evaluating specific experts, there is precedent in the American legal system for inquiring into the degree to which other experts agree with those being evaluated.¹⁶ But precedented or not, just how good is

¹⁴ Of course, in indirect argumentative justification the novice must at least *hear* some of the expert's premises—or intermediate steps between "ultimate" premises and conclusion. But the novice will not share the expert's *justifiedness* in believing those premises.

¹⁵ These items fall under Kitcher's category of "unearned authority" (1993: 315).

¹⁶ Appealing to other experts to validate or underwrite a putative expert's opinion—or, more precisely, the *basis* for his opinion—has a precedent in the legal system's proce-

this appeal to consensus? If a putative expert's opinion is joined by the consensual opinions of other putative experts, how much warrant does that give a hearer for trusting the original opinion? How much evidential worth does consensus or agreement deserve in the doxastic decision-making of a hearer?

If one holds that a person's opinion deserves prima facie credence, despite the absence of any evidence of their reliability on the subject, then numbers would seem to be very weighty, at least in the absence of additional evidence. Each new testifier or opinion-holder on one side of the issue should add weight to that side. So a novice who is otherwise in the dark about the reliability of the various opinion-holders would seem driven to agree with the more numerous body of experts. Is that right?

Here are two examples that pose doubts for "using the numbers" to judge the relative credibility of opposing positions. First is the case of a guru with slavish followers. Whatever the guru believes is slavishly believed by his followers. They fix their opinions wholly and exclusively on the basis of their leader's views. Intellectually speaking, they are merely his clones. Or consider a group of followers who are not led by a single leader but by a small elite of opinion-makers. When the opinion-makers agree, the mass of followers concur in their opinion. Shouldn't a novice consider this kind of scenario as a possibility? Perhaps (putative) expert E_1 belongs to a doctrinal community whose members devoutly and uncritically agree with the opinions of some single leader or leadership cabal. Should the numerosity of the community make their opinion more credible than that of a less numerous group of experts? Another example, which also challenges the probity of

dures for deciding the admissibility of scientific expert testimony. Under the governing test for admitting or excluding such testimony that was applicable from 1923 to 1993, the scientific principle (or methodology) on which a proffered piece of testimony is based must have "gained general acceptance in the particular field in which it belongs". (*Frye v. United States*, 292 F. 1013 D.C. Cir. (1923)). In other words, appeal was made to the scientific community's opinion to decide whether the basis of an expert's testimony is sound enough to allow that testimony into court. This test has been superseded as the uniquely appropriate test in a more recent decision of the Supreme Court (*Daubert v. Merrell Dow Pharmaceuticals*, 509 U.S. 579 (1993)); but the latter decision also appeals to the opinions of other experts. It recommends that judges use a combination of four criteria (none of them necessary or sufficient) in deciding whether proffered scientific expert testimony is admissible. One criterion is the old general acceptance criterion and another is whether the proffered evidence has been subjected to peer review and publication. Peer review, obviously, also introduces the opinions of other experts. Of course, the admissibility of a piece of expert testimony is not the same question as how heavily a hearer—e.g., a juror—should trust such testimony if he hears it. But the two are closely intertwined, since courts make admissibility decisions on the assumption that jurors are likely to be influenced by any expert testimony they hear. Courts do not wish to admit scientific evidence unless it is quite trustworthy. Thus, the idea of ultimately going to the opinions of other experts to assess the trustworthiness of a given expert's proffered testimony is certainly a well-precedented procedure for trying to validate an expert's trustworthiness.

greater numbers, is the example of rumors. Rumors are stories that are widely circulated and accepted though few of the believers have access to the rumored facts. If someone hears a rumor from one source, is that source's credibility enhanced when the same rumor is repeated by a second, third, and fourth source? Presumably not, especially if the hearer knows (or justifiably believes) that these sources are all uncritical recipients of the same rumor.

It will be objected that additional rumor spreaders do not add credibility to an initial rumor monger because the additional ones have no established reliability. The hearer has no reason to think that any of their opinions is worthy of trust. Furthermore, the rumor case doesn't seem to involve "expert" opinions at all and thereby contrasts with the original case. In the original case the hearer has at least some prior reason to think that each new speaker who concurs with one of the original pair has *some* credibility (reliability). Under that scenario, don't additional concurring experts increase the total believability of the one with whom they agree?

It appears, then, that greater numbers should add further credibility, at least when each added opinion-holder has positive initial credibility. This view is certainly presupposed by some approaches to the subject. In the Lehrer-Wagner (1981) model, for example, each new person to whom a subject assigns "respect" or "weight" will provide an extra vector that should push the subject in the direction of that individual's opinion.¹⁷ Unfortunately, this approach has a problem. If two or more opinion-holders are totally *non-independent* of one another, and if the subject knows or is justified in believing this, then the subject's opinion should not be swayed—even a little—by more than one of these opinion-holders. As in the case of a guru and his blind followers, a follower's opinion does not provide any additional grounds for accepting the guru's view (and a second follower does not provide additional grounds for accepting a first follower's view) even if all followers are precisely as reliable as the guru himself (or as one another)—which followers must be, of course, if they believe exactly the same things as the guru (and one another) on the topics in question. Let me demonstrate this through a Bayesian analysis.

Under a simple Bayesian approach, an agent who receives new evidence should update his degree of belief in a hypothesis H by conditioning on that evidence. This means that he should use the ratio (or quotient) of two likelihoods: the likelihood of the evidence occurring if H is true and the likelihood

¹⁷ Lehrer and Wagner say (p. 20) that one should assign somebody else a positive weight if one does not regard his opinion as "worthless" on the topic in question—i.e., if one regards him as better than a random device. So it looks as if every clone of a leader should be given positive weight—arguably, the same weight as the leader himself, since their beliefs always coincide—as long as the leader receives positive weight. In the Lehrer-Wagner model, then, each clone will exert a positive force over one's own revisions of opinion just as a leader's opinion will exert such force; and the more clones there are, the more force in the direction of their collective opinion will be exerted.

of the evidence occurring if H is false. In the present case the evidence in question is the belief in H on the part of one or more putative experts. More precisely, we are interested in comparing (A) the result of conditioning on the evidence of a single putative expert's belief with (B) the result of conditioning on the evidence of concurring beliefs by two putative experts. Call the two putative experts X and Y, and let X(H) be X's believing H and Y(H) be Y's believing H. What we wish to compare, then, is the magnitude of the likelihood quotient expressed in (1) with the magnitude of the likelihood quotient expressed in (2).

$$(1) \frac{P(X(H) / H)}{P(X(H) / \sim H)}$$

$$(2) \frac{P(X(H) \& Y(H) / H)}{P(X(H) \& Y(H) / \sim H)}$$

The principle we are interested in is the principle that the likelihood ratio given in (2) is always larger than the likelihood ratio given in (1), so that an agent who learns that X and Y both believe H will always have grounds for a larger upward revision of his degree of belief in H than if he learns only that X believes H. At least this is so when X and Y are each somewhat credible (reliable). More precisely, such comparative revisions are in order if the agent is *justified* in believing these things in the different scenarios. I am going to show that such comparative revisions are not always in order. Sometimes (2) is not larger than (1); so the agent—if he knows or justifiably believes this—is not justified in making a larger upward revision from the evidence of two concurring believers than from one believer.

First let us note that according to the probability calculus, (2) is equivalent to (3).

$$(3) \frac{P(X(H) / H) P(Y(H) / X(H) \& H)}{P(X(H) / \sim H) P(Y(H) / X(H) \& \sim H)}$$

While looking at (3), return to the case of blind followers. If Y is a blind follower of X, then anything believed by X (including H) will also be believed by Y. And this will hold whether or not H is true. So,

$$(4) P(Y(H) / X(H) \& H) = 1,$$

and

$$(5) P(Y(H) / X(H) \& \sim H) = 1.$$

Substituting these two values into expression (3), (3) reduces to (1). Thus, in the case of a blind follower, (2) (which is equivalent to (3)) is the same as (1), and no larger revision is warranted in the two-concurring-believers case than in the single-believer case.

Suppose that the second concurring believer, Y, is not a *blind* follower of X. Suppose he would sometimes agree with X but not in all circumstances. Under that scenario, does the addition of Y's concurring belief always provide the agent (who possesses this information) with more grounds for believing H? Again the answer is no. The appropriate question is whether Y is more likely to believe H when X believes H and H is true than when X believes H and H is false. If Y is just as likely to follow X's opinion whether H is true or false, then Y's concurring belief adds nothing to the agent's evidential grounds for H (driven by the likelihood quotient). Let us see why this is so.

If Y is just as likely to follow X's opinion when H is false as when it's true, then (6) holds:

$$(6) \quad P(Y(H) / X(H) \ \& \ H) = P(Y(H) / X(H) \ \& \ \sim H)$$

But if (6) holds, then (3) again reduces to (1), because the right-hand sides of both numerator and denominator in (3) are equal and cancel each other out. Since (3) reduces to (1), the agent still gets no extra evidential boost from Y's agreement with X concerning H. Here it is not required that Y is certain to follow X's opinion; the likelihood of his following X might only be 0.80, or 0.40, or whatever. As long as Y is just as likely to follow X's opinion when H is true as when it's false, we get the same result.

Let us describe this last case by saying that Y is a *non-discriminating reflector* of X (with respect to H). When Y is a non-discriminating reflector of X, Y's opinion has no extra evidential worth for the agent above and beyond X's opinion. What is necessary for the novice to get an extra evidential boost from Y's belief in H is that he (the novice) be justified in believing (6')

$$(6') \quad P(Y(H) / X(H) \ \& \ H) > P(Y(H) / X(H) \ \& \ \sim H)$$

If (6') is satisfied, then Y's belief is at least partly *conditionally independent* of X's belief. Full conditional independence is a situation in which any dependency between X and Y's beliefs is accounted for by the dependency of each upon H. Although full conditional independence is not required to boost N's evidence, *partial* conditional independence is required.¹⁸

We may now identify the trouble with the (unqualified) numbers principle. The trouble is that a novice cannot automatically count on his putative

¹⁸ I am indebted here to Richard Jeffrey (1992: 109–10). He points out that it is only conditional independence that is relevant in these kinds of cases, not "simple independence" defined by the condition: $P(Y(H) / X(H)) = P(Y(H))$. If X and Y are even slightly reliable independent sources of information about H, they won't satisfy this latter condition.

experts being (even partially) conditionally independent of one another. He cannot automatically count on the truth of (6'). Y may be a non-discriminating reflector of X, or X may be a non-discriminating reflector of Y, or both may be non-discriminating reflectors of some third party or parties. The same point applies no matter how many additional putative experts share an initial expert's opinion. If they are all non-discriminating reflectors of someone whose opinion has already been taken into account, they add no further weight to the novice's evidence.

What type of evidence can the novice have to justify his acceptance of (or high level of credence in) (6')? N can have reason to believe that Y's *route* to belief in H was such that even in possible cases where X fails to recognize H's falsity (and hence believes it), Y *would* recognize its falsity. There are two types of causal routes to Y's belief of the right sort. First, Y's route to belief in H might entirely *bypass* X's route. This would be exemplified by cases in which X and Y are causally independent eyewitnesses of the occurrence or non-occurrence of H; or by cases in which X and Y base their respective beliefs on independent experiments that bear on H. In the eyewitness scenario X might falsely believe H through misperception of the actual event, whereas Y might perceive the event correctly and avoid belief in H. A second possible route to Y's belief in H might go *partly through* X but not involve uncritical reflection of X's belief. For example, Y might listen to X's reasons for believing H, consider a variety of possible defeaters of these reasons that X never considered, but finally rebut the cogency of these defeaters and concur in accepting H. In either of these scenarios Y's partly "autonomous" causal route made him poised to avoid belief in H even though X believes it (possibly falsely). If N has reason to think that Y used one of these more-or-less autonomous causal routes to belief, rather than a causal route that guarantees agreement with X, then N has reason to accept (6'). In this fashion, N would have good reason to rate Y's belief as increasing his evidence for H even after taking account of X's belief.

Presumably, novices could well be in such an epistemic situation vis-à-vis a group of concurring (putative) experts. Certainly in the case of concurring *scientists*, where a novice might have reason to expect them to be critical of one another's viewpoints, a presumption of partial independence might well be in order. If so, a novice might be warranted in giving greater evidential weight to larger numbers of concurring opinion-holders. According to some theories of scientific opinion formation, however, this warrant could not be sustained. Consider the view that scientists' beliefs are produced entirely by negotiation with other scientists, and in no way reflect reality (or Nature). This view is apparently held by some social constructionists about science, e.g., Bruno Latour and Steve Woolgar (1979/1986); at least this is

Kitcher's (1993: 165–66) interpretation of their view.¹⁹ Now if the social constructionists are right, so interpreted, then nobody (at least nobody knowledgeable of this fact) would be warranted in believing anything like (6'). There would never be reason to think that any scientist is more likely to believe a scientific hypothesis H when it's true (and some other scientist believes it) than when it's false (and the other scientist believes it). Since causal routes to scientific belief never reflect "real" facts—they only reflect the opinions, interests, and so forth of the community of scientists—(6') will never be true. Anybody who accepts or inclines toward the indicated social-constructionist thesis would never be justified in believing (6').²⁰

Setting such extreme views aside, won't a novice normally have reason to expect that different putative experts will have some causal independence or autonomy from one another in their routes to belief? If so, then if a novice is also justified in believing that each putative expert has some slight level of reliability (greater than chance), then won't he be justified in using the numbers of concurring experts to tilt toward one of two initial rivals as opposed to the other? This conclusion might be right when *all* or *almost all* supplementary experts agree with one of the two initial rivals. But this is rarely the case. Vastly more common are scenarios in which the numbers are more evenly balanced, though not exactly equal. What can a novice conclude in those circumstances? Can he legitimately let the greater numbers decide the issue?

This would be unwarranted, especially if we continue to apply the Bayesian approach. The appropriate change in the novice's belief in H should be based on two sets of concurring opinions (one in favor of H and one against it), and it should depend on *how reliable* the members of each set are and on *how (conditionally) independent* of one another they are. If the members of the smaller group are more reliable and more (conditionally) independent of one another than the members of the larger group, that might imply that the evidential weight of the smaller group exceeds that of the larger one. More precisely, it depends on what the novice is *justified* in believing about these matters. Since the novice's justifiedness on these matters may be very weak, there will be many situations in which he has no distinct or robust justification for going by the relative numbers of like-minded opinion-holders.

This conclusion seems perfectly in order. Here is an example that, by my own lights, sits well with this conclusion. If scientific creationists are more numerous than evolutionary scientists, that would not incline me to say that a novice is warranted in putting more credence in the views of the former than

¹⁹ I myself interpret Latour and Woolgar as holding a more radical view, viz., that there is no reality that could causally interact, even indirectly, with scientists' beliefs.

²⁰ This is equally so under the more radical view that there are no truths at all (of a scientific sort) about reality or Nature.

in the views of the latter (on the core issues on which they disagree). At least I am not so inclined on the assumption that the novice has roughly comparable information as most philosophers currently have about the methods of belief formation by evolutionists and creationists respectively.²¹ Certainly the numbers do not *necessarily* outweigh considerations of individual reliability and mutual conditional independence. The latter factors seem more probative, in the present case, than the weight of sheer numbers.²²

5. Evidence from Interests and Biases

I turn now to the fourth source of possible evidence on our original list: evidence of distorting interests and biases that might lie behind a putative expert's claims. If N has excellent evidence for such bias in one expert and no evidence for such bias in her rival, and if N has no other basis for preferential trust, then N is justified in placing greater trust in the unbiased expert. This proposal comes directly from common sense and experience. If two people give contradictory reports, and exactly one of them has a reason to lie, the relative credibility of the latter is seriously compromised.

Lying, of course, is not the only way that interests and biases can reduce an expert's trustworthiness. Interests and biases can exert more subtle distorting influences on experts' opinions, so that their opinions are less likely to be accurate even if sincere. Someone who is regularly hired as an expert witness for the defense in certain types of civil suits has an economic interest

²¹ More specifically, I am assuming that believers in creation science have greater (conditional) dependence on the opinion leaders of their general viewpoint than do believers in evolutionary theory.

²² John Pollock (in a personal communication) suggests a way to bolster support for the use of "the numbers". He says that if one can argue that $P(X(H) / Y(H) \& H) = P(X(H) / H)$, then one can cumulate testimony on each side of an issue by counting experts. He further suggests that, in the absence of countervailing evidence, we should believe that $P(X(H) / Y(H) \& H) = P(X(H) / H)$. He proposes a general principle of probabilistic reasoning, which he calls "the principle of nonclassical direct inference", to the effect that we are defeasibly justified in regarding additional factors about which we know nothing to be irrelevant to the probabilities. In Pollock (2000) (also see Pollock 1990) he formulates the idea as follows. If factor C is irrelevant (he means *probabilistically* irrelevant) to the causal relation between properties B and A, then conjoining C to B should not affect the probability of something's being A. Thus, if we have no reason to think that C is relevant, we can assume defeasibly that $P(Ax / Bx \& Cx) = P(Ax / Bx)$. This principle can be applied, he suggests, to the case of a concurring (putative) expert. But, I ask, is it generally reasonable for us—or for a novice—to assume that the opinion of one expert is probabilistically irrelevant to another expert's holding the same view? I would argue in the negative. Even if neither expert directly influences the opinion of the other, it is extremely common for two people who work in the same intellectual domain to be influenced, directly or indirectly, by some common third expert or group of experts. Interdependence of this sort is widespread, and could be justifiably believed by novices. Thus, probabilistic irrelevance of the sort Pollock postulates as the default case is highly questionable.

in delivering strong testimony in any current trial, because her reputation as a defense witness depends on her present performance.

As a test of expert performance in situations of conflict of interest, consider the results of a study published in the *Journal of American Medical Association* (Friedberg et al., 1999). The study explored the relationship between published research reports on new oncology drugs that had been sponsored by pharmaceutical companies versus those that had been sponsored by nonprofit organizations. It found a statistically significant relationship between the funding source and the qualitative conclusions in the reports. Unfavorable conclusions were reached by 38% of nonprofit-sponsored studies but by only 5% of pharmaceutical company-sponsored studies.

From a practical point of view, information bearing on an expert's interests is often one of the more accessible pieces of relevant information that a novice can glean about an expert. Of course, it often transpires that *both* members of a pair of testifying experts have interests that compromise their credibility. But when there is a non-negligible difference on this dimension, it is certainly legitimate information for a novice to employ.

Pecuniary interests are familiar types of potential distorters of an individual's claims or opinions. Of greater significance, partly because of its greater opacity to the novice, is a bias that might infect a whole discipline, sub-discipline, or research group. If all or most members of a given field are infected by the same bias, the novice will have a difficult time telling the real worth of corroborating testimony from other experts and meta-experts. This makes the numbers game, discussed in the previous section, even trickier for the novice to negotiate.

One class of biases emphasized by feminist epistemologists involves the exclusion or underrepresentation of certain viewpoints or standpoints within a discipline or expert community. This might result in the failure of a community to gather or appreciate the significance of certain types of relevant evidence. A second type of community-wide bias arises from the economics or politics of a sub-discipline, or research community. To advance its funding prospects, practitioners might habitually exaggerate the probativeness of the evidence that allegedly supports their findings, especially to outsiders. In competition with neighboring sciences and research enterprises for both resources and recognition, a given research community might apply comparatively lax standards in reporting its results. Novices will have a difficult time detecting this, or weighing the merit of such an allegation by rival experts outside the field.²³

²³ In a devastating critique of the mental health profession, Robyn Dawes (1994) shows that the real expertise of such professionals is, scientifically, very much in doubt, despite the high level of credentialism in that professional community.

6. Using Past Track Records

The final category in our list may provide the novice's best source of evidence for making credibility choices. This is the use of putative experts' past track records of cognitive success to assess the likelihoods of their having correct answers to the current question. But how can a novice assess past track records? There are several theoretical problems here, harking back to matters discussed earlier.

First, doesn't using past track records amount to using the method of (direct) "calibration" to assess a candidate expert's expertise? Using a past track record means looking at the candidate's past success rate for previous questions in the E-domain to which she offered answers. But in our earlier discussion (section 2), I said that it's in the nature of a novice that he has no opinions, or no confidence in his own opinions, about matters falling within the E-domain. So how can the novice have any (usable) beliefs about past answers in the E-domain by which to assess the candidate's expertise? In other words, how can a novice, *qua* novice, have any opinions at all about past track records of candidate experts?

A possible response to this problem is to revisit the distinction between *esoteric* and *exoteric* statements. Perhaps not every statement in the E-domain is esoteric. There may also be a body of exoteric statements in the E-domain, and they are the statements for which a novice might assess a candidate's expertise. But does this really make sense? If a statement is an exoteric statement, i.e., one that is epistemically accessible to novices, then why should it even be included in the E-domain? One would have thought that the E-domain is precisely the domain of propositions accessible only to experts.

The solution to the problem begins by sharpening our *esoteric/exoteric* distinction. It is natural to think that statements are categorically either esoteric or exoteric, but that is a mistake. A given (timeless) statement is esoteric or exoteric only *relative* to an epistemic standpoint or position. It might be esoteric relative to one epistemic position but exoteric relative to a different position. For example, consider the statement, "There will be an eclipse of the sun on April 22, 2130, in Santa Fe, New Mexico." Relative to the present epistemic standpoint, i.e., the standpoint of people living in the year 2000, this is an esoteric statement. Ordinary people in the year 2000 will not be able to answer this question correctly, except by guessing. On the other hand, on the very day in question, April 22, 2130, ordinary people on the street in Santa Fe, New Mexico will easily be able to answer the question correctly. In that different epistemic position, the question will be an exoteric one, not an esoteric one.²⁴ You won't need specialized training or knowledge

²⁴ In the present discussion only *epistemic* esotericness, not *semantic* esotericness, is in question (see note 10).

to determine the answer to the question. In this way, the epistemic status of a statement can change from one time to another.

There is a significant application of this simple fact to the expert/novice problem. A novice might easily be able to determine the truth-value of a statement after it has become exoteric. He might be able to tell *then* that it is indeed true. Moreover, he might learn that at an earlier time, when the statement was esoteric for the likes of him, another individual managed to believe it and say that it is (or would be) true. Furthermore, the same individual might repeatedly display the capacity to assert statements that are esoteric at the time of assertion but become exoteric later, and she might repeatedly turn out to have been right, as determined under the subsequently exoteric circumstances. When this transpires, novices can infer that this unusual knower must possess some special manner of knowing—some distinctive expertise—that is not available to them. They presumably will not know exactly what this distinctive manner of knowing involves, but presumably it involves some proprietary fund of information and some methodology for deploying that information. In this fashion, a novice can verify somebody else's expertise in a certain domain by verifying their impressive track record within that domain. And this can be done without the novice himself somehow being transformed into an expert.

The astronomical example is just one of many, which are easily proliferated. If an automobile, an air-conditioning system, or an organic system is suffering some malfunction or impairment, untrained people will often be unable to specify any true proposition of the form, "If you apply treatment X to system Y, the system will return to proper functioning." However, there may be people who can repeatedly specify true propositions precisely of this sort.²⁵ Moreover, that these propositions are true can be verified by novices, because novices might be able to "watch" the treatment being applied to the malfunctioning system and see that the system returns to proper functioning (faster than untreated systems do). Although the truth of the proposition is an exoteric matter once the treatment works, it was an esoteric matter before the treatment was applied and produced its result. In such a case the expert has knowledge, and can be determined to have had knowledge, at a time when it was esoteric.²⁶

It should be emphasized that many questions to which experts provide answers, at times when they are esoteric, are not merely yes/no questions that might be answered correctly by lucky guesses. Many of them are questions that admit of innumerable possible answers, sometimes indefinitely many

²⁵ They can not only recognize such propositions as true when others offer them; they can also produce such propositions on their own when asked the question, "What can be done to repair this system?"

²⁶ I have discussed such cases in earlier writings: Goldman 1991 and Goldman 1999 (p. 269).

answers. Simplifying for purposes of illustration, we might say that when a patient with an ailment sees a doctor, he is asking her the question, "Which medicine, among the tens of thousands of available medicines, will cure or alleviate this ailment?" Such a question is unlikely to be answered correctly by mere guesswork. Similarly, when rocket scientists were first trying to land a spaceship on the moon, there were indefinitely many possible answers to the question, "Which series of steps will succeed in landing this (or some) spaceship on the moon?" Choosing a correct answer from among the infinite list of possible answers is unlikely to be a lucky guess. It is feats like this, often involving technological applications, that rightly persuade novices that the people who get the correct answers have a special fund of information and a special methodology for deploying it that jointly yield a superior capacity to get right answers. In this fashion, novices can indeed determine that others are experts in a domain in which they themselves are not.

Of course, this provides no algorithm by which novices can resolve all their two-expert problems. Only occasionally will a novice know, or be able to determine, the track records of the putative experts that dispute an issue before him. A juror in a civil trial has no opportunity to run out and obtain track record information about rival expert witnesses who testify before him. Nonetheless, the fact that novices can verify track records and use them to test a candidate's claims to expertise, at least in principle and in some cases, goes some distance toward dispelling utter skepticism for the novice/2-expert situation. Moreover, the possibility of "directly" determining the expertise of a few experts makes it possible to draw plausible inferences about a much wider class of candidate experts. If certain individuals are shown, by the methods presented above, to have substantial expertise, and if those individuals train others, then it is a plausible inference that the trainees will themselves have comparable funds of information and methodologies, of the same sort that yielded cognitive success for the original experts.²⁷ Furthermore, to the extent that the verified experts are then consulted as "meta-experts" about the expertise of others (even if they didn't train or credential them), the latter can again be inferred to have comparable expertise. Thus, some of the earlier skepticism engendered by the novice/2-expert problem might be mitigated once the foundation of expert verification provided in this section has been established.

²⁷ Of course, some experts may be better than others at transmitting their expertise. Some may devote more effort to it, be more skilled at it, or exercise stricter standards in credentialing their trainees. This is why good information about training programs is certainly relevant to judgments of expertise.

7. Conclusion

My story's ending is decidedly mixed, a cause for neither elation nor gloom. Skeptical clouds loom over many a novice's epistemic horizons when confronted with rival experts bearing competing messages. There are a few silver linings, however. Establishing experts' track-records is not beyond the pale of possibility, or even feasibility. This in turn can bolster the credibility of a wider class of experts, thereby laying the foundation for a legitimate use of numbers when trying to choose between experts. There is no denying, however, that the epistemic situations facing novices are often daunting. There are interesting theoretical questions in the analysis of such situations, and they pose interesting practical challenges for "applied" social epistemology. What kinds of education, for example, could substantially improve the ability of novices to appraise expertise, and what kinds of communicational intermediaries might help make the novice-expert relationship more one of justified credence than blind trust.²⁸

References

- Baron-Cohen, Simon (1995). *Mindblindness*. Cambridge, MA: MIT Press.
- Bezuidenhout, Anne (1998). "Is Verbal Communication a Purely Preservative Process?" *Philosophical Review* 107: 261–88.
- Brewer, Scott (1998). "Scientific Expert Testimony and Intellectual Due Process," *Yale Law Journal* 107: 1535–681.
- Burge, Tyler (1993). "Content Preservation," *Philosophical Review* 102: 457–88.
- Coady, C. A. J. (1992). *Testimony*. Oxford: Clarendon Press.
- Craig, Edward (1990). *Knowledge and the State of Nature—An Essay in Conceptual Synthesis*. Oxford: Clarendon Press.
- Dawes, Robyn (1994). *House of Cards: Psychology and Psychotherapy Built on Myth*. New York: Free Press.
- Foley, Richard (1994). "Egoism in Epistemology," in F. Schmitt, ed., *Socializing Epistemology*. Lanham, MD: Rowman & Littlefield.
- Friedberg, Mark et al. (1999). "Evaluation of Conflict of Interest in Economic Analyses of New Drugs Used in Oncology," *Journal of the American Medical Association* 282: 1453–57.
- Gentzler, J. (1995). "How to Discriminate between Experts and Frauds: Some Problems for Socratic Peirastic," *History of Philosophy Quarterly* 3: 227–46.

²⁸ For helpful comments on earlier drafts, I am indebted to Holly Smith, Don Fallis, Peter Graham, Patrick Rysiew, Alison Wylie, and numerous participants at the 2000 Rutgers Epistemology Conference, the philosophy of social science roundtable in St. Louis, and my 2000 NEH Summer Seminar on "Philosophical Foundations of Social Epistemology".

- Goldman, Alvin (1991). "Epistemic Paternalism: Communication Control in Law and Society," *Journal of Philosophy* 88: 113–31.
- Goldman, Alvin (1999). *Knowledge in a Social World*. Oxford: Clarendon Press.
- Hardwig, John (1985). "Epistemic Dependence," *Journal of Philosophy* 82: 335–49.
- Hardwig, John (1991). "The Role of Trust in Knowledge," *Journal of Philosophy* 88: 693–708.
- Jeffrey, Richard (1992). *Probability and the Art of Judgment*. New York: Cambridge University Press.
- Kitcher, Philip (1993). *The Advancement of Science*. New York: Oxford University Press.
- LaBarge, Scott (1997). "Socrates and the Recognition of Experts," in M. McPherran, ed., *Wisdom, Ignorance and Virtue: New Essays in Socratic Studies*. Edmonton: Academic Printing and Publishing.
- Latour, Bruno and Woolgar, Steve (1979/1986). *Laboratory Life: The Construction of Scientific Facts*. Princeton: Princeton University Press.
- Lehrer, Keith and Wagner, Carl (1981). *Rational Consensus in Science and Society*. Dordrecht: Reidel.
- Pollock, John (1990). *Nomic Probability and the Foundations of Induction*. New York: Oxford University Press.
- Pollock, John (2000). "A Theory of Rational Action." Unpublished manuscript, University of Arizona.