

Living with Ignorance in a World of Experts

Alexander A. Guerrero

forthcoming in

PERSPECTIVES ON IGNORANCE FROM MORAL AND SOCIAL PHILOSOPHY

(ed. by Rik Peels, Routledge, 2016)

please cite to final version

1. Introduction

I only know a tiny fraction of what there is to know. And I know this. And you are no different, whoever you are. Maybe things were not always like this. Maybe we used to live in a less complicated world, in which we all knew, or could easily know, much of what there was to know: things about family, culture, hunting, and gathering. But that is not the world we live in now. In modern societies there is an extensive and thorough division of labor. And with that division of labor comes a division of knowledge. This chapter isn't about the causes or explanation of this phenomenon. It is, however, about a distinctive problem that arises from this phenomenon: the problem of asymmetric ignorance.

We often find ourselves in situations of asymmetric ignorance, where both parties know of the relevant asymmetry.¹ I am a patient and you are a doctor. I am a client and you are a lawyer. I am a citizen and you are a politician. I am a consumer and you are a manufacturer. I am not a geologist (chemist, engineer, plumber, nurse, accountant, soldier, historian, cab driver, judge, etc.) and you are. And you also know that there is this asymmetry between us. Let's refer to these as cases of *known* asymmetric ignorance. Obviously, there could also be cases of unknown or partially unknown asymmetric ignorance. Those cases raise distinct problems. I will leave them aside. Cases of known asymmetric ignorance are common and troublesome enough on their own.

In addition to looking at the subset of asymmetric ignorance cases in which both parties know of the asymmetry, I want to restrict attention to cases that have two additional features.

First, that the asymmetric ignorance is *systematic* over some relatively significant domain. So, Igor is ignorant—and Nora is not ignorant—not just about some fact, *F*, but about a whole host of *F-related* or *close to F* facts, such that those facts comprise a domain of significant and systematic knowledge. I want to concentrate on cases where the asymmetry of ignorance is more extensive, deeper, and more systematic. Spelling this out precisely would require more work; I will leave it rough for now. The basic category I want to capture are those cases in which it would be natural to invoke the idea of asymmetric ignorance borne of asymmetric *expertise*. A very common, and very important kind of known asymmetric ignorance case arises when non-experts with respect to some domain, *D*, interact with people

¹ Much has been written about how exactly to understand ignorance. Two broad distinctions are between ignorance as *lack of knowledge* and ignorance as *lack of true belief*. For a helpful discussion, see Peels (2010). For the purposes of this chapter, I will consider ignorance to be lack of true belief, although little turns on this. One might lack true belief about some matter in more than one way. As I note in Guerrero (2007), one might be to lack true belief in some fact, *F*, in at least three different kinds of cases:

(I1) cases in which a person is ignorant because she has never thought about the issue (and so has no beliefs about *F* or only has unexamined 'implicit' beliefs about *F*)

(I2) cases in which a person is ignorant because, though she has thought about the issue, she has come to have false beliefs about *F* (she believes that not-*F* when in fact, *F*)

(I3) cases in which a person is ignorant because, though she has thought about the issue, she doesn't know what to believe (she doesn't believe that *F* or that not-*F*)

Because my focus will be on cases of known asymmetric ignorance, the ignorance in question will usually be of either the (I1) or (I3) variety.

who are experts with respect to D. I want to restrict attention to those kinds of cases, relying on an intuitive conception of expertise and what it is to be an expert.²

Second, I want to focus on cases in which the experts in question may have or appear to have interests—economic, professional, personal—that are in less than complete alignment with the non-experts, and in which this non-alignment gives rise to at least some reason to be skeptical of the expert. This doesn't mean that the experts are definitely deceiving or manipulating the non-experts; only that this is a possibility that one must take seriously. There are background, imperfectly aligning interests and potential biases of which one is aware. These are not cases in which the expert is your best, most trusted friend, or in which the incentive structures are set up so that what is good for her is always good for you, and vice versa. Nor are they cases in which one can't imagine why the expert might not tell the full story or exactly the story that she most confidently believes. Given even a little cynicism on our part, we can easily imagine this. Nora might not be trying to help us, or at least not *just* trying to do this.

So, my focus is restricted to cases in which (a) there is known, asymmetric ignorance between two or more parties, (b) that asymmetry is the result of asymmetric expertise, (c) there is some measure of non-alignment between the expert's interests and the non-expert's interests, and (d) that non-alignment is such so that one has at least some reason not to completely trust the expert(s). Call contexts in which these four conditions hold “strategic expertise contexts.”

What motivates my focus on cases with these restrictions? The short answer: cases like this are pervasive and present a host of distinct and difficult problems for those of us on the ignorant, non-expert side of things (which will, at times, be all of us). A few salient examples: political representatives and constituents, doctors and patients, lawyers and clients, military officials and the general public, manufacturers and potential consumers, specialist employees and the corporate managers trying to manage them, scientists and the non-scientist public, mechanics and the automotively ignorant.

Additionally, although there has been some discussion of expert testimony, to my knowledge there has been no discussion of cases that involve both asymmetric expertise and some concern about non-aligned incentives. Not that this concern has never arisen for people thinking about experts. For example, Elizabeth Fricker writes: “where I know another to be epistemically expert relative to me on a topic, it is not just rationally permissible, but rationally mandatory for me to accept her judgment in preference to my own, *just so long as I have good ground to trust her sincerity.*”³ But what if we lack good ground to trust her sincerity? (Note that this will not always be the same as having good ground to *distrust* a person; we just lack the background positive reason to think them especially worthy of trust or especially likely to be sincere with us.) What if both of the experts, or all of the experts, are subject to possible biases or have non-aligned interests? Or what if we do not know which experts are subject to biases and which aren't? Is our situation hopeless? But isn't this often, if not mostly, our situation?

In this chapter, I want to focus on two questions that arise in a particularly pronounced way in these contexts of strategic expertise:

² A completely separate paper could be written on what it is for a person to be an expert. I will work with a rough and ready sense which allows for a wide variety of kinds of experts. In general, however, an expert in a domain D is someone who is able and likely (a) to answer a question in domain D correctly, or with considerably greater accuracy than most non-experts; (b) to identify the considerations that are relevant to answering the question; and/or (c) to deploy methods in the future that will help develop answers to or understanding of other questions within D. More demanding, thicker ideas of expertise will require all of (a)-(c). Thinner conceptions will require only some of (a)-(c).

There is an odder sense of “expert” where a person can be an expert in an ultra-narrow domain, so that we might consider things like being in a particular location at a particular moment in time, being in a better perceptual situation, and so on, might be enough to make one a very localized kind of “expert” with respect to some matter. I will have in mind a more regular notion of expertise in which a person is only an expert if the domain of expertise is fairly broad and systematically related. But see, for example, Fricker (2006), who goes a different direction on this, discussing a more “inclusive” sense of expertise.

³ Fricker (2006), p. 243 (emphasis added).

(1) What should we think about the epistemology of testimony in strategic expertise contexts? More specifically: can we rationally or justifiably form beliefs relying on expert testimony in these contexts? Do they constitute knowledge? Under what circumstances?

(2) Are there interpersonal or institutional mechanisms that can improve the reliability of testimony and the appropriateness of trust in strategic expertise contexts? What are they? How might they work?

As suggested above, these contexts are commonplace, including some of the most central aspects of our moral and social lives: law, politics, medicine, science, commerce. Many have commented on the role that expertise plays in modern societies. The other side of this is the great extent to which we all find ourselves in conditions of asymmetric ignorance. The phenomenon of voter ignorance, which is an example of asymmetric ignorance (in relation to the position of elected representatives), is both widespread and much remarked upon.⁴ I want to draw attention to the fact that widespread and significant asymmetric ignorance about matters that are crucially important to our well-being extends not just to politics, but to almost every domain of contemporary social life. It is important for epistemologists and other philosophers to think about how and whether asymmetric ignorance poses a problem, and to consider what solutions might be available.⁵ This chapter is one such effort, focusing on our reliance on testimony in these contexts.

2. The “Standard Picture” of Testimony

There has been a considerable amount of recent philosophical work on the epistemology of testimony—on the epistemic evaluation of our beliefs that are based on what others tell us. This work has focused on the questions of when these beliefs are justified, when they amount to knowledge, and why these beliefs are justified (when they are) and why they amount to knowledge (when they do).

In attempting to answer these questions, almost all parties to the epistemological debates assume what we might call the “standard picture” of testimonial context or what Jonathan Adler calls “core cases” of testimony. This standard picture shifts the focus entirely away from cases of asymmetric expertise or cases involving experts at all.⁶ Drawing from Adler’s entry in the Stanford Encyclopedia of Philosophy on “Epistemological Problems of Testimony” and Jennifer Lackey’s introduction to the Oxford University Press volume edited by her and Ernest Sosa, *The Epistemology of Testimony* (2006), we can identify the features of these “core cases” of testimony or the “standard picture” assumptions. The focus has been on cases with this basic form:

- (1) Testifier to Subject: “ p ”
- (2) Subject forms a belief that p on the basis of Testifier’s utterance in (1)

In addition, the standard picture assumes that:

⁴ The literature here is significant, and all comes down in a similar place. For a characteristic recent example, consider Somin (2013), p. 13: “The sheer depth of most individual voters’ ignorance is shocking to many observers not familiar with the research.”

⁵ Efforts in this regard have been made, particularly with respect to the problem of asymmetric ignorance in politics, law, and science, and in particular with respect to the funding of science and the creation of science policy and scientifically-informed policy. For discussions of problems posed by asymmetric ignorance in these domains, see Kitcher (2001), Douglas (2005), Brewer (1998), Michaels (2008), Whyte and Crease (2010), Guerrero (2014), among others.

⁶ For example, in Sanford Goldberg’s excellent monograph, *Relying on Others: An Essay in Epistemology*, there is exactly one mention of expertise, and no extended discussion of the problem of relying on others when those others are experts.

- (A1) The primary speech-act of testimony is a Testifier's asserting something
- (A2) The assertion is an instance of literal usage
- (A3) Testifier's purpose in asserting is to communicate truths and inform Subject
- (A4) The testimony is the sustaining basis of the corresponding belief in Subject
- (A5) Subject has no special reason to be skeptical of Testifier
- (A6) Testifier is assumed not to have "expert knowledge" on the topic of her assertion
- (A7) Subject has no special knowledge about Testifier
- (A8) Testifier is a stranger to Subject

Additionally, three empirical claims are made about testimony as an actual phenomenon in the world:

- (EC1) Subjects usually accept the assertions of Testifiers
- (EC2) Subjects usually have no way of checking or confirming the Testifier's reliability or sincerity
- (EC3) Testimony is usually truthful

Let me first say a bit about what has been said to motivate the above assumptions about the "standard picture," before turning to consider the empirical claims.

Adler notes that many engaged in investigating the epistemology of testimony are interested in both our far-reaching dependence on testimony, and our apparent vulnerability in relying on testimony. Many have noted that a vast number of our beliefs arise through reliance on testimony and the inferences it justifies (Hume (1748), Reid (1764), Price (1969); Sosa (1994)). Adler (2012) suggests that to properly engage this "Vulnerability Problem" we should use "a class of core cases that isolate our dependence on the word of the speaker and whatever epistemic resources are available in ordinary conversational contexts." We don't want to make reliance on testimony an easier case for us by adding in lots of information about Testifiers that we usually lack.

Regarding the empirical claims: why would testimony usually be truthful? Why might we expect truthful testimony in these "core" cases? Different answers have been offered, often relying on general norms regarding communication and assertion. David Lewis argues that truthfulness is a presupposition of linguistic communication (Lewis 1969, 1975). Paul Grice (1975) suggests that there is pressure to make cooperative contributions in line with what have come to be called the Gricean maxims. Some suggest that detected departure from truthfulness leads to sanctions and damaged reputations, and this creates incentives toward truthfulness. More generally, Tyler Burge (1993) has argued that lying or deception are only rarely well-motivated or rational options. Another suggestion is that Testifiers know that testimony must meet a minimal standard of prior plausibility or face challenge or rejection.

The three key empirical claims—that we usually accept testimony, that testimony is usually truthful, and yet we usually have no way of checking the reliability or credibility of the testimony or the Testifier—make testimony seem to be something of a remarkable phenomenon, despite being a thoroughly unremarkable part of social life. The assumptions of the standard picture are an attempt to focus on some of the features of testimony that seem both prevalent and unremarkable, but also to be something of a worst-case for the reliability of testimony: we don't know the Testifier, we don't know their track record, they aren't any kind of expert, we are relying on them pretty much completely. And despite this—and this is the remarkable part—Testifiers are usually truthful, and testimony usually transmits truth. Taking this as the standard picture allows epistemologists to think about how this could possibly work, and to evaluate the epistemic merit of what is taking place under these 'strangers giving directions' kind of cases.

However, the focus on the difficulties presented by this kind of "core case" has obscured difficulties that are not present in these "core cases," but which are actually central features of a great number of cases of "the actual totality of testimony" in modern life. Namely, cases of testimony that occur in asymmetric expertise contexts, strategic contexts, and strategic expertise contexts more

specifically.⁷ In those cases, we cannot assume that there is no special reason to distrust the testifier, nor should we assume that certain general maxims—Gricean norms regarding communication, Lewisian norms regarding conventional truthfulness of speakers—are applicable: at least not as licensing inferences regarding how speakers are actually behaving or likely to behave. And we cannot assume that the Testifier lacks special “expert knowledge” on the topic of her assertion.

Adler and those who focus on the “standard picture” suggest that the “core cases” are something like the statistically normal cases of actual testimony, and thus an appropriate place to focus on to understand testimony and the epistemological trickiness of testimony. Adler (2012) says that the “relevant reference class is the actual totality of testimony.” But no empirical evidence is provided to support these claims, nor to support the above three empirical claims.

It is easy to overestimate the number of cases that actually exhibit the features of “standard picture” cases, given that in many ordinary contexts, we will not be basing our belief solely or even predominantly on testimony. It is true that we receive a great deal of testimony throughout a normal day, but much of what is testified to is also supported by our perceptual observations, other background evidence, memories, and so on. If we focus on cases in which it is true that testimony really is the sustaining basis of our belief, there are reasons to be skeptical of claims that “testimony overwhelmingly transmits truth.” No empirical evidence is introduced to support these claims of statistical prevalence or what is “usually true” about testimony, Testifiers, and Subjects. But even if such evidence were available and offered, there would still be a concern that this focus ignores a very prevalent kind of testimony—cases of testimony that occur in strategic contexts and strategic expertise contexts in particular.

In ordinary life, we make assertions for many reasons other than to simply provide information—we aim to inform, but also to recommend, encourage, advise, warn, console, convince, manipulate, and deceive. And this is true for experts as well. Some, such as Sperber (2001), have even suggested that communication evolved not from the benefits of providing information, but from the benefits to Testifiers of manipulating the beliefs of Subjects. In some cases, we can do these things just by making (a specifically chosen set of) truthful assertions. But the truth is not always, or not always fully and exclusively, on our side. And it is plausible that testimony in strategic expertise contexts actually presents some of the most difficult epistemological questions, particularly if we are interested in something like the “Vulnerability Problem.”

If one thinks that we should start with the “core cases” in order to learn about the epistemology of testimony in those cases, so as to understand the basics of the epistemology of testimony, there is a worry about the generality and applicability of the lessons learned. The implicit suggestion has been that work on these “core” cases will provide significant insight into the epistemic issues that arise in these other cases—call this the continuity suggestion. In the next several sections, I will show that we should reject the continuity suggestion.

3. Beyond the Standard Picture

If we focus on instances of testimony that occur in strategic expertise contexts, which parts of the standard picture have to be given up? In this section, I will draw attention to those parts of the standard picture that are incompatible with strategic expertise contexts. (Throughout, when I reference “testimony that occurs in strategic expertise contexts” I will be focusing on the testimony that is offered by experts to non-experts.)

3.1 Strategic Communication v. Informational Communication

⁷ There is some work, which I will discuss later, that focuses on the problem of asymmetric expertise and how non-experts or “novices” can rationally or justifiably rely on the testimony of experts. See Goldman (2001), Walton (1997), Fricker (2002). And this assumption is sometimes rejected in the context of work on scientific testimony and legal testimony—see Kitcher (1993); Walton (1997); Brewer (1998). But it is standard in most discussions of the epistemology of testimony.

First, let us consider the assumptions that we might call the “non-strategic” assumptions:

- (A3) Testifier’s purpose in asserting is to communicate truths and inform Subject
- (A5) Subject has no special reason to be skeptical of Testifier

There are clearly many kinds of cases in which these assumptions will not, or may not, hold.

- A salesperson as Testifier, who wants Subject to buy what she is selling.
- A doctor as Testifier, who wants Subject to accept that the adverse outcome was just “one of those things that happens.”
- A political candidate as Testifier, who wants Subject’s vote.
- A scientist as Testifier, who wants Subject to give her a grant.

In all of these cases, we can imagine scenarios in which the Testifier’s purpose in asserting is *not* to communicate truths or inform, but to get the Subject to make a decision or come to have a belief that will be beneficial to the Testifier. Perhaps they will also communicate truths and inform while doing this, perhaps not. But this is not their *purpose* in asserting. Similarly, these may all be cases in which Subjects do have special reasons to be skeptical of Testifiers, precisely because they are aware that Testifiers might be motivated in self-interested ways as just suggested.

These kinds of cases all involve asymmetric expertise. But there are also many other familiar kinds of cases that also involve strategic communication, but not expertise:

- A suitor as Testifier, who wants to win Subject’s affections.
- A university president as Testifier, who wants Subject to donate money.
- A student as Testifier, who wants Subject to accept his late-paper excuse.

I include this latter group to suggest how common it is for the non-strategic assumptions to be inappropriate. Note that these are not all instances in which Testifier is trying to deceive Subject or to take advantage of Subject. The motivations are more diverse, and more complicated. Furthermore, we often know that we are, or may be, in this kind of situation. But we may not always know or realize this.

3.2 Expertise and Asymmetry

An obvious part of the standard picture that is incompatible with strategic expertise contexts is the ‘non-expertise’ assumption:

- (A6) Testifier does not have any special “expert knowledge” on the topic of her assertion.

This assumption is made to focus our attention on run-of-the-mill testimony cases, and the predicament we find ourselves in with respect to those cases. The background thought is something like this: if we knew that the testimony was from an expert, that might change our view about the reliability, knowledge-transmission, justification-transmission, etc. of the testimony in question, but in a way that would depart from ordinary, non-expert testimony.

Importantly, it is a mistake to think that when expert testimony is involved, this simply makes the case *better* for the reliability, knowledge-transmission, justification-transmission, etc., of the testimony. It is true that it makes things in *one* way better: the Testifier is an expert, and on any plausible understanding of expertise, this will mean that the Testifier is either reliable and accurate or at least better than average with respect to reliability and accuracy of her beliefs in the domain of her expertise. On the other hand, it makes things in at least one way *worse*, at least in many cases: it may be much harder for a Subject to have or trust any independent judgment of the plausibility of the testimony offered by the expert Testifier, and it may be much harder for a Subject to get independent confirmation of the truth of the testimony from a non-expert source. We will discuss this at length later, but there is a way in which

our epistemic dependence on the testimony of experts is more complete or total or, as Hardwig (1985) puts it, “blind.” And we can encounter distinctive problems in this regard when the fact of asymmetric expertise is a matter of mutual knowledge between a Testifier and Subject. In contexts of known asymmetric expertise, non-experts may be relatively helpless, and this helplessness is a matter of mutual knowledge.

4. The Epistemology of Testimony in Strategic Expertise Contexts

Those working on “core cases” with the standard picture in the background have tended to argue about three main epistemological issues regarding testimony: the appropriateness of default rules; whether, when, and how testimony transmits knowledge; and whether one should accept reductionism or non-reductionism about testimony. In this section, I show that we should reject the continuity suggestion by presenting the debates as conducted regarding “core cases” arising under the standard picture, and then showing how moving to strategic expertise contexts dramatically alters the epistemological issues. We should be wary of claims suggesting that what is true of testimony in “core cases” is true of testimony more generally. Methodologically speaking, if one attempts to make general claims about the epistemology of testimony, then these claims must be tested against a broad set of testimonial contexts. Many theses about the epistemology of testimony look much more, or much less, plausible if we think about the broader—and more realistic—world of testimony which includes strategic expertise contexts. Instead of attempting to conclude that “testimony” in general has this or that epistemic feature, claims should be offered in a much more restricted way: “testimony in cases with background features X, Y, and Z has epistemic properties 1, 2, and 3.”

4.1 Default Rules for Testimony

Some have argued for a default rule for testimony:

If the Testifier asserts that p to the Subject, then, under normal conditions, it is correct (or epistemically appropriate, or permissible, or warranted) for Subject to accept (believe) Testifier’s assertion.

Features of the standard picture sketched above are used to suggest that such a default rule is appropriate. One argument in this vein concerns the reliability of testimony, where “testimony” is treated as a unified category. Shogenji (2006) argues that the omnipresence of testimonially-based beliefs—and therefore the ubiquity of reliance on the reliability of testimony—can be used to give greater confirmation for reliability of testimony: “[e]ven if the degree of tacit confirmation by a single observation is small, there are plenty of such observations... Their cumulative effect is substantial and should be sufficient for justifying our trust in testimony.” But as with reliability arguments in general, there is a worry that at most what we get from this kind of reliability argument for a default rule is something bordering on the tautological: in reliable contexts, testimony is reliable.

If we include in the category of “testimony” all of the testimony that occurs in strategic expertise contexts, or strategic contexts in general, it seems much more of an open question whether “testimony” in general is reliable. So, too, with the defense of the default rule in general. The default rule regarding the justifiability of believing testimony covers testimony that occurs “under normal conditions” and for which there is no “special” reason not to accept the testimony. Presumably, defenders of this rule would say that strategic contexts are not “normal” or that they introduce “special” reasons. Are these just descriptive claims? If so, are they correct?

One might defend such claims by arguing that strategic contexts are a statistically small number of the total testimonial contexts, so that the general claims about the overall reliability of testimony or what conditions are “normal” might still hold. But I think we should be skeptical of this claim. It seems that “stranger giving directions” contexts are actually the more unusual kinds of cases. As Graham (2004) points out, testimony is different than perception in that it is the product of individuals with free will, with

communicative-intentions that do not operate in law-like fashion, and which are compatible with intentional untruthfulness. I think we are often in strategic contexts. Maybe mostly.

Even leaving that aside (as there is no sophisticated empirical data available to support either case), when we know that we are in a strategic context, as we often will, we will not be able to rely on these general claims about the reliability of testimony or our default justification in accepting testimony. At least not without showing or arguing that testimony is reliable or is default justified even in strategic contexts, or even in strategic expertise contexts. And I think this will be very hard to do.

So, perhaps the default rule survives, but its application will be significantly limited. Many contexts will be such that it does not apply.

4.2 Transmitting Knowledge

Another central question in the epistemology of testimony concerns when a Testifier can “transmit” knowledge to a Subject simply via his or her testimony. Many⁸ have suggested a principle like this one:

Knowledge transmission: If Testifier knows that p and Testifier asserts that p to Subject, and Subject accepts p on the basis of Testifier’s testimony, then Subject knows that p .

This principle might be defended on the basis of the following pair of claims:

- (a) knowledge is a norm of assertion: one should assert that p only if one knows that p ⁹
- (b) speakers generally conform to the knowledge norm of assertion

If we give up the standard picture and reject the claim that speakers generally conform to the knowledge norm of assertion, there are significant worries about the above knowledge transmission principle. As before, there is a question of what kind of evidence there is for claim (b). Is this evidence sufficient to ground the claim that speakers “generally” conform to this norm? If we suppose that the norm is often violated in strategic contexts, does this undermine claim (b)?

Leaving that question aside, it at least seems that in strategic contexts it is an open question whether claim (b) is true. Speakers in such contexts will often, or at least sometimes, assert things they don’t even believe, things they know are false, things they doubt, things they know are not the whole truth of the matter, things they only hope are true, and so on. So, in a strategic context, when Testifiers do know that p , and assert that p to a Subject, and the Subject accepts p on that basis, this may seem to be something of a special case, rather than anything to be expected or relied upon.

If that’s right, then there is a real worry about the ability for knowledge to transmit via testimony in strategic contexts, and perhaps in particular in strategic expertise contexts, even in cases in which the Testifier does know that p . It would just be a matter of luck that, for a particular Subject, in a particular case, the Testifier testifies that p and actually knows that p . Even if a Subject comes to believe some true proposition p on the basis of testimony from a Testifier who knows p , it could have easily have been otherwise—the Subject could have easily falsely believed that p . All it would have taken is that the Testifier in the particular case have been one who did not actually know that p . Furthermore, in cases of asymmetric expertise, it would be particularly difficult for a Subject to tell if she were in a “good” case or a “bad” case. Our ignorance, as non-experts, means that we will not be in a good position to do this.

So, if we accept a safety condition on knowledge, as has been defended by Duncan Pritchard (2007) and Ernest Sosa (1999) among many others, then we should reject the claim that knowledge transmits via testimony in strategic contexts. Although there are many debates about exactly how to understand safety conditions on knowledge, the basic thought is that an agent knows a true proposition p only if that agent could not easily have falsely believed p . Or, as Williamson (2000) puts it, “if one knows, then one could not easily have been wrong in a similar case.” In strategic expertise contexts, we,

⁸ For example: Hintikka (1962); Evans (1982); Adler (1996); Audi (1997); Reynolds (2002); Sutton (2007).

⁹ This claim is suggested and defended by many, including Unger (1975); Williamson (1996), (2000); Adler (2009).

as the ignorant party, could very easily have come to falsely believe p ; we could easily have been wrong in a similar case.¹⁰

If we think that there are many cases in which it is obscure whether we are in a strategic context or not, these concerns about safety and luck may have even wider implications for the claim that testimony transmits knowledge. Put another way, the extent to which people do not conform to the knowledge norm of assertion has a direct effect on the plausibility or generality of the claim that testimony is a means by which to transmit knowledge.

If we are in situations of known, widespread asymmetric ignorance, we should be more skeptical of testimony as a source of knowledge—given the many other motivations Testifiers may have for making assertions, and given our inability to discern the good cases from the bad due to our ignorance. Robust levels of political ignorance, scientific ignorance, medical ignorance, and so on, have left us vulnerable to classes of both real experts and pseudo-experts, both of whom can be enlisted to mislead and manipulate us into accepting all manner of false propositions, while pretending to believe, or in some cases genuinely believing, that they know these propositions. Michaels (2008) sets out in fascinating detail some of the strategies invoking expert testimony of various forms that were employed by the tobacco, beryllium, and phenylpropanolamine industries, but other examples abound.¹¹

4.3 Reductionism vs. Non-Reductionism

A third main issue in epistemological debates regarding testimony is the kind of source of justified belief that testimony is, if it is a source of justified belief. Some, such as Coady (1973), defend “non-reductionism,” arguing that testimony is a *basic* source of justification (warrant, entitlement, knowledge), like sense perception, memory, and (perhaps) inference. Others defend “reductionism,” arguing that testimony is only a *derivative* source of justification, dependent on the justification we have for sense perception, memory, and inductive inference. In that sense, it can be “reduced” to these other sources of justified belief. A related question is how much ‘positive epistemic work’ a Subject has to do in order to rationally or justifiably form a belief on the basis of testimony. Can we just passively receive testimony, relying on it unless we have special or unusual evidence against accepting the testimony in this particular case? Or do we always need to have positive reasons for accepting testimony or for relying on this particular Testifier? The reductionist picture holds that testimony is like inference and places a significant burden on the recipient of testimony. Non-reductionist views hold that testimony is closer to perception or memory and places a relatively light burden on the recipient of testimony.

If we move outside of the standard picture, the non-reductionist account may begin to look implausible. If someone is in a strategic expertise testimonial context, it is plausible that more work must be done on the part of the Subject than simply passively receiving the input, the testimony. And if we consider that many testimonial contexts will involve some dimension of strategy or asymmetric expertise, the non-reductionist picture may look implausible across the board. Instead, acceptance of testimony in strategic expertise contexts will only be justified a posteriori, upon having done some investigation into the source of the testimony, the incentives of the Testifier, and other evidence that bears on the Testifier’s expertise, credibility, reliability, and honesty. Gathering this evidence will require familiar methods: perception, memory, induction, and so on.

Thinking, then, of the standing of testimony more generally, we should be skeptical of non-reductionism. If strategic contexts abound—as I have suggested that they do—then we should be warier of testimony than we often are. At any rate, the argument for non-reductionism cannot rely on a simple

¹⁰ There are complexities here depending on how the safety condition is developed. In particular, if it is developed in a more externalist direction, it may not matter that we are unable to discern “good” experts from “bad” experts, as long as we in fact happen to be dealing with “good” experts. It strikes me as dubious to credit such an individual with knowledge, and that this dubiousness is explained by safety-like reasons, but there is not space to make that argument in full here. Thanks to Karl Schafer for raising this concern.

¹¹ Stanley (2015) offers a great number of examples in the political case.

appeal to “normal” cases.¹² Rather than being like perception in the regular world, testimony is like perception in a fun house, a house of mirrors, or at least a world in which a significant portion of the houses are fun houses. Perhaps we can tell when we are in a fun house, a strategic context, and so reductionism is only true of testimony in those contexts, and non-reductionism may be true elsewhere. But this won’t always be the case, and it is an empirical question the relative proportion of contexts that are like the standard picture would have them as opposed to being more strategic. In this, I side with those like Elizabeth Fricker (1994) who claims that hearers can and ought to obtain independent evidence to confirm the belief that a speaker is trustworthy. This seems particularly right in cases in which there are strategic considerations in the background, in which the testimony is or would be the sustaining basis for our belief, and in which there is known asymmetric expertise (so that the expert knows the non-expert is unlikely to be able to assess the plausibility of the testimony directly).

In this section, I have argued that in strategic expertise contexts, testimony is not default justified, testimony is not knowledge transmitting, and non-reductionism about testimony is implausible. If we think that these contexts either (a) constitute a large portion of the total testimonial contexts in which testimony is the sustaining basis of a Subject’s belief or (b) are a large enough portion to “taint” the more standard picture kinds of cases due to our inability to screen for them, then we should also think that testimony in general is not default justified, is not knowledge-transmitting (at least not without something more), and that non-reductionism about testimony is false. Perhaps we should think some of these things anyway.

Where does this leave us? Well, what I have suggested only applies to the case of bare testimony—testimony we receive and do nothing to investigate or question. But we needn’t see ourselves as purely helpless in all of this. In contexts of strategic expertise, constituted in part by asymmetric ignorance, we will need some way of checking on the accuracy, reliability, and sincerity of Testifiers and particular items of testimony for us to rationally form beliefs based on that testimony (or for that testimony to be knowledge- or justification-conferring). The difficult thing about these kinds of cases is it may seem that we won’t be in a position to do this, because of the ignorance that gives rise to the problem in the first place.

In the final section of the paper, I discuss methods and mechanisms that might be used or put in place to aid non-experts in contexts of strategic expertise. The effort here can be seen as helping to develop what Sperber et al. (2010) have referred to as institutions of “epistemic vigilance.” One hope is that, if these methods can be used successfully, even under conditions of ignorance, then in those cases in which they are used, belief based on testimony will be justified, knowledge may be transmitted through testimony, and testimony is likely to prove reliable. That’s a reason for optimism—we needn’t see ourselves as in the deep dependence well that some such as Hardwig (1985) would suggest that we are in. On the other hand, the success of these methods is far from assured.

5. How to Respond to Strategic Expertise Testimony?

The problem of how non-experts can rely on or trust experts is not a new one. Plato struggles with the problem in *Charmides*. More recently, Alvin Goldman (2001) has focused on the problem of a non-expert trying to determine which of two purported experts to believe in cases in which they disagree. Douglas Walton (1997), Harry Collins and Robert Evans (2007), and Elizabeth Anderson (2011) all offer criteria or mechanisms which they suggest can be helpful for non-experts when assessing and relying on experts and expert testimony. Throughout, I will assume that one of the strategies that is not available is simply to

¹² There may be a more elaborate argument stemming from the need for some sort of default entitlement to get the whole structure of justification off of the ground, but I think such an argument will be less plausible in the testimonial case than in, say, the perceptual case. But fuller discussion is required.

have the non-experts *become* experts. That is, I will assume that addressing our relative ignorance of the domain directly is not an option. This seems essential, given the reasons we have to think that asymmetric ignorance and asymmetric expertise will remain pervasive features of modern life.

There are at least four questions that non-experts might need to answer; it is important to keep them distinct (not all suggestions will help with all of them):

Expertise

(Q1) Is this particular person, E, an expert in this particular domain, D?

Comparative Expertise

(Q2) Is purported expert E1 more expert in domain D than some other purported expert, E2?

Sincerity of Testifier

(Q3) Does the Expert Testifier believe her testimony is the truest thing she is able to say on the topic in question?

Quality of Testimony

(Q4) Is this particular piece of expert testimony, T, likely to be reliable (accurate, true)?

Obviously, these can come apart. One might not be an expert, but might be perfectly sincere. One might be an expert but have reasons to deceive or mislead. That is a central case we are considering. One might be an expert and sincere, but subject to some kind of unconscious bias in this particular case. Importantly, the answers to questions (Q1) and (Q2) may be completely distinct from (Q3) and (Q4). This is the reason for focusing not just on expertise—as all of the above authors do—but also on strategic dimensions of communication and testimony. These are two distinct problems that create testimonial difficulties for us, but they compound when combined. Goldman (2001), Walton (1997), Collins and Evans (2007), and Anderson (2011) all discuss the need to screen for potential conflicts of interests or biases. It is not that they are unaware of this concern. But I think that they may understate it. I will say more about this in a moment.

Additionally, one needn't always have an answer to all of them in a given case. For example, with respect to (Q2), it might be enough to know that the person is an expert in the domain, even if they are less of an expert than other experts. But there might be other cases in which two experts disagree, in which case an answer to (Q2) might be helpful.

My aim in this section is to describe the systemic elements needed to convert strategic expertise contexts into “reliable expert” contexts—contexts in which expert testimony might be the grounds for justified, reliable belief and in which expert testimony might be knowledge-transmitting. These elements, if realizable, would be sufficient to help non-experts arrive at reasonable (although of course, fallible) answers to questions (Q1)-(Q4) in particular cases. My argument so far has been that ordinary contexts are not reliable expert contexts in this way, and that “doing nothing” is not sufficient to result in a reliable expert context of this sort. Now we consider: what else is needed?

5.1 Broad Appearances

Collins and Evans (2007) discuss what they call “external” expertises that we all have and can deploy to help judge whether some individual is an expert or not. These are “external” in that they do not require expertise in the domain. They focus on what they call “ubiquitous discrimination” which is “what we have all been learning since we could speak” and which is “just a particular application of our regular judgments about friends, acquaintances, neighbors, relations, politicians, salespersons, and strangers” (2007, p. 45). They also discuss “local discrimination” which is similar, but relies on more specific local knowledge of people or institutions. Both of these are supposed to rely on wide-ranging knowledge about institutions, social practices, social networks, and common sense judgments. They are somewhat vague about what exactly is included here. I will discuss things like knowledge of credentials and institutions

later. But there are clearly some things that do some work, even falling short of these more sophisticated things. So, for example, you wouldn't go to a doctor who operated behind a dumpster. You wouldn't take investment advice from an unkempt teenager on the bus. We think we can tell when people are trustworthy, shady, smart, and so on, just by interacting with them for a little bit, by noticing their manner, their dress, their apparent social position, their cadence and vocabulary, and so on. In principle, this kind of external, informal, ubiquitous discernment capacity might help us with all of (Q1)-(Q4).

An obvious worry about this kind of discrimination, however, is that it can be highly unreliable and is easy to exploit. Collins and Evans note this, mentioning the "white coat" phenomenon, which leads people wearing white coats to be seen as experts on a wide range of subjects relating to science and medicine, even if they lack any expertise. In general, these broad appearance-tracking mechanisms will work as stereotypes do—they will get some broad range of cases correct, perhaps, but will go woefully astray in other cases, and are easy to take advantage of, since they are commonly known and predictable. Salespeople and high-quality schmoozers in all lines of work exploit them as a matter of professional practice. Furthermore, they often align with pernicious and unreliable forms of bias relating to gender, race, ethnicity, religion, nationality, class, age, disability, and so on.

Goldman (2001) suggests that another, more focused kind of "external" observation—able to be made even by those who are completely "outside" of the domain—might help us with discerning comparative expertise in a way that would help answer (Q2). He notes that, at least in some cases, non-experts are afforded an opportunity to observe the argumentative back and forth between two purported experts. As non-experts, they are not in a position to judge the argument directly, on the merits. But non-experts can attend to what Goldman calls relative "dialectical superiority." This is not supposed to be just about one having greater "debater skill" than the other. Instead, non-experts are to attend to how two experts, E1 and E2, respond to each other. Does E1 seem rattled by the objections E2 is raising? Does E2 always have a ready response to E1, whereas the reverse is not true? Does E2 speak fluidly and clearly in response to challenges and questions? All of these things are supposed to serve not just as broad discriminating features as those identified above, but as proxies or heuristics, suggesting greater relative preparation, examination, and thought about the issue in question. This "indirect argumentative justification" is supposed to function as a kind of inference to the best explanation: the best explanation of why E1 doesn't have anything to say in response or speaks haltingly or seems confused or rattled is that E1 is not that much of an expert or is less of an expert than E2.

As before, the main concern is obvious: appearances can be misleading. It is entirely possible to train or coach individuals so that they will seem more polished, prepared, and so on, without actually being in any non-superficial way more of an expert than they were before. And indeed, this is exactly what is done when people are brought to testify in court or before political bodies. This is what is done to prepare salespeople to do their job. This is perhaps the main skill that political candidates and political representatives have to possess and improve upon in order to succeed in electoral politics. This is how people are prepared to be in television commercials or other forms of media. Additionally, since, by hypothesis, non-experts are not attending to the expert content detail of the argumentative back and forth, there is virtually no limit to the effectiveness of the coaching. A good coach would train her expert to always *make* a confident response to every potential objection, even if the expert *had* no response whatsoever at the level of content. That is perhaps why we end up with movie stars as politicians more often than might be expected.

Another concern is that non-experts may not always have an opportunity to observe experts engaged in debate in this way.

Ultimately, without more, strategies that focus on appearances in this way will serve us poorly in addressing any of the (Q1)-(Q4). Perhaps if there were no background strategic interests, so that we didn't have to worry about experts (or their handlers) exploiting these tendencies of ours, we might be in a better situation. But we know that this gamesmanship takes place—indeed, these are precisely contexts in which we should expect it.

5.2 Disclosures of Interest and Bias

Most writing about the problem of relying on experts or expert testimony, including Goldman, Collins and Evans, Anderson, and Walton, note that it is important to have expert testifiers disclose potential or actual conflicts of interest or possible sources of bias. The hope is some combination of (a) discouraging allowing experts to testify if they are hopelessly biased and (b) ensuring that all hearing the testimony can discount it appropriately given the existence of a possible bias or conflict of interest. (Sometimes learning about bias or sources of funding is really a way of learning about the invalidity of an individual's credentials or claims to expertise—they can only get funding from Institution with Ideological Agenda or Industry Attempting to Evade Regulation—something I will discuss in the section on credentials below.)

It is plausible that before relying on testimony in any context, and in particular in a strategic expertise context, it is important to be aware of possible biases or conflicts of interest on the part of the speaker. Of course, the strategic expertise contexts are defined as ones in which the expert will have some interest that is out of alignment with the interests of the hearer. So it won't be any kind of news. But, it still seems like a good minimal practice to aid in improving the reliability of expert testimony that actually ends up being relied upon. If one knows that the Testifier is being paid by X to say Y, that the *only* reason the Testifier is saying Y is because they are being paid to do so, and one knows that X would want us all to accept Y regardless of the truth of Y, then one should entirely discount the Testifier's testimony as a reason to accept Y. And one certainly shouldn't accept Y on the *basis* of the Testifier's testimony.

Unfortunately, most cases are considerably muddier than this. Here is an initial problem: almost all experts will have biases and interests to disclose, if we expand our focus beyond the usual and unduly limited assessment of sources of financial funding and compensation. People have interests in making a name for themselves, in obtaining and keeping power and prestige, in having their views more widely known, and sometimes just in getting you to believe what they want you to believe (and not because they believe it, but because it is useful for them if you do). These other biases and interests will not always raise epistemic concerns—perhaps the best way to make a name for oneself is to be the best and most reliable source of knowledge in some domain. The concern is that these biases and interests may in some cases diverge from purely good epistemic practices.

Still, perhaps it is useful to know that E1 has *even more* interest in possibly misleading me than E2 does. Perhaps. But I think the helpfulness can be overstated, particularly if we are focused less on the comparative question of E1 v. E2, and more on the question of whether reliance on the testimony of either E1 or E2 is justified or knowledge-transmitting.

Consider, for example, the great many cases in which we all know that the expert—and indeed all of the experts like the expert—has an interest in getting people to believe what he or she says. This will often be the case. Anyone trying to sell us something, when they know more about what they are selling than we do (as will almost always be the case). Lawyers or doctors or dentists or mechanics or consultants trying to get us to pay for their services. Politicians trying to get us to vote for them (where they are experts perhaps about policy but certainly about what they want to do with political power) or to accept their explanation for why some good or bad thing happened. It will do us little good to learn that they have an interest in the matter. We already know this. So, at most, this seems like a necessary element of a reliable expert context. But it is clearly far from sufficient. (Not that anyone would have thought it would be, since it really engages at most (Q2)-(Q4)—it is no help at all with (Q1), the question of whether a particular person is an expert.)

5.3 Track Records

Another thing that might help non-experts answer questions (Q1)-(Q4) is the *track record* of the expert under consideration. Collins and Evans, Goldman, and Anderson all mention track record as a possible aid.

One initial difficulty is that the track record itself can be hard for non-experts to assess. After all, if we are not experts, we will be hard-placed to evaluate the merits of various results and accomplishments obtained in the past by the expert we are considering. Goldman makes the nice point that in some cases what might have at one point been *esoteric*, expert knowledge can become *exoteric*,

common knowledge or easily obtainable knowledge. So, perhaps in 1998, the expert predicted, on the basis of some elaborate astronomical theory, that there will be an unusual lunar eclipse on April 30, 2018. Prior to April 30, 2018, it might be hard for a non-expert to evaluate the model and calculations involved. But after April 30, 2018, presumably anyone could evaluate whether the unusual eclipse actually transpired. So, perhaps that will help in some cases. And in other cases the record will be relatively easily observable. You have 10 friends all of whom have used the same mechanic for years, and all of whom always found both that the mechanic gave them the best price of any mechanic they consulted, and that the problem always was resolved after one visit to the mechanic and never recurred. Even as a non-expert, that should give you some confidence that the person is an expert, that she gives sincere recommendations, and that in a particular case she is likely to have the correct answer.

Let us call this kind of track record an *exoteric record of expert success*. This kind of track record is in principle observable, understandable, and knowable by non-experts, and is such that it provides evidence that a particular individual is actually an expert in a particular domain.

Another kind of track record that we might care about is not about the expert's performance, *qua* expert, but about her *record of epistemic integrity*. This will be important for helping provide reasonably well supported answers to (Q3) and (Q4) in particular. Anderson (2011) offers some related ideas in her discussion of "criteria for judging honesty" and "criteria for judging epistemic responsibility." Things we might be interested include the following:

- evidence of previous expert-related dishonesty (e.g. plagiarism, faking data)
- evidence of a record of misleading statements (e.g. cherry-picking data, quotations out of context)
- evidence of a record of misrepresenting views of expert opponents
- evidence of evasion of peer-review or refusal to allow other experts to assess work
- evidence of refusal to disclose data, methodology, or detailed results
- evidence of refusal to disclose results contrary to the expert's own views
- evidence of "dialogic irrationality": repeating claims after they have been publicly refuted, without responding to the refutations
- evidence of a record of "over-claiming" of expertise: claiming expertise beyond the expert's domain of expertise
- evidence of a record of "lending" one's expertise to support other individuals or institutions that themselves lack epistemic integrity in some of the above ways
- evidence of being an "opinion for hire"—offering expert testimony for pay, perhaps particularly if that testimony conflicts with other things the expert has said

These are all ways in which an expert might lack "epistemic integrity." Of course, they won't all be relevant for all of the different kinds of experts we might encounter. There is nothing quite like "peer reviewed journals" for the work of mechanics and politicians. That said, there might be similar institutions in those domains that aim to assess the work and veracity of even these kinds of experts. As with records of expert success, there will be both esoteric and exoteric records of epistemic integrity. Some of these records will be difficult or impossible for non-experts to access or understand; others will be relatively straightforward to understand, if they are encountered.

It seems that in cases in which they are *available*, both exoteric records of expert success and records of epistemic integrity will be of significant value in creating a "reliable expert" context. That is, when they are available, and when they strongly support a particular expert as both expert and of high epistemic integrity, testimony from that expert can be the basis for justified belief and can transmit knowledge. Unfortunately, one ideally needs evidence of both kinds of records—one or the other usually won't be sufficient, since one might reasonably believe the person is an expert but have no evidence of trustworthiness or sincerity in particular cases, or vice versa.

The more general problem with track records is related to this difficulty: they either won't exist at all; they won't exist in an easily, publicly available form; or they will be difficult for the general public to understand or put in appropriate context. Unlike with credentials—which we will discuss next—track records do not come easily attached to particular experts. It is true that those with some training and motivation will, in some cases, be able to uncover relevant evidence. But this will often be difficult outside of very high profile cases of explicit and detected academic fraud. For some kinds of decisions about trusting experts, such as hiring a local doctor or mechanic, one can use local knowledge to help develop and uncover a useful track record for that particular person. But that will be much harder in many cases, at least if one is just trying to do it oneself. Anderson (2011) suggests that the internet and sources like Wikipedia can help non-experts in this regard. But I am skeptical of this as a general resource for obtaining information about particular experts and their exoteric records of expert success or epistemic integrity. The information currently available for many experts or purported experts is very patchy, very unreliable (often the result of self-interested efforts at image management), and too esoteric for most non-experts. And of course there will be some topics or some kinds of expertise for which a track record of success is simply not available—either because the field is too young, or because there is not an adequately broadly accepted view about what counts as “success” in the field.

5.4 Experts Assessing Experts: Credentials, Exams, Licenses, Institutions, Peer-Reviewed Publications, Rankings, Awards, Expert Reviews

A natural complement or supplement for the patchy, possibly unreliable, and irregular track records that might be available to non-experts are the relatively more systematic ways in which experts evaluate and mark other experts. There are many different ways in which experts assess each other or allow or restrict entry into the community of experts. Universities, professional schools, and other educational institutions involve experts—professors and various academic instructors—assessing the work of future experts or experts-in-training. (These institutions themselves are subject to regular expert review, ranking, and accreditation.) Work of those in training must be above a certain standard to earn various grades and credentials. In some cases, exams (written by experts, to test for expertise) have to be passed. For career progression in many expert careers, one has to receive favorable evaluations by recognized experts—in the form of various “letters of recommendation,” financial grant support for research, and/or satisfactory review of one’s work through various forms of peer-review. For many professions, one has to obtain a license from a professional association or governmental body (or both) before one is allowed to work as an expert. In some fields, experts regularly review each other’s work in public settings (think newspaper book reviews) and contribute to published reputational ratings and other forms of rankings.

These mechanisms do much of the heavy-lifting in sorting experts from non-experts in modern societies. There will be some kinds of expertise—particularly those based on non-professional life experience—for which these formal credentials and markers are not the main evidence of at least some threshold level of expertise. But for most kinds of expertise, credentials of some form are at least a minimum qualification. So, are these enough to help us reasonably rely on experts and answer our central questions (Q1)-(Q4)?

Unfortunately, no. Again, as with disclosures of interest, knowledge of credentials, licenses, institutional affiliations, and so on might be necessary to help us answer (Q1). Is this person an expert at all? But this knowledge won't get us much further—if it even gets us that far.

For one thing, the details of the credentials, institutions, degrees, rankings, reputation ratings, and so on all matter a lot for assessing both whether the person really is an expert, and questions of comparative expertise when looking at more than one expert, but non-experts are often in a bad position to assess those details. Our ignorance regarding these matters is what generates the problem in the first place. This makes it possible to make a person look like a well-credentialed expert without them being an expert. Institutions can be created that purport to provide expert credentials but are in fact shams. Professions have reputation ratings that confer “excellent” ratings on everyone, nearly everyone, or everyone willing to pay to be included. Or a person might have an expert credential in X but suggest that it provides evidence of expertise in Y, with X and Y being difficult for non-experts to distinguish. In this

way, the evaluations made by experts (or even whether the evaluations really are made by experts) in this category can amount to esoteric knowledge effectively masquerading, troublingly, as exoteric knowledge.

A second difficulty is that some of these entities are overly promiscuous in bestowing credentials, overly reluctant to revoke or challenge claims to expert status, or are even essentially institutions generating pseudo-expertise or helping to support the credentials of pseudo-experts. For example, educational institutions have powerful financial incentives to enroll and graduate enough students, particularly in the world of for-profit educational programs. Professional organizations often self-regulate and set standards for membership (that was the idea of having experts judging experts!) and for predictable reasons have a hard time disciplining or expelling those who fall short of professional standards. Enough financial or political support can create credentialing institutions even when they have no or only inadequate expertise attached to them. And various commercial entities and industries have even made the effort to create institutions and “peer-reviewed” journals so as to bolster the credentials of those pushing the line they prefer.¹³

Another problem with this kind of evidence—judgments about experts by experts—is that there are some cases in which, for strategic competition reasons, we are unlikely to get unbiased rankings or ratings. Think, for example, of the difficulty of getting unbiased ratings of politicians—either in terms of the quality of policy they help bring about, or in terms of the extent to which they try to do what they said they would do. There are various entities that “fact check” assertions of politicians, and those are helpful, certainly.¹⁴ But so much happens behind the scenes that makes it difficult to make more detailed judgments about those actions and statements that occur outside of easy public view. And it would take far more time and effort than most individuals have or are willing to expend to keep apprised of all that they might need to know.

So, even for addressing (Q1), these mechanisms have significant limitations. That said, they are at least sometimes important and reliable, and there are institutions that have better and worse track records themselves in terms of both expert success and epistemic integrity. With respect to (Q2), then, it may be possible to form reliable judgments about answers to this kind of comparative question if one is dealing with a case in which one of the two experts has obviously better credentials (more of them, from a broader set of institutions, which themselves have a better track record).

A related possibility, which is the subject of much discussion in the law, is the extent to which an expert is aligned with the “expert consensus” in his or her field.¹⁵ We might want to know this either in general or with respect to the particular testimony we are considering. As with these other mechanisms, knowing that the person is broadly supported in their view by other experts is a way of strengthening our foundation for thinking that the expert is reliable in general. And knowing that the person’s testimony is in line with the expert consensus on a topic also improves our epistemic position with respect to our justification for accepting the testimony and seeing it as knowledge transmitting. On the other hand, as Goldman (2001) points out, there are worries with looking to expert consensus. One worry is that it is a situation in which there is a guru with slavish followers. Additional supporters only lend support to our epistemic position if those supporters have some credibility on the topic and are at least partially conditionally independent of each other. As Goldman puts it, they must not be mere non-discriminating “reflectors” of each other’s view, or of some third-party’s view.

So, in some cases, credentials and other forms of expert judgments about experts can be helpful. Unfortunately, they help very little with the issues raised by (Q3) regarding trustworthiness and sincerity and epistemic integrity. Of course, academic institutions and professional organizations all attempt to train individuals how to behave ethically and with integrity, and they require members to attest to their

¹³ The volume *Agnotology: The Making and Unmaking of Ignorance*, edited by Proctor and Schiebinger (2008), offers a host of real world examples of these efforts.

¹⁴ Both FactCheck.org (www.factcheck.org) and the Pulitzer-prize winning PolitiFact (www.politifact.com) are particularly successful entities in this regard.

¹⁵ For extensive discussion of the standards of admissibility of expert testimony in law, and the role played by alignment with expert consensus, see Haack (2008) and Brewer (1998).

work being faultless in this regard, but it would be naïve to think this ensures the epistemic integrity of the members of these communities or the holders of these credentials. For one thing, only the most flagrant violations are sanctioned. And, as noted above, these are often far from exclusive communities with any kind of centralized oversight. Finally, the financial interests are often considerable—that is often why people pursue the degrees, licenses, ratings, and other credentials—and those interests set up potential conflicts that raise distinct problems. This is actually an additional reason to take “expert consensus” seriously, at least if it is not the case that all experts in the field have the same incentives: this kind of consensus suggests that the testimony is not unduly motivated simply by financial or other interests and incentives. It is worth repeating: this only helps if not all experts in the field have the same incentives in this regard.

5.5 Incentive Alignment Strategies

In some contexts, it might be possible to improve the alignment of the interests of experts and the interests of non-experts. If there are ways for the non-experts to gauge whether experts have successfully answered some question, made some prediction, or solved some problem—if these things are or can be made exoteric in the right way—it can be possible for non-experts to offer rewards or impose penalties or for experts to make credible commitments of one form or another. This happens, in an obvious way, in various marketplaces for experts: a lawyer working on a contingency fee, investors who are required to disclose their investment success rates, and various forms of legal or financial liability for experts who fall below some standard of reasonable expert conduct. Think of medical malpractice, or other forms of tort liability, for example. Or think of the story of the Chinese airline officials—charged with overseeing efforts to ensure Y2K compliance of China Southern Airlines—being in the air on an airplane as the year 2000 began.¹⁶ They would be very interested in making sure that the computer systems actually would operate without flaw as the New Year turned.

Of course, these forms of alignment and reward and explicit bonding won’t always work. There needs to be some relatively exoteric, observable fact that can be the basis for reward or punishment. And there are dangers of setting the benchmark in the wrong place. For example, consider the way in which attaching executive compensation to various financial targets and stock prices can lead executives to engage in fraud and other forms of misrepresentation in order to boost stock prices or meet various financial goals.¹⁷ Or consider the difficulty in ascertaining whether some physician conduct fell below the standard of care so as to constitute negligence. It will often require expert judgment to make these assessments, and then we must again confront the problems of experts regulating themselves.

One nice feature of these strategies, when they can be devised effectively, is that they can help both with ensuring *expert* performance and ensuring *reliable* performance, since the dimension of trust drops out if the target is placed appropriately and is sufficiently exoteric and observable. Of course, for that same reason, these do not so much *solve* the problem of relying on expert testimony as *obviate* the need for a solution to that problem.

Let me summarize the basic situation we are in as ignorant, non-experts trying to think about how we might reasonably rely on expert testimony. We often use broad appearances, but these are unreliable and unlikely to be helpful in any detailed way with answering any of the four core questions: (Q1)-(Q4). Disclosure of biases and interests may be necessary, but will be far from sufficient, particularly in cases in which all experts have biases and interests and we know this. Track records—particularly exoteric records of expert success and records of epistemic integrity—can also be helpful, but they are not easily available, if available at all. Finally, expert judgments about experts can also be helpful, but are often more esoteric than exoteric, can be significantly misleading in cases that we might call “false credentials,” and are

¹⁶ See <http://www.propertyandcasualty.com/doc/airline-execs-to-fly-at-midnight-january-1-to-0001>.

¹⁷ For discussion, see D. Guerrero (2004).

subject to a host of biases and other problems stemming from the fact that these are communities that have their own internal reasons for producing the evaluations that they do. All of this suggests that I think we are often not in “reliable expert” contexts—contexts in which expert testimony might be the grounds for justified, reliable belief and in which expert testimony might be knowledge-transmitting.

What should we think about this? Here is one place where institutional reform might greatly improve our epistemic situation. Consider this sketch of an idea: have all academic experts, all licensed experts in various fields (law, medicine, financial advising, automotive repair, etc.), all experts who wish to testify before legal or political bodies, and so on, register with a “Database of Experts” as part of membership in what we could call the “Society of Experts.” There could be different sections in the society for different kinds of experts. And decisions could be made about which kinds of credentials/institutions were such that they could identify a person as an expert. This could be more inclusively—allowing almost anyone who met some minimal criteria and wanted to be listed to be listed—or exclusively—having some group of experts (expertise experts, perhaps) determine the bar to inclusion in the database for different categories of expertise. On the former, inclusive route, the idea would be that mere inclusion wouldn’t be of much significance; what would matter is what then goes in your ‘file’—as discussed below. On the latter route, being included itself would be some reason to think that a person was an expert, but there would then be hard questions (as discussed above) about exactly who should be involved in determining what credentials, experience, and so on are required to be listed in the database. One possibility would be to allow there to be different databases, used for different purposes, perhaps, or to let the market sort out which of these approaches was most useful for different purposes.

This database would include all relevant credential information, along with work experience, but it would also have as full a history as possible of both esoteric and exoteric records of expert success and records of evidence relevant to the individual’s epistemic integrity. Entries in this database could be submitted by the experts themselves, users, other experts, and non-experts, but would have to be themselves screened for veracity and could be marked as “contested” or “attested to” by various people (perhaps with the names of the persons contesting or attesting attached). Existing institutions and bodies could submit their own contributions regarding ratings, professional credentials, and so on, so that the database would not replace these things if they already existed, but put them in broader context along with other sources of information about particular experts. For the system to work, it would have to itself be run and maintained by some neutral entity—in something of the way that Consumer Reports or the Better Business Bureau work for products and businesses, or in the way that Wikipedia works with respect to some entries.

There are non-trivial issues about maintaining the integrity of the information in the database and finding funds to support the enterprise that do not lead to its corruption. This has at times been a concern raised about, for example, Wikipedia and the Better Business Bureau network, and other related kinds of entities: everything from Yelp to Angie’s List to Rate My Professors.¹⁸ There are also issues about confidentiality that might arise concerning previous instances in which an expert worked with someone or testified in a particular case. But it seems there would be ways in which steps could be taken to keep this information appropriately confidential. Even if it were imperfect in various ways, this kind of institution would help generate far more “reliable expert” contexts, and would help to highlight potentially problematic experts or purported experts. And refusal of an expert to be included or a decision by an expert to contest all the entries about him- or herself would also be epistemically useful: suggesting that significant skepticism about that expert’s testimony is probably warranted.

This database would be useful for non-experts in ordinary life, but also for actors within legal and political systems attempting to discern whether an individual can be relied upon for information or even appointment to important technocratic or bureaucratic posts. It would also be useful for journalists and news media attempting to decide which experts to use as expert sources.

¹⁸ See, for example, Mokey (2010), Fleming (2011), Pinsker (2015).

Indeed, since having the database itself doesn't introduce a mechanism by which to ensure that the potential expert Testifier's purposes are to communicate the truth, one possibility would be to use the database along with *random selection* of experts from that database who satisfied various desiderata in terms of domain of expertise, credentials, track record, and so on. This would help to prevent "cherry picking" of experts for purposes of legal or political testimony, or for use as experts in media stories. It would also allow for assessment of whether a particular expert had outlier views or whether there was something close to a consensus within some domain of expertise. For example, one could choose 100 experts at random and see what the balance of opinion was on some question.

This Database of Experts wouldn't be perfect, particularly in contexts in which all experts in a domain might have similar biases or epistemically unproductive purposes, but it might put many of us non-experts in a relatively better place when we are faced with the question of whether or not to rely on expert testimony. Perhaps this solution isn't the best one, or will only work when supplemented with other mechanisms. Still, I think that Hardwig (1985) and Goldman (1999) and others who have railed against the limits of individualist epistemology are definitely on to something. In the modern world, ignorant as we all are, we have to rely on the testimony of others, and we have to rely on the testimony of experts. To do this in a way that is reliable and epistemically justifiable, we will need help. And that help, I contend, will be institutional and must be socially built and supported.

Bibliography

- Adler, Jonathan. 1990. "Conservatism and Tacit Confirmation." *Mind* 99: 559–570.
 —1994. "Testimony, Trust, Knowing." *Journal of Philosophy* XCI: 264–275.
 —1996. "Transmitting Knowledge." *Noûs* 30: 99–111.
 —1997. "Lying, Deceiving, or Falsely Implicating." *Journal of Philosophy* XCIV: 1–18.
 —2002. *Belief's Own Ethics*. Cambridge: The MIT Press.
 —2009. "Another Argument for the Knowledge Norm." *Analysis* 69: 407–412.
 Anderson, Elizabeth. 2011. "Democracy, Public Policy, and Lay Assessments of Scientific Testimony." *Episteme* 8.2: 144–164.
 Audi, Robert. 1997. "The Place of Testimony in the Fabric of Justification and Knowledge." *American Philosophical Quarterly* 34: 405–422.
 —2006. "Testimony, Credulity, and Veracity." In Lackey, J. and E. Sosa (eds.), *The Epistemology of Testimony*. Oxford: Oxford University Press. 25–48.
 Brewer, Scott. 1998. "Scientific Expert Testimony and Intellectual Due Process." *Yale Law Journal* 107: 1535–1681.
 Burge, Tyler. 1993. "Content Preservation." *The Philosophical Review* 102: 457–488.
 Coady, C.A.J. 1973. "Testimony and Observation." *American Philosophical Quarterly* 10:149-155.
 Collins, Harry and Robert Evans. 2007. *Rethinking Expertise*. Chicago: University of Chicago Press
 Collins, Harry and Martin Weinel. 2011. "Transmuted Expertise: How Technical Non-Experts Can Assess Experts and Expertise." *Argumentation* 25(3): 401-413.
 Douglas, Heather. 2005. "Inserting the public into science." In Sabine Maasen and Peter Weingart eds., *Democratization of Expertise?* Netherlands: Springer, 153-169.
 Dretske, Fred. 1982. "A Cognitive Cul-de-sac." *Mind* 91: 109–111.
 Evans, Gareth. 1982. *The Varieties of Reference*. John McDowell (ed.), Oxford: Oxford University Press.
 Fleming, Troy. 2011. "'Pay for Play' Scandal at the Better Business Bureau Leads to Consumer Mistrust of the Business Rating Organization." *Loyola Consumer Law Review* 23: 445-458.
 Fricker, Elizabeth. 1987. "The Epistemology of Testimony." *Proceedings of the Aristotelian Society Supplementary* 61: 57–83.
 —1994. "Against Gullibility", in Matilal and Chakrabarti (eds.) *Knowing from Words: Western and Indian Philosophical Analysis of Understanding and Testimony*: 125–161.
 —2002. "Trusting others in the sciences: a priori or empirical warrant?" *Studies in History and Philosophy of Science Part A* 33.2: 373-383.

- 2006. “Testimony and Epistemic Autonomy.” In Lackey, J. and E. Sosa (eds.), 2006, *The Epistemology of Testimony*. Oxford: Oxford University Press.
- Gelfert, Axel. 2006. “Kant on Testimony.” *British Journal for the History of Philosophy* 14: 627–642.
- Goldberg, Sanford C. 2010. *Relying on others: An essay in epistemology*. OUP Oxford.
- Goldman, Alvin. 1999. *Knowledge in a Social World*. Oxford: Oxford University Press.
- 2001. “Experts: Which Ones Should You Trust?” *Philosophy and Phenomenological Research* 63(1): 85-110.
- Graham, Peter J. 2004. “Metaphysical Libertarianism and the Epistemology of Testimony.” *American Philosophical Quarterly* 41: 37–50.
- Grice, H. P. 1975. “Logic and conversation.” In P. Cole and J. P. Morgan (eds), *Syntax and Semantics, Vol. 3: Speech Acts*. New York: Seminar Press.
- Guerrero, Alexander. 2007. “Don’t Know, Don’t Kill: Moral Ignorance, Culpability, and Caution.” *Philosophical Studies* 136: 59-97.
- 2014. “Against Elections: The Lottocratic Alternative.” *Philosophy and Public Affairs* 42: 135-178.
- Guerrero, Douglas. 2004. “The Root of Corporate Evil.” *The Internal Auditor* 61(6): 37-40.
- Haack, Susan. 2008. “Of Truth, In Science and Law.” *Brooklyn Law Review* 73: 563-586.
- Hardwig, John. 1985. “Epistemic Dependence.” *Journal of Philosophy* 82: 335-349.
- Hintikka, Jaakko. 1962. *Knowledge and Belief*. Ithaca: Cornell University Press.
- Hume, David. 1739. *A Treatise of Human Nature*. 1888 edition, L.A. Selby-Bigge, ed., Oxford: Clarendon Press.
- 1748. *An Enquiry Concerning Human Understanding*. 1977 edition, Hackett Publishing Company.
- Kitcher, Phillip. 1993. *The Advancement of Science*. Oxford: Oxford University Press.
- 2001. *Science, Truth, and Democracy*. Oxford: Oxford University Press.
- Lackey, J. and E. Sosa (eds.). 2006. *The Epistemology of Testimony*. Oxford: Oxford University Press.
- Lewis, David. 1969. *Convention*. Cambridge: Harvard University Press.
- 1975/1983. “Languages and Language.” *Philosophical Papers Vol. I*. Oxford: Oxford University Press, 163–188.
- McDowell, John. 1980. “Meaning, Communication, and Knowledge”, in *Philosophical Subjects*, Z. van Straaten (ed.), Oxford: Oxford University Press, 117–139.
- 1994. “Knowledge by hearsay”, in *Meaning, Knowledge and Reality*. Cambridge: Harvard University Press, 195–224.
- Michaels, David. 2008. “‘Manufactured Uncertainty,’ Contested Science and the Protection of the Public’s Health and Environment,” in *Agnotology: The Making and Unmaking of Ignorance* (Ed. by Robert Proctor & Londa Schiebinger).
- Mokey, Nick. 2010. “Yelp Faces Extortion Claims in Cash-For-Ratings Scandal.” *Digital Trends* available at <http://www.digitaltrends.com/computing/yelp-faces-extortion-claims-in-cash-for-ratings-scandal/>
- Peels, Rik. 2010. “What is Ignorance?” *Philosophia* 38(1): 57-67.
- Pinsker, Joe. 2015. “The Covert World of People Trying to Edit Wikipedia—For Pay,” *The Atlantic*, available at <http://www.theatlantic.com/business/archive/2015/08/wikipedia-editors-for-pay/393926/>
- Price, H. H. 1969. *Belief*. New York: Humanities Press.
- Pritchard, Duncan. 2007. “Anti-Luck Epistemology.” *Synthese* 158: 277-98.
- Proctor, Robert & Londa Schiebinger, (eds). 2008. *Agnotology: The Making and Unmaking of Ignorance*. Stanford: Stanford University Press.
- Reid, Thomas. 1764 (1983). *Inquiry and Essays*. R. Beanblossom and K. Lehrer (eds.), Indianapolis: Hackett.
- Reynolds, S. 2002. “Testimony, Knowledge, and Epistemic Goals.” *Philosophical Studies* 110: 131–161
- Shogenji, T. 2006. “A Defense of Reductionism about Testimonial Justification of Beliefs”, *Noûs* 40: 331–346
- Somin, Ilya. 2013. *Democracy and Political Ignorance*, Stanford: Stanford University Press.

- Sosa, Ernest. 1994. "Testimony and Coherence", in Matilal and Chakrabarti (eds.) *Knowing from Words: Western and Indian Philosophical Analysis of Understanding and Testimony*: 59–67.
- 1999. "How to Defeat Opposition to Moore." *Philosophical Perspectives* 13: 141-54.
- Sperber, Dan. 2001. "An Evolutionary Perspective on Testimony and Argumentation", *Philosophical Topics* 29: 401–413.
- Sperber, Dan, Fabrice Clement, Chrisophe Heintz, Olivier Mascaro, Hugo Mercier, Gloria Origgi, and Deidre Wilson. 2010. "Epistemic Vigilance." *Mind and Language* 25(4): 359-393.
- Stanley, Jason. 2015. *How Propaganda Works*, Princeton: Princeton University Press.
- Sutton, J.. 2007. *Without Justification*, Cambridge: The MIT Press.
- Unger, Peter. 1975. *Ignorance*, Oxford: Oxford University Press.
- Walton, Douglas. 1997. *Appeal to Expert Opinion*, University Park: Pennsylvania State Press.
- 2004. *Relevance in Argumentation*, Mahwah, N. J.: Erlbaum
- Whyte, Kyle Powys & Robert Crease. 2010. "Trust, Expertise and the Philosophy of Science," *Synthese* 177 (3): 411-425.
- Williamson, Timothy. 1996. "Knowing and Asserting." *The Philosophical Review* 105: 489–523.
- 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.
- Wolterstorff, Nicholas. 2001. *Thomas Reid and the Story of Epistemology*. Cambridge: Cambridge University Press.