

11

Experts and Peer Disagreement

Jennifer Lackey

11.1 Introduction

It is often argued that widespread disagreement among epistemic peers in a domain threatens expertise in that domain. In this chapter, I will sketch two different conceptions of expertise: what I call the *expert-as-authority* and the *expert-as-advisor* models. While it is standard for philosophers to understand expertise as authoritative, such an approach renders the problem posed by widespread peer disagreement intractable. I will argue, however, that there are independent reasons to reject both this model of expertise and the central argument offered on its behalf. I will then develop an alternative approach—one that understands expertise in terms of advice—that not only avoids the problems afflicting the expert-as-authority model, but also has the resources for a much more satisfying response to the problem of widespread peer disagreement.

11.2 Peer Disagreement and the Threat to Experts

There are two kinds of conclusions that are said to follow about expertise from disagreement. The first is a *metaphysical result*: widespread disagreement among epistemic peers in a domain, *D*, *threatens the existence of experts* in *D*. This is because it is often accepted that expertise requires knowledge. For instance, according to Alvin Goldman, “an expert . . . 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” (Goldman 2001, p. 92). But peer disagreement in *D* threatens knowledge in *D*. In particular, peer disagreement is said to provide experts with counterevidence that defeats their knowledge in *D*.

Sarah McGrath (2008) formulates the following argument of this sort specifically with respect to moral beliefs, where a belief is **CONTROVERSIAL** “if and only if it is denied by another person of whom it is true that: you have no more reason to think that he or she is in error than you are” (McGrath 2008, p. 91):

- P1. Our controversial moral beliefs are **CONTROVERSIAL**.

P2. CONTROVERSIAL beliefs do not amount to knowledge.

C. Therefore, our controversial moral beliefs do not amount to knowledge.¹

A general version of McGrath's argument can be fleshed out as follows: my epistemic peer relative to a question is someone whom I regard as equally likely to be mistaken about this question.² If you and I are experts and we regard one another as equally likely to be mistaken about a matter, then it seems that we ought to withhold belief when we disagree about this matter. This is because I have no reason to prefer my own belief to yours if I think we're equally likely to be mistaken about it. And if I have no good reason to prefer my own belief to yours, then it is irrational for me to steadfastly cling to it in the face of our disagreement. Since such irrationality is incompatible with knowledge, widespread peer disagreement in a domain undermines expertise. Moreover, even for those who endorse weaker requirements for expertise, such as justified belief, it is plausible that widespread peer disagreement threatens such states, for irrationality is similarly incompatible with them. The upshot of these considerations, then, is a fairly robust form of skepticism: there are no experts in areas rife with peer disagreement, such as religion, politics, and philosophy.

By way of avoiding this metaphysical result, there are at least two responses, the first being to distinguish between *primary* and *secondary* questions with respect to expertise. According to Goldman:

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. (Goldman 2001, p. 92)

Even if there is widespread disagreement over primary questions in a given domain, there is often lots of agreement regarding secondary matters. For instance, while there is no consensus whatsoever about the morality of abortion, most accept that the central relevant issues are the rights of the mother, whether fetuses are persons, the extent of our moral obligations to dependents, and so on. It might be argued, then, that it is

¹ Similarly, Patrick Hurley writes:

There are some areas in which practically no one can be considered an authority. Such areas include politics, morality, and religion. For example, if someone were to argue that abortion is immoral because a certain philosopher or religious leader has said so, the argument would be weak regardless of the authority's qualifications. Many questions in these areas are so hotly contested that there is no conventional wisdom an authority can depend on. (Hurley 2008, p. 133)

² See Elga (2007) for this conception of being an epistemic peer. See Kelly (2005) for an alternative view.

knowledge of primary questions that is threatened by widespread peer disagreement, and thus expertise, when understood as involving knowledge of only secondary questions, remains.

Moreover, with respect to primary questions, it is not clear that expertise requires knowledge. So here is where the second response to the metaphysical result comes in: it might be sufficient for being an expert that one reliably offer true testimony, regardless of what one knows oneself. A very modest doctor, for instance, might have doubts that prevent her from believing and, therewith, from knowing the diagnoses of her patients, but she might nonetheless be very reliable when she testifies about them. Or a Creationist teacher's religious faith might prevent her from believing and, thus, knowing that *Homo sapiens* evolved from *Homo erectus*, but she might reliably convey this fact to her students. Perhaps, then, we can save the existence of experts in controversial domains by loosening the requirements on expertise to involve only the offering of reliable testimony of primary questions. Indeed, this alternative conception to the one proposed by Goldman already has support in the literature. According to Michael Cholbi, for instance, "Within the extant philosophical literature, there is general agreement that a moral expert is someone who very reliably, though not necessarily infallibly, provides correct moral advice in response to moral situations and quandaries" (Cholbi 2007, p. 324). If this is the case, then it is certainly possible that experts persist in the face of widespread peer disagreement since one side of the debate might actually be right, and thus be reliably providing correct moral testimony.

There is, however, an *epistemological result* that is also relevant here: widespread disagreement among epistemic peers in *D* threatens our ability to identify experts in *D*. McGrath makes this point when she writes, "in general, identifying those with genuine expertise in some domain will be most straightforward when we have some kind of *independent check*, one not itself subject to significant controversy, by which we can tell who is (and who is not) getting things right. . . . But significantly, we possess no . . . independent check for moral expertise" (McGrath 2008, p. 97). To put this point generally, in domains where there is widespread peer disagreement, there is also often the same degree of disagreement about not only a host of surrounding issues, but also about who the experts are. When this is the case, there aren't resources to appeal to that are not themselves in dispute to help us pick out experts in a given domain. This has the consequence that even if experts exist, we cannot rely on them in controversial areas, such as religion, politics, and philosophy, because we don't know who they are. In terms of our ability to benefit from expertise, then, this epistemological result leaves us no better off than the metaphysical one.

11.3 The Expert-as-Authority Model

In this section, I want to examine the epistemological result from widespread peer disagreement in relation to what I call the *expert-as-authority* model. A helpful place

to begin understanding this concept is with the notion of authority itself. Perhaps most influential here is the preemptive view of authority developed and defended by Joseph Raz in the political domain and recently adapted for broadly epistemological use by Linda Zagzebski.³ Since I will be interested primarily in the epistemological issues surrounding authority, I will follow Zagzebski's work most closely here, with occasional references to Raz and others.

According to Zagzebski, beliefs formed on the basis of the testimony of an authority can be epistemically rational, where "what is essential to authority is that it is a normative power that generates reasons for others to do or to believe something preemptively" (Zagzebski 2012, p. 102). Modeling her conception of authority on Raz's, Zagzebski holds that a preemptive reason is "a reason that replaces other reasons the subject has" (Zagzebski 2012, p. 102). What this means is that a subject should not treat the testimony of an authority as evidence to be weighed against, or aggregated with, other relevant evidence that she might have. Rather, she should let the authority "stand in for [her] in [her] attempt to get the truth in that domain, and to adopt his belief" without deliberation (Zagzebski 2012, p. 105).

Now when I replace my car or winter coat, I toss out the old one to make room for the new one. So Zagzebski's talk of a subject replacing her reasons with those provided by an authority invites us to think of preemption along these lines: a subject tosses out all of her old reasons on the topic when she takes on those provided by the authority. But this is a mistake, Zagzebski tells us. Letting an authority's testimony replace one's other reasons does not eliminate the latter from one's doxastic framework. They are there; they are just crucially not available for being weighed or aggregated with those provided by the authority. In this way, the authority's testimony normatively screens off one's other relevant reasons or evidence. To illustrate this, Zagzebski provides the example of stopping at a red light because this is what the law requires. When one takes the law as an authority on this matter, one does not weigh stopping at a red light with one's other relevant reasons on the topic, such as one's preference to not stop because one is in a hurry or one's concern about getting a costly ticket. Instead, the authority of the law preempts all of these other reasons. This does not mean, however, that such considerations cease to exist or that one is no longer aware of them. They are just screened off from interacting with the reason provided by the authority.

On Zagzebski's view, authority understood in this sense can be justified in one of two different ways: by a subject conscientiously judging either (i) that she is more likely to form a true belief and avoid a false belief, or (ii) that she is more likely to form a belief that survives her conscientious reflection, if she believes what an authority believes than if she tries to figure out what to believe

³ See Raz (1988) and Zagzebski (2012). See also Keren (2014).

herself.⁴ Conscientious reflection is “using our faculties to the best of our ability in order to get the truth” (Zagzebski 2012, p. 48). This is not an externalist notion, where one can strive to be as conscientious as possible but still fall radically short. It is doing the best that one can epistemically, where this is grounded in a natural trust that Zagzebski argues we all have in our own faculties—a trust that cannot be supported with a non-circular defense of the reliability of these faculties. Let us call this the *Authority View* (AV) of the rationality of beliefs.

The first model of expertise that I would like to consider, then, is that of an *expert-as-authority* in the sense found in the AV. On this model, the testimony of experts is not one piece of evidence to be weighed with other evidence the hearer might have; rather, experts provide preemptive reasons for belief in the domain of their expertise. Moreover, beliefs formed on the basis of expertise on this view can be justified via the sort of conscientious reflection discussed above. Understanding experts in this way provides a framework for capturing many of our ordinary beliefs and practices involving expertise. When we are looking to identify an expert, it is commonly thought that what we want is someone whose testimony reliably puts us in touch with the truth in a given domain. And if we judge someone to be better, perhaps even significantly so, at doing this than we would be were we to form beliefs on our own, then it arguably makes sense to treat the expert’s reasons as preemptive.

Indeed, the most powerful support for this model of expertise comes from what we might call the *Track Record Argument*, which emphasizes precisely this point: if the testimony of an authority is taken, not as providing a preemptive reason for belief, but as simply one piece of evidence to be weighed with one’s other evidence, then it is said that one’s record in getting at the truth will be worsened. Since acquiring true belief is a key epistemic goal, preemptively trusting the testimony of an authority is epistemically justified. Joseph Raz puts this point as follows:

Suppose I can identify a range of cases in which I am wrong more than the putative authority. Suppose I decide because of this to tilt the balance in all those cases in favour of its solution. That is, in every case I will first make up my own mind independently of the “authority’s” verdict, and then, in those cases in which my judgment differs from its, I will add a certain weight to the solution favoured by it, on the ground that it, the authority, knows better than I. This procedure will reverse my independent judgment in a certain proportion of the cases. Sometimes even after giving the argument favoured by the authority an extra weight it will not win. On other

⁴ More precisely:

Justification Thesis 1 for the Authority of Belief (JAB 1)

The authority of another person’s belief for me is justified by my conscientious judgment that I am more likely to form a true belief and avoid a false belief if I believe what the authority believes than if I try to figure out what to believe myself.

Justification Thesis 2 for the Authority of Belief (JAB 2)

The authority of another person’s belief for me is justified by my conscientious judgment that I am more likely to form a belief that survives my conscientious self-reflection if I believe what the authority believes than if I try to figure out what to believe myself. (Zagzebski 2012, pp. 110–11)

occasions the additional weight will make all the difference. How will I fare under this procedure? If, as we are assuming, there is no other relevant information available, then we can expect that in the cases in which I endorse the authority's judgment my rate of mistakes declines and equals that of the authority. In the cases in which even now I contradict the authority's judgment the rate of my mistakes remains unchanged, i.e., greater than that of the authority. This shows that only by allowing the authority's judgment to pre-empt mine altogether will I succeed in improving my performance and bringing it to the level of the authority. (Raz 1988, pp. 68–9)

In addition to endorsing this argument from Raz, Zagzebski discusses empirical studies that show not only that preemption leads to better results, but also that rats are better at preempting than humans are. For instance, suppose that a green light has flashed 75 percent of the time and one's goal is to predict whether a red or a green light will flash. Humans proportion their verdicts to the probabilities and thus predict green 75 percent of the time, while rats always choose the option that appears 75 percent of the time—thus always predicting green. The result is that rats end up right 75 percent of the time while humans end up worse off. According to Raz and Zagzebski, this shows not only that we ought to preemptively trust an authority if our goal is true belief, but also that “once we identify... an expert, we should follow that person's... advice all of the time” (Zagzebski 2012, p. 115, note 11).

The spirit of the Track Record Argument is highly intuitive: if I recognize that someone is better than I am at getting at the truth in a given area, wouldn't it be better epistemically for me to just wholly defer to her rather than aggregating her input with my own views on the topic? As John Hardwig writes, “because the layman is the epistemic inferior of the expert (in matters in which the expert is expert), rationality sometimes consists in refusing to think for oneself... appeals to epistemic authority are [thus] essentially ingredient in much of our knowledge” (Hardwig 1985, p. 336). To make this point vivid, suppose that I recognize that I am unreliable at identifying instances of sexism and I also know that my colleague is excellent at doing so. If I take my colleague's reliable testimony that someone's comment was sexist as simply a piece of evidence to be combined with all of my unreliable beliefs on the matter, then I would surely be worse off than if I just took her word for it. This is not only due to my unreliable beliefs possibly outweighing my colleague's for me, but also because the very beliefs and processes that I am using in the weighing process might themselves be riddled with sexism. Paulina Sliwa makes this point when she writes, “if [an] agent is worried about her moral judgment being biased, there is no guarantee that her further deliberation won't be biased or overly impressed with her self-interested reasons as well. And once she starts reflecting further, she opens the way to succumbing to temptation and rationalizing the testimony away... Moral advice is so useful precisely because it's a means to put an end to one's reflections about what the right thing to do is” (Sliwa 2012, p. 183).

While the Track Record Argument provides support for understanding experts-as-authorities, the problem of peer disagreement is particularly salient on this view. For if I am looking for an expert whose testimony is going to give me preemptive reasons for belief, then I had better make sure that I find a reliable one. This is because once I identify an expert, her testimony is authoritative for me: all of my evidence on a given

topic is normatively screened off, replaced by only what the expert tells me. Given this degree of reliance, picking reliable experts is of paramount importance epistemically. But as noted earlier, widespread disagreement in a domain seriously calls into question our ability to identify reliable sources in that domain. If I am trying to choose between two purported experts who radically disagree on a topic, and I myself am not an expert on the topic, what independent resources do I have to draw on to figure out which is the more reliable one? Indeed, the problem is exacerbated by the fact that local disagreements over a given question often involve disagreement regarding many other issues that bear on it. With two radically different frameworks, each being championed by a purported expert, having the resources for identifying which is the reliable one seems hopeless.

But the situation is not as bleak for expertise as it might seem. For despite the fact that disagreement poses a significant obstacle to picking out experts when they are taken to be authorities, there are serious problems both with the AV that provides the basis for this model and with the Track Record Argument. After developing these problems, I will propose an alternative conception of expertise that fares better than the expert-as-authority model not only by avoiding the objections facing this view, but also in dealing with peer disagreement.

Let's begin with the AV itself. The first and most obvious problem with this view is that *it provides all of the resources for rendering rational the beliefs of paradigmatically irrational communities*, such as white supremacists, cults, and terrorists. To see this, notice that it is surely possible for a member of a white supremacist group to conscientiously judge that if she believes the teachings of her group, the result will survive her conscientious reflection better than if she tries to figure out what to believe on her own. This is especially clear when the beliefs in one's doxastic framework that are relevant to one's conscientious judging are themselves shaped and guided by one's membership in the community in question. If a person has been raised among white supremacists, for instance, then it is quite natural for her to judge that she is more likely to form beliefs that survive conscientious reflection if she believes what her fellow white supremacists believe since it is the very beliefs of her community that provide the framework through which she is so conscientiously judging. Indeed, the more insular a community is, the more likely it is for beliefs of its members to survive conscientious reflection.

The second problem with the AV is that *it fails to provide the resources for rationally rejecting an authority's testimony when what is offered is obviously false or otherwise outrageous*. Suppose, for instance, that I conscientiously judge that the pastor of my church is an authority on moral matters and he testifies to me that women are morally inferior to men. According to this conception of expertise, this instance of testimony is not one piece of evidence to be weighed against all of the other relevant evidence I have about the moral capacities of men and women; instead, it replaces all of the evidence I have on the topic. It is thus fully rational for me to now believe that women are morally inferior to men, despite the massive amounts of compelling evidence I have to the contrary. But why should one person's testimony—even when it is from a recognized

authority—swamp all of my other relevant evidence on the question, especially when the proffered report is clearly false?

Zagzebski is aware of the general worry that the AV might lead to the rational acceptance of crazy claims from authorities and she offers the following by way of response:

Suppose the authority's belief is something outrageous. Can't that count as a defeater of your belief that it is an authority? Yes it can, but that does not count against preemption. Suppose your physician tells you to take 4,000 pills an hour for the rest of your life. I assume that you trust your belief that you should not take so many pills more than you trust your judgment that your physician is an authoritative guide to your health... But as long as you conscientiously think the physician is a better guide, you have reason to take the physician's directive as one that preempts your own decision about what you should do in that domain.

(Zagzebski 2012, p. 116)

This response, however, seems to be little more than simply biting the bullet. In particular, Zagzebski admits that so long as I continue to regard someone as an authority, then it is rational for me to accept his testimony as a preemptive reason for sharing his belief, no matter how outrageous or morally perverse it is. Of course, it might be thought that very few people will be such that they will continue to regard someone as an authority in the face of such outrageous or morally perverse claims. But, sadly, this is simply not the case. Many groups have recognized authorities who teach their members precisely the sorts of claims at issue, such as that women are morally inferior to men or that certain racial groups are intellectually superior to others. And given that the reasoning employed by the members of these groups is often largely influenced by the beliefs of the other members, it is not at all surprising that many of these sorts of claims fail to raise red flags about the credentials of the authorities at issue. In these cases, then, Zagzebski simply grants that one can rationally accept the testimony in question, such as that women are morally inferior to men.

The third problem with the AV is that *it is unclear how the testimony of an authority can even strike one as clearly false or outrageous, given that all of one's other relevant evidence has been normatively screened off*. For recall that the testimony of an authority replaces all of one's other reasons on the topic. Given this, in virtue of what can the testimony of an authority strike one as crazy? In the above case, for instance, if all of my non-authority based reasons about the moral properties of men and women are normatively screened off, how exactly does the testimony from my pastor that women are morally inferior to men strike me as outrageous? The standard explanation of what happens here is that the substantial amount of background evidence that I have about women's moral capacities, when weighed against the testimony of my pastor, enables me to recognize the absurdity of the proffered testimony. But this background evidence is simply not available on the AV for interaction with my pastor's testimony. Thus, there do not seem to be resources to draw upon on this view for explaining how the outrageous testimony of authorities strikes hearers as outrageous in the first place. Here is another way to put this point: it is not the *content* of the claim that women

are morally inferior to men that is wildly implausible. It is, rather, that it is wildly implausible *given* all of the other information I have on the matter. This, however, just sounds like weighing.

Somewhat relatedly, the AV will be hard pressed to explain what a subject ought to do epistemically when two purported authorities disagree on a question. If the testimony of each authority provides preemptive reasons, then presumably they cannot be weighed against one another or with any of one's other background evidence. So should one accept both instances of testimony, even if they are contradictory? If not, should one withhold belief with respect to both, even if one is far more epistemically supported than the other? Neither option seems epistemically appropriate. But choosing one instance of testimony over another because of the support it enjoys would seem to involve the weighing of evidence, which is clearly at odds with the AV. It is, then, entirely unclear how the AV has the resources for providing a plausible answer here.

The upshot of these considerations is that non-evidential and preemptive authority-based reasons are simply incommensurable with evidence. Indeed, Zagzebski admits as much: she argues that preemptive reasons are deliberative, first-personal reasons while evidence-based reasons are theoretical. She then writes that "theoretical reasons aggregate with each other, but they do not aggregate with first-person reasons" (Zagzebski 2012, p. 114). This renders beliefs based entirely on authority immune to defeaters, as defeaters fundamentally involve evidential relations. Authority-based reasons, then, not only swamp past reasons, as we saw above, they screen off future reasons. If I come to believe through my pastor's testimony that women are morally inferior to men, then future evidence to the contrary is screened off since it is unable to interact with my preemptive reasons. Relatedly, deliberation involving authority-based reasons is closed off, as such a process involves the weighing of evidence. If I am deliberating about whether to join the army, for instance, I would weigh the pros and cons of such a commitment and try to determine whether one side has more in favor of it than the other. But authority-based reasons cannot figure into this sort of deliberation. Thus if I take my pastor to be an authority with respect to making major decisions in life, then when he tells me to join the army, this testimony simply cannot interact with the evidentially grounded reasons I have for not joining.

This shows that there are no rational beliefs grounded in authority, where authority is understood as set out by the proponent of the AV. This is because all beliefs are subject to the possibility of defeat⁵ and the only way to rationally retain beliefs when confronted with a defeater is by having a defeater-defeater. Forming beliefs on the basis of the testimony of experts-as-authorities, then, does not result in rationality or epistemic justification. Thus, there are no rational beliefs that are grounded in preemptive reasons.

There is one final problem that I would like to raise for the AV and, therewith, the conception of expertise on which it is based: what happens if I ultimately come to reject

⁵ This is true at least with respect to all of the beliefs at issue here.

as an authority someone who once had this status for me? The likely answer on the AV is *epistemic catastrophe*. If I have been relying on the testimony of an authority, then I have been normatively screening off all other evidence relevant to the question at issue and forming beliefs about this question purely on the basis of what the authority tells me. But then this also means that I most likely have missed out on numerous opportunities to gather relevant evidence. For instance, there is certainly no need for me to *take in, consider, or seek out*, evidence relative to the question at issue beyond what the authority tells me. This can lead to my missing out on many opportunities to hear the opposing views of others and the reasons for their beliefs. There is no need for me to *deliberate* about what the authority tells me or to consider it in light of either past, present, or incoming evidence on the matter. Thus, any connections that I might have made, or inferences that I might have drawn, from such a process of deliberation will be lost. And overall, these losses can have a crippling effect on my epistemic life, especially should I ultimately come to reject my authority. For I am then left without a great deal of evidence on the question at issue, evidence that I ignored in favor of the testimony of my authority.⁶

We have seen, then, that there are serious problems afflicting the expert-as-authority model grounded in the AV. Let us now turn to the central support for this model of expertise: the Track Record Argument. There are at least three problems with this argument. The first is that it is not at all clear how it is compatible with the response offered above to the problem of an authority offering outrageous testimony. For, on the one hand, both Raz and Zagzebski claim that an authority's testimony can be rejected in virtue of its being outrageous and, on the other hand, that the preemption model should be favored because it will better one's track record at getting at the truth, but only if the authority's advice is followed "all of the time." Doesn't accepting an authority's testimony "all of the time" mean doing so even when it is crazy?

Both Raz and Zagzebski suggest that an acknowledged authority is simply not an authority in cases where one has reason to doubt the proffered testimony, such as when it is outrageous. So presumably the idea is that the problem is avoided because one should follow the advice of a *genuine authority* all of the time, and the status of being an authority does not survive doubt. But surely this is not going to save the preemption model of authority-based reasons. For the view suggested here is that one should be evaluating authorities, comparing what they say with one's background information to determine whether their testimony permits the retention of their status as authorities. Otherwise put, the recommendation of the proponent of the AV seems to be this: always follow the advice of an authority, except when you have strong evidence that the person is no longer an authority on the matter at issue. However, isn't this not only just what the proponent of a model of authority that

⁶ In addition to the epistemic consequences, this can have a profound psychological impact. After putting all of my doxastic eggs in one basket, I might be left feeling completely lost as an epistemic agent once my authority is rejected.

involves weighing says, but also what the humans, rather than the rats, are doing in the empirical studies cited above?

The second problem with the Raz/Zagzebski argument above is this: suppose that it is granted that in order to avoid worsening one's track record for getting the truth, one should *always* follow the advice of an authority. Nothing that has been said above, however, shows that this should be done *preemptively*. For humans would do just as well as rats if they accepted the testimony of an authority all of the time, but without the authority-based reasons replacing their other reasons. Instead, the authority-based reasons might simply be regarded as so epistemically powerful that they always outweigh the other reasons. The result would be the same: both the rats and the humans end up being right 75 percent of the time.

The third problem with the Raz/Zagzebski move is that it is not the case that in order to avoid worsening one's track record for getting the truth, one should *always* follow the advice of an authority. Here are some alternative policies that would have even better epistemic results: follow the advice of an authority, except when one is certain that the authority is wrong; follow the advice of an authority, except when one knows that the authority is wrong; follow the advice of an authority, except when what the authority says is highly doubtful. If humans adopted any of these policies, they would end up faring better than the rats and, moreover, following them relies directly on not screening off the normative force of background evidence, as the AV requires. For instance, something would strike one as highly doubtful only against the background of one's other relevant information. The upshot of these considerations, then, is that there is simply no connection between the goal of true beliefs and either preemption or the requirement to always follow the advice of an authority. Given that this is the central argument offered for favoring the preemption model of authority to the weighing one, there is still further reason to reject the AV.

While we have seen that the problem of peer disagreement poses a particularly difficult problem for experts when they are understood as authoritative in the sense found in the AV, we have also seen that there are independent reasons to reject the AV and the Track Record Argument that is used to support it. Let us now turn to an alternative conception of expertise.

11.4 The Expert-as-Advisor Model

I now want to turn to a different conception of expertise, what I call the *expert-as-advisor* model, and show that it not only avoids the problems afflicting the expert-as-authority model, but also that it has the resources for a much more satisfying response to the problem of peer disagreement.

An expert that is an advisor does not give authoritative testimony or preemptive reasons for belief; rather, her testimony provides evidence for believing a given proposition and, in this way, offers guidance. Consider the testimony provided by an

expert at a trial: no one would tell the jurors that the testimony of a given expert is authoritative or provides preemptive reasons for belief. Indeed, jurors themselves would be superfluous in many ways if experts functioned authoritatively. Instead, competing expert testimony is often presented from both sides—the prosecution and the defense—with jurors needing to evaluate the full body of evidence in reaching a verdict. The experts here are, then, advising the jurors rather than dictating to them what they ought to believe.

To be clear, though, I am not suggesting that from an epistemic point of view, sometimes we ought to treat experts as authorities and other times we ought to treat them as advisors. The considerations in section 11.3 reveal that there are serious epistemic problems with both the authority model of expertise and the central argument on its behalf and, thus, we should never aim to form beliefs as the rats do in the empirical studies. What I am proposing here is that the testimony of experts should always be regarded as a piece of evidence to be weighed with the other relevant evidence we have on the matter. Epistemically, then, experts should always be regarded as advisors.

When offering guidance or advice, certain features of the proffered testimony become far more important than when the testimony is authoritative. For instance, if I am advising you about what to believe, I am not asking you to simply trust me or to take my word for it, and thus it is helpful that I provide reasons for holding the belief. I might explain the matter to you or offer arguments or other support for the belief at issue. I may be sensitive to your particular epistemic needs or concerns. I might be patient and open-minded in my treatment of you as an epistemic agent. In general, though, I am far more likely to cultivate understanding in you about the matter if I am your advisor rather than your authority.

Of course, this is not to say that reliability doesn't matter to being an expert-as-advisor. If, for instance, I am wholly unreliable in a given domain with respect to both primary and secondary questions, then no matter how able I am at presenting support for holding a belief, I am not an expert in the relevant domain. But reliability is simply one dimension of being an expert-as-advisor, whereas it is the sole, or at least the dominant, feature of being an expert-as-authority.

Given this, we have far more resources available to us for identifying experts, even in the face of widespread disagreement. We can evaluate the arguments proffered on behalf of a particular view, we can assess how able the expert in question is at enhancing our understanding of the matter, we can determine how effective the expert is at explaining matters in clear and accessible ways, and so on. This far richer framework for evaluating experts focuses not only on their testimony and beliefs, but also on their broader character traits. An ethics consultant serving at a hospital will be effective largely by helping doctors, patients, and their families navigate through difficult medical decisions. Sure, her reliably offering true testimony is important, but equally important are her abilities to clearly explain the terrain, to listen attentively and receptively to the concerns and values of those around her, and to answer questions in a thoughtful and constructive way. This is a place where the intellectual and moral

virtues are relevant. For we might not be interested only in experts helping us reach the truth but in *how* they help us reach the truth. A dogmatic and close-minded ethics consultant might offer lots of true statements about a medical decision but she might do so in a way that fails to move a patient and her family. They might feel ignored or bullied, which might result in the expert's testimony failing to gain any traction with them, thereby leaving them with fewer true beliefs.

What I especially wish to emphasize here, though, is that being able to identify experts in the face of widespread disagreement is not hopeless when they are assessed as advisors rather than as authorities. For even if I cannot tell who is more reliable in a given domain, I surely can tell who is better at effectively clarifying the terrain, or at being a sensitive listener, or at being open-minded to the issues that are of concern to me. These skills and virtues can be available even when who is right is not. Moreover, disagreement between experts can be beneficial when what is being sought is advice rather than an authoritative directive since there is more evidence available to be weighed when arriving at a conclusion.

This conception of expert-as-advisor accords well with arguments I have given elsewhere to show that knowledge is not sufficient for epistemically proper assertion. In particular, I identify a specific kind of knowledge, what I call *isolated second-hand knowledge*, and argue that it is an epistemically insufficient basis for certain classes of assertions. "There are two components to this phenomenon: the subject in question knows that p solely on the basis of another speaker's testimony that p —hence the knowledge is second-hand; and, second, the subject knows nothing (or very little) relevant about the matter other than that p —hence the knowledge is isolated" (Lackey 2011, p. 254). One class of assertions where I show that isolated second-hand knowledge is an insufficient epistemic basis is that of experts offering flat-out assertions in the domain of their expertise. The following is one of the examples I use to illustrate this:

DOCTOR: Matilda is an oncologist at a teaching hospital who has been diagnosing and treating various kinds of cancer for the past fifteen years. One of her patients, Derek, was recently referred to her office because he has been experiencing intense abdominal pain for a couple of weeks. Matilda requested an ultrasound and MRI, but the results of the tests arrived on her day off; consequently, all the relevant data were reviewed by Nancy, a very competent colleague in oncology.⁷ Being able to confer for only a very brief period of time prior to Derek's appointment today, Nancy communicated to Matilda simply that her diagnosis is pancreatic cancer, without offering any of the details of the test results or the reasons underlying her conclusion. Shortly thereafter, Matilda had her appointment with Derek, where she truly asserts to him purely on the basis of Nancy's reliable testimony, "I am very sorry to tell you this, but you have pancreatic cancer." (Lackey 2011, pp. 34–5)

Despite the fact that there is no epistemic barrier to Matilda knowing that Derek has pancreatic cancer, I argue that her isolated second-hand basis for this belief is an

⁷ In the original case, Nancy was a competent medical student but I follow Benton (2016) in modifying the details for the purposes of clarity.

epistemically inadequate one for her flat-out asserting that this is the case. This is because expert assertion requires that the asserter be in a position to respond to questions or challenges regarding the expert judgment, and isolated second-hand knowledge does not enable this. Thus, I conclude that knowledge is not sufficient for epistemically proper assertion, not because more or less justification is needed for epistemically proper assertion, but because sometimes—such as in the case of expert assertion—a certain *kind* of epistemic basis is needed for proper assertion.

The point that is of import for our purposes here is that cases involving isolated second-hand knowledge show that testimony offered by experts requires that they at least have the capacity to function as an advisor. This is because they need to be in a position to answer questions or challenges regarding their testimony. But nothing in the expert-as-authority model captures this. It would be perfectly acceptable on this view to not have anything to offer to a challenge or to simply respond, “Because I said so” or “Because I’m an authority on the matter.” After all, that the testimony is from an authority just *is* the reason for believing on this view, so there is no requirement that the expert have further reasons to offer beyond this. The problematic nature of isolated second-hand knowledge, then, provides further support for conceiving of experts as advisors.

The expert-as-advisor model also avoids the problems afflicting the authoritative one. First, there are no particular worries that the advisor view will render rational the beliefs of paradigmatically irrational communities, such as white supremacists, cults, and terrorists. For rationality on this model fundamentally involves the weighing of evidence and so beliefs will only be as rational as the evidential support grounding them. Second, it is very clear how one could reject an authority’s testimony when what is offered is obviously false or otherwise outrageous: such testimony will end up being outweighed by the vast amounts of background evidence one has on the topic. Finally, since one’s other relevant evidence on the topic in question will always be available for deliberation, there are no problems understanding how the testimony of an authority can strike one as clearly false or outrageous.

It is also worth noting that the expert-as-advisor model provides the resources for capturing many of our ordinary ascriptions and practices involving expertise. For instance, nearly every hospital in the United States has an ethics consultant on staff. The University of Washington’s School of Medicine describes this role as follows: “An ethics consultant is an expert in ethics who provides ethics consultations and may also serve as an educator to the committee or program. In some health care institutions an ethics consultant provides ethics expertise to workgroups that are addressing systems issues and have the need to better understand the ethics and preferred practices from an ethics perspective.” But calling ethics consultants “experts” and having patients and committees consult with them in this role is at serious odds with many philosophical views on the matter. Recall, for instance, the earlier claims by Hurley and McGrath that widespread disagreement calls into question moral expertise or at least our ability to pick out experts. And, according to Bernard Williams:



There are, notoriously, no ethical experts. . . . Anyone who is tempted to take up the idea of there being a theoretical science of ethics should be discouraged by reflecting on what would be involved in taking seriously the idea that there were experts in it. It would imply, for instance, that a student who had not followed the professor's reasoning but had understood his moral conclusion might have some reason, on the strength of his professional authority, to accept it. . . . These Platonic implications are presumably not accepted by anyone.

(Williams 1995, p. 205)

Williams here takes the stronger view that there aren't any ethical experts rather than that we simply cannot reliably pick them out as such. And while he might not have in mind all that is involved in the expert-as-authority model, he clearly finds problematic what is at the core of the view: forming beliefs purely on the basis of an expert's *authority*, at least when ethical matters are at issue.

But the expert-as-advisor model offers a way of explaining our reliance on experts, even in domains where there is widespread disagreement. Consider, again, ethics consultants at hospitals. Typically, what is being sought from such experts is not a settled view, but advice and guidance navigating a complex and morally fraught matter. Indeed, in a 2008 article in *American Medical News*, Lainie Friedman Ross, M.D., Ph.D., Associate Director of the University of Chicago's MacLean Center for Clinical Ethics, seems to have in mind precisely this conception of expertise when he says, "Rarely do ethicists tell you—the physician—what to do, but rather they help you think about options and ways to negotiate compromise." In this way, experts-as-advisors function more as personal trainers than as oracles.

This model of expertise also fits the physician-patient relationship that exists for many. Given our heavy reliance on the Internet and the accessibility of vast quantities of medical information via sites like WebMD, many patients already possess a great deal of knowledge about their conditions and options when consulting with their doctors. What they are often looking for, then, is direction in exploring the best way to move forward, where such a relationship is collaborative or advisory, rather than authoritative.

Even in domains where testimony is sometimes taken to be authoritative for various reasons, the advisor model can be easily adapted. Consider religious contexts: while it may be taught that one's religious leaders are authorities, this doesn't mean that, from an epistemic point of view, this is how they ought to be treated. Indeed, the general version of this point is precisely what was argued in the previous section. But it needn't follow from this that there aren't any religious experts. If one consults with a pastor about one's struggles with belief in God or whether to stay in one's marriage in the face of infidelity, for instance, one can certainly regard the pastor as a wise and deeply rich source of guidance. One may talk through one's doubts or concerns, ask questions, listen to the pastor's own personal struggles, inquire about the grounds for the church's teachings, and so on. One may also weigh the testimony of the pastor very heavily—far more so than one would weigh the input of even a highly respected friend. But none of this requires treating the pastor's reasons preemptively. In fact, putting aside the

epistemic problems with doing so, it would also be extremely odd psychologically to consult with one's pastor about struggling with one's faith and then regarding her as an oracle. How would this even look? In response to grappling with whether the immense suffering in the world is compatible with there being a God, one would then go on to just accept that God does in fact exist simply because the pastor says so.

It is also interesting to note that the Track Record Argument offered in support of the expert-as-authority model makes it such that the arguments and explanations that the expert-as-advisor model says render one a *better* expert make one a *worse* expert on the authority model. This is because in order to avoid one's own novice input reducing the number of true beliefs acquired, the Track Record Argument requires that experts be trusted all of the time on all relevant issues. But an expert offering arguments and explanations on behalf of her testimony encourages deliberation on the part of the recipient, which, in turn, invites the latter's own novice input to enter the picture. And such "meddling," we are told, ultimately decreases truth-conduciveness.

There is, however, a significant objection to the expert-as-advisor model that needs to be addressed. Many people are novices or laypersons when it comes to the questions regarding which they are seeking the input of experts and hence they might be utterly ill equipped to take what they are told as mere "advice." Indeed, John Hardwig has argued extensively for the thesis that rationality requires sometimes believing propositions for which one has no evidence at all, and thus when a layperson relies on an expert, that reliance is necessarily "blind."⁸ For instance, if I am consulting with a forensic scientist about gunshot residue, I might have absolutely no framework with which to evaluate what I am told. Because of this, my dependence on her expertise is not properly understood as advice or guidance; rather, given my absolute novice status, she is an authoritative expert on the matter relative to me. And this, it is argued, reveals a deep limitation with the expert-as-advisor model.

By way of response to this point, notice that my lacking evidence for believing something beyond what is offered by the testimony of an expert does not at all call into question whether believing on the basis of expertise involves the weighing of evidence. I might have lots of evidence on behalf of the reliability of the expert herself, for instance, even if I have none whatsoever regarding the content of what she reports. Indeed, an expert's testimony might be weighed so heavily that it virtually guarantees my corresponding belief, either because it is just about all of the evidence I have on the matter or because the expert is so epistemically superior to me. In this way, my readiness to accept whatever the expert tells me might be misinterpreted as evidence-less believing, when in fact it is not.

This point can be developed by extending our earlier analogy. As should be clear, personal trainers can shoulder more or less of the burden for those whom they train. One helping with fitness, for instance, may at times be almost literally doing the exercises for her client, while at other times she may be sitting back watching him take them on

⁸ See Hardwig (1991).

almost entirely on his own. The same is true in the epistemic domain. The mere fact that experts are advisors does not mean that we cannot be highly dependent on them. Sometimes, we may be so out of our depth that deference is nearly guaranteed, while at other times we may be simply looking for a bit of guidance. The crucial point is that from an epistemic point, we should never, ever behave like the rats—we should never entirely screen off our own reasons in a domain when relying on experts, no matter how good they are.

11.5 Conclusion

We have seen that while widespread disagreement among epistemic peers in a domain threatens our ability to identify experts when they are understood as authorities, it does not do so when they are regarded as advisors. Even if I cannot pick out who is right in every dispute, I can still determine who is better at presenting the relevant issues clearly, who can help me navigate complex terrains with insight and nuance, who can guide me toward drawing conclusions from other things that I accept, and so on. Such features not only have obvious intellectual value, they also govern our actual practices involving expertise. There is, then, no reason to be driven to radical forms of skepticism about expertise in controversial areas, such as religion, politics, and philosophy, as experts can advise, even when they are not authoritative.⁹

References

- Benton, Matthew. 2016. "Expert Opinion and Second-Hand Knowledge." *Philosophy and Phenomenological Research* 92: 492–508.
- Cholbi, Michael. 2007. "Moral Expertise and the Credentials Problem." *Ethical Theory and Moral Practice* 10: 323–34.
- Elga, Adam. 2007. "Reflection and Disagreement." *Noûs* 41: 478–502.
- Goldman, Alvin. 2001. "Experts: Which Ones Should You Trust?" *Philosophy and Phenomenological Research* 63: 85–110.
- Hardwig, John. 1985. "Epistemic Dependence." *The Journal of Philosophy* 82: 335–49.
- Hardwig, John. 1991. "The Role of Trust in Knowledge." *The Journal of Philosophy* 88: 693–720.
- Hurley, Patrick. 2008. *A Concise Introduction to Logic*, 10th edn. Belmont, CA: Wadsworth.
- Kelly, Thomas. 2005. "The Epistemic Significance of Disagreement." In John Hawthorne and Tamar Szabo Gendler, eds, 167–96. *Oxford Studies in Epistemology, volume 1*. Oxford: Oxford University Press.
- Keren, Arnon. 2014. "Trust and Belief: A Preemptive Reasons Account." *Synthese* 191: 2593–615.

⁹ I am grateful to Matthew Benton, Baron Reed, an anonymous referee, and audience members at the Expertise, Humility, and Disagreement Conference at Monash University, at the GAP.9 Conference (Osnabrück, Germany), at the Northwestern-Cologne Disagreement Conference, and at Ryerson University.

- Lackey, Jennifer. 2011. "Assertion and Isolated Second-Hand Knowledge." In Jessica Brown and Herman Cappelen, eds, 251–75. *Assertion: New Philosophical Essays*. Oxford: Oxford University Press.
- McGrath, Sarah. 2008. "Moral Disagreement and Moral Expertise." In Russ Shafer-Landau, ed., *Oxford Studies in Metaethics* 3: 87–107. Oxford: Oxford University Press.
- Raz, Joseph. 1988. *The Morality of Freedom*. Oxford: Clarendon Press.
- Sliwa, Paulina. 2012. "In Defense of Moral Testimony." *Philosophical Studies* 158: 175–95.
- Williams, Bernard. 1995. *Making Sense of Humanity*. Cambridge: Cambridge University Press.
- Zagzebski, Linda. 2012. *Epistemic Authority: A Theory of Trust, Authority, and Autonomy in Belief*. Oxford: Oxford University Press.