

'We're Not Very Good at This': Dealing with Evidence of Unreliability

Nathan L. King (University of Notre Dame)

[Early draft; some revisions remain. Do not cite or quote. Comments welcome.]

We know a lot. And we have rational beliefs about many topics: our immediate environment, our own thoughts, the mental states of others, matters of common sense, the past, the future, scientific facts, mathematics, and – perhaps more controversially – morality. For the purposes of this paper, I'll add beliefs about religious matters to the list. I'll assume, for instance, that Christian, Islamic and Jewish belief is rational – at least *prima facie* rational. Indeed, I'll assume that such belief often enjoys *prima facie* rational status independent of arguments or other publicly available evidence. In this respect, my starting point accords with those who take religious belief to be properly basic – that is, rational in a way that does not depend on further beliefs of the relevant subject. On the picture I'll assume here, religious beliefs, if rational, are typically rationally grounded in religious experience rather than in arguments. It is the rational status of such experience-based beliefs that will take center stage in this paper.

Prima facie rational beliefs lose their positive epistemic status when the relevant subject acquires a defeater. Where one's initial grounds G render a belief B *prima facie* rational, a *defeater* is a reason R such that (R & G) does not render B rational. A central way in which one's belief might be defeated is when one acquires evidence that this belief was not formed in a reliable way (evidence, e.g., that one has heeded misleading evidence or has made a mistake in reasoning).

Much recent work in epistemology has focused on two ways in which evidence of fairly widespread unreliability may be revealed: (i) empirical studies; and (ii) disagreement. Richard Foley has produced significant work on both topics, chiefly in his *Intellectual Trust in Oneself and Others*. In this paper, my aim is to apply Foley's work to the epistemology of religious belief, in particular, the epistemology of religious disagreement. The plan of the paper is as follows. In section I, I summarize the evidence of unreliability to be gleaned from empirical studies of certain sorts of judgment (e.g., judgments about probability). I then explain a handful of Foley-inspired strategies for coping with this unreliability. Section II explores the extent to which these strategies can be applied to the epistemology of experience-based religious beliefs.¹ There I'll suggest

¹ Why not simply apply what Foley says about *disagreement* to the religious case? Why model the epistemology of religious disagreement on these empirical studies, which don't involve disagreement at all? Why not focus on the Foleyan account of peer disagreement? Fair questions. By way of response: First, it is not at all clear that typical religious disagreements are *peer* disagreements, so the Foleyan account of peer disagreement may not apply in this context (More on this point below). Second, the beliefs targeted in the empirical studies share important features in common with controversial religious beliefs. It is plausible that these factors are doing crucial epistemic work both in the empirical cases and in the disagreement cases. In both cases, higher-order evidence about the unreliability of our belief-forming practices is doing crucial epistemic work. Inasmuch as Foley deploys several strategies for dealing with this higher-order evidence in the empirical cases, it is worth considering whether the strategies can be applied to the religious case as well.

that the Foleyian religious epistemologist will give *public evidence* (e.g., apologetic arguments) a key role in solving the epistemological problem of religious disagreement.

Before we get to the studies, one caveat. The following material marks an attempt to draw together several threads from the epistemology literature. I'll attempt to relate studies in empirical psychology to recent work in religious epistemology, and to relate both of these to Foley's work in *Intellectual Trust*. Discussions in these sub-sections of the literature often trade in talk of different epistemic statuses. For example, there's Alstonian justification, Plantingean warrant, and Foley rationality. One might worry that simultaneous discussion of all these topics and statuses carries the risk of conceptual slippage. I share this worry. However, here's an *initial* reason to think that the discussion isn't simply doomed to confusion: all of the statuses just mentioned are defeasible. And it is in the capacity of a prospective defeater that both religious disagreement and empirical studies are alleged to have their epistemic significance. The language of defeat is common currency in discussion of the different topics and statuses that will concern us. So there is at least some reason to think that these matters can be brought together after all. With this hopeful thought in mind, let's give it a try.

I. Evidence of unreliability

A. The empirical studies

I'll start with a topic that is near (though not dear) to my heart – the academic job market. As any candidate knows, the selection process typically depends on a candidate's performance in short interviews. In light of this, it is worrisome that a wide range of studies seems to indicate that interviews do not enable interviewers to reliably judge the likelihood of candidates' future performance. In judging the likelihood of future success given interview performance, human subjects are typically unreliable, even if not irrational. This result is consistent in studies ranging over interview-based predictions regarding the future success of academic job candidates, the future success of graduate and medical students, and even recidivism rates among prisoners. With respect to the studies of student success predictions, committees that used only impersonally gathered data (e.g., MCAT scores and grade point averages) more accurately predicted student success than did committees that used the objective data *plus* the interviews. Despite their central role in institutional selection processes, short interviews do not provide a reliable guide to future success. The most commonly accepted explanations of this phenomenon are that:

- (i) Interviewers are exposed to *irrelevant information* in the course of the interview (e.g., information about the applicant's appearance or mannerisms), and this information skews judgments about a candidate's merits;
- (ii) Due to the brevity of the interviews, interviewers are exposed to a narrow range of data, which may not be *representative* of the applicant's total set of abilities.

As Foley notes, these explanations are complementary. (We'll consider similar explanations below in connection with the epistemology of religious belief. For now, note that these explanations, though illuminating, are not necessary to get us worried about our interview-based judgments – for *that* purpose, the evidence of unreliability is enough. I'll argue below that a similar point applies to analogous explanations of unreliability in religious belief formation.)

Another widely studied phenomenon is *overconfidence bias*. As Thomas Gilovich notes, "One of the most documented findings in psychology is that the average person purports to believe extremely flattering things about him or herself – beliefs that do not stand up to objective analysis."² A disproportionate number of us believe that we are above average in intelligence, fair-mindedness, and driving skill. That is, far more of us believe ourselves to be above average in these respects than can really *be* above average.

In one study, a group of high school seniors was asked to rate their leadership abilities. Over 70% thought that they were above average, while only 2% thought that they were below average. In another study – this one concerning ability to get along with others – all students surveyed thought that they were above average, 60% thought that they were in the top 10%, and 25% thought that they were in the top 1% of their peer group. One might think that age and experience make one less prone to these self-serving errors. However, a survey of university professors found that 94% believe they are better than average professors. Psychologists have explored several explanations of this self-serving behavior; among them is the notion that such beliefs yield practical benefits (e.g., they foster self-esteem). What is crucial for our purposes is that the beliefs of the relevant subjects can't all be *true* – if one takes the studies at face value, they show that human subjects are unreliable in certain domains of self-assessment.

Finally, consider the results of some studies in *probabilistic reasoning*. Repeated studies on both statistically naïve and sophisticated subjects have indicated a tendency to assign a higher probability to certain conjunctions than to their individual conjuncts.³ In one well-known study, subjects were given the following information:

Linda is thirty-one years old, single, and very bright. She majored in philosophy. As a student, she was very concerned with issues of discrimination and social justice, and also participated in antinuclear demonstrations.

Given this information, students were asked to rank the probability of, among others, the proposition (T) Linda is a bank teller; (F) Linda is active in the feminist movement; and (T & F) Linda is a bank teller who is active in the feminist movement. Even in tests whose format highlighted the relation between a conjunction and its conjuncts, 85% of

² Gilovich, *How We Know What Isn't So* (New York: The Free Press, 1991), 77.

³ Daniel Kahneman and Amos Tversky, "Extensional versus Intuitive Reasoning: The Conjunction Fallacy in Probability Judgment." *Psychological Review* 90:4 (1983), 293-315. For a helpful discussion of the epistemic significance of these studies, see Alvin Goldman, *Epistemology and Cognition* (Cambridge, MA: Harvard, 1986) chapter 15.

subjects rated (T & F) more probable than T – thus violating the rules of the probability calculus.

In other studies, subjects routinely ignored prior probabilities and base-rate information. To take an example discussed by Foley, consider a study in which participants are asked to solve the following problem:

Suppose that 1% of the population has colon cancer. Suppose also that when an individual with colon cancer is given the standard diagnostic test, there is an 80% probability of the test being positive. On the other hand, when an individual without colon cancer is tested, there is a 10% probability of the test being positive. John Smith has tested positive. On the basis of this evidence, what are the chances that Smith has colon cancer?

Subjects commonly answer that the probability is slightly less than 80%; most subjects say that the probability is greater than 50%. The correct answer is about 7.5%. The standard explanation of this mistake is that subjects fail to pay adequate attention to the very low (1%) base-rate of colon cancer within the general population.⁴

The above studies apparently show that, in making certain kinds of judgments, we humans are just not very good. Before we consider the implications of the studies, here's a caveat: I'm not committed to the claim that all of the above-discussed studies ultimately show human subjects to be irrational (or even unreliable) with respect to the judgments at issue. I have selected these cases for illustrative purposes. If you have reservations about some of the studies I have mentioned, you are free to substitute different examples – they aren't far to find. In any case, I'll assume that *some* such studies can be taken to show that human thinkers are unreliable with respect to certain subjects. What interests me, and what interests Foley, are the egocentric epistemic questions that arise in light of this. If human subjects are not in general reliable in making certain kinds of judgments, does this defeat the beliefs *I* form when making such judgments? If so, is this defeat partial or complete? And what might I do to overcome it?

B. Strategies for dealing with unreliability

Foley views the above studies (or the unreliability they reveal) as a threat to intellectual self-trust. In the face of the studies, self-trust may be undermined, and the relevant beliefs defeated. Foley thus surveys several strategies for reducing this threat to self-trust, and for defeating the defeaters one acquires when one becomes aware of the studies.

He begins by noting that, when challenged with evidence showing that human subjects are not in general reliable when making judgments of type X, one might respond by asserting that one is a member of a "protected class" of people who are invulnerable to the relevant kind of error. Call this the *protected class strategy*. This response, says Foley, is "mere bluster" unless one has reason to think that one is a

⁴ Foley, *Intellectual Trust*, 64.

member of the protected class. For example, training in economics or probability theory may make one less vulnerable to ignoring base-rates than the average member of the population. But without such training – or some other reason to believe that one is a protected class member – one lacks adequate reason to think that one is not prone to the errors to which similar subjects are vulnerable.

Even when one is not a member of a protected class, one can sometimes employ strategies that dramatically reduce one's tendency to make unreliable judgments. With respect to probability judgments, subjects whose problems were stated in terms of relative frequencies rather than single case probabilities tended to perform quite well. For instance, when subjects were given a reformatted version of the cancer diagnosis problem discussed above, most gave an answer of under 10% - a dramatic improvement in accuracy.⁵ Simply reformatting the problems yielded a significant improvement in reliability. This suggests that even untrained subjects, when presented with such problems, can employ the *reformatting strategy* in order to improve their epistemic position. Subjects equipped to use this strategy, Foley says, are thereby equipped to restore much of the intellectual self-trust eroded by the studies.

With respect to some of the other studies, improving one's epistemic position is not so straightforward. Recall the interviewing studies. In contrast to some of the probability studies, there is apparently no simple reformatting strategy to employ here. And in Foley's view, without some independent evidence, I cannot rationally take myself to be a member of some protected class of those unaffected by the irrelevant, unrepresentative evidence I encounter while interviewing a candidate. The protected class strategy won't work – at least not by itself. A "pessimistic" response to this situation, Foley says, is to completely discount the interview, assign it no evidential weight, and revert to my pre-interview opinion. (If I had no pre-interview opinion, on this response, I should suspend judgment about the candidate's credentials.)

Foley does not think such a response is mandatory in all cases. Accordingly, he seeks a middle way between skeptical pessimism and the sort of naïve optimism that would ignore the studies altogether. The strategy of *self-monitoring* plays a key role in this *via media*. The studies "put me on notice" because they show that subjects similar to me are unreliable when forming opinions on the basis of short interviews. I can't rationally ignore this fact. But, says Foley, even if I lack independent evidence showing that I am a member of a protected class, I need not discount the interview entirely. I can instead monitor the way I formed my opinion using two key sources of information: (i) *introspection*; and (ii) *publicly available evidence*.

With respect to the first source of information, I can look inward, trying to discern whether I gave undue weight to such irrelevant factors as the candidate's appearance or mannerisms. I can try to recall whether I was attentive during the interview, or whether I was in a particularly good or bad mood. If these introspective episodes fail to reveal epistemically suspect tendencies, I receive at least some evidence that I haven't made the same mistakes to which others are prone.

⁵ Recall: the correct answer was 7.5%; most subjects gave an answer of greater than 50%.

As Foley emphasizes, however, effective self-monitoring typically involves more than just careful introspection.⁶ In addition, it usually requires attention to *publicly available evidence*. Applied to the interview case, this might involve comparing notes with other interviewers who were present, asking whether these interviewers thought my questions and responses revealed any sort of bias, or whether my non-verbal cues indicated undue irritation (or undue enthusiasm). I might also compare my interview-based impression with the objective data in the candidate's file. By pooling together the resources of public and introspective evidence, I can limit the epistemically bad tendencies to which I might otherwise fall prey.

If I cannot limit these tendencies, I can *recalibrate* my opinion to account for my biases (e.g., if I am aware that I tend to prefer candidates whose views accord with my own, and the present candidate's views conflict with mine, I can lower my confidence in my interview-based judgment).⁷ If I conduct the process of self-monitoring or recalibration carefully enough, the resulting opinion may be immune to further self-criticism in light of my deepest epistemic standards. Given the goal of having accurate and comprehensive beliefs, immunity to such criticism would then render my interview-based opinion rational in Foley's sense – empirical studies notwithstanding. In Foley's view, then, the empirical studies needn't be an immovable obstacle to self-trust, nor need they ultimately defeat the relevant beliefs.

II. Unreliability and religious belief

How is all this relevant to the epistemic status of religious belief? To begin to address this question, I'll spell out some respects in which religious beliefs are similar to those discussed in the empirical studies. These relevant similarities raise problems for the all-things-considered rationality of experience-based religious beliefs, problems that mirror those discussed above.

I'll start by granting that none of the empirical studies target beliefs that are exactly analogous to religious beliefs. For instance, some of the studies target inferential beliefs, while we are concerned with religious beliefs whose rational status is non-inferentially grounded. And of course, the beliefs in the empirical studies differ in *content* from their religious counterparts. I'm not sure that these differences matter, given the salient features that the two sorts of beliefs share in common. For it is hard to see how, in these cases anyway, differences in a belief's topical content will make epistemic differences. And – differences between inferential and non-inferential beliefs notwithstanding – in neither case is the target belief regarded as *ungrounded*. Reformed epistemologists deny that inferential grounding is the only variety. And in both cases the target belief's rational status is called into question by evidence indicating that human subjects are not reliable in the relevant field of inquiry. The supporting link

⁶ For arguments to the effect that introspection alone is not a particularly effective self-monitoring mechanism, see Daniel Kahneman, "A Perspective on Judgment and Choice: Mapping Bounded Rationality," *American Psychologist* 58 (2003): 697-720.

⁷ Alternatively – on the assumption that I have no direct voluntary control over my confidence levels – I can attempt to exercise indirect control by reflecting on the fact that my belief is likely to have been formed unreliably.

between the relevant grounds and the beliefs they support is allegedly weakened, even if not severed. Let's spell out some similarities between the cases in further detail.

First, like the beliefs of subjects in the study on overconfidence bias, the beliefs set out as orthodox in the world's diverse religions can't all be *true*. At least on a straightforward reading, the core doctrines of the world's great religions are incompatible.⁸ Further, on the assumption that religious persons typically form their beliefs on the basis of religious experience (perhaps coupled with background beliefs gleaned from their religious community), it would appear that the practice⁹ of forming beliefs in this way is not reliable. Indeed, as far as reliability is concerned, one might regard religious believers as worse-off than overconfident drivers. The latter believe only that they are above average, and *lots* of drivers are above average. Though overconfidence bias leads to a fair amount of unreliability via overestimation, several members of the group turn out to be correct. When it comes to religious beliefs, however, we often have a situation in which at most *one* among a set of incompatible doctrines can be correct. In such cases, a higher percentage of subjects hold false beliefs than in the overconfidence bias cases.

Second, though this point isn't required for my argument, some plausible explanations of unreliability in the empirical studies also seem plausible with respect to religious belief. Recall that in the interview studies, one explanation of unreliability concerned the irrelevant, misleading evidence that short interviews often provide. Similarly, a plausible explanation of unreliability in experience-based religious beliefs is the apparent fact that a lot of evidence gleaned from religious experience is *misleading* evidence. (Or at any rate, the experiential evidence is misleading when conjoined with the diverse subjects' background belief systems. Where these two items together constitute the *inputs* of religious belief forming processes, incompatibility between the *outputs* of these processes tells us that something has gone awry.¹⁰) A second, complementary explanation of this unreliability is that experience-based religious beliefs are formed on the basis of *narrow* grounds.¹¹ As in the interview case, the worry is that one's experiential evidence may not be representative of the total evidence available. Indeed, inasmuch as the typical religious subject will have had religious

⁸ There are subtle issues here. Just what generates the incompatibility between the doctrines of the world's great religions? Are these doctrines explicitly contradictory or only implicitly so? Just how much incompatibility is there? I won't address these questions here. For a careful treatment see Alston, *Perceiving God* (Ithaca: Cornell University Press, 1991), chapter 7. For present purposes, I will assume a straightforward (and fairly literal) reading of religious doctrines, and will assume that there is a significant amount of incompatibility between these doctrines. For a radical alternative reading of such doctrines see John Hick, *An Interpretation of Religion* (New Haven: Yale, 2004).

⁹ Or *practices*. Should we think of the problem of religious diversity as a problem that concerns the reliability of a *single* doxastic practice – say, *forming-beliefs-via-religious experience*? Or do we have a problem of *inter-practice* incompatibility? I am strongly inclined toward the latter construal, though either one would yield a pressing problem about religious disagreement. For more on this point see Alston, *Perceiving God*, chapter 7.

¹⁰ In discussion of the interview cases, much of the evidence (e.g., the candidate's appearance) is often characterized as *irrelevant*. But I take it that a similar point can be made with respect to evidence that is characterized as *misleading*.

¹¹ Though not narrow *inferential* grounds.

experiences only within her own tradition, it is doubtful that these experiences reflect the total range of extant experiential data.

We could add to these explanations. Hypotheses concerning practical benefits often figure in explanations of unreliability in the empirical case studies. And anyone familiar with the psychology of religious belief will know that similar explanations (e.g., the Freudian explanation in terms of *wish fulfillment*) are salient here. For present purposes, I'll set such explanations to the side. Even those strongly inclined to doubt the explanations we have discussed may be inclined to admit that religious beliefs seem to mirror those addressed in the studies in relevant respects. After all, even if all the *explanations* of unreliability are false, the *phenomenon* of unreliability remains. And this is enough for our purposes. Suppose I form my religious beliefs on the basis of religious experience. Similar subjects who formed their beliefs in a way similar to the way in which I formed my beliefs are *unreliable*. They have either heeded misleading evidence, or have assessed non-misleading evidence inappropriately. This situation raises egocentric epistemic questions of central interest to the Foley religious epistemologist – questions of the same shape as those that arose in light of the empirical studies discussed above. If subjects like me are not generally reliable in forming beliefs about religious matters on the basis of religious experience, does this defeat the beliefs I form on this basis? If so, is this defeat partial or complete? And what might I do to overcome it? These are the questions that will concern us in the remainder.

A. Defeat?

Let's start with the questions about defeat. Remember: I'm concerned subjects who form their beliefs on the basis of religious experience. I'm granting for the sake of argument that such beliefs are *prima facie* rational. I won't defend the strong thesis that such beliefs (or the rational efficacy of their grounds) are *completely* defeated by the evidence of unreliability in religious matters. In the face of defeat, a belief's *ultima facie* rational status depends in part on the extent to which the subject's grounds render her belief *prima facie* rational. All else being equal, the greater support the grounds provide, the harder it will be to reduce their epistemic efficacy to the point where suspending judgment becomes more rational than retaining belief. This provides reason to doubt that extremely well supported religious beliefs will fall prey to complete defeat.¹²

It does seem, however, that the evidence of unreliability provides at least a partial defeater for experience-based religious beliefs. Where one's initial grounds G render a belief B *prima facie* rational, a *partial defeater* is a reason R such that (R & G) does not render B as rational as G alone does. (That is, the conjunction (R & G) supports B to

¹² By *complete defeat* I mean simply that the defeater makes withholding or disbelief more rational than belief. Here is another way to make the point developed in the above paragraph. Experience-based religious beliefs need not be similar to each other in all epistemically relevant respects. A number of factors may be epistemically relevant to the *prima facie* status of such beliefs, including (i) the character of the religious experience itself; and (ii) the degree of self-support enjoyed by the subject's experiential doxastic practice. Experience-based religious beliefs often differ along these dimensions. Those who wish to defend the universal claim that awareness of religious disagreement *always* provides a complete defeater must show that these differences make *no* difference to the all-things-considered status of religious beliefs.

a lesser extent than G alone.) And the evidence from religious disagreement seems to fit this description. To see this, consider a subject, Bob, who forms the belief that <God was in Christ, reconciling the world to himself> via religious experience (call his belief “B*”, and his experiential grounds “G*”). Bob then gets good evidence of massive unreliability among similar subjects who formed their beliefs – which are incompatible with his – via similar methods. Call this new bit of evidence “R*”. Does the conjunction (R*&G*) support Bob’s belief, B*, as strongly as G* itself? It seems pretty clear that it does not.

If you doubt this, suppose that Bob subsequently obtained evidence against R*: all the “proponents” of religions that reject B* sign a sworn affidavit indicating that they were “just kidding” when they claimed to form beliefs incompatible with B* via religious experience. All of these subjects have actually refrained from giving B* – and any beliefs that conflict with it – any serious thought. They are instead concerned with other matters, like spoofing Bob. Would this revelation make Bob’s belief, B*, more rational than it is when his evidence consists in (G* & R*)? It would certainly seem to. But if it does, then R* fits the definition of a partial defeater – for in that case, Bob’s grounds for B* would in effect revert to G*. So, G* supports B* to a greater degree than (G* & R*). The upshot of this discussion is that awareness of widespread unreliability in experiential religious belief formation is at least a partial defeater for beliefs formed by way of religious experience.¹³ (To my mind, the only candidate exception to this rule is the believer who rationally takes her experiential grounds to be indefeasible. This believer would be analogous to the believer who rationally takes herself to have a *rational demonstration* of her religious beliefs.¹⁴ I doubt whether there are many believers of either sort.)

[*An important aside:* the above argument need not assume that parties to religious disagreements are *epistemic peers* in the sense most often discussed in the literature. In that sense of the term, subjects count as peers if and only if they (a) have the same evidence; and (b) are equally well-disposed to respond appropriately to shared evidence. At least if we are counting religious experiences as evidence, there is good reason to doubt that parties to typical inter-religious disagreements have the same evidence. So such subjects aren’t peers. Epistemological problems about religious disagreement are not typically problems of *peer* disagreement. However, it would be a mistake to infer from this that the epistemic challenge of religious diversity has been met. Indeed, the argument sketched above is an attempt to formulate a diversity-based challenge without the assumption of peerhood. Suppose you and I are party to a religious disagreement. Even if our experiential evidence differs – thus preventing us from being peers – there remain the higher-order questions of whose evidence is more likely to lead to the truth, and whose faculty of judgment is more likely to assess

¹³ Alston seems to agree with this assessment: “it can hardly be denied that the fact of religious diversity reduces the rationality of engaging in [Christian Mystical Practices] below what it would be if the problem did not exist.” *Perceiving God*, 275.

¹⁴ On this see Plantinga, *Warranted Christian Belief*, 440.

evidence appropriately. In the face of disagreement, these questions suffice to get defeat-related worries off the ground.]¹⁵

B. Defeater-Defeaters and Improving Our Epistemic Lot

Of course, defeaters – including partial ones – can often be defeated. Defeat itself is sometimes a *prima facie* matter. This raises the question of what kinds of evidence and methods are admissible as defeater-defeaters with respect to experience-based religious beliefs. It also raises the related question of what we can do to improve our epistemic situation. What might make human subjects more reliable assessors of religious experiential evidence? Our task here will be to consider whether any of the Foley strategies discussed above can be of use in the religious case.

Reformatting and Recalibration

Perhaps unsurprisingly, there is no simple reformatting strategy (as in certain probability cases) that can be applied in order to make human subjects reliable in forming beliefs on the basis of religious experience. Or if there is such a strategy, we are not aware of it. With respect to the Foley strategies discussed above, our options are down to *self-monitoring*, *recalibration*, and the *protected class strategy*. Recalibration in this case amounts to a concession that one's belief is at least partially defeated. In the face of such a defeater, rationality requires reduced confidence in the target belief. As I have indicated above, the extent to which the target belief is defeated depends on, among other things, the rational efficacy of one's original grounds. There may therefore be large differences across subjects and religious traditions in the extent to which experience-based religious beliefs are defeated in the way we have suggested. The extent to which recalibration is appropriate depends on these factors. So, I have nothing general to say here except that I don't see that there's a correct "one-size-fits-all" rule. The point just made about partial defeat tells us that recalibration is in order; it does not tell us how much recalibration is in order.

Self-monitoring and the Protected Class Strategy

At any rate, recalibration is not a strategy for defeating defeaters. As a means of avoiding defeat, some religious epistemologists suggest the *protected class strategy*. Plantinga, for instance, says this:

[L]et's suppose that rationality and epistemic duty do, indeed, require treating similar cases similarly. Clearly you do not violate this requirement if the beliefs in question are *not* on a par. And the Christian believer thinks that they are *not* on a par... [S]he must think that there is an important epistemic difference: she thinks that the other person has made a mistake, or has a blind spot, or hasn't

¹⁵ I argue for this sort of point at greater length in "Disagreement: What's the Problem? or A Good Peer is Hard to Find," *Philosophy and Phenomenological Research* (forthcoming).

been wholly attentive, or hasn't received some grace she has, or is blinded by ambition or pride or mother love or something else; she must think that she has access to a source of warranted belief [her dissenter] lacks...[T]he serious believer will not take it that we are all, believers and unbelievers alike, epistemic peers on the topic of Christian belief.¹⁶

The Christian, Plantinga thinks, will take herself to be a member of a protected class. But does she have *reason* to think this? In addressing *this* question, Plantinga considers only the case of the believer who carefully rethinks her Christian belief in the face of religious disagreement. He concludes that the subject who embarks on such reflection and still finds herself with Christian beliefs can thereby retain her beliefs without being properly subject to epistemic censure. This suggests – though it does not entail – that Plantinga's protected class strategy is closely tied to a kind of self-monitoring strategy. The protected class strategy is not alleged to work on its own. In light of this, let's consider how self-monitoring might look in the religious case.¹⁷

As we saw above, Foley's self-monitoring is one strategy for improving one's epistemic position in the face of evidence of unreliability. Perhaps this strategy can be of use in the case of experience-based religious belief. The Foley's religious epistemologist may suggest *introspection* as a starting point on this score. Applied to the religious case, this might involve looking inward to determine whether one was thinking clearly when one formed the target belief on the basis of religious experience. Was the experience highly emotional? Might this have clouded one's thought? Did one have especially strong psychological motives for forming the belief? Did one *want* the belief to be true? One might supplement this sort of introspection by formulating the epistemological problems that arise from religious disagreement in their clearest, most menacing form. Having done this, one might further introspect: in light of these problems, does the target belief still *seem* true? If so, perhaps this is at least some evidence that one is not among the benighted subjects who have been led into falsehood by misleading evidence or mistakes in reasoning.

In the religious case, as with the empirical cases discussed above, it would be unwise to stop with introspection. In the face of the interview studies, Foley argues, the best sort of self-monitoring requires one to seek publicly available evidence as a way to confirm or disconfirm one's opinion. This is epistemically wise in part because in the interview studies, one has excellent reason to think that the information gleaned from the interview is unrepresentative of the total available evidence regarding the applicant's abilities. And it is epistemically unwise to form (or retain) extremely confident beliefs on the basis of evidence one has reason to think is skewed. Something similar seems to apply to the religious case. One's religious experiences – undergone within one's own religious tradition – represent a narrow range of the total available religious experiences. It is wise to seek further evidence under such circumstances. Indeed, the Foley's epistemologist will insist on it, at least if his epistemic standards caution against forming beliefs on the basis of an unrepresentative body of evidence.

¹⁶ Plantinga, *Warranted Christian Belief*, 451, 53, 54.

¹⁷ Indeed, Plantinga takes the protected class strategy to be legitimate *downstream* of self-monitoring.

But where to look? Seeking further evidence by way of religious experiences *outside* one's own tradition is, in certain cases, regarded as *idolatrous*. Further, even if this weren't an obstacle, given the pervasive inconsistencies between systems of religious doctrine, it would seem that much of the extant evidence from religious experience is misleading – at least given the assumption that subjects aren't routinely making errors in evidence assessment. There is a kind of dilemma here. On the one hand, one's own religious experience provides an evidential base that is unlikely to be representative of the total available experiential evidence – and this provides reason to be wary of beliefs formed on an experiential basis. On the other hand, the total available experiential evidence is a very noisy data set; it points in all sorts of incompatible directions. This means that a lot of the evidence is misleading, which provides a different sort of reason for thinking that one is in poor epistemic conditions.

In light of these problems, *publicly available evidence* of the sort one finds in natural theology and historical apologetic arguments may play a crucial role in improving our epistemic situation. That is, even if epistemologists like Alston, Plantinga, and Moser have a correct picture of the *prima facie* rational status of religious belief, traditional evidences (philosophical, scientific, or historical) may play an important epistemological role as defeater-defeaters.¹⁸ In a way that is analogous to consulting public evidence in the interview cases, rational arguments for religious claims – if successful – can provide a kind of independent check on experience-based beliefs.

The above suggestion is likely to meet with opposition. Won't the relevant apologetic arguments contain premises whose rational status is subject to a disagreement-based objection of precisely the sort we have been considering? And won't that mean that these premises (and the beliefs they support) are again subject to at least partial defeat? Perhaps so. Those of us who hold religious beliefs may have to get used to the idea that our beliefs are partially defeated. But even if correct, this would not show that the pursuit of argumentative evidence is useless when applied to the problem of religious diversity. The relevant arguments may help the believer avoid complete defeat, and they may lessen the degree of partial defeat to which her beliefs are subjected. At any rate, these are possibilities whose exploration the Foleyian religious epistemologist will happily commend.

Conclusion

We began by asking whether experience-based religious beliefs can remain rational in the face of religious disagreement. In Foleyian terms: Can such beliefs survive the scrutiny of one's deepest epistemic standards, given the goal of having accurate and comprehensive beliefs? In addressing this question we saw that the Foleyian epistemologist will advise believers faced with the problem of religious diversity to seek further evidence – including *public evidence*. More modestly, the Foleyian will give such advice to subjects whose epistemic standards include the imperative to seek a body of evidence that is representative of the total available evidence. Whether their

¹⁸ Alston himself is open to this. See *Perceiving God*, 270.

beliefs would survive such scrutiny depends, in some measure, on the character of the evidence they would find if they heeded this imperative. *Ex hypothesi*, the believers in question don't yet have this evidence, so *they* are ill equipped to make judgments about its probative force. If such subjects want to determine whether their beliefs are Foley-rational, they may have a good deal of evidence-seeking to do. Perhaps I do, too.¹⁹

¹⁹ Thanks to Alex Arnold, Andrew Bailey, Nathan Ballantyne, Robert Garcia, Matthew Lee, Jeff Snapper, and Chris Tucker for helpful comments and discussion.