

# Testimony Amidst Diversity

*Max Baker-Hytch*

That testimony is one of the principle bases on which vastly many people hold their religious beliefs is hard to dispute. Equally hard to dispute is that our world contains an array of mutually incompatible religious traditions each of which has been transmitted down the centuries chiefly by way of testimony. In light of this latter it is quite natural to think that there is something defective about holding religious beliefs primarily or solely on the basis of testimony from a particular tradition. The present chapter takes up the question of what that defect consists in. I first consider whether religious diversity entails that a religious believer's testimony-based beliefs are not formed in a suitably epistemically reliable manner even conditional upon the truth of her religion. After casting doubt on this thought I turn to look at the idea that testimony-based beliefs are subject to defeaters in light of awareness of religious diversity, and I suggest that many such beliefs are not obviously so. According to my diagnosis the problem, rather, is that believers who base their religious beliefs just on testimony will be very unlikely to have reflective (that is, second-order) knowledge even if they possess first-order knowledge, and I explain why this is a notable shortcoming.

## 1. The facts of religious diversity

Before we proceed any further it will be helpful if we can characterize the facts of religious diversity in more detail. By 'religious belief-system' I shall mean a set of beliefs concerning such interconnected topics as what ultimate reality is like, who or what (if anything or anyone) is worthy of worship, and what makes for a good

human life. What we seem to find in our world is that there are a number of distinct religious belief-systems each of which is believed by very large numbers of people, and furthermore, that the contents of these various belief-systems (at least when taken at face-value) are mutually inconsistent with one another, in the sense that at most only one of those systems could be wholly true. To be sure, certain pairs of these belief-systems—for example, Judaism and Christianity—overlap quite considerably with one another, so that significant portions of the pair could simultaneously be true. At the same time, other pairs of these belief-systems—Christianity and Buddhism, for instance—overlap very little with one another.

What's more, we seem to find that there is a rather tight correlation between living in a certain part of the globe at a certain time and adhering to a certain religious belief-system. That is, rather than an even smattering of adherents of various systems throughout the earth's populations, we see something of a 'clustering' of particular religious belief-systems among the inhabitants of particular societies. An obvious explanation for this observation is that the adoption of religious beliefs by individuals very often occurs by way of social belief-forming practices in which the testimony of others, especially the testimony of adults to children, plays a central role.

It is important, finally, to distinguish between two sorts of religious beliefs that are held by the adherents of most of the world's religions. On the one hand there are those settled doctrinal commitments such as the belief that God is three persons in one or that the scriptures of one's religion were revealed to a certain individual at a certain historical time and place. On the other hand, there are those beliefs concerning, say, how God is acting in one's life or the world more generally at a given time or what God's will for one's life is in a given situation. (Notably, not all religious belief-systems make room for the having of beliefs of the latter sort.) It is the settled doctrinal commitments rather than the latter sort of beliefs that are frequently acquired via testimony and socialisation. Beliefs of the latter sort, I take it, are generally acquired by way of certain experiences rather than via testimony, though it is plausible that the interpretation (and perhaps even the initial content) of such experiences can be shaped by an individual's prior doctrinal commitments.

## 2. Religious diversity and epistemic (un)reliability

Do the facts of religious diversity show that someone's testimony-based religious beliefs are not formed in a reliable manner, even supposing those beliefs to be true? If so, given that reliability of some sort is very plausibly a necessary condition for knowledge,<sup>1</sup> it follows that testimony-based religious beliefs fail to be knowledge even

---

<sup>1</sup> I intend 'reliability' here to cover the sort of modal reliability that is lacking in Gettier cases.

if true—a fairly damning indictment. Something along these lines may well be what J. S. Mill had in mind when he wrote that

[M]ere accident has decided which of these numerous worlds is the object of [a religious person's] reliance... [T]he same causes which make him a churchman in London would have made him a Buddhist or a Confucian in Peking. (1991 [1859]: 229-30)

Mill's thought here might be that given that the very same (type of) belief-forming process would lead people to hold different and mutually incompatible sets of religious beliefs depending on the place-time they inhabit, it would be a matter of sheer luck were such an unreliable process to lead someone to acquire a set of true religious beliefs.<sup>2</sup>

Epistemologists have identified a number of distinct kinds of epistemic reliability and there remains some disagreement over which kinds (if any) are necessary for knowledge. Rather than take a stand on this issue I'll consider the three kinds of reliability that have attracted most of the discussion in recent decades. Note that the notion of a belief's 'basis' which features in the following definitions is to be understood in terms of a token causal process which generates and sustains the belief.<sup>3</sup>

SENSITIVITY. S's belief that  $p$  is sensitive IFF: if  $p$  had been false then S wouldn't have believed  $p$  on the basis on which she actually believes  $p$ .<sup>4</sup>

SAFETY. S's belief that  $p$  is safe IFF: S couldn't easily have formed a false belief in a proposition relevantly similar to  $p$  on a basis relevantly similar to the one on which she actually believes  $p$ .<sup>5</sup>

STATISTICAL RELIABILITY. S's belief that  $p$  is statistically reliable IFF: the salient belief-forming process type instantiated by S's belief that  $p$  is such as to

---

<sup>2</sup> This is how Tim Mawson (2009) interprets Mill's argument. Alvin Plantinga suggests that Mill's argument, with its invocation of accidental truth, 'seems to be designed to appeal to reliabilist intuitions' (1995: 211).

<sup>3</sup> For arguments in favour of construing a belief-forming basis (or equivalently, a belief-forming 'method') in terms of the actual facts about what causes the belief, rather than in terms of how it appears to the agent 'from the inside' that her belief arose, see Williamson (2000: 155-56).

<sup>4</sup> Fred Dretske (1971) and Robert Nozick (1981: 172-79) have argued for accounts of knowledge that require sensitivity.

<sup>5</sup> Timothy Williamson (2000) and Duncan Pritchard (2012) have argued for accounts of knowledge that require safety.

yield a preponderance of truth over falsehood across its instances in the actual world and nearby possible worlds.<sup>6</sup>

It's worthy of note that sensitivity and safety are factive, in the sense that, necessarily, if a belief is false then it does not possess either of these properties.<sup>7</sup> Statistical reliability, by contrast, isn't factive, for it can be the case both that a belief held on a given basis is false and that the type of belief-forming process which the belief's basis instantiates is one that leads to true beliefs in the great majority of cases.

In considering whether testimony-based religious beliefs can be reliable in any of the three senses just outlined, given the facts of religious diversity, we can consider the fictional case of a fairly typical religious believer—we'll call her Jane—who has come to hold her beliefs in the central doctrines of Christianity primarily on the basis of the testimony of her parents and religious teachers and on the basis of the testimony of the Bible.

### *Sensitivity*

Let's begin by considering whether Jane's testimony-based doctrinal beliefs are sensitive. Since sensitivity is factive, no such belief is sensitive if false. If true, whether her testimony-based beliefs are sensitive depends on whether she would still have had the testimonial basis she actually has for those beliefs in the nearest possible world(s) in which the contents of those beliefs are false.

Now, it's a slightly tricky matter to think about the sensitivity of beliefs that are *necessarily* true if true at all, and plausibly some of Jane's doctrinal beliefs will be of this character. Her belief that God exists is one such; her belief that God is triune is another. If these propositions are true, then the nearest worlds in which they are false will be metaphysically impossible worlds. Still, some philosophers are inclined to think that we can sensibly talk about what goes on in metaphysically impossible worlds—indeed, that we *need* to be able to do so—and that the usual similarity metric can be deployed in order to determine whether one metaphysically impossible world is closer than another to the actual world.<sup>8</sup> Running with this thought, let's consider Jane's belief that God exists. Supposing for the sake of the argument that God exists in the actual world, in the nearest world in which God doesn't exist, would

---

<sup>6</sup> Alvin Goldman (1979) has defended an account of justification that requires statistical reliability.

<sup>7</sup> It is impossible for a belief that  $p$  to be both sensitive and false for the following reason: if it is *actually* the case that not- $p$  and that  $S$  believes  $p$  on a basis  $B$ , then it follows (given the usual Lewis-Stalnaker semantics for subjunctives, according to which  $[p \& q] \rightarrow [p \square \rightarrow q]$ ) that *if it were the case* that not- $p$  then  $S$  would believe  $p$  on the basis of  $B$ , hence the belief is insensitive. It is impossible for a belief that  $p$  to be both safe and false for the following reason: if it is *actually* the case that  $p$  is false and that  $S$  believes  $p$  on a basis  $B$ , then it follows that  $S$  *could easily* have believed a false proposition relevantly similar to  $p$  (namely,  $p$  itself) on a basis relevantly similar to  $B$  (namely,  $B$  itself), hence the belief is unsafe.

<sup>8</sup> See Nolan (1997).

Jane still have her testimonial basis for believing God exists? It is tempting to think that she wouldn't, and the reason she wouldn't is that she, and everyone and everything else, wouldn't exist either. Note that this thought doesn't assume that a world like ours could only exist if God exists. It simply assumes that the worlds in which God doesn't exist but which are most similar to a world in which God exists and created the physical universe (similar in as many ways as are consistent with God not existing) will be worlds in which there isn't a physical universe. If in the actual world the coming into being of the physical universe was caused by God and nothing else, then subtracting God whilst holding fixed as much else as possible should yield the result that no physical universe comes into being. We should conclude, then, at least tentatively, that Jane's testimony-based belief that God exists is sensitive if true.

Let's consider another of Jane's beliefs that I suggested is necessarily true if true at all, namely, the belief that God is triune. Assuming for the sake of the argument that in the actual world there exists a triune God, how do things go in the nearest world in which God is not triune? Well, plausibly, that world isn't an atheistic one. That is, it seems that a world in which a triune God exists is much closer to a world in which there exists a deity of some sort but who lacks the property of being triune than it is to a world in which no deity exists at all. Would Jane still have her testimonial basis for believing that God is triune in the world in question? Well, that depends whether the doctrine of the trinity would still have arisen and been testimonially transmitted down through the ages in the worlds in which God is not triune but which is otherwise most similar to a world in which God exists and is triune (similar in as many ways as are consistent with God not being triune).

If in the actual world the proximate cause which brought it about that the doctrine of the trinity arose among the early Christians and was propagated down through history to people like Jane was the special activity of a triune God, then it is plausible that the nearest world in which God doesn't have the property of being triune but as much else as possible remains the same will be a world in which the doctrine of the trinity doesn't arise (because God refrains from communicating it to people) and hence it doesn't get propagated down through history to people like Jane. If that is how things are, then Jane's belief in the trinity is sensitive if true.

On the other hand, if in the actual world a triune God exists but engaged in no special causal activity in order to give rise to the belief in the doctrine of the trinity, but instead simply looked on with approval as a certain group of first-century Jews correctly hit upon that doctrine, then it would seem that in the nearest world in which God isn't triune, the doctrine of the trinity does still arise and get propagated down through history to people like Jane. If that is how things are, then Jane's belief in the trinity is not sensitive even if it is true. Similar points will apply to Jane's beliefs in other doctrines: whether those beliefs are sensitive if true will depend on the causal

mechanism by which those doctrines came to be initially believed and subsequently transmitted down through history, and in particular, on whether God was causally involved in some sense that goes beyond merely sustaining the world. And it does seem rather plausible that given the truth of a monotheistic belief-system like Jane's, God would indeed have been directly causally involved in the dissemination of important truths about himself.

### *Safety*

Again, safety is a factive property of a belief and so no belief of Jane's that is false is safe. If a testimony-based religious belief of Jane's is true, then whether it is safe depends on whether the testimonial basis for Jane's belief (or a basis relevantly similar to it) could easily have led Jane to hold a false belief with relevantly similar content to the belief she actually holds.

Now, one might think that Jane's testimony-based beliefs are unsafe (even if true) in virtue of the very plausible counterfactual according to which, if Jane had been raised in a significantly different religious environment then she would have held religious beliefs that conflict with her actual ones, and would have held those beliefs on a basis that is relevantly similar to her actual basis. We might add that the contents of the beliefs Jane would have held in those counterfactual circumstances are relevantly similar to the contents of her actual religious beliefs. This sort of counterfactual has been thought by many a religious sceptic to show that it would, in some epistemically troubling sense, be merely lucky if a religious believer were to end up with true religious beliefs.<sup>9</sup> Safety is a property, according to many epistemologists, which is lacked by beliefs that are merely luckily true. One might think, then, that we can build an argument for the unsafety of Jane's testimony-based religious beliefs that takes the aforementioned counterfactual claim as a key premise. Such an argument would run like so:

- (1) If Jane had been raised in a significantly different religious environment, she would have held religious beliefs that conflict with her actual beliefs, and would have done so on a basis relevantly similar to her actual basis.
- (2) Jane could easily have been raised in a significantly different religious environment.
- (3) IF: if  $p$  were true then  $q$  would be true, and it could easily have been the case that  $p$ , THEN: it could easily have been the case that  $q$ .

---

<sup>9</sup> Philip Kitcher, for instance, writes that 'Had the Christians been born among the aboriginal Australians, they would believe, in just the same ways, on just the same bases, and with just the same convictions, doctrines about Dreamtime instead of about the Resurrection' (2011: 26).

- (4) Jane could easily have held religious beliefs that conflict with her actual beliefs, on a basis relevantly similar to her actual basis.
- (5) If (4), then Jane's actual religious beliefs are not safe (even if true).
- (6) Jane's actual religious beliefs are not safe (even if true).<sup>10</sup>

Such an argument faces a serious difficulty, however. The difficulty concerns the liability of the argument to over-generalize in ways that lead to an implausible scepticism about testimonial knowledge in non-religious contexts.<sup>11</sup>

If we deem (2) to be true then for the sake of consistency we ought to think that it is very often true of recipients of mundane, non-religious testimony that they *could easily* have been in somewhat different circumstances in which they would have received misleading testimony and would have accepted it. (2) can be read as claiming that the world in which Jane is raised in a significantly different religious environment, and hence receives significantly different religious testimony, is a fairly close world—close enough to be relevant to the safety of her beliefs. But if we affirm this claim, then we should affirm a parallel claim about many ordinary, non-religious testimony cases. Take a case in which a tourist asks for directions from Waterloo train station to Westminster Abbey from a knowledgeable-looking local and receives correct directions.<sup>12</sup> Plausibly such cases often result in knowledge. But if we are saying that the world where Jane was raised in a wholly different cultural and religious outlook is a *close* world, then we must surely judge to be *very close* the world in which the tourist asks a different, equally knowledgeable-looking local for directions but winds up with misleading directions.

The closeness of a world in which one falls into error about a similar proposition isn't enough by itself to undermine the safety of one's belief, however. It also needs to be the case that the basis on which one would fall into error in that close world is relevantly similar to one's actual basis. Now, (1) claims that in the world in which Jane receives a significantly different religious upbringing, she comes to hold religious beliefs that conflict with her actual beliefs and does so on a basis that is relevantly similar to her actual basis. But once again, if we are willing to count as relevantly similar the basis that Jane employs in that world—a basis that involves (inter alia) receiving testimony from a different chain of individuals and a different set of religious scriptures—then for the sake of consistency we ought to count as *very similar* the basis that the tourist would have employed in the very close world in

---

<sup>10</sup> Thanks to John Hawthorne for suggesting this formulation.

<sup>11</sup> For a rather different sort of response to a safety-based construal of a sceptical argument from religious diversity, see Bogardus (2013: 379-82).

<sup>12</sup> This is very similar to a case discussed by Jennifer Lackey (2007: 352), which she judges to be a case of knowledge.

which she asks a different, equally knowledgeable-looking local and gets wrong directions to Westminster Abbey.

The case of the tourist asking for directions is hardly unique among ordinary testimony cases in being such that the recipient of testimony need not have been in a greatly different environment than the one she was in in order to have received misleading but seemingly equally credible testimony. The upshot, then, is that the above argument requires interpretations of modal closeness and of relevant similarity of basis which, if applied consistently, will result in excessive scepticism concerning a range of ordinary testimony cases. In short, if we want to avoid such scepticism, then we should deny either (1) or (2) of the above argument.

With all that said, perhaps there is a rather different way to show that testimony-based religious beliefs are unsafe even if true. Thus far we have been examining whether Jane's testimony-based religious beliefs are unsafe given the supposition that they are *all* true. Presumably, though, virtually no religious believer will have a set of testimony-based religious beliefs that are *all* true, even supposing that the believer in question is an adherent of a religious belief-system whose central doctrines are true. As we have been construing it, safety requires that one couldn't easily, on a relevantly similar basis, have held a false belief whose content is relevantly similar to the content of one's actual belief. One fails to satisfy this requirement if one *actually* holds a false belief that is relevantly similar in content to the target belief and does so on a very similar basis. For example, suppose one hears a particular news anchor testify that *p* and that *q* and that *r*, where these are all propositions about a certain plane crash that has just occurred. Suppose that *p* and *q* are true but *r* is false, and suppose that one believes all three propositions on the basis of the news anchor's testimony. According to the way of construing safety with which we are working, one's belief in *p* is unsafe despite being true: one *actually* believes (and therefore, could easily have believed) a false proposition similar in content to *p* and on the same (and hence, a very similar) basis to the basis on which one believes *p*. Analogously, the thought might be that even a religious believer who has arrived at a largely but not entirely true set of religious beliefs on the basis of testimony will not enjoy the status of safety for those true beliefs, in virtue of the fact that her testimonial source has also given her a few false religious beliefs—a few 'bad companions', as it were—which count as having content that is relevantly similar to that of her true beliefs. Presumably this line of thought will apply *a fortiori* to religious believers who don't enjoy nearly such a high ratio of truth to falsehood among their testimony-based religious beliefs.

Again, though, it seems to me that the way that safety needs to be interpreted for this line of thought to succeed will lead to implausible scepticism. Just consider cases in which parents impart 'general knowledge' to their children. Bill, for instance,

loves to teach his daughter Gracie about world geography and is for the most part an excellent guide on such topics as names of capital cities, oceans, active volcanoes, mountains, and so on. Even the greatest parents just occasionally err, though, like when Bill told Gracie that the state capital of Kentucky is Louisville (rather than Frankfort). Will anyone really want to say that Gracie therefore gets to know *nothing* about world geography by way of her father's testimony? It appears that if we want to avoid such excessive scepticism, safety will have to be interpreted in a way that doesn't require one not to have *any* false beliefs about the same topic as the target belief.

In sum, if they are mostly true, then Jane's testimony-based religious beliefs should be deemed safe. Securing the opposite verdict requires interpretations of safety's crucial moving parts which lead to implausible scepticism about many ordinary, non-religious cases of testimony.

### *Statistical reliability*

Statistical reliability, recall, is a property a belief possesses in virtue of exemplifying a type of belief-forming process that generates true beliefs in the majority of cases in which it is exemplified in the actual and nearby worlds. Let's consider again Jane, and whether her testimony-based religious beliefs possess this property. The following is an argument for the conclusion that they don't:

- (7) The salient belief-forming process type exemplified by Jane's testimony-based religious beliefs is one that is also exemplified by adherents of religious belief-systems which are mutually inconsistent with Jane's belief-system.
- (8) If a process type is exemplified by many mutually inconsistent token beliefs, then that process type is not statistically reliable.
- (9) The salient process type exemplified by Jane's testimony-based religious beliefs is not statistically reliable.

What should we make of such an argument? (8) is hard to argue with. (7), on the other hand, seems to require that we adopt a very coarse-grained approach to classifying the process type to which Jane's beliefs belong. Specifically, in order for (7) to come out true it needs to be the case that the salient process type in the case of Jane's beliefs is a process type whose description involves no mention of the particular testimony chains that led to Jane's belief or the particular religious texts she consulted. In short, the description will have to be about as coarse-grained as 'trusting religious testimony'. This is a process type that is indeed exemplified by testimony-based beliefs of adherents of all manner of mutually inconsistent religious belief-systems across the globe. If this is the right way to individuate process types for the purposes of ascertaining statistical reliability, then the above argument is sound.

But such a coarse-grained approach to individuating process types suffers from a number of problems. An especially difficult one is known as the ‘no distinction’ problem. A pair of cases will serve to illustrate the problem. Tom looks at a nearby tree through a clear windowpane in good lighting and without being under the influence of any impairing substances, and correctly believes that there is a tree before him. Tim looks at a nearby tree through a filthy windowpane in poor lighting whilst heavily intoxicated, and correctly believes that there is a tree before him. It is plainly obvious that Tom forms his belief in a much more reliable fashion than Tim. But if belief-forming process types are individuated in so coarse-grained a manner—coarse enough to count the adherents of a wide variety of religious belief-systems as employing the same type of belief-forming process as one another—then we won’t be able to capture the intuitively obvious difference in statistical reliability between Tom’s and Tim’s beliefs. It looks like such an approach will count Tom and Tim as employing the same type of belief-forming process as one another, something like ‘visually perceiving nearby objects’, and so they will mistakenly be counted as exhibiting the same degree of statistical reliability as one another.<sup>13</sup>

For this and other reasons coarse-grained approaches to typing processes, of the sort that would be needed in order to get (7) to come out true, are out of favour. Now, the consensus among epistemologists seems to be that the generality problem for reliabilism (the problem of how to non-arbitrarily assign token beliefs to process types) hasn’t yet been fully resolved. But one recurring suggestion that does seem plausible is that our assigning of tokens to types should be constrained by the following consideration: in assigning a token belief to a particular process type we should pick a type whose description makes reference to those factors that are causally operative in producing the token belief in question, and only to those factors.<sup>14</sup> One reason that this suggestion may not be fully sufficient to solve the generality problem is that there remains a difficulty about how narrowly or broadly we are to describe the causal factors that are operative. Still, it seems that this suggestion should at least help to narrow down significantly the range of process types under which we might classify a given belief. So, for instance, if the fact that the sun is shining outside has no causal bearing at all on Emma’s coming to believe that there is a mug on her desk, then we shouldn’t assign Emma’s belief to a process type whose description includes anything to do with whether the sun is shining. On the other hand, if Emma has a rare disorder whereby her vision only works well when she looks at objects that are

---

<sup>13</sup> As Earl Conee and Richard Feldman note, ‘the reliabilist needs the relevant type for each belief to be sufficiently narrow to include only beliefs that are equally well justified’ (2002: 102).

<sup>14</sup> A number of epistemologists who claim to have solved the generality problem have been guided by this thought. See for instance Becker (2008), Alston (1995), and Goldman (1986: 50).

illuminated by sunlight rather than by artificial light, then we should assign her belief to a process type whose description makes mention of whether the sun is shining.

Another parameter that is relevant to how we classify beliefs under process types is that of a process type's *depth*. A type's depth is a matter of how long a stage of the causal process terminating in the production of a belief gets taken into account in the description of the type.<sup>15</sup> For example, the type 'receiving testimony from a friend who got his or her information from a national news broadcast' is a deeper type than merely 'receiving testimony from a friend'. It is plausible that when we are considering the statistical reliability of beliefs that are the result of accepting testimony that has passed along a chain of testifiers, we need the description of the relevant process type to be sufficiently deep so as to mention something about the sorts of causal factors that generated belief in the individual (or individuals) at the beginning of the testimony chain. Here is a plausible contrast: the salient testimonial process type exemplified by South Korean schoolchildren's beliefs about twentieth-century history is much more statistically reliable than the salient testimonial process type exemplified by North Korean schoolchildren's beliefs about twentieth-century history. But unless we individuate the testimonial process types with reference to the sorts of causal factors that are operative at the beginning of the respective testimony chains, we won't have any grounds for counting the South Korean schoolchildren's beliefs as exemplifying a different (salient) process type than that which is exemplified by the North Korean schoolchildren's beliefs, and hence won't have any grounds for counting one set of beliefs as exemplifying a much higher degree of statistical reliability than the other. In the North Korean case, the testimony chains that eventually result in the schoolchildren's twentieth-century history beliefs trace back in significant part to fabrications by senior figures in the political establishment, motivated (*inter alia*) by a desire to instil in the populace a hatred of other nations with whom North Korea is at war. In the South Korean case, the testimony chains that eventually result in the schoolchildren's twentieth-century history beliefs trace back (by and large) to veridical first-hand experiences of the historical events in question. The process types in these cases thus ought to be individuated by reference to these originating events.

Guided by the foregoing considerations about typing belief-forming processes, let's consider again Jane's testimony-based religious beliefs. What happened at the beginning of the testimony chain that resulted in Jane's belief is crucial. In particular, were the doctrinal beliefs generated in the individuals at the beginning of the testimony chain the proximate result of genuinely revelatory events brought about by special divine action? Or were those beliefs the result of fabrication or confusion or the like? If they were the result of direct divine activity, then given what

---

<sup>15</sup> For more on the issue of a type's depth, see Swinburne (2001: 13-20).

we have observed about the importance of individuating testimonial process types with reference to the causal factors operating at the beginning of a testimonial chain, it seems that the relevant process type here should be one whose description makes mention of such special divine action. And if the salient process type exemplified by Jane's testimony-based doctrinal beliefs is something like 'accepting religious testimony from a chain of testifiers tracing back to revelatory events that were the product of special divine action', then it *won't* be the case—contrary to (7)—that the salient belief-forming process type exemplified by Jane's testimony-based religious beliefs is one that is also exemplified by adherents of religious belief-systems which are mutually inconsistent with Jane's belief-system. Again, it is quite plausible that given the truth of a monotheistic belief-system like Jane's, God would indeed have been directly causally involved in initiating those testimony chains that eventually resulted in Jane's beliefs.

In sum, I have suggested that if Jane's testimony-based religious beliefs are mostly true, then contrary to the allegation under consideration, they likely do exemplify reliability of various important kinds. Let's now turn to consider a rather different way in which such beliefs might be thought to be wanting in light of the facts of religious diversity.

### 3. Defeaters arising from the facts of religious diversity

A natural thought is that religious belief held on the basis of testimony alone is deficient because awareness of the facts of religious diversity can so readily generate a defeater for such beliefs. Following John Pollock (1986), it is usual to distinguish between two sorts of defeaters: rebutters and undercutters. A rebutting defeater for one's belief that  $p$  consists of some reason one acquires for thinking that  $p$  is false. An undercutting defeater for one's belief that  $p$  consists in some reason one acquires for thinking that one's belief is not reliably connected to the truth—that its basis is unsafe or insensitive or statistically unreliable or perhaps something else still.

#### *Rebutting defeat*

Let's consider first whether learning of the facts of religious diversity might generate a rebutting defeater for testimony-based religious beliefs. Do the facts of religious diversity constitute evidence for the falsity of some religious belief-system or other? Plausibly those facts do indeed disconfirm some such systems.<sup>16</sup> In particular, the facts of religious diversity appear to constitute strong disconfirmatory evidence

---

<sup>16</sup> Stephen Maitzen (2006) has argued that that these facts strongly disconfirm theistic religious belief-systems. For responses, see Baker-Hytch (2015) and Mawson (2012).

against religious belief-systems which involve the conjunction of the following two claims: (i) that there exists a perfectly loving God, and (ii) that God will mete out post-mortem punishment upon any human individual (or at least any human individual cognitively capable of religious belief) who fails to have a fairly comprehensive set of true religious beliefs by the end of his or her earthly life. The facts of religious diversity seem quite improbable conditional on (i) and (ii) because it is highly expectable that a God of the sort described by these two propositions would arrange the world in such a way as to give more-or-less everyone a realistic opportunity to acquire a comprehensive set of true religious beliefs during their earthly lifetimes, contrary to the facts that actually obtain. On the other hand, it is considerably less obvious that the facts of religious diversity have disconfirmatory force against religious belief-systems that deny either (i) or (ii).

As for belief-systems that affirm (i) but deny (ii), there is significant scope for appealing to a wide array of reasons God might plausibly have for permitting significant religious diversity (for a time, at least). God might well value such diversity intrinsically, or he might value the spreading of the correct religious belief-system by way of the sorts of social belief-forming practices that humans rely on in so many spheres of life and whose employment, plausibly, has some considerable intrinsic value.<sup>17</sup> Belief-systems that deny (i) but affirm (ii), if any there be—that is, belief-systems that affirm the existence of a non-loving God who damns all individuals who die lacking a comprehensive set of correct religious beliefs—are not disconfirmed in virtue of the facts of religious diversity. The reason is simply that if God is not loving then there isn't any reason to think that he would be likely to give everyone an equal shot at salvation. And as for belief-systems that deny both (i) and (ii), neither are these systems disconfirmed by the facts of religious diversity, for if ultimate reality is neither loving nor such as to punish people who fail to have comprehensively correct religious beliefs at the time of their deaths then there is no reason to expect the world to be arranged so as to ensure that most people end up with correct religious beliefs. All in all, then, I find it doubtful that the facts of religious diversity constitute a rebutting defeater for adherents of religious belief-systems that deny either (i) or (ii).

### *Undercutting defeat*

Let's now consider whether learning of the facts of religious diversity might generate an undercutting defeater for a testimony-based believer. Undercutting defeaters, recall, are reasons for thinking that one's belief that  $p$  is held on an unreliable basis. If one learns that the red-looking widget is being illuminated by trick lighting that would make a widget of any colour appear red, one learns, among other things, that one could very easily have formed a false belief about the widget's colour given the

---

<sup>17</sup> See Baker-Hytch (2015).

way that one formed one's belief that the widget is red. If one learns that the passerby from whom one obtained directions to Westminster Abbey is in fact a notorious prankster who takes delight in misleading tourists, one thereby learns, among other things, that the testimonial source upon which one based one's belief that Westminster Abbey is on the next street is a statistically unreliable source.

But how exactly is it that becoming apprised of the facts of religious diversity could amount to learning that one's testimony-based religious beliefs are likely not to be reliably connected to the truth? I suggest in the following way: in learning that there exist a plurality of other testimonially-transmitted religious belief-systems that conflict with one's own, one thereby learns that there are testimonial religious belief-forming processes that are relevantly similar to one's own—relevantly similar, that is, when we consider things just from the internal perspectives of agents—and yet which yield beliefs that contradict one's own. Put another way, for all the testimony-based religious believer is able to tell 'from the inside', her epistemic situation is fully symmetrical to that of testimony-based believers in other religious traditions whose beliefs contradict hers, and so she has no right to think that her testimonial process is any more reliable than those used by adherents of conflicting religious traditions.<sup>18</sup>

One worry about this line of thought concerns its potential to generate scepticism beyond the religious realm. The underlying principle which is at work here appears to be something like this: if one has good reason to think that one's own epistemic situation is relevantly *internally* similar to the situation of an agent who is forming her beliefs in an unreliable manner (or so one believes), then one is not entitled to think that one's own beliefs are being formed in any more reliable a manner. Call this 'Principle I'. A worry about Principle I is that radically sceptical possibilities involving deceiving demons or scientists who envat people's brains and feed them illusory experiences also exhibit this very feature of involving an agent whose beliefs are being formed in a highly unreliable fashion but whose situation is internally indistinguishable from ours. In short, there is a concern that if Principle I holds then we will not be entitled to think that we ourselves are not subject to the sorts of radical sceptical scenarios to which I just alluded. Someone might respond that there is an important difference between radically sceptical scenarios and the facts of religious diversity: namely, that the former are mere *possibilities* (remote ones at that) whereas the latter are actual. Accordingly, the thought might go, Principle I comes into force only when one has reason to think that there are *actual* agents who are in relevantly internally similar situations to oneself and who are forming their

---

<sup>18</sup> Sandy Goldberg seems to be thinking along these lines when he characterizes the problem posed by religious diversity in the following way: '[N]o one is in a position to reliably determine, of the process-type they use on a given occasion, whether it is reliable: everyone regards her own process as reliable, even though a majority are wrong about this, and no one can tell the difference "from the inside" between the various types (or their reliability profiles)' (2014: 290).

beliefs unreliably. But this response falters if we do in fact have some reason to suppose that there are actual agents with mental lives internally indistinguishable from our own but who are radically in error, for example, the so-called ‘Boltzmann brains’ postulated by some multiverse cosmological theories: lonely, free-floating minds which emerge temporarily from fluctuations out of cosmic chaos. In any case, note that this response depends on the assumption that we ourselves are not victims of radically sceptical scenarios. How are we entitled to this assumption? It is not very satisfying to reply that we just are. A better answer is that we have evidence that we are not such victims, evidence that is lacked by victims of evil demons or mad scientists or the like.<sup>19</sup> But if this is right then an externalist notion of evidence must be in play: a notion of evidence according to which one’s evidence does not merely consist of that which is internally accessible to an agent through introspection but additionally of facts about the external world to which the agent is reliably connected. And if such externalism can be invoked in order to buttress anti-sceptical assumptions, then it isn’t clear why it can’t also be invoked on behalf of a testimony-based believer whose beliefs are true and are in fact reliably formed: she has evidence that her counterparts who follow conflicting religions do not.

But even if we set aside concerns about principle I, there are at least a couple of ways in which a testimony-based religious believer might justifiably judge that her epistemic situation is not symmetrical with that of believers in conflicting traditions. For one thing, it is very often the case that people who base their religious beliefs principally upon testimony have received that testimony from individuals whom they know personally and trust. Why is this significant? Well, when one personally knows and trusts someone else, one typically has opportunities to observe significant portions of the track record of the other person’s testimony. By contrast, one does not typically have such opportunities when it comes to individuals whom one does not personally know and trust. It would seem, then, that many religious believers have at least some degree of entitlement to prefer the testimony of those whom they personally know and trust to the testimony of those whom they don’t.<sup>20</sup>

Secondly, consider the following conditional: if religion R were true then the testimony chains on which (many) adherents of R base their beliefs would be reliable (safe, sensitive, statistically reliable). A claim of this form is much more likely to be true of some religious traditions than others, and this is something that can be ascertained independently of knowing which religion (if any) is true. The way it can be ascertained is by reflecting on how probable it is, conditional upon the truth of a given religion R, that there has been some kind of special divine action which ensured

---

<sup>19</sup> Williamson (2000: Ch. 8) offers such a response to the problem of radical scepticism.

<sup>20</sup> For a defence of the thought that personal knowing can put one in a privileged position with respect to gaining knowledge from the testimony of another, see Benton (ms.).

the reliable testimonial dissemination of the core tenets of R. Such a thing will be probable only given that there is an agency who is sufficiently powerful as to be able to engage in such action and ensure its success, who is sufficiently knowledgeable about how to engage in such action, and who is sufficiently morally good as to want to try to communicate truths about itself among humans. Given the truth of a religion which entails only the existence of a deity who is significantly limited in one or more of these respects or no deity at all, it isn't probable that there has been special divine action which has succeeded in ensuring the reliable dissemination of the core tenets of that religion. It looks, then, as though the only religions whose truth makes it significantly likely that there has been divine action of the aforementioned sort are those religions whose truth entails the existence of an extremely powerful, knowledgeable, and morally good deity—in other words, monotheistic religions. Testimony-based believers in monotheistic religions have available to them a significant symmetry breaker when they compare their situations to those of testimony-based believers in religious traditions whose truth entails at most the existence of a limited deity (or deities), a symmetry-breaker of the following form: if my belief that  $p$  is true then it is likely that my belief that  $p$  was reliably formed; (even) if your belief that not- $p$  is true, it is not likely that your belief that not- $p$  was reliably formed. Admittedly, this latter sort of consideration doesn't apply between adherents of traditions that make it equally likely their respective testimony chains are reliable, and so insofar as such believers really are required to regard their situations as symmetrical and insofar as they are therefore required to think of themselves as no more likely to be using a reliable process than are those who appear to be in a situation symmetrical to their own, they ought to become agnostic as to whether the portions of their belief-systems over which they differ have been reliably transmitted to them.

#### 4. Reflective knowledge

So what (if anything) is deficient about religious belief based on testimony alone? I've argued that if true, the testimony-based beliefs of at least some religious belief-systems can satisfy a range of reliability conditions. I've suggested, moreover, that defeat is not inevitable for many religious believers who base their beliefs just on testimony and come to be aware of the facts of religious diversity. In my view, the deficiency lies principally in this: a person's religious beliefs which are based just on testimony may amount to knowledge if they are true, but even if they do she very likely won't know that they do. That is, despite knowing, she won't know that she

knows.<sup>21</sup> Ernest Sosa (2009) has distinguished between mere animal knowledge, which is essentially reliably true belief, and reflective knowledge, which involves a knowledgeable perspective on one's reliability.<sup>22</sup> My contention is that religious believers who base their first-order religious beliefs on testimony alone—and by 'basing their religious beliefs on testimony alone' I mean that they do not employ any (non-circular) arguments for the reliability of their testimonial sources—such believers will not be able to rise to the level of reflective (i.e. second-order) knowledge. First allow me to explain why I think that that this is so, before going on to outline why a lack of reflective knowledge in the religious domain should strike us as problematic.

Consider the belief-forming process by which someone who bases her first-order religious beliefs on testimony alone might come to believe that those first-order beliefs amount to knowledge. Let's consider Jane once again. Jane bases her first-order religious beliefs on testimony alone and doesn't possess any arguments for the reliability of her testimonial sources. Her second-order beliefs concerning the epistemic status of her first-order beliefs will then presumably be the product of some sort of default disposition to assume that her testimonial sources are reliably connected to the truth. Let's suppose that Jane's first-order beliefs are indeed reliably formed (safe, sensitive, statistically reliable), as I argued they likely would be if true, and that given their truth they amount to first-order knowledge. Are there any good reasons for assigning the token processes responsible for Jane's second-order beliefs to a type that is *not* also exemplified by the adherents of conflicting religious belief-systems who also come to believe that their testimony-based religious beliefs amount to knowledge by simply relying on a default disposition to assume that their testimonial sources are reliably connected to the truth? I would suggest not.

Recall how I suggested that there are good reasons for typing testimonial belief-forming processes in a way that takes account of the sorts of causal factors that were operative at the origins of a testimony chain. The upshot of this suggestion was that those religious believers (if any there be) who are in receipt of testimony that traces back to revelatory events caused by special divine action should be counted as employing a different testimonial process type from that which is employed by believers who are at the receiving end of a testimony chain with no such special origins—after all, the causal factors operative at the beginning of the former sort of testimony chain are very different from those that are operative at the beginning of the latter sort of chain. This sort of consideration, however, doesn't seem to apply to

---

<sup>21</sup> Defenders of the KK principle will deny the possibility of knowing whilst not knowing whether one knows. For a general argument against the KK principle, see Williamson (2000: Chap. 5).

<sup>22</sup> Sosa writes that 'we can more generally distinguish animal knowledge, which requires only that one track reality, on one hand, and reflective knowledge, on the other, which in addition requires awareness of how one knows, in a way that precludes the unreliability of one's faculties' (1997: 427).

Jane's second-order beliefs about the epistemic status of her first-order beliefs. The causal factors operative in producing Jane's second-order beliefs—namely, some sort of default disposition to think that her testimonial sources are reliably connected to the truth—are roughly the same sorts of causal factors as are operative in generating the second-order beliefs of those religious believers whose first-order beliefs are not in fact knowledge. There seem to be no grounds for assigning Jane's process to a different type than the one to which we assign the processes responsible for the second-order beliefs of believers in conflicting religious traditions who base their first-order beliefs on testimony alone. In light of this, despite having first-order knowledge by way of testimony, Jane lacks reflective (i.e., second-order) knowledge.

Now, someone might object to the foregoing claim by pointing out that given the truth of Jane's religious belief-system, God would be rather likely to be involved not just in setting up the testimony chain that led to Jane's first-order beliefs but also in prompting Jane directly to form the second-order belief that her first-order beliefs amount to knowledge.<sup>23</sup> If this thought is correct, though, it looks as though we have a case of causal overdetermination. For Jane presumably *also* possesses that default disposition to think that her testimonial sources are reliably connected to the truth. If God had failed to prompt Jane to form the second-order belief that her first-order beliefs are knowledge, Jane would still have formed *that very belief*, and would have done so as a result of her default disposition to do so (the same sort of disposition that many adherents of false religious traditions employ to form the second-order belief that their first-order beliefs are knowledge). This makes it at best very dubious that we should classify Jane's second-order belief under a type whose description includes divine promptings. And note that this is quite different to the situation with her first-order beliefs. While it is plausibly true that had Jane not been on the receiving end of a divinely-initiated testimony chain she would nevertheless have formed some religious beliefs by way of some *other* religious testimony chain, one with *non*-divine origins. But those beliefs would not have had the same content as the beliefs Jane in fact formed.<sup>24</sup> I would add that while there is good reason to think that the truth of any one of the major monotheisms makes it overwhelmingly likely that God acted specially in order to bring about revelatory events at the beginnings of the relevant testimony chains, it is far from clear that the truth of any of these belief-

---

<sup>23</sup> Plantinga's (2000) 'Aquinas-Calvin' model of how Christian belief is likely to be warranted if true involves the idea that God would engage in such prompting by way of what Plantinga calls 'the internal instigation of the Holy Spirit', though I would note that Plantinga doesn't specifically claim that such promptings would be intended to result in second-order beliefs.

<sup>24</sup> This does, of course, assume that the nearest world in which Jane's religious beliefs are the result of a testimony chain with *non*-divine origins is a world in which Jane is a follower of a different religion, rather than a world in which Jane is a follower of the same religion but that religion is instead false. This assumption seems reasonable when we consider that facts about which religion Jane follows make much less of a contribution to the overall similarity of a pair of worlds than the facts about which religious belief-system is true.

systems makes it especially likely that God would engage in special prompting intended to result in second-order beliefs.

But why should a lack of second-order religious knowledge be so troubling? Isn't it enough that Jane's testimony-based religious beliefs amount to knowledge at the first-order? There are a couple of ways of approaching the answer to this.

For one thing, a lack of second-order knowledge shows up as a deficiency when we consider that it is actually very easy to obtain second-order knowledge in various other domains. Typically it is very easy to know that a perceptual belief is an item of knowledge, for instance. What is the process by which we form second-order beliefs about whether our perceptual beliefs are knowledge? Presumably it is some kind of default disposition to assume that things are going normally with our perceptual systems unless there are internally accessible indicators that things are going awry—indicators such as a failure of various sensory impressions to cohere fully with one another, reports of an event by other people that conflict with the deliverances of one's own senses concerning that event, memories of having been exposed to substances that one knows to have an adverse affect on perception, and so on.<sup>25</sup> Assuming that we are not subject to persistent and systematic deceptions in the perceptual realm (such as those foisted by Cartesian evil demons or the like), this belief-forming process type is presumably highly statistically reliable: the vast majority of the time when we employ it to form a second-order belief about whether things are going normally with our perceptual systems, and hence whether a given perceptual belief is in the market for first-order knowledge, we get it right. Moreover, the second-order beliefs we form in this way will typically be safe and sensitive: in the typical case, when one forms a second-order belief in the described manner concerning whether a first-order perceptual belief is an item of knowledge, one could not easily have been mistaken, and furthermore, if that first order belief had failed to be an item of knowledge, one typically would not have formed the second-order judgment to the effect that such a belief amounts to knowledge.<sup>26</sup> Similar points could plausibly be made about the processes we use to form second-order beliefs about the epistemic status of our first-order logical and arithmetical beliefs, memorial beliefs, introspective beliefs, and so on.

Another way to highlight the deficiency inhering in a lack of reflective knowledge in the religious domain is to consider the relation between knowledge and permissible assertion. Suppose, as a number of epistemologists have recently argued,

---

<sup>25</sup> As Sosa writes, 'What, for example, is the competence we exercise in taking the light to be normal when we trust our colour vision in an ordinary case? It seems a kind of default competence, whereby one automatically takes the light to be normal absent some special indication to the contrary' (2007: 32).

<sup>26</sup> This will typically be because one would not even have formed the first-order belief, aware that it wouldn't have amounted to knowledge had one formed it in the circumstances.

that knowledge is the norm of assertion.<sup>27</sup> Suppose, that is, that one is entitled to assert that  $p$  only if one knows that  $p$ . Suppose that Jane knows various religious propositions as a result of trusting testimony but that for the sorts of reasons just outlined she doesn't know whether she knows these propositions. In that case, Jane does not know whether she is entitled to assert the contents of her religious beliefs. Even supposing she in fact is entitled to assert those propositions in virtue of knowing them, there may yet be an important sense in which Jane is criticisable for asserting them, for if she routinely asserts propositions whilst not knowing whether she knows them, she will be liable in a wide range of situations to make assertions that she is not entitled to make.

Now, it may be true that in other controversial domains, notably philosophy, it is similarly hard to know whether one has knowledge and hence hard to know whether one is entitled to assert the contents of one's beliefs. But in philosophy, at least, there is ample room for hedging assertions in a way that exempts a speaker from the strictures of the knowledge norm. By contrast, Jane's religion exhorts her—commands her indeed—to *proclaim* the contents of her beliefs. And proclamation, I take it, involves outright assertion rather than assertion which is hedged or somehow softened so as to be exempt from the knowledge norm.<sup>28</sup> Given that she doesn't know whether her first-order religious beliefs are knowledge, Jane doesn't know whether she will be conforming with the knowledge norm when she proclaims the contents of her religious beliefs.

## 5. Conclusion

This chapter has sought to identify what, if anything, is wrong with religious beliefs that are based upon testimony, given the facts of religious diversity. The epistemic reliability of such beliefs isn't principally what is at issue, I've suggested, because at least some religious belief-systems are such that conditional upon their truth it is likely that beliefs in their core tenets which are based upon testimony will be safe, sensitive, and statistically reliable. Nor is it inevitable that such beliefs face defeaters in light of the facts of religious diversity. Rather, the problem with basing one's religious beliefs just on testimony, so I have claimed, is that one will typically be unable to rise to the level of reflective knowledge, the having of which requires a knowledgeable perspective on one's reliability.

---

<sup>27</sup> See, for instance, Sosa (2011: Ch. 3), Williamson (2000: Ch. 11), DeRose (2002).

<sup>28</sup> Softening might involve asserting things like 'if my religion turns out to be true then  $p$ ', which clearly doesn't amount to asserting that  $p$ .

## References

- Alston, William P. 1995. "How to Think About Reliability". *Philosophical Topics* 23 (1): 1-29.
- Baker-Hytch, Max. 2015. "Mutual Epistemic Dependence and the Demographic Divine Hiddenness Problem". *Religious Studies* forthcoming.
- Becker, Kelly. 2008. "Epistemic Luck and the Generality Problem". *Philosophical Studies* 139: 353-66.
- Benton, Matthew. ms. "Epistemology Personalised".
- Bogardus, Tomas. 2013. "The Problem of Contingency for Religious Belief". *Faith and Philosophy* 30 (4): 371-92.
- DeRose, Keith. 2002. "Assertion, Knowledge, and Context". *The Philosophical Review* 111: 167-203.
- Dretske, Fred. 1971. "Conclusive Reasons". *Australasian Journal of Philosophy* 49 (1): 1-22.
- Goldberg, Sanford. 2014. "Does Externalist Epistemology Rationalize Religious Commitment?". In *Religious Faith and Intellectual Virtue*, edited by Laura Frances Callahan and Timothy O'Connor. 279-98. Oxford: Oxford University Press.
- Goldman, Alvin. 1979. "What Is Justified Belief?". In *Justification and Knowledge*, edited by George S. Pappas. 1-23. Dordrecht: D. Reidel.
- . 1986. *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.
- Kitcher, Philip. 2011. "Challenges for Secularism". In *The Joy of Secularism: 11 Essays for How We Live Now*, edited by George Levine. 24-56. Princeton, NJ: Princeton University Press.
- Lackey, Jennifer. 2007. "Why We Don't Deserve Credit for Everything We Know". *Synthese* 158: 345-61.
- Mawson, T. J. 2009. "Mill's Argument against Religious Knowledge". *Religious Studies* 45 (4): 417-34.
- . 2012. "The Rationality of Classical Theism and Its Demographics". In *Scientific Approaches to the Philosophy of Religion*, edited by Yujin Nagasawa. 184-204. London: Palgrave Macmillan.
- Mill, John Stuart. 1991 [1859]. "On Liberty". In *The Collected Works of John Stuart Mill*, edited by J.M. Robson. Toronto, ON: Toronto University Press.

- Nolan, Daniel. 1997. "Impossible Worlds: A Modest Approach". *Notre Dame Journal of Formal Logic* 38 (4): 535-72.
- Nozick, Robert. 1981. *Philosophical Explanations*. Cambridge, MA: Harvard University Press.
- Plantinga, Alvin. 1995. "Pluralism: A Defence of Religious Exclusivism". In *The Rationality of Belief and the Plurality of Faith: Essays in Honor of William P. Alston*, edited by Thomas D. Senior. 191-215. London: Cornell University Press.
- . 2000. *Warranted Christian Belief*. Oxford: Oxford University Press.
- Pollock, John L. 1986. *Contemporary Theories of Knowledge*. Totowa, NJ: Rowman & Littlefield.
- Pritchard, Duncan. 2012. "Anti-Luck Virtue Epistemology". *Journal of Philosophy* 109 (3): 247-79.
- Sosa, Ernest. 1997. "Reflective Knowledge in the Best Circles". *Journal of Philosophy* 94 (8): 410-30.
- . 2007. *A Virtue Epistemology. Vol. 1, Apt Belief and Reflective Knowledge*. Oxford: Oxford University Press.
- . 2009. *Reflective Knowledge. Vol. 2, Apt Belief and Reflective Knowledge*. Oxford: Oxford University Press.
- Swinburne, Richard. 2001. *Epistemic Justification*. Oxford: Clarendon Press.
- Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.