

## Religious Diversity and Epistemic Luck<sup>1</sup>

Max Baker-Hytch

**Abstract** A familiar criticism of religious belief starts from the claim that a typical religious believer holds the particular religious beliefs she does just because she happened to be raised in a certain cultural setting rather than some other. This claim is commonly thought to have damaging epistemological consequences for religious beliefs, and one can find statements of an argument in this vicinity in the writings of John Stuart Mill and more recently Philip Kitcher, although the argument is seldom spelled out very precisely. This paper begins by offering a reconstruction of an argument against religious beliefs from cultural contingency, which proceeds by way of an initial argument to the unreliability of the processes by which religious beliefs are formed, whose conclusion is then used to derive two further conclusions, one which targets knowledge and the other, rationality. Drawing upon recent work in analytic epistemology, I explore a number of possible ways of spelling out the closely related notions of accidental truth, epistemic luck, and reliability upon which the argument turns. I try to show that the renderings of the argument that succeed in securing the sceptical conclusion against religious beliefs also threaten scepticism about various sorts of beliefs besides religious beliefs.

**Keywords** Religious diversity · Cultural contingency · Epistemic luck · Reliability · Undercutting Defeater · Knowledge

### §1. Introduction

A familiar criticism of religious belief starts from the claim that a typical religious believer holds the particular religious beliefs she does just because she happened to be raised in a certain cultural setting rather than some other. John Stuart Mill once wrote that “mere accident has decided which of these numerous worlds is the object of [a religious person’s] reliance... the same causes which make him a churchman in London would have made him a Buddhist or a Confucian in Peking” (1991, pp.

---

<sup>1</sup> I am very grateful for the comments of Christopher Kyle, Wes Skolits, Brian Leftow, and Tim Mawson on earlier drafts of this paper, and for conversations with members of Oxford University’s New Insights and Directions in Religious Epistemology project, which helped me to arrive at my formulations of various key epistemological notions. I am also grateful to Oxford’s philosophy of religion reading group for an extremely helpful critical discussion of the paper, and for the questions and comments of audience members at a lecture in Oxford’s philosophy faculty on 5<sup>th</sup> February 2014 in which this paper was presented. And I am grateful for the comments of an anonymous referee for this journal on the paper’s penultimate draft.

229-230). And in much the same vein, Philip Kitcher has suggested 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... Since they are all on a par, we should trust none of them” (2011, p. 26). Even supposing that someone was exposed to the right religious influences and thus ended up with true religious beliefs, the thought seems to be that he or she would simply have been lucky to have done so, in a sense that is incompatible with her beliefs having a positive epistemic status of a certain sort.

My aim in this paper is to try to reconstruct what I take to be the strongest version of an argument from cultural contingency, and then to subject it to scrutiny. Very significant work has been done by epistemologists in recent years in clarifying the interrelated notions of accidental truth, epistemic luck, and reliability—work which largely has yet to be applied to the perennial issue of the cultural contingency of religious beliefs and its epistemological consequences. Applying the latest and most promising analyses of these concepts, I contend that the argument from cultural contingency against religious belief threatens significantly more widespread scepticism than its advocates intend.

## §2. Reconstructing an Argument from Cultural Contingency

The first thing to note about the sort of argument we are considering, as it is typically presented, and as it is presented by Mill and Kitcher, is that it involves a counterfactual conditional: “You believe that  $p$ , but *if you had been raised in different cultural circumstances, you would instead have believed  $q$* ” (where  $q$  is incompatible with  $p$ ). Of course, as Alvin Plantinga has pointed out, a counterfactual of this sort is true not just of many religious beliefs, but also of all sorts of other beliefs, including scientific beliefs (1995, p. 212). Had I been raised in first-century Athens, I would have believed that *the earth is flat*, which is obviously incompatible with my actual belief that *the earth is round*. There is, however, a standard fix for this issue: namely, to emphasise the similarity of belief aetiology between the actual and the counterfactual cases.<sup>2</sup> Whereas my belief that *the earth is round* is the causal result of seeing countless images of the earth taken from space, the belief I would have held if I were raised in ancient Athens—the belief that *the earth is flat*—would have been the causal result of something very different: accepting mere speculation about the earth’s shape. On the other hand, the thought goes, when it comes to most religious beliefs,

---

<sup>2</sup> For more on this, see Bogardus (2013, p. 377-78).

the difference between the causal history of the actual belief and that of the counterfactual belief is not nearly so marked. Mill and Kitcher seem implicitly to be aware of this point. Thus, Mill: “the *same causes* which make [a man] a churchman in London would have made him a Buddhist or a Confucian in Peking” [emphasis mine] (1991, pp. 229-230). And Kitcher: “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” [emphasis mine] (2011, p. 26). In our reconstruction of the argument from cultural contingency, then, the counterfactual conditional we employ should have the following form:

*CF*: If *S* had been raised under different cultural circumstances, then *S* would have held a belief incompatible with her actual belief, whose causal history would have been *relevantly similar* to that of her actual belief.

Another point to note before we proceed to our reconstruction is that there is more than one epistemic property that an argument from cultural contingency might seek to target. On the one hand, the target might be *knowledge*.<sup>3</sup> Knowledge very plausibly requires that belief be formed in a reliable manner,<sup>4</sup> and one might try to show that the cultural contingency involved in the causal origins of many religious beliefs implies that those beliefs are formed in an *unreliable* manner, and hence, that those beliefs do not amount to knowledge, even supposing they are true. Indeed, Plantinga notes that Mill’s argument, with its invocation of accidental truth, “seems to be designed to appeal to reliabilist intuitions” (1995, p. 211). On the other hand, one might wish to target *rationality*, where that notion is understood roughly in terms of one’s epistemic entitlement to believe something.<sup>5</sup> Epistemologists often speak of “defeaters”, where a defeater for one’s belief that *p* consists of a reason for giving up that belief; a reason, the having of which renders one *irrational* in continuing to believe that *p*.<sup>6</sup> John Pollock distinguished two sorts of defeaters: rebutting and undercutting (1986, pp. 37-39). A rebutting defeater for one’s belief that *p* consists in evidence for the falsity of *p*. Suppose one believes that the first Europeans to set foot on the North American continent were Spaniards, but that one acquires evidence that the first Europeans to do so were Icelanders. In that case, one has acquired a rebutting defeater for one’s original belief. An undercutting defeater, on the other hand, is a reason for thinking that the source of one’s belief that *p* is not

<sup>3</sup> This is how Mawson (2009) interprets Mill’s argument.

<sup>4</sup> I intend “reliable” here to cover the sort of modal reliability that appears to be lacking in Gettier cases.

<sup>5</sup> The notion I have in mind is equivalent to what is sometimes termed “deontic justification”.

<sup>6</sup> For a recent and extensive treatment of defeaters, see Bergmann (2006, ch. 6).

trustworthy or reliable. If, for instance, one acquires evidence that the Ordinance Survey map one is following is a joke edition with fake landmarks included, then one acquires an undercutting defeater for one's belief that there is an ancient well in the next village (a belief based solely on reading the map); one is irrational in continuing to hold that belief on the basis of that map. And if the map one is using is *in fact* unreliable, then the undercutting defeater that one acquires is non-misleading.<sup>7</sup>

It looks as though the evidence concerning cultural contingency is liable to generate an undercutting rather than rebutting defeater, if it generates a defeater at all. And in order to show that that evidence generates a non-misleading undercutting defeater, one will need to show that the cultural contingency at work in the causal histories of religious beliefs does *in fact* imply that those beliefs are unreliably formed. I suggest, then, that an argument from cultural contingency is best formulated in terms of a core argument for the unreliability of the processes by which religious beliefs are formed, whose conclusion is subsequently used to support two further conclusions, one which targets knowledge and the other, rationality. I shall shortly suggest that the core argument—in particular, its second premise—can be spelled out in a couple of different ways depending on which notion of reliability one employs, but for now, allow me to state that argument in a schematic way, and clarify it later.

*Core argument*

- (1) For many religious believers, *R*, it is true that *CF*: If *R* had been raised under different cultural circumstances, then *R* would have held different and incompatible religious beliefs whose causal histories would have been relevantly similar to those of her actual religious beliefs.
- (2) If *CF* is true for *R*, then *R*'s religious beliefs are unreliably formed.
- (3) For many *R*, *R*'s religious beliefs are unreliably formed [from (1) and (2)].

*Further argument I: Knowledge*

- (4) If one's belief is unreliably formed, then that belief is not knowledge.
- (5) For many *R*, *R*'s religious beliefs are not knowledge [from (3) and (4)].

*Further argument II: Rationality defeat*

- (6) If one correctly believes that one's belief is unreliably formed, then one has a non-misleading undercutting defeater for one's belief.
- (7) If a religious believer *R* of whom *CF* is true believes (3), then *R* has a non-misleading undercutting defeater for her religious beliefs [from (3) and (6)].

---

<sup>7</sup> Misleading evidence is evidence that renders a true proposition less likely. Hence, one's evidence for the unreliability of the map is non-misleading evidence just in case the map in fact is unreliable.

In what follows, my focus will be upon the core argument. Before I can begin to scrutinise that argument, however, I must distinguish two kinds of reliability.

### §3. Modal Reliability and Process Reliability

On the one hand, there is what might be termed *modal reliability*. Modal reliability is what is lacking in a Gettier case, in which a subject is lucky, in a certain sense, to wind up with a true rather than false belief. Farmer Jones, for example, looks into his field and forms the true belief that his prize cow Daisy is in the field, though what he is in fact looking at is a black and white sheet caught up in some trees that looks just like Daisy, obscuring the real Daisy from view. Given how he forms his belief, Farmer Jones appears to be very lucky to get a true rather than false belief. Put in terms of the framework of possible worlds, there are some very close worlds in which Farmer Jones gets a false belief: worlds in which Farmer Jones still forms his belief by looking at the black and white sheet caught up in the trees, but in which Daisy has in fact wandered into the neighboring field. Duncan Pritchard has termed this kind of luck “veritic luck” (2005, p. 146). A modal reliability condition on knowledge aims to exclude precisely this sort of luck. Two candidates for such a condition have been given serious consideration in the recent epistemological literature: the sensitivity condition and the safety condition. Sensitivity can be formulated in the following way:

*Sensitivity:*  $S$  sensitively believes that  $p$  in a world  $w$  if, and only if, in the nearest world(s) to  $w$  in which  $p$  is false and  $S$  forms her belief about  $p$  in the same way as in  $w$ ,  $S$  doesn't believe that  $p$ .<sup>8</sup>

There are several notorious problems with saying that sensitivity is necessary for knowledge, including sensitivity's failure to preserve deductive closure and to allow for inductive knowledge,<sup>9</sup> as well as a further problem that is especially relevant for our present purposes: it cannot give sensible results concerning beliefs in necessary truths, because where a belief in a necessary proposition is concerned, there simply *isn't* a possible world in which the proposition in question is false.<sup>10</sup> Given that some of the most important tenets of major religious systems seem to be the sorts of things

---

<sup>8</sup> Dretske (1971) and Nozick (1981, pp. 172-179) have both defended a condition along these lines.

<sup>9</sup> For more on these issues, see Hawthorne (2004, pp. 31-46) and Pritchard (2005, pp. 34-66).

<sup>10</sup> On the usual Lewis-Stalnaker semantics for counterfactuals, a counterfactual with a necessarily false antecedent is trivially true (Lewis 1973, p. 24; Stalnaker 1975, p. 170). Hence, for any proposition  $p$  that is necessarily false, the counterfactual conditional  $p > q$  is true, but so too is  $p > \neg q$ .

are necessarily true if true at all,<sup>11</sup> this will pose a serious obstacle for an argument from cultural contingency. For the reasons alluded to, then, it wouldn't be wise to utilize the sensitivity condition in such an argument. On the other hand, the much more popular safety condition appears capable of excluding veritic luck whilst avoiding the aforementioned difficulties:<sup>12</sup>

*Safety:*  $S$  safely believes that  $p$  in a world  $w$  if, and only if, there is no nearby world  $w^*$  in which  $S$  arrives at a false belief with relevantly similar propositional content and with a relevantly similar causal history.<sup>13</sup>

Clearly enough, this safety condition can handle cases involving a badly formed belief in a necessary truth. Consider a case in which John forms a belief about the square root of 81, by way of consulting a calculator that is (unbeknownst to John) broken and is issuing answers more or less randomly. Although John gets the correct answer in the actual world, there will be some very close worlds in which John forms a distinct but similar belief about the square root of 81 via a very similar causal chain, and ends up in error. Hence, John's belief is unsafe.

At the same time, some epistemologists have suggested that modal reliability may not be sufficient for knowledge by itself, and that there is a distinct kind of reliability that is also necessary for knowledge.<sup>14</sup> One reason for thinking this is that there is a certain sort of case in which a subject forms a modally reliable belief, despite her belief having been formed by way of a *type* of process that intuitively seems not to be capable of producing knowledge. Suppose, for instance, that Alfie forms various beliefs out of wish-fulfillment, but that unbeknownst to him, he has an epistemic guardian angel who was assigned to him at birth and who has been tasked with making Alfie's every wish come true. Alfie's belief that *Santa Claus is real*, formed by way of wish-fulfillment, is safe, given that the guardian angel acts to ensure its truth: there is no nearby world in which Alfie forms a similar belief in a similar manner and yet ends up with a false belief. But intuitively there is something wrong with Alfie's forming his beliefs on the basis of wish-fulfillment. What's wrong, perhaps, is that he forms his beliefs by way of a type of process that is *generally* unreliable—that is, a

---

<sup>11</sup> I'm thinking, for instance, of such tenets as the Christian doctrine of the trinity, or the claim shared by all the major monotheistic traditions that God is morally perfect, omniscient, omnipotent, and so on, or the claims of many Eastern traditions that there is a fundamental metaphysical unity of all things.

<sup>12</sup> Bogardus (2013) interprets the argument from cultural contingency in terms of safety.

<sup>13</sup> It is true that cruder versions of safety, which merely require that  $S$  doesn't falsely believe  $p$  in a nearby world, cannot give sensible results concerning badly formed beliefs in necessary truths. The more sophisticated version of safety that I have offered is close to formulations presented by Pritchard (2012, pp. 256-257) and Williamson (2009, p. 235).

<sup>14</sup> See, for example, Pritchard (2012) and Becker (2008).

process that generates an insufficiently high proportion of true beliefs when used under normal circumstances.<sup>15</sup> The process he uses—*wish-fulfillment*—is one that, when used under normal circumstances in the actual world (and in worlds with broadly similar physical setups to the actual world), generates considerably more false than true beliefs. It has a low *truth-ratio*. Alfie’s failure to know might thus be explained in terms of his failure to satisfy the following condition:

*Process reliability:* *S*’s belief that *p* is process reliable if, and only if, the salient process type that produced *S*’s belief has a high truth-ratio in the actual world and in worlds with similar physical laws.

Of course, appealing to process reliability to explain why Alfie’s beliefs aren’t knowledge depends on its being the case that the process he uses is rightly classified as *wish-fulfillment* rather than *wish-fulfillment in the presence of an epistemic guardian angel who ensures the truth of one’s wishes*. The latter process presumably has a very high truth-ratio, while the former doesn’t. Alfie’s token belief-forming episode is of course an instance of *both* the aforementioned process types (and of indefinitely many other process types at greater and lesser levels of generality), and so, if we are to appeal to the unreliability of the process type that Alfie uses, there needs to be a fact concerning *which* process type is the salient one. The notorious difficulty of determining the salient process type in a principled manner is known as the generality problem, and it is a problem that will become relevant later.

In what follows, I shall examine the core argument interpreted first as targeting modal reliability (specifically, safety), and then, having tried to show that the argument thus understood leads to unwarranted scepticism, I shall examine its prospects when the argument is understood in terms of process reliability.

#### §4. Arguing that Cultural Contingency Implies Unsafety

Let’s now try to restate our core argument so as to secure the conclusion that religious beliefs formed by way of culturally contingent causal processes are not *safe*. For the purposes of illustration, let’s consider Mill’s Victorian churchman, who believes that *God is triune*, and who, by way of similar causal processes as led him to

---

<sup>15</sup> It should be noted that although many epistemologists consider knowledge to be incompatible with belief that is the product of a generally unreliable process, Lasonen-Aarnio (2010) has defended the view that cases like that of Alfie are in fact cases of knowledge (knowledge being mere safe belief), but that our reluctance as onlookers to *attribute* knowledge in such cases is explained by our wish not to reward the subjects in such cases for their employment of an unreasonable belief-formation policy—a policy that *in general* does not yield knowledge.

his Trinitarian belief, would have believed that *God is non-triune* had he been raised among Confucians in Peking. We should assume for the sake of the argument that the churchman's actual belief happens to be true.<sup>16</sup> To say that the churchman's Trinitarian belief fails to be safe is to say that there are nearby worlds—including some of those worlds where he was raised in Peking—in which the churchman has a false belief whose propositional content and causal history are relevantly similar to those of his actual-world belief. In general, then, a proposition of the following form needs to be true of a person's religious beliefs if those beliefs are to be deemed unsafe:

*CF\**: There are nearby possible worlds in which *S* has a false belief with relevantly similar propositional content and a relevantly similar causal history to her actual-world belief.

Accordingly, our core argument, understood in terms of a failure of safety, can be stated as follows:

- (1\*) For many religious believers, *R*, it is true that *CF\**: there are nearby worlds in which *R* has different and (*ex hypothesi*) false religious beliefs with relevantly similar propositional content and relevantly similar causal histories to her actual-world religious beliefs.
- (2\*) If a person safely believes that *p* in a world *w*, then there is no nearby world *w\** in which she has a false belief with relevantly similar propositional content and a relevantly similar causal history.
- (3\*) Therefore, for many religious believers *R*, *R*'s religious beliefs are not safe [from (1\*) and (2\*)].

From there, the rest of the argumentation—the knowledge argument and the rationality defeat argument—can proceed as before, at least, given the plausible assumptions that knowledge requires safety and that correctly believing that one's belief is unsafe gives one a non-misleading undercutting defeater for that belief.

What I want to argue now is that premise (1\*) of the above argument is false, or at any rate, that unless we deem it false we will be committed to an unwarranted degree of scepticism about many ordinary cases of testimony. In general, when presented with the claim that *S*'s belief that *p* is unsafe, there are three ways in which one might try to challenge that claim: (i) one could try to show that the possible worlds in which *S* has a false belief with relevantly similar content and a relevantly similar

---

<sup>16</sup> The argument needs to be able to succeed even on this assumption, because it is trivial and uninteresting to show that what is assumedly a false belief is not an instance of knowledge.

causal history to her actual belief are not in fact *close* worlds; (ii) one could try to show that, even supposing there are close worlds in which *S* has a false belief whose propositional content is relevantly similar to *p*, the causal history of *S*'s belief in those worlds is not relevantly similar to the causal history of *S*'s actual belief; (iii) one could try to show that, even supposing that there are close worlds in which *S* has a false belief whose causal history is relevantly similar to that of her actual belief, the propositional content of *S*'s belief in those worlds is not relevantly similar to *p*. I am going to attempt both (i) and (ii), though not (iii). I doubt it is worth trying to show, for instance, that the content of the churchman's actual belief that *God is triune* is not relevantly similar to the content of his counterfactual belief that *God is non-triune*.

Now, clearly, the terms "close" and "relevantly similar" are significantly vague.<sup>17</sup> In disputed cases, such as cases of religious belief, it won't really do to appeal to intuitive judgments about the similarity of pairs of cases of religious beliefs that are mutually incompatible with one another. The intuitive judgments of interlocutors in the atheist-theist debate are not likely to form a consensus on this matter. But what we can do, it seems to me, is use our intuitive judgments about cases on which all parties to the debate are agreed as a kind of anchor, in a manner that will become clear shortly. The anchor I propose to use is a sort of ordinary (i.e. non-religious) case of testimony-based belief, which I think nearly everyone will grant is a case of knowledge.

Let's term coin a term of art: an *error world* is a possible world in which a subject *S*, who correctly believes *p* in the actual world, instead believes a false proposition whose content is relevantly similar to *p* (the limiting case being *p*). Assuming the truth of the churchman's Trinitarian belief in the actual world, the world in which the same individual was instead raised in Peking among Confucians is an error world: in that world he (*ex hypothesi*) falsely believes that *God is non-triune*. My contention, then, is that for some ordinary testimony cases that virtually everyone will deem knowledge, the closest error worlds are *at least as close* to the actual world as the closest error worlds are to the actual world for fairly typical cases of religious belief, such as that of Mill's Victorian churchman. What's more, I contend that for such testimony cases, the causal history of belief in the closest error worlds is *at least as similar* to the belief's causal history in the actual world as is true for many cases of religious belief. It's worth noting that the closeness of an error world to the actual world is in part, though by no means wholly, determined by the similarity of the belief's causal history across worlds. The two variables are not completely

---

<sup>17</sup> The seeming inescapability of such vagueness in formulating central epistemological notions is apparently part of the motivation for Timothy Williamson's knowledge-first approach to epistemology, which eschews the project of locating non-circular necessary and sufficient conditions for knowledge, reliability, or justification (see his 2000, pp. 2-5; 2014).

independent of one another. In particular, a difference between two worlds with respect to the environment in which the believer forms her belief will entail differences with respect to the causal history of her belief, as well as entailing other differences between the two worlds that don't directly bear upon the causal history of her belief (e.g., the way the wind is blowing her hair).

So, consider the following testimony case. Jenny gets off the train in Truthland and obtains a correct belief about the directions to the British embassy by way of the testimony of someone who seems to Jenny, and in fact is, both honest and knowledgeable about the whereabouts of the embassy. What's more, everyone else in Jenny's nearby vicinity is both honest and knowledgeable about the whereabouts of the embassy. I take it that many of us will be inclined to attribute knowledge to Jenny here. But it seems to me that for Jenny in this case, the nearest error world is *at least as close* to actuality as the nearest error world is to the actual world for Mill's churchman. And relatedly, the causal history of Jenny's belief in the closest error worlds is *at least as similar* to her belief's causal history in the actual world as the churchman's belief's causal history is to the causal history of his counterfactual belief in the world in which he is raised in Peking. Put another way, the degree to which Jenny's environment must be counterfactually altered so that she instead winds up with a false belief seems to be *no greater than* the degree of counterfactual alteration involved in switching the churchman's entire cultural setting. One way to change Jenny's environment so that she gets a false rather than true belief is by switching her location to Falsityland, where she would be surrounded largely by people who are mistaken (though sincere) about the directions to the British Embassy. But if we want to maintain that Jenny acquires knowledge in the actual world then we'll of course have to say that it is *isn't* epistemically relevant that she would have fallen into error had she been located in Falsityland; we'll have to say that the world in which Jenny is located in Falsityland is *not* a sufficiently close world, or that the causal history of her belief in Falsityland is *not* relevantly similar to the causal history of the belief she forms in Truthland. And yet the argument against religious beliefs from cultural contingency, cast in terms of safety, will have to maintain that it *is* epistemically relevant that a religious believer would have had false religious beliefs had she been located in religious "Falsityland" rather than in religious "Truthland". And so, if one insists that such a believer's beliefs aren't safe even if they're true, then one will also be committed to saying that Jenny (and people in situations like that of Jenny) has an unsafe belief. But that seems implausibly sceptical.<sup>18</sup>

---

<sup>18</sup> Note that this is different from Tomas Bogardus's recent objection to a safety-based contingency argument (2013, pp. 379-82). Bogardus's objection is twofold: (i) safety is not necessary for knowledge; (ii) in any case, the epistemic luck that would be involved in happening to be born in a culture with true religious beliefs is a benign

An important caveat to the above must be entered at this point. In using the example of Mill's churchman, I assumed that the individual in question was fairly far removed from the possible influence of a different religious tradition, meaning that a considerable counterfactual alteration of his environment would be required in order to place him under such influence. Now, I take it that the situation of Mill's Victorian churchman is fairly typical of believers who live in largely religiously homogeneous societies. In some sections of contemporary Western societies, though, there are presumably at least some individuals whose circumstances are such that very little would have to have been different in order for them to have ended up with a different set of religious beliefs.<sup>19</sup> That is, the nearest error worlds for the religious beliefs of such individuals are very close indeed. A child growing up in a multicultural district of a Western European city, say, might very easily have been exposed to the influence of a different religious tradition than the one to which she was actually exposed. Consider Aaliyah, who would have been sent to a school with a wholly different religious orientation, but for the whimsical decision of an educational administrator. Supposing (*arguendo*) that Aaliyah's Islamic beliefs are true, it looks as though not very much *at all* needed to go differently in order for Aaliyah to wind up with (*ex hypothesi*) false religious beliefs—indeed, less than had to go differently in order for Jenny, who was fairly geographically far from Falsityland, to have ended up with a false belief. So it looks as though we can say that Aaliyah's belief is unsafe, seemingly *without* risking the kind of epistemic “collateral damage” that we faced if we said the same of Mill's churchman.

But this isn't much of a consolation to the advocate of the argument from cultural contingency. Mill and Kitcher and others who advocate the argument from cultural contingency are presumably thinking of their arguments as applying in full force to religious believers who come to their beliefs by way of religious enculturation, *regardless* of whether their circumstances are quite as modally precarious as those of Aaliyah. Now, to be clear, the extent to which a religious believer is *modally* distant from the influence of a different religious tradition is not straightforwardly determined by her *geographical* remoteness from a different tradition, though geographical remoteness will typically contribute to modal distance. The two will come apart in the case in which a believer has been raised in a location that is geographically far removed from the influence of a different religious tradition but in which it is true that if (say) the believer's parents had died or had not managed to stay financially above water (states of affairs that could fairly easily have come about), she would have been sent to be raised by relatives living in a far-flung land and thus

---

kind of luck. That sort of epistemic luck is analogous, he claims, to safely crossing a bridge that was nearly rendered unsafe: first-order safety is compatible with being lucky to have used a safe method.

<sup>19</sup> Thanks to Chris Kyle for drawing my attention to this consideration.

exposed to a wholly different religious outlook. Nonetheless, for many—perhaps most—individuals who are geographically far removed from the influence of other religious traditions, it is true that not easily could they have been subject to the influence of another tradition. Furthermore, even for a good many believers who are living in fairly close physical proximity to adherents of another tradition, a significant amount would have had to go differently in order for them to have been subject to those other religious influences.<sup>20</sup> Religious believers in a situation as precarious as Aaliyah's, vis-à-vis which religious beliefs they end up holding, are presumably going to be in a considerable minority, even among inhabitants of moderately pluralistic Western societies. But in order to avoid the sorts of unintended sceptical consequences sketched above (i.e., regarding Jenny), the argument from cultural contingency would have to restrict its focus exclusively to cases like Aaliyah's; exclusively to those religious beliefs that are less modally reliable than any beliefs that we take to be uncontroversial cases of knowledge. Such a narrowing of the argument's scope would be seriously disappointing as far as Mill and Kitcher are concerned. Let's see, then, how the argument from cultural contingency fares if conducted in terms of the reliability of belief-forming process types.

### §5. Arguing from the Unreliability of Process Types

If one wants to run the core argument in terms of process type reliability without assuming the falsehood of any particular set of religious beliefs—an assumption that would render the argument rather uninteresting—then one would think the best way to proceed is to try to show that the salient process type is one that is responsible for a variety of mutually incompatible religious beliefs. After all, a process type is not going to have a truth-ratio high enough to satisfy a process reliability condition on knowledge if too many of the beliefs for which it is responsible contradict one another. Even without making any assumptions about the truth-values of any *particular* beliefs that are the outputs of a given process type, we can still infer that the truth-ratio of that process type is no greater than the largest number of the type's outputs that could all be true at once divided by the total number of the type's outputs. I suggest, then, that if construed in terms of process types, the core argument ought to run in the following way:

---

<sup>20</sup> I am grateful to an anonymous referee for this journal for encouraging me to clarify the relation between modal distance and geographical remoteness.

- (1\*\*) For many religious believers, *R*, the salient process type that is responsible for producing *R*'s religious beliefs is such that a significant proportion of its outputs are mutually incompatible.<sup>21</sup>
- (2\*\*) If the salient process type that is responsible for producing one's belief is such that a significant proportion of its outputs are mutually incompatible, then that process type does not have a truth-ratio sufficiently high to count as reliable.
- (3\*\*) For many religious believers, *R*, the salient process type that is responsible for producing *R*'s religious beliefs does not have a truth-ratio sufficiently high to count as reliable [from (1\*\*) and (2\*\*)].

My focus once again will be the first premise, (1\*\*). The issue that now rears its ugly head is that of picking out the *salient* process type. As I mentioned earlier, any particular belief will be the result of a token causal process that can be classified under indefinitely many process types, and these types will vary widely in their reliability. But rather than try to resolve the thorny problem of how the salient process type ought to be selected (the generality problem), the approach I'll take instead is to explore the prospects for two broad approaches to this issue that could be taken by the advocate of the argument from cultural contingency. One approach is to pick out a fairly narrow, though nonetheless repeatable, process type. I'll suggest, however, that this approach faces a serious difficulty if monotheistic religious beliefs are the target of the argument. Another approach is to select a rather wider (i.e. more general) process type, something like *accepting the testimony of a religious tradition*. I shall argue that this latter approach may in fact succeed, but only if the defender of the argument from cultural contingency is prepared to accept some rather costly consequences.

## §6. Narrower Process Type Individuation

Consider that what determines the narrowness of a process type is the extent to which the specification of that type involves particular individuals—including particular objects, times, and events. The process type *forming a belief by way of accepting the testimony of Huw Edwards via the TV set in my lounge between 6 and*

---

<sup>21</sup> Note that when the core argument is run in terms of process reliability, the argument's first premise doesn't need to appeal to a counterfactual claim, though it *could*. The relevant counterfactual would be something like: if *R* had been raised under different cultural circumstances, then *R* would have had different and (*ex hypothesi*) false religious beliefs that would have been the product of the same salient process type as produced her actual religious beliefs.

*6.30PM on Saturday 8<sup>th</sup> March 2014* is a very narrow process type compared with the type *forming a belief by way of accepting the testimony of a TV news anchor*, which mentions no particular objects, events, or times. The latter type has vastly many more instances than the former. What I want to suggest is that the argument from cultural contingency faces a serious difficulty if the process type picked out as the relevant one is specified in a way that involves particular individuals. As I said, the argument from cultural contingency needs to assume for the sake of the argument the truth of the religious beliefs that are being targeted—it is otherwise uninteresting to show that those beliefs are not knowledge. What the defender of the argument from cultural contingency is then looking to show is that the relevant process type, as a result of which a given believer formed her religious beliefs, is such that it yields many beliefs that are incompatible with the believer’s actual beliefs; the defender of the argument is trying to show, that is, that the relevant process type has an insufficiently high truth-ratio. But if the process type in question is specified in a way that involves the *particular* chain of individuals whose testimony was actually causally efficacious in bringing about a person’s religious beliefs, or in a way that involves the *particular* religious texts or rituals that actually contributed to her belief-formation, then it looks as though that process type will *not*, on the whole, yield very many beliefs that differ from her actual beliefs. Now, it is worth emphasising at this point that we need not restrict our attention just to the *actual world* truth-ratio of process types (i.e. the ratio of true beliefs to false beliefs produced in the actual world by a given process type); rather, we can—and according to Kelly Becker, *should*—also take into account the truth-ratio of a given process type across a nearby region of possible worlds centred on the actual world.<sup>22</sup> Given this, one might think that even if we do mention in our specification of the process type things like a particular set of religious scriptures or a particular testimony chain, we will nonetheless get to take into account worlds where, say, that set of scriptures or that testimony chain yields many claims which fail to match the facts.<sup>23</sup> One might think that when all these worlds are factored in, it could well turn out that even a narrow process type, one whose specification includes a particular set of religious scriptures or a particular testimony chain, nonetheless yields a significant proportion of mutually inconsistent outputs and thus a significant proportion of false beliefs.

---

<sup>22</sup> The reason for this is that on a narrow type-selection approach such as the one we are currently considering (and which Becker advocates), a given process type may not have many actual-world instances, and so in order to adequately assess its truth-ratio, we need to consider to take into account the process type’s possible outputs throughout a nearby region of worlds; worlds that are fairly similar to the actual world. See Becker (2007, p. 89).

<sup>23</sup> Issues about modal fragility may threaten the notion that a particular set of religious scriptures could have included substantially different propositional content and yet have remained *the very same* set of scriptures. The argument I am about to offer, though, doesn’t rely on assuming the modal fragility of scriptures or testimony chains.

I think this line of reasoning is flawed, however. Suppose that the beliefs of a certain monotheistic believer are the target of the argument from cultural contingency on this occasion, so that the truth of her various religious beliefs is granted for the purposes of the argument. In that case, the defender of the argument from cultural contingency must reckon with the fact that given the truth of a monotheistic belief-system (according to which God created humans and wishes to relate personally to them), it is pretty epistemically probable that God has providentially guided or overseen the processes by which the scriptures of her religion have been written, assembled, and subsequently transmitted down through history, such that not easily could those scriptures have ended up asserting more than a small proportion of false claims. If God has indeed guided these transmission processes, and did so as a result of a settled intention rather than a whim—something which also seems highly epistemically probable given the truth of a monotheistic belief-system on which God wishes to relate to humans—then a possible world where a large proportion of scriptural assertions (as they are in that world) fail to match the facts (as they are in that world) will be extremely far from the actual world, and so they will barely, if at all, register in our evaluation of the process type's truth-ratio.

Now, there is an obvious objection to the argument I just offered: there exist disagreements (hence, mutually incompatible beliefs) not only at the *inter*-religious level, but also, frequently, at the *intra*-religious level.<sup>24</sup> There is diversity of belief even among believers who formed their beliefs as a result of the *same* processes of historical transmission, by way of the *same* scriptures and traditions. What can be said about this? Well, to begin with, it is important to note that we implicitly relativize our reliability judgements about process types to a particular domain and a particular level of specificity. Here's what I mean. Suppose *The Daily Post* prints a fair number of errors with respect to the precise times at which events occurred. Nonetheless, that, by itself, doesn't seem to licence the conclusion that reading *The Daily Post* is an unreliable way to form the belief *that* some particular event has occurred at all, just that it's an unreliable way to form a belief with more specific content concerning the precise time at which some event occurred. Again, suppose *The Daily Post* frequently prints mistakes regarding the sports scores. This doesn't appear, by itself, to licence the conclusion that *The Daily Post* is an unreliable way to form beliefs about some domain other than sports scores, say, political current affairs. Analogously, then, the fact that some religious process type (say, the type *forming beliefs by way of reading the Qur'an*) leads to a good deal of mutually incompatible beliefs concerning the various *specific ways* in which the core tenets of a religious system could be true, doesn't licence the conclusion that the process type in question

---

<sup>24</sup> Thanks to Brian Leftow and Chris Kyle for drawing my attention to this objection.

is an unreliable way of forming the mere belief *that* those tenets are true. And it does look like a great deal of intra-religious disagreement—not *all* instances of it, to be sure—occurs at the level of beliefs with more specific content; beliefs not merely *that* the core tenets of a religion are true, but concerning the *specific ways* in which those core tenets are true. For instance: about exactly how it is that Jesus' death achieved the reconciliation of God with humans; about the precise nature and scope of the commandment to love one's neighbour; about the exact manner in which the Hajj is to be undertaken; about the means by which God created the world; and so on. The claim can still stand, then, that conditional on the truth of a given monotheistic belief system, it is very epistemically likely that God has guided the processes by which at least the core tenets of that system have been transmitted down through history, such that not at all easily could those processes have yielded more than a small proportion of false beliefs whose content is no more specific than that the core tenets of that system are true.

With that said, this won't be a claim that can be made on behalf of religious belief-systems that don't avow belief in a deity capable of such supernatural guidance, and so perhaps a version of the argument from cultural contingency cast in terms of narrower process types can still succeed against *those* sorts of religious beliefs. I take it, though, that the argument from cultural contingency has usually been thought of as threatening *all* varieties of culturally instilled religious beliefs. That is certainly how Mill and Kitcher seem to see things. Accordingly, let us consider an approach on which a wider (more general) process type is held to be the relevant one.

### §7. Wider Process Type Individuation

An appropriately wide process type whose specification involves no mention of particular testimony chains or particular religious texts, traditions, or rituals might be something like *accepting the testimony of a religious tradition*. A process type like *this* will surely be wide enough so as to be instantiated by a large variety of mutually incompatible religious belief tokens, and thus have an insufficiently high truth-ratio. Indeed, consider again Mill's claim that "*the same causes* which make [a man] a churchman in London would have made him a Buddhist or a Confucian in Peking [emphasis mine]" (1991, pp. 229-230). As far as I can tell, *accepting the testimony of a religious tradition* is about as narrow a process type as can be selected here without rendering Mill's claim false; any more detail—any mention of *particular* religious teachers, texts, rituals, or experiences—and the pair of cases Mill has in mind won't rightly be described as involving the *same type* of causal process. Similarly, Kitcher

would have to have in mind something like this fairly wide process type in order for it to be true that, as he claims, “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 [emphasis mine]” (2011, p. 26). Again, Aboriginal and Christian ways of coming to believe their respective doctrines seem to have in common nothing more specific than that they both involve a tendency to accept the passing on of religious teachings and practices from one generation to the next. This strategy we are now considering, then, does indeed appear to fulfil the desideratum of picking out a process type that encompasses many mutually incompatible religious beliefs. But I now want to draw out three implications of this approach—implications at which the advocate of the argument from cultural contingency might understandably balk.

### *§7.1. Implications for Moral Epistemology*

One implication of advancing the argument from cultural contingency on the basis of a wide type-selection approach is that *moral* belief-forming process types at an equivalent level of generality, some of whose outputs we tend to be very reluctant to deem non-knowledge, have truth-ratios that are barely higher than that of *accepting the testimony of a religious tradition*. If the type *accepting the testimony of a religious tradition* doesn’t have a high enough truth-ratio to yield knowledge, then nor do these other process types. I won’t dwell long on this point, since it has already received significant discussion. Plantinga, for instance, gives the example of someone’s belief that *racial discrimination is wrong*, which has been instilled by her cultural upbringing (1997, pp. 297-298). Peter van Inwagen and Jerome Gellman make similar points concerning political beliefs, which typically one has arrived at partly by way of upbringing, such as the belief that *democracy is a fairer system of government than totalitarianism*.<sup>25</sup> The (wide) process type at issue here, something like *accepting cultural moral/political values*, looks to be responsible for nearly as many mutually incompatible belief outputs as does *accepting the testimony of a religious tradition*. By parity reasoning, then, it looks like the defender of the argument from cultural contingency should conclude that one doesn’t know, on that basis, that (say) racial discrimination is wrong. Of course, one might bite the bullet regarding culturally instilled moral beliefs, and yet maintain that moral beliefs formed by way of rational reflection are in good shape, nonetheless. The problem is, the equivalently wide process type here, *rational moral/political reflection*, appears to issue about as many mutually incompatible outputs as the aforementioned types.

---

<sup>25</sup> See van Inwagen (1995, p. 238) and Gellman (1993, p. 350).

A way around this problem might be to adopt some form of moral anti-realism, according to which moral propositions are true or false just in virtue of facts about the moral beliefs of the relevant moral community. On such a view, one could plausibly get the result that moral belief-formation by way of one's upbringing can often result in knowledge: given social constructivism, process type reliability would require only that the type *accepting cultural moral/political values* (and, indeed, the type *rational moral/political reflection*) doesn't produce more than a small proportion of mutually conflicting beliefs *within* a given moral community; it wouldn't require such consistency *across* moral communities. There is the worry with this constructivist approach that one will be committed to saying that an appalling (say, fascistic) moral belief that is held by the majority of a certain community will both be true and constitute knowledge for members of that community. Such a consequence could perhaps be avoided by enlarging moral communities both spatially and temporally to the point that such beliefs are held only by a minority, thus rendering them false and so non-knowledge. This enlargement of moral communities will, however, presumably also have the effect of increasing the overall degree of disagreement occurring within them, thus decreasing the overall truth-ratio of the process types *accepting cultural moral/political values* and *rational moral/political reflection*, perhaps to the point that they won't be counted reliable.

### *§7.2. Implications for the Theory of Knowledge*

Another implication is this. If one picks out a wide process type like *accepting the testimony of a religious tradition* as the salient type for the purpose of criticising religious beliefs, then it seems that on pain of unjustified arbitrariness, one is committed to selecting the relevant process types at an *equivalent* width when non-religious beliefs are concerned. But a wide type-selection approach is problematic, as reliabilists well know. Becker notes that "The most telling cases occur when an agent forms two (or more) true beliefs from what we would [on this wide type approach] deem identical process types, but intuitively one of the beliefs constitutes knowledge and the other does not" (2008, p. 362). This is known as the *no distinction* problem. Consider two cases of true belief: In case A, I look at a nearby tree through a clear windowpane in good lighting and without being under the influence of any impairing substances, and correctly believe that *there is a tree before me*. In case B, I look at a nearby tree through a filthy windowpane in poor lighting whilst heavily intoxicated and correctly believe that *there is a tree before me*. What is the *salient* process type that is instantiated in each case? Well, as I said, the defender of the argument from cultural contingency is committed to saying that the salient process

type in each of these cases is a process type of approximately the same width as *accepting the testimony of a religious tradition*. And so, unfortunately, it looks like the defender of the argument from cultural contingency is committed to saying that the relevant type for case A *is the same as* the relevant process type for case B—something like *visually perceiving nearby objects*. These two belief tokens both fall under the wide process type *visually perceiving nearby objects*, and yet they clearly differ enormously with respect to their reliability.<sup>26</sup> Even given the truth of the beliefs in both cases, we intuitively want to say that case A is a case of knowledge and hence, reliably formed belief, but that case B isn't. Thus, the difference between the epistemic statuses of these two cases will have to be accounted for in terms of something other than the reliability of process types. The situation seems to be as follows:

- (i) Given the wide type-selection approach needed by the defender of the argument from cultural contingency, there are pairs of cases (where case A is intuitively a case of knowledge, and hence, reliably formed belief, and case B is not) that both fall under the same salient process type.
- (ii) Given (i), the salient process type in case A is reliable *if and only if* the salient process type in case B is reliable.
- (iii) Given (ii), denying that the salient process type in case B is reliable entails denying that the salient process type in case A is reliable.
- (iv) Denying that the salient process type in case A is reliable is implausible, since it implies A isn't a case of knowledge (on the assumption that type reliability is required for knowledge).
- (v) So, the difference in epistemic status between case A and case B must be explained with reference to something other than process type reliability.

One natural thought at this point is that we could look to modal reliability (i.e. safety) to explain why A is a case of knowledge and B is not. In case A my belief couldn't easily have been false, whilst in case B it could. That way, we can allow that in both A and B the salient process *type* is reliable, but still say, as we intuitively want to say, that case A is a case of knowledge and B is not.

Unfortunately, this strategy of placing the burden of explanation entirely upon the safety condition will not work. There are cases in which a subject has a true belief that is safe (and, indeed, internally justified too), but which is intuitively *not* an

---

<sup>26</sup> Of course, there are other, narrower types, such that these two beliefs fall under different types. But because the argument from cultural contingency requires the *salient* type when it comes to religious beliefs to be something as wide as *accepting the testimony of a religious tradition*, the defender of the argument from cultural contingency is in the position of having to say that the *salient* type here is something equivalently wide, namely, *visually perceiving nearby objects*.

instance of knowledge. One much discussed case of this sort is due to Duncan Pritchard, and involves a man, “Temp”, who forms a belief about the temperature of the room by looking at a thermometer that he justifiably believes to be working properly, but which is in fact broken and fluctuating randomly (2009, pp. 40-41; 2012, p. 260). Also unbeknownst to Temp, an unseen agent is constantly adjusting the temperature of the room to ensure that it always matches Temp’s belief. It appears, then, that Temp’s true belief is safe; that there is no nearby world in which Temp forms a false belief, whose content and causal history are relevantly similar to those of his actual belief.<sup>27</sup> And yet it appears very doubtful that Temp really *knows* the temperature.<sup>28</sup> In order to explain why cases like this are not knowledge, Pritchard and others have been led to add to their analyses of knowledge a special sort of process type reliability condition, a *cognitive virtue/achievement* requirement, according to which the truth of one’s belief must be creditable to one’s use of a reliable cognitive process.<sup>29</sup> But if one is pursuing a *wide* type-selection policy—as I have suggested the defender of the argument from cultural contingency needs to do—it doesn’t look like one can offer any explanation for Temp’s lack of knowledge in terms of process *type* reliability either. This is because Temp’s process type, picked out at a level of generality roughly equivalent to *accepting the testimony of a religious tradition*, will be something like *consulting a measurement device*.<sup>30</sup> This latter is a process type that we had better be in a position to count as reliable or else face some rather implausible sceptical consequences—i.e., that people *never* gain knowledge by consulting measurement devices. In short, if the defender of the argument from cultural contingency pursues the wide type-selection approach that she seems to need, she will be committed to an account of knowledge that faces counterexamples in the form of Temp-style cases *unless* she can locate an additional necessary

---

<sup>27</sup> This assumes, of course, that not easily could the unseen agent have failed to be present to fit the temperature of the room to Temp’s belief about it—but there is no reason why that modal fact couldn’t be built into the case.

<sup>28</sup> The recipe for generating such cases is straightforward: have a person form a true belief on the basis of some indicator, which she justifiably believes to be reliable; have that indicator be defective, such that it intuitively cannot lead to knowledge; and have the situation feature (again, unbeknownst to the person) a modally stable causal agency that consistently alters the truth-making facts to fit the defective indicator’s deliverances.

<sup>29</sup> Pritchard argues that these cases call for a cognitive ability condition, a requirement that the truth of one’s belief be creditable to a truth-conducive cognitive ability, where a cognitive ability is a special sort of belief-forming process type (2012, pp. 260-264). Becker claims that such cases involve process luck—that is, luck that one was subject to a reliable process token, given that many of the other tokens belonging to the same (relevant) type are unreliable (2007, pp. 32-36). He argues that “Eliminating egregious cases of process luck entails a process [type] reliabilism condition” (2007, p. 77).

<sup>30</sup> If one wants to say that in Temp’s case, the salient process type is *consulting a broken measurement device*, hence subdividing the type *consulting a measurement device* into narrower types based on truth-conduciveness, then one ought similarly to subdivide the type *accepting cultural religious testimony* into such narrower types: *accepting truth-conducive cultural religious testimony* and *accepting falsity-conducive cultural religious testimony*. But, of course, an argument from cultural contingency depends for its success upon *not* individuating process types in this manner.

condition on knowledge: one which is distinct from either an internal justification condition, a modal reliability condition, or a process type reliability condition, and which both explains such cases and is not obviously too restrictive. That is a tall order indeed.

### *§7.3. Implications for Science*

There is a third, seemingly even tougher problem for the defender of the argument from cultural contingency, which concerns the deliverances of science. The problem is roughly that two process types at a level of generality approximately equivalent to *accepting the testimony of a religious tradition*, namely, *constructing and testing scientific hypotheses* (for the scientists themselves) and *accepting expert scientific testimony* (for the laypeople), encompass scientific belief tokens from *all* eras of scientific research. Given that many scientific beliefs from bygone eras are incompatible with modern scientific beliefs,<sup>31</sup> the overall truth-ratio of the types *constructing and testing scientific hypotheses* and *accepting expert scientific testimony* is not going to be all that high; that truth-ratio can *at most* be equal to the greatest number of scientific beliefs that could all be true at once divided by the total number of scientific beliefs. Even without assuming the falsity of any particular scientific beliefs, we can see that, given such a wide type-selection policy, the truth-ratio of the aforementioned scientific process types will get dragged down due to the mutual inconsistencies between scientific beliefs of different eras. Of course, a much more fine-grained individuation of scientific belief-forming process types—which, for instance, includes in the description of a process type the specific experimental methods, techniques, background theory, and equipment used—wouldn't lump all scientific beliefs together under one type, and thus wouldn't have the upshot that scientific beliefs all stand or fall together with respect to the process type reliability condition. But were the defender of the argument from cultural contingency to adopt this fine-grained approach to type selection, she would have to do likewise for religious beliefs, thus mentioning specific religious texts and testimony chains in her description of the relevant process type, and I have already suggested that that approach won't work (not against monotheistic religious beliefs, at any rate), given the needed assumption of the truth of the targeted religious beliefs.

With that said, one fairly striking difference between religious diversity, on the one hand, and the diversity of expert scientific opinion, on the other, is that the latter phenomenon is principally *diachronic*; that is to say, disagreement between scientific experts, though to some degree *synchronic* (occurring between contemporaries *at a*

---

<sup>31</sup> For a colourful list of now abandoned scientific theories, see Laudan (1981).

given time), is much more pervasive *across* historical time-periods. On the other hand, religious diversity is clearly both diachronic and synchronic; it is both the case that modern people have substantially different religious beliefs to peoples living thousands of years ago, *and* that various peoples hold substantially different religious beliefs from one another at any given point in history. Accordingly, one move the defender of the argument from cultural contingency might make is to index process types to *times*, or time periods. The details would need some working out, in particular, concerning how times are to be individuated. At any rate, the basic thought is that by indexing process types to time periods, the epistemically salient process type employed by Ptolemaic astronomers—*constructing and testing scientific hypotheses during  $t_x$* —can be distinguished from the epistemically salient process type used by modern day cosmologists—*constructing and testing scientific hypotheses during  $t_y$* —and thus the unreliability of the former need not spread to the latter. Of course, this would mean that during what Thomas Kuhn has termed “revolutions” in the history of science, during which a once-dominant theoretical paradigm is overthrown and various theoretical frameworks compete to become the new dominant paradigm,<sup>32</sup> it will be harder for a scientific practitioner to know the truth of a given hypothesis (assuming the hypothesis in question is even true at all). After all, a process type such as *constructing and testing scientific hypotheses during  $t_{\text{some particular scientific revolution}}$*  would be responsible for a significant proportion of mutually inconsistent outputs and thus would *at best* have only a moderate truth-ratio. But this seems in fact to be a fitting consequence. Now, consistency demands that religious belief-forming process types also be indexed to times, but since religious diversity is a synchronic phenomenon to a much greater degree than is scientific expert disagreement, this shouldn’t undermine the argument from cultural contingency; the type *accepting cultural religious testimony during  $t_{\text{the early twenty-first-century}}$*  will still yield the large proportion of mutually inconsistent outputs that the defender of the argument from cultural contingency needs it to.

Still, one might rightly worry that even with process types indexed to time periods, scientific beliefs whose epistemic credentials are seemingly very strong will get unfairly dragged down by the fact that, at any given time, some areas of scientific research lag considerably behind others, and yet on this wide type-selection approach they are all lumped together under one process type. Why, for instance, should the epistemic status of Charles Darwin’s (intuitively, well-justified) beliefs about biological heredity be affected by the fact that at the time he was working, some fairly far removed areas of science such as astronomy had produced many false beliefs? One solution to this worry would be to index process types to individual *propositions*; to

---

<sup>32</sup> See Kuhn (1996).

specify types according to the following schema: *forming a belief about whether p by way of a causal process type C during t*.<sup>33</sup> In this way, the process type reliability pertaining to a given scientific belief will be determined only by the truth-ratio among beliefs held during the same time period concerning just that one proposition. This move would go some way to alleviating the concern about degenerate fields of scientific research dragging down progressive fields. Again, *ideally* one would have recourse in one's specification of the causal process type to the specific experimental methods and theoretical background employed, but this isn't available on the wide type-selection approach that the defender of the argument from cultural contingency needs. Also, note that this move of indexing process types to particular propositions has the following cost for the defender of the argument from cultural contingency. It's admittedly true that a process type such as *forming a belief about whether God is triune by way of accepting the testimony of a religious tradition during t<sub>n</sub>* will yield a pretty high proportion of mutually inconsistent outputs and so at best have an insufficiently high truth-ratio for knowledge, and it's true that this will in general be the case whenever the proposition is one believed only by some but not other religious traditions. But when it comes to propositions that are believed by virtually all religious traditions, it looks as though the proportion of mutually inconsistent outputs of the relevant process type will be very low, and so given the truth of such a proposition, the type reliability condition will be satisfied.

## §8. Conclusion

To summarise the foregoing, I began by arguing that the core argument, in my reconstruction of an argument from cultural contingency, cannot—without either hugely restricting its scope or giving rise to unintended sceptical consequences—be framed in terms of the unsafety of religious beliefs. Rather, the core argument must be conducted in terms of the unreliability of religious belief-forming process types. But I suggested that the argument does not significantly threaten monotheistic religious beliefs—not beliefs of core tenets of a religious system, at any rate—if the salient process type is held to be one whose specification involves particular religious texts or particular testimony chains. Selecting a wider process type, whose specification includes no such particulars, can indeed achieve the desired result: locating a process type a significant proportion of whose outputs are mutually inconsistent, thus whose overall truth-ratio cannot, even given the truth of one or

---

<sup>33</sup> One might try to index a process type to something broader than a single proposition, such as a “subject matter”, but it is very hard to see how the individuation of subject matters won't simply end up being *ad hoc*.

other consistent subset of the process's outputs, be high enough to satisfy a process type reliability condition. But I claimed that the problems generated by this wide approach to selecting process types can be solved only if one is committed to the following: (1) with respect to moral beliefs, either the claim that humans have unreliably formed moral beliefs and thus lack moral knowledge, or the constructivist claim that the truth conditions for moral beliefs are determined by the consensus of moral communities; (2) an analysis of knowledge that faces counter-examples (in the form of Temp-style cases) *unless* a genuinely necessary condition for knowledge can be identified which is distinct from either a modal reliability condition, a process type reliability condition, or an internalist justification condition (all three of these conditions being met in Temp-style cases, given a wide type-selection approach); (3) with respect to scientific beliefs, either the view that scientific beliefs formed at different periods in history and resulting from widely disparate fields of research are all equivalent in terms of their reliability, or an approach that indexes process types to times and propositions and so concedes that on the assumption that they are true, beliefs of religious propositions that are held by virtually all religious traditions can satisfy the process type reliability condition. Thus, I conclude that the argument from cultural contingency succeeds only to the extent that these philosophical positions are plausible.

## References

- Becker, K. (2007). *Epistemology modalized*. London: Routledge.
- Becker, K. (2008). Epistemic luck and the generality problem. *Philosophical Studies*, 139, 353-366.
- Bergmann, M. (2006). *Justification without awareness: a defense of epistemic externalism*. Oxford: Oxford University Press.
- Bogardus, T. (2013). The problem of contingency for religious belief. *Faith and Philosophy*, 30(4), 371-392.
- Brueckner, A., & Buford, C. (2013). Becker on epistemic luck. *Philosophical Studies*, 163, 171-175.
- Dretske, F. (1971). Conclusive reasons. *Australasian Journal of Philosophy*, 49(1), 1-22.
- Gellman, J. (1993). Religious diversity and the epistemic justification of religious belief. *Faith and Philosophy*, 10(3), 345-364.
- Hawthorne, J. (2004). *Knowledge and lotteries*. Oxford: Oxford University Press.
- Kitcher, P. (2011). Challenges for secularism. In G. Levine (Ed.), *The joy of secularism: 11 essays for how we live now* (pp. 24-56). Princeton, NJ: Princeton University Press.
- Kuhn, T. S. (1996). *The structure of scientific revolutions* (3rd ed.). Chicago, IL: University of Chicago Press.

- Lasonen-Aarnio, M. (2010). Unreasonable knowledge. *Philosophical Perspectives*, 24(1), 1-21.
- Laudan, L. (1981). A confutation of convergent realism. *Philosophy of Science*, 48(1), 19-49.
- Lewis, D. (1973). *Counterfactuals*. Cambridge, MA: Harvard University Press.
- Mawson, T. J. (2009). Mill's argument against religious knowledge. *Religious Studies*, 45(4), 417-434.
- Mill, J. S. (1991). On liberty. In J. M. Robson (Ed.), *The collected works of John Stuart Mill* (Vol. XVIII). Toronto, ON: Toronto University Press.
- Nozick, R. (1981). *Philosophical explanations*. Cambridge, MA: Harvard University Press.
- Plantinga, A. (1995). Pluralism: a defence of religious exclusivism. In T. D. Senior (Ed.), *The rationality of belief and the plurality of faith: essays in honor of William P. Alston* (pp. 191-215). London: Cornell University Press.
- Plantinga, A. (1997). Ad Hick. *Faith and Philosophy*, 14(3), 295-298.
- Pollock, J. L. (1986). *Contemporary theories of knowledge*. Totowa, NJ: Rowman & Littlefield.
- Pritchard, D. (2005). *Epistemic luck*. Oxford: Oxford University Press.
- Pritchard, D. (2009). Safety-based epistemology: wither now? *Journal of Philosophical Research*, 34, 33-45.
- Pritchard, D. (2012). Anti-luck virtue epistemology. *Journal of Philosophy*, 109(3), 247-279.
- Stalnaker, R. C. (1975). Causation and conditionals. In E. Sosa (Ed.), *Oxford readings in philosophy* (pp. 165-179). New York, NY: Oxford University Press.
- van Inwagen, P. (1995). Non est Hick. In T. D. Senior (Ed.), *The rationality of belief and the plurality of faith: essays in honor of William P. Alston* (pp. 216-241). London: Cornell University Press.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford: Oxford University Press.
- Williamson, T. (2009). Reply to John Hawthorne and Maria Lasonen-Aarnio. In P. Greenough, & D. Pritchard (Eds.), *Williamson on knowledge* (pp. 313-329). Oxford: Oxford University Press.
- Williamson, T. (2014). Knowledge first. In M. Steup, J. Turri, & E. Sosa (Eds.), *Contemporary debates in epistemology* (pp. 1-10). Oxford: Wiley Blackwell.