10-1-2013

The Problem Of Contingency For Religious Belief

Tomas Bogardus

THE PROBLEM OF CONTINGENCY
FOR RELIGIOUS BELIEF

Tomas Bogardus

In this paper, I hope to solve a problem that's as old as the hills: *the problem of contingency for religious belief.* Paradigmatic examples of this argument begin with a counterfactual premise: had we been born at a different time or in a difference place, we easily could have held different beliefs on religious topics. Ultimately, and perhaps by additional steps, we're meant to reach the skeptical conclusion that very many of our religious beliefs do not amount to knowledge. I survey some historical examples of this argument, and I try to fill the gap between the counterfactual premise and the skeptical conclusion as forcefully as possible. I consider the following possibilities: there are no additional steps in the argument; or there are and they concern the alleged safety condition on knowledge, or the alleged non-accidentality condition on knowledge, or the unclarity produced by disagreement. On every possibility, the argument from the counterfactual premise to the conclusion of widespread skepticism is invalid. It seems, then, that there is no serious problem of contingency for religious belief.

The Problem "In the Wild"

"*But you only believe that because . . .*" begins a common objection to a belief's rationality, and it ends by pointing out a factor that is irrelevant to the truth of the belief. Typically, this is an origin story that highlights the historical contingency of the belief: "you were raised Catholic," or "you studied at Oxford," or "you read only the *Daily Worker*." Epistemologists have recently discovered a burning interest in this type of objection, which we might call "the problem of contingency" for our beliefs, or "the problem of irrelevant causal factors," or "the problem of historical variability."[1] Our

---

[1]See for example Nathan Ballantyne, "The Problem of Historical Variability," in *Disagreement and Skepticism*, ed. Diego Machuca (Routledge Studies in Contemporary Philosophy, 2013); Adam Elga, "Lucky to Be Rational," unpublished manuscript; Joshua Schechter, "Luck, Rationality, and Explanation: A Reply to Elga's 'Lucky to Be Rational,'" unpublished manuscript; Miriam Schoenfield, "Permission to Believe: Why Permissivism Is True and What It Tells Us about Irrelevant Causal Influences on Belief," *Noûs* (forthcoming); and Roger White, "You Just Believe that Because . . . ," *Philosophical Perspectives* 24 (2010), 573–615.

There has also be a surge of interest in what are often called "evolutionary debunking arguments." Does our evolutionary history constitute a problematic "irrelevant causal factor" for some of our beliefs, especially our moral beliefs? See for example Richard Joyce, *The*

371

discussion below will bear on much of this new work, but my primary goal in this paper is to solve the problem of contingency specifically *for religious belief*, and that problem is as old as philosophy itself. One finds an embryonic statement of it in the pre-Socratic philosopher Xenophanes (DK 21 B15):

> If oxen and horses and lions had hands and were able to draw with their hands and do the same things as men, horses would draw the shapes of gods to look like horses and oxen to look like oxen, and each would make the gods' bodies have the same shape as they themselves had.[2]

Elsewhere (DK 21 B16), Xenophanes points out, with wry amusement, that Ethiopians say the gods are "snub-nosed and black," whereas the Thracians say the gods are "blue-eyed and red-haired." These statements were a part of his larger attack on Greek popular religion, meant to support skepticism toward popular beliefs about the gods. Xenophanes's idea, I take it, is that our opinions on religious matters are shaped—to an embarrassing and perhaps undermining degree—by matters that are completely on the wrong side of the question. Change certain facts about the religious *believer*—facts untethered to the objects of the belief—and, he claims, the content of the religious belief will change in predictable ways, ways that may easily alter the belief's truth-value. We should therefore be skeptical of these religious beliefs, he thinks.

Breezing along the timeline of philosophy, we find the thought maturing in John Stuart Mill:

> And the world, to each individual, means the part of it with which he comes in contact; his party, his sect, his church, his class of society . . . [I]t never troubles him that mere accident has decided which of these numerous worlds is the object of his reliance, and that the same causes which make him a Churchman in London, would have made him a Buddhist or a Confucian in Pekin.[3]

Mill's concern is that we arrive at our religious beliefs accidentally, driven by causes which are—for Mill as for Xenophanes before him—in a disturbing way independent of the truth of our beliefs. Change the setting of one's biography from London to Beijing, and the causes which actually drove one toward Christianity (say), would have driven him towards Buddhism

---

Evolution of Morality (Cambridge, MA: MIT Press, 2006); Justin Clarke-Doane, "Morality and Mathematics: The Evolutionary Challenge," *Ethics* 122 (2012), 313–340; Guy Kahane, "Evolutionary Debunking Arguments," *Noûs* 45 (2011), 103–125; and Sharon Street, "A Darwinian Dilemma for Realist Theories of Value," *Philosophical Studies* 127 (2006), 109–166. Below, I'll say nothing about the import of our discussion for those arguments. Yet evolutionary debunking arguments are structurally quite similar to the problem of contingency for religious belief. And so perhaps the defense I give here of the rationality of religious belief could be deployed *mutatis mutandis* as a defense of the rationality of moral realism in the face of evolutionary debunking arguments. Or perhaps not. Working out those connections will be left to the reader, or to a future paper.

[2]H. Diels and W. Kranz. *Die Fragmente der Vorsokratiker*, 6th ed. (Zürich: Weidmann, 1951).

[3]John Stuart Mill, *On Liberty* (London: Longmans, Green, and Co., 1913), 10–11.

or Confucianism. That we arrive where we actually do is, therefore, a mere accident. And this should trouble us, according to Mill, though it too rarely does.

Finally, reflect on two contemporary arguments, where our problem finds its fullest flower.[4] The first is from John Hick:

> [R]eligious allegiance depends in the great majority of cases on the accident of birth: someone born into a devout Muslim family in Pakistan is very likely to be a Muslim, someone born into a devout Hindu family in India to be a Hindu, someone born into a devout Christian family in Spain or Mexico to be a Catholic Christian; and so on. The conclusion I have drawn is that a "hermeneutic of suspicion" is appropriate in relation to beliefs that have been instilled into one by the surrounding religious culture.[5]

As with Mill, Hick appeals to a troubling accident of birth. But, as a development from Mill, Hick tells us a bit more about the causes which drive us to our religious beliefs: they are "instilled" in us by the surrounding religious culture. This is why, according to Hick, changing the culture changes the resulting religious belief: different inputs, fed into the same belief-producing method, yield predictably different outputs. And for Hick—as for Mill and Xenophanes before him—this accidentality, this historical contingency, rationally requires us to be suspicious of our religious beliefs. In Hick's own case, this suspicion resulted in his abandoning the religious beliefs with which he was instilled. But it's unclear whether he thinks the suspicion requires that course in all cases.[6]

And witness an even more recent statement of the problem of contingency for religious belief, from Philip Kitcher:

> Most Christians have adopted their doctrines much as polytheists and the ancestor-worshipers have acquired theirs: through early teaching and socialization. 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

---

[4]For another recent statement of the problem, see Antony Flew, "The Presumption of Atheism," in *The Presumption of Atheism and Other Philosophical Essays on God, Freedom, and Immortality* (New York: Barnes & Noble, 1976): "One positive reason for being especially leery towards religious opinions is that these vary so very much from society to society; being, it seems, mainly determined, as Descartes has it, 'by custom and example.' The phrase occurs, in Part II of his Discourse on the Method, almost immediately after the observation: 'I took into account also the very different character which a person brought up from infancy in France or Germany exhibits, from that which . . . he would have possessed had he lived among the Chinese or with savages.'"

[5]John Hick, "The Epistemological Challenge of Religious Pluralism," *Faith and Philosophy* 14 (1997), 281.

[6]He later says (ibid.) that the "relativity of religious belief to the circumstances of birth" should "warn us to look critically at such claims." Presumably this requires that we take a second look at our religious beliefs, to reevaluate our grounds. But exactly what Hick recommends should happen to our confidence during this second look—perhaps nothing, perhaps significant downgrade, perhaps suspension of belief—is unclear.

Resurrection. The symmetry is complete. . . . Given that they are all on a par, we should trust none of them.[7]

Kitcher adds detail to Hick's story, detail concerning how religious beliefs are meant to be instilled by the surrounding culture. According to Kitcher, it's through "early teaching and socialization." And, again, we reach the same conclusion: reflection on the way in which these religious beliefs are formed should lead us to doubt them. Noticing the origin of these religious beliefs provides a *defeater* for them, a reason to substantially lower our confidence in our religious beliefs, and perhaps give them up entirely.

No doubt there is a broad thicket of thorny issues here, and many distinct arguments each with a claim to be *the* problem of contingency for religious belief. My project below is to state the problem as forcefully as possible, to find a version of the argument with the greatest promise to defeat religious knowledge.

The tie that binds these sundry arguments is something like this: the connection between many of our religious beliefs and the truth isn't good enough for knowledge, even if those beliefs are true. And, when we are made aware of this frail connection, we acquire an *undermining defeater* for the rationality of many if not all of our religious beliefs. In the face of that defeater, we can no longer sensibly claim to *know* that these religious beliefs are true. In what follows, then, I will construe the problem of contingency for religious belief as a problem for religious *knowledge*, i.e., as an argument for the conclusion that religious beliefs meeting certain widespread conditions do not rise to the level of knowledge, even if they are true. The question for us is, on its strongest interpretations, does the argument succeed? I will show that it doesn't.

A Bare Counterfactual Argument

Let's make a first attempt to lay out the problem of contingency. One finds explicit in Mill's and Kitcher's statements—and implicit in Xenophanes's and Hick's statements—certain counterfactuals to the effect that, if you had been born at a different time and in a different place, your religious beliefs would have been different. Had Christians been born among the Aboriginal Australians, Kitcher says, these Christians would have had radically different religious beliefs, beliefs concerning Dreamtime rather than the Resurrection.

Upon reading Mill, Hick, and Kitcher, one could be forgiven for thinking that the argument hurries from just that counterfactual premise directly to the skeptical conclusion. That, in fact, is how Alvin Plantinga reads Hick, attributing to Hick something like this argument:

1) If you had been born and raised elsewhere, else when, you would have had different religious beliefs.

---

[7]Philip Kitcher, "Challenges for Secularism," in *The Joy of Secularism: 11 Essays for How We Live Now*, ed. George Levine (Princeton, NJ: Princeton University Press, 2011), 26.

Therefore,

2) Your religious beliefs don't count as genuine knowledge.[8]

Let's dub this "the Bare Counterfactual Argument." No doubt the argu-
ment produces, in Plantinga's words, the "sense of intellectual vertigo"
we're looking for, while also grasping tightly the virtue of simplicity. But
the argument does so at the expense of validity. In his response, Plantinga
points out that he believes he was born in Michigan. Yet had he been born
elsewhere and else when—in Madagascar, say—he would not have be-
lieved he was born in Michigan. So, by the general principle that is meant to
carry us from (1) to (2), it seems to follow that Plantinga doesn't genuinely
know he was born in Michigan. Since that's absurd, we should conclude
that the inference from (1) to (2) is invalid.

And of course one may multiply instances of this counterexample:
there are many beliefs that seem to satisfy the general form of the argu-
ment, but for which we would judge the conclusion false. As Plantinga
points out, that seems to be the case for Hick with respect to his own plu-
ralistic religious beliefs. Says Plantinga: "Pluralism isn't and hasn't been
widely popular in the world at large; if the pluralist had been born in
Madagascar . . . he probably wouldn't have been a pluralist." And so it
seems that Hick's position is self-defeating. On the one hand, premise (1)
seems true of Hick's religious beliefs, beliefs which he takes to constitute
knowledge. On the other hand, he thinks the above skeptical argument
is valid. That's trouble. And what goes for Hick seems to go for religious
skeptics generally.

These seem to me compelling objections to this natural "bare counter-
factual" interpretation of the problem of contingency. Yet, in his reply to
Plantinga, Hick offered a defense that is subtle, underappreciated, and—I
believe—successful. There, Hick clarifies a reading of the problem of
contingency different from (1)–(2) above, a reading which allows him to
evade Plantinga's objections. But, I will argue, Hick evades these objec-
tions at a high cost.

Here is Hick's response to Plantinga's objection concerning self-defeat:
"This is true; but how relevant is it? One is not usually a religious plu-
ralist as a result of having been raised from childhood to be one, as (in
most cases) one is raised from childhood to be a Christian or a Muslim
or a Hindu, etc."[9] Thus does Hick run a highlighter over his original con-
clusion that we should be skeptical of "beliefs that have been instilled
into one by the surrounding religious culture," emphasis on *instilled*.
And this seems in keeping with the quotations from Mill and especially
Kitcher, whose counterfactual says that had Christians been born among

---

[8]Alvin Plantinga, "Pluralism: A Defense of Religious Exclusivism," in *The Rationality
of Belief and the Plurality of Faith: Essays in Honor of William P. Alston*, ed. Thomas D. Senor
(Ithaca, NY: Cornell University Press), 191–215.

[9]Hick, "The Epistemological Challenge of Religious Pluralism," 277–286.

aboriginal Australians, they would believe *in just the same ways* doctrines incompatible with Christianity. Evidently, the problem for our religious beliefs is meant to be the *way* in which they've been produced, the *method* by which we formed them: early teaching and socialization, as Kitcher puts it. And since Hick's pluralist religious beliefs were not formed by that questionable method, those beliefs are off the hook. To my mind, by this exemption Hick evades Plantinga's self-defeat objection, and in an acceptably non-ad-hoc way.

But what of Plantinga's belief that he was born in Michigan? Wasn't that belief instilled into him since childhood? And consider Plantinga's final response to Hick, in which he notes that he also holds moral beliefs (e.g., the belief that racial intolerance is wrong).[10] These beliefs seem to count as genuine knowledge, and yet they also seem to have been instilled into him since childhood. These are serious objections, to be sure. But I believe that Hick's appeal to methods can snatch the problem of contingency from the teeth of both of these apparent counterexamples.

The problem Hick points to, if there is one, is with belief-producing methods that could easily go awry. The charge is that religious believers commonly use such methods. To test whether a method could easily go awry, we imagine changing certain inputs to the method—facts about us and our environs—that shouldn't bear on the question that the method is employed to answer. In the religious case, we imagine changing where and when we were born, for instance. If such changes to the input of the method change the content of the output belief *in ways that may easily alter the output's truth-value*, then we have a problem. Even assuming one's religious beliefs are true, one must admit that her method of arriving at these beliefs could easily have led her into falsehood. That's the idea, anyway.

But it's clear that the method by which Plantinga formed his belief that he was born in Michigan is *not* suspicious in this way. For changing the inputs in the relevant way—where and when he was born—does indeed change the content of the output belief, but *not* in a way that could easily alter the output's truth-value. Imagine Plantinga's having been born and raised in Madagascar, trusting his parents on the question of his birthplace. The content of his belief changes—actually, Michigan; counterfactually, Madagascar—but the method doesn't *go awry*: it would still produce a true belief in the counterfactual scenario. Contrast these with the religious beliefs Plantinga would likely have held had he been born in Madagascar.[11] By Plantinga's own lights—that is, assuming the truth of his actual religious beliefs—had he been born in Madagascar, a great number of his religious beliefs could easily have been false. And that's the rub.

---

[10]Alvin Plantinga, "Ad Hick," *Faith and Philosophy* 14 (1997), 295–298.

[11]Plantinga is a particular kind of Protestant Christian. According to Wikipedia, the Malagasy are roughly 20 percent Protestant and 40 percent Christian. Of course, if Madagascar's demographics prove not to provide a good example for Plantinga here, another could be easily supplied.

The skeptic's question is not, then, as Plantinga thinks: "Could your religious-belief-producing methods easily have produced *different* religious beliefs?" The skeptic's question is rather: "Could they easily have produced beliefs that would be *false* by your own lights?" It's no problem if changing certain inputs to a method—facts about you and your environs, for example—changes the outputs of the method. For those facts may bear on the question that the method has been employed to answer, and so the method's varying outputs may track the truth. This is what happens in Plantinga's Michigan/Madagascar case. Rather, the alleged problem opens its steely maw when the method you've employed to answer a question might easily lead you into error, were certain facts that do not bear on the question to change. And that's the case, according to Hick, Kitcher, et al., with religious belief. So here too a slight modification of (1) will avoid Plantinga's counterexample.

Now, what of Plantinga's belief that racial intolerance is wrong? One might run this Plantingian objection to the bare counterfactual argument above like so: Plantinga believes that racial intolerance is wrong. That belief counts as genuine knowledge. But the methods used to acquire this belief, namely early teaching and socialization, are "modally fragile" in the problematic way: had Plantinga been born elsewhere, else when, he easily might have believed falsely, by his own lights. And the time and place of his birth do not bear on the question of whether racial intolerance is wrong. But then the bare counterfactual argument drives us to the conclusion that Plantinga doesn't really know that racial intolerance is wrong, which contradicts our earlier assumption. But that assumption is true: Plantinga clearly does know. The bare counterfactual argument must, therefore, be invalid. So ends the objection.

Again, I believe that an appeal to methods can rescue Hick from the alleged counterexample. If Plantinga really does believe that racial intolerance is wrong merely on the basis of having been raised from childhood to believe it—i.e., merely on the basis of passively, unreflectively receiving that testimony—then it's not implausible that Plantinga—now aware of the diversity of opinion on the question and the fragility of his method—doesn't genuinely know that racial intolerance is wrong. *That's* the assumption that should be jettisoned, Hick could plausibly say. But, more realistically, Plantinga believes that racial intolerance is wrong by means of a different, better method: rational reflection and intuition. He can *just see* and thereby know that it's wrong, in the way that we can *just see* and thereby know that Gettier's Smith lacks knowledge.[12] But then it's

---

[12]To refute the claim that knowledge is equivalent to justified true belief, Edmund Gettier ("Is Justified True Belief Knowledge?" *Analysis* 23 [1963], 121–123) famously described two cases in which a subject, Smith, has a justified belief but lacks knowledge even though, due to an lucky twist of fate, the belief is also true. In one case, job-candidate Smith believes, on the employer's word, that Jones will get the job and also, based on a count, that Jones has ten coins in his pocket. Smith concludes that *p*: the man who will get the job has ten coins in his pocket. Here's the twist: Jones won't get the job but Smith will, and unknown to Smith he himself has ten coins in his pocket. Though there are thin wisps of an argument offered in

false that this method is modally fragile in the allegedly problematic way. This sort of rational insight cannot easily go awry, and so *that* is the assumption that should be jettisoned. Either way, then, the objection to the bare counterfactual argument fails.

This parallels Hick's own response in the case of his pluralist religious beliefs. He is not, as he says, a pluralist because he was raised from childhood to be one. Rather, his pluralism is the output of rational reflection. And that method has no problem with modal fragility. In this way, Hick survives the charge of self-defeat, and the bare counterfactual argument might, suitably modified, survive Plantinga's objections.[13]

To sum up, the bare counterfactual argument is, as written, subject to decisive objections. But there's reason to believe that the bare counterfactual argument doesn't fully capture the problem of contingency for religious belief. That problem is meant to afflict beliefs that have been formed in a particular way, by a particular method: unreflective "early teaching and socialization," as Kitcher says. And the problem with this method is that it could easily have delivered what would be, by the believer's own lights, falsehoods.

Focusing on methods in this way might allow the problem of contingency to side-step Plantinga's objections. We may carve a loophole to let slip Hick's reflective belief in pluralism, other philosophers' reflective atheism, etc. But we carve this loophole at a cost: reflective religious beliefs will also slip through. After all, reflective religious believers use the same method—rational reflection—to arrive at their religious views as Hick used to arrive at his. Their religious beliefs are not merely the result of "early teaching and socialization," so they are as exempt as Hick is from the problem of contingency.

And that's a cost. We might have thought that the problem of contingency casts a wide net, threatening religious believers generally. But in reality, as we have just seen, to avoid trivial counterexamples and self-defeat, the problem of contingency must target only unreflective religious belief, belief formed genuinely e.g., on the basis of passive receipt of testimony

---

favor of the claim that Smith doesn't know *p*, most readers find that claim *obvious* and not in need of any further support. These readers take themselves to *just see* that Smith lacks knowledge of *p* in the scenario as described. And the suggestion here is that Plantinga knows that racial intolerance is wrong in the same way: direct acquaintance with—"just seeing"—the truth of that proposition.

[13]This response should also be given, it seems to me, to Peter Van Inwagen's criticism of Hick. See Peter Van Inwagen, "Non Est Hick," in *The Rationality of Religious Belief and the Plurality of Faith*. Van Inwagen asks us to consider our political beliefs. Surely the methods we used to acquire these could easily have gone awry: just witness the widespread disagreement on these matters. But then, the objection concludes, Hick's argument requires too much skepticism—or, as was Van Inwagen's concern, too much disapprobation—since it rules that we do not know much at all about politics. Those sympathetic to Hick may say, in response, that if the method used really is modally fragile—something like passive, unreflective receipt of testimony—skepticism is plausibly called for. But for many of us, our political beliefs are the result of a modally stable method: rational reflection. And such beliefs are immune to Hick's skeptical argument (and to his charges of arrogance, which was Van Inwagen's primary concern).

during childhood. Reflective religious belief is in the clear.[14] And that, I should think, exonerates the religious beliefs of anyone reading this essay, as well as the religious beliefs of anyone who has ever reflected on anything remotely like this essay. So, as I say, an appeal to methods comes with a cost.

But it also comes with benefits: it avoids the decisive objections to the bare counterfactual argument. So, holding the lessons of this section in hand, let's attempt to carefully state the problem of contingency for religious belief in a way that attends properly to *methods*.

The Safety Argument

These days, when epistemologists speak of the virtues of belief-producing methods, they often have in mind either *sensitivity* or *safety*. Suppose Smith believes truly that p (e.g., it's 70°F in here) on the basis of some method (e.g., checking her thermometer). To say that Smith's belief is *sensitive* is to say that, had p been false, Smith would not have believed via this method that p. To say that Smith's belief is *safe* is to say that, were Smith to believe that p via this method, p would be true.[15] Or, alternatively, that not easily would Smith have believed falsely via that method.[16]

We'd like to sharpen the problem of contingency for religious belief by focusing on methods. The original quotations from Hick, Kitcher, et al. seem to support a reading of the argument on which the problem is that the *method* whereby religious beliefs are formed lacks some virtue

---

[14]How reflective must a belief be to avoid the net? It's hard to say. If we construe "reflective" too strictly, so that only philosophers like Hick meet it, then the problem of contingency begins to look like it calls for an implausible degree of skepticism, and catches in its net even many less-than-fully-reflective religious skeptics. On the other hand, if we construe "reflective" too loosely, so that nearly everyone's beliefs are reflective, then the problem of contingency loses its sting, casting doubt on virtually nothing. In earlier work, Hick says that "in some ninety-nine percent of cases the religion which an individual professes and to which he or she adheres depends upon the accidents of birth." See John Hick, *An Interpretation of Religion: Human Responses to the Transcendent* (New Haven, CT: Yale University Press, 1989), 2. This is an empirical question that begs for good data, but my sense is that to maintain such a wide net for the problem of contingency requires a quite high—implausibly high, I say—standard for "reflective."

[15]See, for example, Duncan Pritchard, *Epistemic Luck* (Oxford: Oxford University Press, 2005), 163: "If a believer knows that p, then in nearly all, if not all, nearby possible worlds in which the believer forms the belief that p in the same way as she does in the actual world, that belief is true." And Steven Luper, "Restorative Rigging and the Safe Indication Account," *Synthese* 153 (2006), 161–170: "at time t, S knows p by arriving at the belief p through some method M only if: M would, at t, indicate that p was true only if p were true."

[16]See, for example, Ernest Sosa, "How to Defeat Opposition to Moore," in *Philosophical Perspectives 13: Epistemology*, ed. J. Tomberlin (Oxford: Blackwell, 1999), 142: "[A] belief by S is 'safe' iff: as a matter of fact, though perhaps not as a matter of strict necessity, not easily would S believe that p without it being the case that p." And John Hawthorne, *Knowledge and Lotteries* (Oxford: Oxford University Press, 2004), 56n17: "Insofar as we withhold knowledge in Gettier cases, it seems likely that 'ease of mistake' reasoning is at work, since there is a very natural sense, in such cases, in which the true believer forms a belief in a way that could very easily have delivered error." And R. M. Sainsbury, "Easy Possibilities," *Philosophy and Phenomenological Research* 57 (1997), 907: "If you know, you couldn't easily have been wrong."

required for genuine knowledge. And learning of this deficiency is a de-feater for our religious beliefs, one might think.

But which virtue shall we focus on? Sensitivity or safety? There are two reasons not to favor sensitivity. First, favoring sensitivity would re-quire an unnatural reading of the original statements of the argument. For example, we find this in Mill: "the same causes which make him a Churchman in London, would have made him a Buddhist or a Confucian in Pekin." This sounds to my ears more like a violation of the safety con-dition rather than of the sensitivity condition. For Mill, there are nearby possibilities in which one forms religious beliefs via the same method she actually used, and yet in which she would believe something which is, by her own lights, false. And likewise with the counterfactual we get from 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." This too sounds more like a violation of safety than of sensitivity.

So, construing the problem of contingency as an attack on the sensitivity of religious beliefs is unfaithful to the texts. But it is also uncharitable, since recent work in epistemology has cast serious doubt on the claim that sensitivity is required for knowledge, or that learning that one's belief is not sensitive should count as a defeater for that belief. For a nice sum-mary of the case against sensitivity, the reader may consult Roger White's recent discussion.[17] We have two reasons, then, to construe the problem of contingency as an attack on the safety of religious beliefs.

Let's call this "the safety argument":

3) If you had been born and raised elsewhere, else when, and formed religious beliefs using the same method you actually used, then, by your own lights, you easily might have believed falsely.

Therefore,

4) Your religious beliefs were not formed safely.

Therefore,

5) Your religious beliefs don't count as genuine knowledge.

In my estimation, something like this argument is most likely what Xeno-phanes, Mill, Hick, and Kitcher had in mind. As a condition on knowl-edge, safety has much to be said for it, and many adherents in the current literature. And counterfactuals like (3) do have a certain glow of truth, and appear to menace the safety of religious belief.

---

[17]White, "You Just Believe that Because . . . ," section 4.1. Though for a few spirited recent defenses of sensitivity, the reader may consult *The Sensitivity Principle in Epistemology*, ed. Kelly Becker and Tim Black (Cambridge: Cambridge University Press, 2012).

Nevertheless, I believe this argument is doubly unsound, since both in-ferences are invalid. Before I explain, though, let's recall from above that, in order to avoid self-defeat, talk of methods in premise (3) of the skeptic's argument will be taken to cover only *unreflective* religious believers. Re-flective believers—including the skeptics who wield this argument—have nothing to fear.

And neither, I will argue, do unreflective religious believers. For con-sider the inference from (3) to (4). (3) tells me that, had something gone different in my past, then, using the same method I actually used to ac-quire religious beliefs, I easily might have believed falsely. But it does not follow from this that I actually believed unsafely. To see this, consider a non-epistemic analogy. On your drive home from work, you always use the Homeward Bridge. Today it's as sturdy as ever, and you safely cross it. Earlier that day, however, a pack of unprincipled teenagers tried to deto-nate some small-scale explosives under Homeward Bridge, which would have rendered it unsound and hazardous by the time of your approach. Fortunately for you, their matches were wet and they skulked home in low spirits. But had those kids succeeded, as they easily might have, the bridge would have been unsafe. Nevertheless, Homeward Bridge was a safe method for you to get home today. In this case, an analogue of premise (3) does not entail an analogue of premise (4): the fact that something nearly happened in the past that would have made your method unsafe doesn't entail that the method was unsafe when you actually used it.

Here's an epistemic case to seal the point. Suppose that the infamous Evil Epistemologist has poisoned the world's water supplies with a drug that radically impairs human cognitive faculties. Once exposed to the drug, all of one's faculties become completely unreliable. However, a be-nevolent nurse used his only dose of antidote to immunize you in the maternity ward. Your faculties are therefore safe from this poison, while everyone else's faculties are in serious danger. Yet it is true that, had you been born in a different time and a different place, and used the same faculties and methods you actually used, you easily might have believed things that would be, by your own lights, false. So we have here a counter-example to the general form of the inference that is meant to carry us from (3) to (4). The fact that something might have happened which would have rendered my faculties unsafe does not entail that my faculties are actually unsafe. The first inference of the argument is, therefore, invalid.[18]

---

[18]It is worth noting that if the inference from (3) to (4) were valid, then we would plau-sibly have counterexamples to the inference from (4) to (5), i.e., we could plausibly construct cases of unsafe knowledge. See Juan Comesaña, "Unsafe Knowledge," *Synthese* 146 (2005), 395–404; Ram Neta and Guy Rohrbaugh, "Luminosity and the Safety of Knowledge," *Pacific Philosophical Quarterly* 85 (2004), 396–406; and Christoph Kelp, "Knowledge and Safety," *Journal of Philosophical Research* 34 (2009), 21–31. These authors all develop alleged cases of unsafe knowledge that depend on the move in question: something easily might have happened in the past which *would* have rendered my beliefs now unsafe, therefore they *are* now unsafe. If you think that's a good inference—i.e., if you think the move from (3) to (4) is valid—then these philosophers provide a major obstacle to the move from (4) to (5).

And so is the second inference, for it assumes that safety is required for knowledge. That is, the move from (4) to (5) assumes that a subject can know that p via some method only if that method could not easily have delivered error. Only if, that is, not easily would the subject believe that p via this method without it being the case that p. But—as I will now argue—safety is *not* required for knowledge.

For suppose that the world's most accurate clock hangs in Smith's office, and Smith knows this. This is an atomic clock: its accuracy is due to a clever radiation sensor, which keeps time by detecting the transition between two energy levels in cesium-133 atoms. This radiation sensor is very sensitive, however, and could easily malfunction if a radioactive isotope were to decay in the vicinity.

This morning, against the odds, someone did in fact leave a small amount of a radioactive isotope near the world's most accurate clock in Smith's office. This alien isotope has a relatively short half-life, but–quite improbably–it has not yet decayed at all. It is 8:20 am. The alien isotope will decay at any moment, but it is indeterminate when exactly it will decay. Whenever it does, it will disrupt the clock's sensor, and—for complicated sciencey reasons that I can't explain here—it will freeze the clock on the reading "8:22."

The clock is running normally at 8:22 am when Smith enters her office. Smith takes a good hard look at the world's most accurate clock—what she knows is an extremely well-designed clock that has never been tampered with—and forms the true belief that it is 8:22 am.

Many classic and widely-accepted theories of knowledge rule that Smith knows it is 8:22 am.[19] Further, it strikes me as obvious that Smith knows it is 8:22 am, and perhaps the same goes for you. The threat to her clock remains purely counterfactual, and everything is functioning properly when she forms her belief. Yet, since the isotope could easily have decayed and frozen the clock at "8:22," and since Smith may easily have checked the clock a moment earlier or later, Smith might easily have believed it is 8:22 am without it's being 8:22 am. Smith formed her belief in a way that could easily have delivered error.

Therefore, Smith knows, and yet her belief was not formed safely. Safety, then, is not required for knowledge. And so the inference from (4) to (5) above is invalid, as was its partner inference from (3) to (4). And so we have a surfeit of reasons to think that the safety argument does not threaten the rationality of religious belief, even if the belief is unreflective. Let's turn elsewhere, then, to see if we might find a more intimidating version of the problem of contingency for religious belief.

However, I deny the inference from (3) to (4), and I believe this is a fatal shortcoming for the proposed cases of unsafe knowledge from Comesaña, Neta and Rohrbaugh, and Kelp. (For a fuller defense of this claim, see my "Knowledge Under Threat," forthcoming in *Philosophy and Phenomenological Research*.) But there are clear examples of unsafe knowledge that do not rely on this hasty inference, as we'll see below.

[19]For a defense of this claim, see Bogardus, "Knowledge Under Threat."

The Accidentality Argument

Many epistemologists share an *anti-luck intuition*—knowledge excludes a certain kind of luck—and they typically offer Gettier cases as paradigmatic examples of this type of toxic luck. Gettier himself seems to have the intuition, when he pointed out that it is only by the "sheerest coincidence" that Smith's belief turned out to be true, in support of his claim that Smith lacked knowledge in that second case offered in his famous 1963 paper. As conditions on knowledge, safety and sensitivity are often offered as analyses of this sort of knowledge-precluding luck.

But suppose that, as we've argued above, neither sensitivity nor safety is able to sharpen the problem of contingency into a stinging attack on religious belief. It still may seem as though the problem of contingency for religious belief concerns how lucky one would have to have been to get her religious beliefs right. Think on all the ways one might easily have been raised, all the religious beliefs one might easily have received through early teaching and socialization: if one happened to get it right on religious matters, one might think it could only be due to that sort of knowledge-precluding luck. Though safety and sensitivity look unpromising, there is still hope of developing the problem of contingency along these lines. For we may pursue another strand of thought within anti-luck epistemology: an analysis of luck in terms of *accidentality*.

Peter Unger, for example, offers this analysis of knowledge: S knows that p just in case it is not at all accidental that S is right about its being the case that p.[20] Unger is careful to point out that some types of accidentality are compatible with knowledge. For example, suppose that a disgruntled student tries to run me down with his car, but at the last moment he loses control and slows his roll harmlessly into a ditch. "I can't believe you crashed your car trying to kill me!" I exclaim, though I really can and do. Indeed, I know it. Now, *that the car crashed* was an accident, and so the content of my belief *that it crashed* is true by accident. Yet I know that it crashed. So that sort of accidentality is compatible with knowledge. And it's an accident that I'm alive and able to believe that the car crashed, since the driver failed to kill me only by accident. And so this kind of accidentality too is compatible with knowledge. What's incompatible with knowledge, for Unger, is that there should be any accidentality "between the man and the fact," as he says. That is, *given* my evidence—which may exist or have been acquired by accident—and *given* that I'm there to consider the evidence—this too may be the result of an accident—it cannot be an accident that I form a true belief on my evidence. Or, as Unger says, in believing that the car crashed, it can't be an accident that I'm right about its being the case that the car crashed.

---

[20]Peter Unger, "An Analysis of Factual Knowledge," *Journal of Philosophy* 65 (1968), 159. See also William Harper, "Knowledge and Luck," *The Southern Journal of Philosophy* 34 (1996), 273–283; and Masahiro Yamada, "Getting It Right by Accident," *Philosophy and Phenomenological Research* 83 (2011), 72–105.

Casting the problem of contingency as an attack on the accidentality of religious belief rather than the safety or sensitivity of religious belief will avoid the objections raised in the previous section. In addition, this strategy has the virtue of being faithful to certain heretofore underappreciated commonalities in the original quotations. Mill, for example, says the trouble for religious belief is that "mere accident" has determined where we'd land on religious questions. And Hick stresses that religious allegiance depends in the great majority of cases on an "accident of birth." We have ample textual reason, then, to pursue accidentality as the culprit.

And so perhaps the argument is meant to go like this: when we reflect on certain counterfactuals, we realize that we easily might have ended up with religious belief that we would consider false. And these counterfactuals should convince us that our religious beliefs, if true at all, are true by sheer accident. And by this route we should conclude that our religious beliefs do not amount to knowledge. A bit more carefully, we begin as we did before with the safety argument:

6) If you had been born and raised elsewhere, else when, and formed religious beliefs using the same method you actually used, then, by your own lights, you easily might have believed falsely.

But now we conclude:

7) It's *accidental* that you're right about your religious beliefs (or you're not right at all).

And, hence,

8) Your religious beliefs aren't knowledge.

As with the safety argument above, this argument threatens him who wields it. To avoid that problem of self-defeat, we ought to understand the first premise as true only in the case of *unreflective* methods, such as those allegedly used by many religious believers. Those who believe this argument is sound do so, presumably, on the basis of rational reflection, a method that could not easily go awry. Therefore, as with the safety argument above, reflective religious believers have nothing to fear from this argument.

And neither do unreflective religious believers, since both inferences are invalid. For consider first the move from (6) to (7). (6) tells us that, had something accidental and out of your control gone differently in the past, then the method you used to acquire religious beliefs could easily have led you astray. But this premise is blind to Unger's distinction between benign and toxic kinds of accidentality discussed above. For it may be that, had something accidental and out of your control gone differently in the past, then you would have lacked good evidence that you actually had, and so the method you used could easily have led you astray. But that sort of accidentality is not the toxic luck Unger tried to capture with the phrase

featured in (7). And so it's possible to construct a counterexample to the inference from (6) to (7).

Here's one. Suppose you are an elderly friend of Copernicus in Fromburk in 1515. Copernicus has collected enough data to justify belief in heliocentrism, but, for fear of ecclesiastical reprisals, he has not and will not share these data with the world for over a decade, long after your death. He has, however, shared his findings with you. You can come to know on the basis of Copernicus's testimony that heliocentrism is true, and it's no accident that you're right. Yet had you been born and raised elsewhere, else when—far away from your friend Copernicus—and formed astronomical beliefs using the same method you actually used (namely, expert testimony), you easily could have been led astray, into thinking that heliocentrism is false. Only by a mere accident of birth did you acquire Copernicus's solid gold testimony. But this species of accidentality is not the sort described in premise (7); *given* the evidence you've acquired, it's not an accident that you're right about heliocentrism's being true. So we have here a counterexample to the general form of the inference that is meant to swing us from (6) to (7). The accidentality captured in (6) may not infect the relationship "between the man and the fact" required by (7), and so (6) doesn't entail (7).

More controversially, I believe the inference from (7) to (8) also fails. That inference is valid only if knowledge excludes the kind of accidentality captured by (7) and described by Unger. But I will now argue that, contra Unger, a subject can know that p even if it is accidental that she is right about its being the case that p. First, let's admit to one another that Unger's analysans—"it's not at all accidental that S is right about its being the case that p"—is not exactly a common coinage in ordinary English, and so it's hard to get a grip on it. But he helpfully provides what he takes to be a case in which it's accidental that a subject is right about its being the case that p. Unger imagines an estate owner who employs a servant to tend to the estate. The servant doesn't believe there is a rosebush on the estate, but, for the sheer fun of it, he tells the estate owner that there is one in order to get him to have a false belief. A hilarious prank. Unbeknownst to the servant, there really is a rosebush on the estate, and so the estate owner ends up with a true belief on the basis of the servant's testimony. Unger says that the estate owner does not *know* that there is a rosebush on the estate, since "it is entirely accidental that the estate owner has been right about there being a rosebush on his estate."[21]

From Unger's example, I believe we can extract one sufficient condition for its being accidental that a subject is right about p's being the case. What made the estate owner's belief accidentally true was that it was based on a report that was not aimed at the truth nor at anything connected with the truth. The servant's report hit the truth by the sheerest coincidence;

---

[21]Unger, "An Analysis of Factual Knowledge," 158.

that it should hit the truth without being so aimed could happen only by accident. And so I believe Unger is committed to this proposition:

   9) Necessarily, if a subject believes truly that p only on the basis of a report from some source, and that report was not aimed at the truth nor aimed at anything connected with the truth, then it's accidental that the subject is right about p's being the case.

Now for a moment set aside Unger's non-accidentality analysis of knowledge. Think instead about *swampman*, that bizarre creature of philosophical fiction. Swampmen (like swampwomen) are the results of lightning striking swamps and randomly rearranging bits of organic matter into microphysical duplicates of ordinary people. Astronomically improbable, sure, but metaphysically possible. Your swamp duplicate will look and act just like you. On the basis of misleading memories, he or she will issue utterances like "Of course I'm not a swamp-person! I've been around for decades, I have a mother and a father, etc." His or her eyes, ears, nose, brain—in short all his or her cognitive faculties—will eventually issue reports about the world to him or her just as yours do to you. But, of course, your swamp duplicate's faculties are the result of lightning striking a swamp: they weren't designed by God or nature or anything else, they weren't aimed at anything, let alone at the truth. And so we can be confident of the following proposition:

   10) Necessarily, the reports of swamp-people's cognitive faculties are not aimed at the truth or anything else.

And from (9) and (10) it follows that if a swamp-person ever believes truly that p, it is accidental that he or she is right about p's being the case. And so we may conclude:

   11) Necessarily, if there is a swamp-person that knows something, then it's possible for a subject to know that p despite its being accidental that the subject is right about p's being the case.

Now, if we prove that a swamp-person could have knowledge, we could infer from (11) that Unger's analysis of knowledge is wrong, and so that (8) does not follow from (7) in the above version of the problem of contingency. That would be an important result.

   And we can prove that a swamp-person could have knowledge. For:

   12) *That I am a swamp-person* is consistent with everything I know from the armchair.

By "everything I know from the armchair," I mean everything I can know *a priori*, everything I can know even on the supposition that I am the victim of Descartes' evil demon. Let *a* (for "armchair") pick out the long conjunction of those propositions. If you're like me, you'll think that nothing you know from the armchair—no part of *a*—rules out the proposition that

you're a swamp-person. You cannot introspect or intuit your way to refuting that skeptical hypothesis; empirical work is required.

But notice that we can intuit or introspect our way to the knowledge *that we know at least something*.[22] Even if I am the victim of an evil demon, I can know that I exist, and *I can know that I know* that I exist. So, some conjuncts of proposition *a* clearly entail that I know. And so we can safely endorse this proposition:

13) That *I know nothing* is inconsistent with *a*, i.e., everything I know from the armchair.

And from (12) and (13) it follows that:

14) It's possible that I am a swamp-person and yet I know something.[23]

And now we may combine (14) with (11) to yield:

15) It's possible for a subject to know that p despite it's being accidental that the subject is right about p's being the case.

But then the inference from (7) to (8) in the above statement of the problem of contingency for religious belief is invalid—just as, I've argued, the inference from (6) to (7) is invalid.[24] We have here, then, two reasons to reject this statement of the problem of contingency for religious knowledge.[25]

---

[22]This thought goes back at least to Augustine, who makes the same move in book XI of chapter 26 in his *City of God*. Augustine makes similar moves in *Contra Academicos* 3.10.23 and also in *De Trinitate* 10.10.14.

[23]The form of this modal argument from (12) and (13) to (14) is, where *s* names the proposition that *I am a swampman*, *a* picks out the long conjunction of propositions that I can know from the armchair, and *k* picks out the proposition that *I know something*: ◊(s & a), ¬◊(¬k & a), therefore, ◊(s & k). That's valid, at least on widely accepted systems of modal logic.

[24]A simpler counterexample to Unger's analysis of knowledge might borrow John Hawthorne's swampwatch, a particle-for-particle duplicate of the world's most accurate wristwatch "created by a fortuitous coming together of molecules." See John Hawthorne, "Deeply Contingent A Priori Knowledge," *Philosophy and Phenomenological Research* 65 (2002), 247–269. Now one might think that, given its origins, while swampwatch reports the time its reports are not aimed at the truth or anything else. And so, if a subject were to believe truly that *p* on the basis of swampwatch's reports, it would be an accident—in a familiar and legitimate sense of the word "accident"—that the subject is right about *p*'s being the case. And yet, Hawthorne says, when a subject "uses his swampwatch to inform him about the time . . . we are intuitively ready to say that [she] knows what time it is." If you share Hawthorne's intuition, we have here a case in which S knows that *p* and yet it is accidental that S is right about *p*'s being the case. I'm doubtful, though, that many will share Hawthorne's intuition here. And so I include the above argument concerning swamp-people.

[25]Objection: Unger and friends need not accept (9). They may claim that it fails to identify the knowledge-precluding sort of accidentality that Unger had in mind when offering his estate owner case. Rather, the estate owner's belief was true by accident because the servant's report failed to be "truth-conducive"—as Yamada ("Getting It Right by Accident") puts it on Unger's behalf—since it's false that if one trusts the servant, one will end up with a true belief. But the reports of swamp-people's cognitive faculties *are* truth-conducive (as are the reports of swampwatch from the previous note). And so if a swamp-person gains a true belief from her faculties (or if we gain a true belief from swampwatch), it need not be true by accident. Response: Let's call the kind of accidentality that Unger's analysans picks out "toxic accidentality." It's toxic because it's the type of accidentality that is incompatible with knowledge. Now, there may well be several conditions that are sufficient for toxic ac-

The Argument from Symmetry

We have investigated statements of the problem of contingency for religious belief that center on safety and also on accidentality. Neither avenue leads to a damning indictment of religious belief. I will now offer one final attempt to state the problem of contingency for religious belief, this time focusing on Kitcher's "symmetry argument."

Recall Kitcher's statement of the argument:

> 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. The symmetry is complete. . . . Given that they are all on a par, we should trust none of them. [26]

Perhaps we should understand Kitcher this way: the counterfactual ought to convince us of a troubling epistemic symmetry between adherents of different religions, to convince us that we have no independent, non-question-begging reason to believe that we are in a better epistemic position on religious questions than those who disagree with us. After all, adherents of different religious views used the same method to arrive at their religious beliefs as we did to arrive at ours.

This argument seems to map out terrain that has already been extensively explored (and colonized) by those working on the epistemic significance of peer disagreement. Analogies one finds in that literature suggest themselves here: if my thermometer says it's 70 degrees in here while your thermometer says it's 80 degrees, and I have no reason independent of this disagreement to favor mine over yours, it would be unacceptably

---

cidentality, several ways for a belief to be toxically accidental. And the servant's report—or a report from a swamp-person's cognitive faculty, or a report from swampwatch—may be toxically accidental in one way but not in another. *Even if* a report is truth-conducive and so not toxically accidental in that way, it may be toxically accidental in some other way. For example, a report may be toxically accidental in hitting the truth despite not being aimed at the truth. It seems to me that this is so. But in Unger's analysis, as he himself says (ibid., 159), "a *complete absence* of the accidental is claimed" (emphasis mine). There can't be *any* sort of accidentality "between the man and the fact"; there must be no toxic accidentality whatsoever. And so the objector must prove that (9) is false, and not simply that Unger and his friends are not committed to it by their theory. But (9) seems quite plausible, and so Unger and friends are still in trouble. Furthermore, truth-conduciveness is not required for knowledge. And so it is a dead-end for those wishing to diagnose what's gone wrong in the estate owner case, and it will not serve religious skeptics who wish to wield the problem of contingency against religious belief. For Yamada (ibid.), "to say that an epistemic method is truth-conducive is to say that if one does as the method says, one will end up with a true belief." This indicative conditional sounds fatally close to the proposed safety condition on knowledge: close enough that the atomic clock case from the previous section seems to me to refute Yamada's proposed condition on knowledge just as it refutes the proposed safety condition on knowledge. In that case, it's false that, if you trust the imperiled atomic clock, you will gain a true belief about the time. The isotope may easily decay, giving you a false belief. And so Yamada's truth-conduciveness condition is not required for knowledge, and so will be of no avail to formulating the problem of contingency for religious belief. Even if religious beliefs were not formed via a truth-conducive method, they may nevertheless rise to the level of knowledge.

[26]Kitcher, "Challenges for Secularism," 26.

question-begging or arbitrary to nevertheless stick with the report of my thermometer over yours. The fact that *this thermometer is mine* is not an epistemically respectable tie-breaker. So too with the deliverances of early teaching and socialization, one might think. If I accept Christianity as a result of early teaching and socialization, but I come to learn that many others accept many different religious beliefs on that same basis (or even that this easily *could* have happened), wouldn't it be unacceptably arbitrary or question-begging to nevertheless stick with the deliverances of *my* early teaching and socialization? The fact that *this early teaching and socialization are mine* cannot properly break the epistemic symmetry, and so I ought to give up my religious beliefs. Perhaps this is what Kitcher had in mind.

If so, this version of the problem of contingency would be on rather firmer ground than the previous versions we've discussed. For it seems to rely on a principle very much like a conciliatory view of peer disagreement, a view that has become prominent among those epistemologists working on peer disagreement. And yet the argument *need not* rely on anything quite so strong as a conciliatory view. Kitcher's skepticism could take flight without the burden of the dubious claim that adherents of various religions really are epistemic peers. The threat of skepticism seems more acute if we notice merely that it's very often *unclear* which of the various religionists are in a better epistemic position.[27] And it's unclear whether we would have been in a better epistemic position had we been born and raised elsewhere, else when. That far-reaching unclarity—both counterfactual and actual—should, I believe, be taken to constitute the core of Kitcher's symmetry argument.

And so I'd recommend that the argument be formulated like this:

16) If you had been born and raised elsewhere, else when, and formed religious beliefs using the same method you actually used, then you easily might have had different religious beliefs.

17) Many who *actually have* been born and raised elsewhere, else when, and formed religious beliefs using the same method you actually used, have formed different religious beliefs.

18) It's unclear whether you are in a better epistemic position than either your counterfactual self or those actual people who disagree with you.

Therefore,

19) Your religious beliefs aren't knowledge.

---

[27]For nice developments of this suggestion for those looking toward peer disagreement as fuel for skepticism, see Ballantyne "The Problem of Historical Variability," as well as Nathan King "Disagreement: What's the Problem? *Or* A Good Peer Is Hard to Find," *Philosophy and Phenomenological Research* 85 (2012), 249–272.

The conclusion, as before, is that you have a defeater for your religious beliefs. You can't sensibly go on holding them as before. And so they don't constitute knowledge.

This, I believe, is the strongest statement we have yet looked at of the problem of contingency for religious belief. Yet even this version of the problem does not throw down a successful defeater for religious belief. For consider again the problem of self-defeat, which has limited the scope of our previous statements of the problem of contingency. Kitcher means to shake this argument like a can of pennies in the face of the religious believer, while at the same time preserving his own atheistic answers to religious questions, as well as his belief that the above argument is sound. But what have religious skeptics like Kitcher to say about the premises of this argument? Aren't those premises true of religious skeptics? Suppose we appeal to a now-familiar distinction between reflective and non-reflective belief, between those beliefs gained "through early teaching and socialization" as Kitcher says, and those beliefs gained via rational reflection on the evidence. Which premise, if any, is false of Kitcher's reflective beliefs on religious questions? (And which premise is false with respect to Kitcher's belief that the symmetry argument is sound?)

Perhaps Kitcher would deny (16), and insist it's unlikely that, had he really used the same method of rational reflection, he easily might have arrived at different religious beliefs. Or perhaps he would deny (17), and reject that those who disagree with him do so on the basis of rational reflection. Or, perhaps most plausibly, he may deny (18) and claim that, in the case of his beliefs formed via rational reflection, it is not unclear to him who is in a better epistemic position. He can *just see* the truth of his view, and the falsity of his opponents' views.

You may find all of these responses unsatisfactory. If so, then the argument looks self-defeating. Its premises apply to all of us, even to those like Kitcher who brandish the argument. These premises apply equally well even to Kitcher's belief that the argument is sound: even that belief would be undercut by this argument. On the other hand, perhaps you think one of the above strategies will save religious skeptics the indignity of self-defeat. Perhaps the fact that their beliefs are *reflective* will allow them to plausibly deny one of the premises in their own cases. If so, then—as with previous statements of the problem of contingency—reflective religious believers will also be off the hook: they may, with equal plausibility, deny whichever premise reflective skeptics deny. And then, again, the scope of the argument will be narrower than the religious skeptic might have hoped. It will target, at most, unreflective religious believers.

Yet even many unreflective religious believers will be exempt from this argument. For notice that the essay so far has given us good reason to believe that the "elsewhere, else when" premise isn't enough, on its own, to defeat any external, anti-luck condition on knowledge. Premise (16) doesn't look to guarantee any epistemic trouble *upstream* of experience, neither for the safety of our religious beliefs nor for their non-accidentality. For

similar reasons, I'd think, the same goes for (17). So any pressure we feel from this argument more likely originates from the norms governing how we *respond* to experience, i.e., rationality *downstream* of experience. The premises make relevant an alternative that the believer cannot rule out: that those who disagree with him—including his counterfactual self—are better positioned than him on the question. And one might think that, once this alternative becomes salient, rationality requires conciliation. But that alternative can rise to salience only if one is *made aware* of it. That's why, in the epistemology of disagreement literature, much is made about the "full disclosure" condition: one feels the weight of conciliatory views only after the subject in question becomes aware of the disagreement.

After all, rationality is silent with respect to considerations that one is non-culpably unaware of. We cannot be expected to respond to considerations that are, through no fault of our own, beyond our ken. An alternative we're non-culpably unaware of cannot be relevant to us in a way that threatens knowledge all by itself. So, if an unreflective believer is non-culpably unaware of the premises of this argument, his religious beliefs may well be fully rational in spite of the truth of those premises. But then those premises don't by themselves constitute a defeater—undermining or otherwise—for his religious beliefs: (16)–(18) don't guarantee any epistemic problem upstream of experience, and there may well be no irrationality downstream of experience either, so long as the believer is non-culpably unaware of (16)–(18). To bridge the gap between (16)–(18) and (19), we would need to fasten in a "full disclosure" premise: that the believer in question is aware of the premises (or, at least, *culpably* ignorant of them). Without a premise to that effect, the inference is invalid.

And there may be many such religious believers, quietly warming the pews, cheerfully unaware that (16)–(18) are true of them. If their ignorance is due to any of a thousand exculpatory reasons—simplicity, cultural isolation, occupation with the more pressing matters of life, etc.— the conclusion may well be false of them even if the premises are true. And, furthermore, if an otherwise unreflective religious believer becomes aware of the premises of this argument, rationally reflects on them, and maintains his belief, he thereby becomes eligible for the exemption offered above to reflective religious believers. So here too we find reason to narrow the scope of this argument, reason to fence in the boundaries of the skepticism that it calls for.

Conclusion

Throughout the history of philosophy, many thinkers have held that the historical contingency of our religious beliefs constitutes a potent argument for the conclusion that those beliefs do not amount to knowledge. However, this alleged problem is more often stated enthusiastically than carefully. Here, we have attempted to lay out the argument as forcefully as possible, and we have found no compelling version of the problem of contingency for religious belief. The contingency of our religious beliefs

does not show that they were formed unsafely; and even if it did, safety is not required for knowledge. And the contingency of our religious beliefs does not show that they are accidentally true if true at all; and even if it did, non-accidentality is not required for knowledge. Finally, Kitcher's argument from symmetry, suitably understood, is either self-defeating or applies only to that fraction of the population that holds religious beliefs unreflectively and is unaware of Kitcher's premises due to blameworthy negligence or irrationality. Reflective religious believers—and even many unreflective religious believers—can rest easy. It may well be, then, that there simply is no problem of contingency for religious belief. Or, if there is, that it needn't worry many people.[28]

Pepperdine University