Episteme, 14, 4 (2017) 441–461 © Cambridge University Press
doi:10.1017/epi.2016.13

LUCK: EVOLUTIONARY AND EPISTEMIC

BILLY DUNAWAY
dunawayw@umsl.edu

ABSTRACT

This paper advances two theses about evolutionary debunking arguments in ethics. The first is that, while such arguments are often motivated with the rhetoric of 'luck', proponents of these arguments have not distinguished between the kinds of luck that might lead to the formation of a true belief. Once we make the needed distinctions, the relevance of the kind of luck which can be derived from broadly evolutionary explanations to the epistemological conclusions debunkers draw is suspect. The second thesis is that debunkers might successfully show that epistemic luck is relevant to their concerns. But they will need to include specific premises about the kind of evolutionary mechanism that explains the beliefs that the debunker targets. Proponents of debunking arguments, if they continue to insist on using luck-based considerations, are then hostage to empirical fortune in a way not previously recognized.

1. PRELIMINARIES

Recently a quite large literature has emerged on the topic of evolutionary 'debunking' arguments in ethics.[1] Broadly, these arguments aim to show that, owing to some general features of the evolutionary origins of certain kinds of belief, those beliefs are epistemically suspect. This is just a rough schema for an argument, and within the burgeoning literature there are a variety of ways of filling out the schema. (In particular, which kinds of belief are problematic, how evolutionary origins bear on formation of the relevant beliefs, and what 'epistemically suspect' amounts to are all questions any respectable version will need to answer.) One thought that often motivates the idea that some version of the argument is compelling deploys the notion of *luck*. It is a common theme in the debunking arguments literature that the effects of evolutionary forces which shape our ethical beliefs show that the beliefs in question are, if true, only so because of an accident of luck. And many epistemic notions – knowledge, justification, and the like – plausibly require the absence of a certain kind of luck. The question this paper addresses is whether the luck-based motivations for worrying about the epistemic effects of evolution are apt.

It is very natural to speak in terms of luck when considering the question of whether evolutionary forces might have produced true ethical beliefs. Street, for instance, uses such language in the following passage:

1 See for example Joyce (2006), Street (2006) for early examples, and for further developments Schafer (2010), Kahane (2011), Clarke-Doane (2012), Setiya (2012: 65 *ff.*), Wilkins and Griffiths (2013), and Vavova (forthcoming).

> [I]t turns out that we have been evolving towards affirming whatever evaluative content tends to promote reproductive success. We have thus been guided by the wrong sort of influence from the very outset of our evaluative history, and so, more likely than not, most of our evaluative judgments have nothing to do with the truth. Of course it's possible that as a matter of sheer chance, some large portion of our evaluative judgements ended up true, due to a happy coincidence between the realist's independent evaluative truths and the evaluative directions in which natural selection tended to push us, but this would require a fluke of luck that's not only extremely unlikely, in view of the huge universe of logically possible evaluative judgements and truths, but also astoundingly convenient to the realist. Barring such a coincidence, the only conclusion remaining is that many or most of our evaluative judgements are off track.[2]

Richard Joyce motivates a similar thought by analogy with a case that clearly involves epistemically problematic luck. Joyce imagines a case where you are slipped a pill that causes belief in the claim that Napoleon lost Waterloo. It is extremely natural, given his set-up, to view your Napoleon-beliefs as only true owing to an accident of luck; one could easily have been slipped a pill which caused one to believe that Napoleon won instead. Joyce highlights the pernicious epistemic effects of such luck in the following passage:

> Suppose that there were a pill that makes you believe that Napoleon won Waterloo, and another one that makes you believe that he lost … Now imagine that you are proceeding through life happily believing that Napoleon lost Waterloo (as, indeed, you are), and then you discover that at some point in your past someone slipped you a 'Napoleon lost Waterloo' belief pill. It is not a matter of your learning of the existence of such pills and having no way of knowing whether you have ever taken one; rather, we are imagining that you somehow discover beyond any shred of doubt that your belief is the product of such a pill. Should this undermine your faith in your belief that Napoleon lost Waterloo? Of course it should.[3]

It is important to distinguish the rhetorical effect of these passages from the official arguments Street, Joyce, and other debunkers offer. There are complicated issues of interpretation concerning exactly which premises individual debunkers use to support their conclusions. One possible interpretation of these arguments is that they explicitly rely on a premise about epistemic luck. That is, some debunkers might be arguing along the following lines:

> *Argument from Luck*
> Evolution influences ethical belief-formation in such-and-such way.
> If evolution influences ethical belief-formation in such-and-such way, then ethical beliefs are epistemically suspect.
> Therefore, ethical beliefs are epistemically suspect.

2 Street (2006: 122, emphasis added). See also Street (2008: 208–9). In these and other passages Street is careful to limit the scope of her debunking argument belief in evaluative facts, construed realistically. In what follows I will assume that the argument can be extended to belief in ethical facts.

I will also assume that ethical facts are to be understood in a paradigmatically realist fashion. Street uses her debunking argument against the realist's ethical facts as a lemma in an argument for her anti-realist Constructivism on which, she claims, ethical facts are not debunked. I will not be addressing the second half of Street's argument here. For more on whether Street's favored version of Constructivism avoids epistemological challenges, see Dunaway (ms).

3 Joyce (2006: 179). Similar rhetoric is found in the presentation of evolutionary arguments in other domains in Bloom (2009: 129) and Wilkins and Griffiths (2013: 143).

https://doi.org/10.1017/epi.2016.13 Published online by Cambridge University Press

But as debunking arguments have inspired a massive literature dealing with the question of how best to interpret the arguments, any claim that debunkers rely on a version of the Argument from Luck is bound to be controversial. Fortunately we can address the relevance of luck-based considerations in debunking arguments without taking an interpretive stand on whether any debunkers have explicitly endorsed the Argument from Luck.

This is because luck-based considerations might play a less formal, but no less important, role in motivating the debunker's conclusion. When Street makes claims about a 'happy coincidence' and the effects of 'sheer chance', the vocabulary bears on one natural interpretation a very close connection to luck-based considerations. Roughly: it is very natural to hear the claim that a belief's truth relies on a 'happy coincidence' as claiming that, if the belief is true, then its truth is a matter of luck. The same point goes for other proprietary terms Street uses. And Joyce claims an 'analogy' between ethical beliefs shaped by evolution and beliefs formed by ingesting a Napoleon Pill. (He says, 'The intention of this make-believe scenario is to prime us for an analogical epistemological conclusion regarding an evolved moral sense.') As before, there is a natural reading on which one central feature of the analogy is luck-based: just as the Napoleon Pill makes one's Napoleon-beliefs rely on an accident of luck, evolution shows a similar reliance on luck in the case of ethical belief.

The upshot is that even if debunkers do not explicitly accept the Argument from Luck, it is nonetheless very natural to hear their central terminology and analogies as relying on the epistemic consequences of luck. We can ask whether the intuitive force of this connection is warranted. Should we think of beliefs that are true owing to a 'happy coincidence' as epistemically suspect because reliance on a coincidence entails the presence of epistemically problematic luck? And should we think of ethical beliefs formed by evolution as epistemically suspect in the same way as the Napoleon Pill beliefs, which are a clear example of epistemic luck? Debunkers who do not rely on the Argument from Luck are not forced into a 'yes' answer here. But since much of the intuitive force of their presentation plays up the presence of luck that arises out of the evolutionary influences on ethical belief, it is an important question whether these connections are deserved. I will focus almost exclusively on this question; nothing I say here addresses debunking arguments that explicitly disavow connections with epistemic luck.

My conclusion is a largely pessimistic one – barring some additional non-trivial empirical premises, debunkers' attempts to evoke problematic epistemic luck are misplaced.

The largely pessimistic conclusion I will argue for consists in two theses. The first is that evolutionary debunking arguments are quick to play up very natural connections with luck, but systematically ignore important distinctions between different *kinds* of luck. Not all types of luck in processes leading up to formation of a true belief are epistemically problematic – in fact, some are perfectly benign. Proponents of debunking arguments, in using the language of luck to motivate central premises, have systematically ignored these distinctions. Absent further premises, the evolution-induced presence of luck in ethical belief implies nothing about the epistemic status of these beliefs.

This conclusion isn't *entirely* pessimistic, however, as the second thesis is that there may well be additional premises available to the debunker which do establish the needed connection between evolutionary luck and the epistemic status of ethical beliefs. In particular, if debunkers can provide additional evidence about the specific way in which evolutionary forces have shaped our ethical beliefs, they can rule out the possibility that the luck involved in such belief-formation is epistemically benign. I will outline what the specific evolutionary processes would have to look like in order to yield this conclusion. While

https://doi.org/10.1017/epi.2016.13 Published online by Cambridge University Press

I will not attempt to settle whether these processes have in fact taken the required shape, one lesson from this second thesis is that the relevance of luck to evolutionary debunking is hostage to empirical fortune. Not every evolutionary explanation that proceeds broadly by a mechanism of natural selection will do the job. It is an open (and under-explored) question, to be settled by naturalistic scientific investigation – and not philosophical debate alone – whether luck is the right tool for casting doubt on our ethical beliefs.

A few preliminary methodological points are in order before proceeding. The primary focus of the following sections will be on whether a focus on luck supports the epistemological conclusion debunking arguments. For simplicity I will focus on one epistemological notion, namely knowledge. Thus instead of asking whether the debunker's argument succeeds in showing that ethical beliefs are epistemically suspect, I will ask the more pointed question of whether luck is relevant to the possibility of ethical knowledge. This is not entirely in keeping with the terms used in the debunking literature, where questions about knowledge are not explicitly addressed. However for present purposes this simplification is a harmless one, and the lessons from the focus on knowledge should carry over to other familiar epistemological notions as well. A central claim of this paper is that there are certain kinds of luck that are incompatible with knowledge, while others are epistemically benign. Analogously, other epistemically significant notions will not be compatible with certain kinds of luck while tolerating other epistemically benign forms of luck.

For example, some will want to think of debunking arguments as showing that ethical beliefs are unjustified, or that once one knows of the evolutionary origins for one's ethical beliefs, one acquires a 'defeater' for them. Again I will not focus on the notions of justification or defeat here, but the lessons below carry over: one way of a belief's being unjustified is for it to rely for its truth on an objectionable kind of luck (for example the kind of luck in the Napoleon Pill case), and one way of acquiring a defeater is to learn that one's belief, even if true, was formed in a way that relies on the relevant kind of luck. Thus proponents of justification- and defeat-based arguments will find the discussion of epistemic luck here to be of the first importance when deciding whether to invoke luck-based considerations when making their case.

One final dialectical point is in order as well. Some debunkers might insist that their argument is not to be couched in familiar epistemological terminology at all, and instead can only be stated using their own proprietary vocabulary. (One could imagine – and might even encounter in the literature – a debunker insisting that the relevant notion of 'off-trackness' cannot be understood by its relationship to knowledge, justification, defeat, and the like.) The versions of the debunking argument will not find the discussion of epistemic luck here to be of much interest. But such versions of the argument are also much less forceful dialectically, since most philosophers will accept that ethical belief is epistemically suspect only if it can be shown that it bears substantial and non-trivial relations to familiar epistemic notions: that it is not knowledge, is unjustified, or is susceptible to defeat. It is the relevance of luck-based considerations to these dialectically worrying versions of the debunking argument that I will focus on in what follows.

2. LUCK AND RISK OF ERROR

Luck-based considerations provide some intuitive motivation for thinking that ethical belief is debunked by evolution. Does this intuitive thought survive scrutiny? Reflection

https://doi.org/10.1017/epi.2016.13 Published online by Cambridge University Press

on some simple cases suggests some natural parameters on when a belief is epistemically suspect due to luck. I will outline some of these parameters in the present section, but nothing in the constraints I mention here will surprise; the epistemology literature has worked over these in more detail than I can hope to cover here. All I will do is outline enough of what I take to be the most plausible aspects that have been noted by this literature, briefly mention why they are plausible, and develop them in enough detail to facilitate a closer look at the relationship between luck and evolutionary debunking.

As an entry point, consider a traditional case from Gettier (1963): Jones knows that Smith has 10 coins in her pocket, falsely believes Smith will get the job, and deduces that the person who will get the job has 10 coins in her pocket. As it turns out, Jones's belief happens to be true: Jones herself has 10 coins in her pocket, and it is she who is the one who will get the job. But while Jones truly believes that the person who will get the job has 10 coins in her pocket, it is natural to say that she doesn't know this, and one diagnosis of this is that this is because her belief is true owing to a kind of luck.[4]

Of course not just any kind of luck leading up to the formation of a true belief is problematic in the same way. Sometimes one's belief will owe its truth to an accident of luck but not in a way that makes the belief fail to be knowledge. If I happen to turn my head just in time to see that John is smiling, I can know that John is smiling, even though I am in a sense lucky to have turned my head and formed this belief at all.[5] The Gettier case and the head-turning case involve different kinds of luck, only one of which is relevant to knowledge. One might use a simple 'safety' condition to capture this distinction, according to which a nearby false belief that is formed by a similar method prevents an actual true belief from being knowledge. This is Simple Safety:

> Simple Safety If one knows p in a world w, then for any world w^* close to w where one believes p in w^* by a similar method, p is true in w^*.

Simple Safety might provide a nice first-pass gloss on why true belief in the Gettier case isn't knowledge (on the grounds that it could easily have been false that Jones has 10 coins in her pocket). And it distinguishes this case from the head-turning case (where, although one was lucky to have turned one's head, in the nearby worlds where one hasn't, one forms a false belief by a very dissimilar method). However we need to refine Simple Safety in a number of ways. For instance it cannot deliver the same verdict in similar cases where the belief in question is a belief in a necessarily true proposition, since these propositions are guaranteed to be true in all close worlds. (This problem is especially relevant when evolutionary debunking of 'basic' ethical beliefs is at issue in the following sections, as the contents of these beliefs are plausible candidates for propositions that are necessarily true if true at all.) Also the appeal to beliefs that are formed by the 'same' method will prove unsatisfactory in cases where issues of method-individuation arise. Here is a revision of Simple Safety that avoids some of these difficulties:

4 Dancy (1985: 134).

5 See similar examples in Pritchard (2004: 133 *ff.*); see also Kahane (2011: 105–6) for a preliminary discussion of kinds of luck as they pertain to debunking arguments.

https://doi.org/10.1017/epi.2016.13 Published online by Cambridge University Press

> SAFETY If one knows p in a world w, then for any world w^* close to w where in w^* one believes a proposition p^* similar to p, and the token causal process which produces the belief in p^* in w^* is similar to the token causal process which produces the belief in p in w, the belief in p^* in w^* is true in w^*.[6]

According to SAFETY, true beliefs in necessary propositions can fail to be knowledge. This is because of the requirement that similar beliefs in close worlds be true. If I correctly guess that 154 is the sum of 81 and 73, I still won't count as knowing that 154 is the sum: given that I was guessing, I could easily have guessed a different number and believed (say) that 183 is the sum.[7]

Not everyone will be happy with the luck-based approach to these cases, and will reject SAFETY and any natural refinement to deal with recalcitrant cases. While this is not the place to present a full defense of the relevance of SAFETY to knowledge and epistemic success, a few preliminary points about SAFETY as stated here are warranted.

There are two similarity-relations in SAFETY: a similarity-relation between counterfactual beliefs and the actual belief that is a candidate for knowledge, and a similarity-relation between the token causal process that produces the actual belief, and the token causal processes that produce counterfactual beliefs. We have already noted, in the move away from SIMPLE SAFETY, how the first similarity-relation accounts for the possibility of lucky true beliefs in necessary truths that are not knowledge. The second similarity-relation is in some ways more important, both in accounting for some alleged counterexamples to SAFETY-like principles in the literature, and in explaining when some causal influences that are irrelevant to the truth of a belief destroy knowledge.

First, some token causal processes leading up to a belief are distinguished by virtue of containing a knowledge-state that plays an explanatory role. One can learn the location of a nearby restaurant by the testimony of a local who knows where the restaurant is, and this knowledge is not destroyed if there is a nearby liar who would have fibbed if asked. In cases like these there is a nearby false belief – one formed on the basis of the testimony of the liar – but it is not produced by a sufficiently similar token causal process. If one actually talked to the honest and knowledgeable local, one's true belief is produced by a causal chain that in part includes the testifier's knowledge. But the nearby false belief produced by the liar is not the product of a chain that involves the speaker's knowledge in the same way. So that token chain is plausibly not similar to a sufficient degree.[8]

---

6 Compare Sosa (1999: 142), Williamson (2000: 325), Hawthorne (2004: 55–6), and Pritchard (2004: 163). One might wish to modify SAFETY to account for cases where one's belief-forming process fails to produce any belief at all (perhaps owing to reference failure – see Manley (2007)). I will omit these complications from the main discussion.

7 Terminological note: here and in what follows, I will speak of beliefs satisfying or failing to satisfy SAFETY. By this, I mean that they do not satisfy the consequent of the conditional labeled SAFETY, and hence on present assumptions do not constitute knowledge.

8 This kind of approach also accommodates the alleged 'Halloween' counterexample to a SAFETY-like principle from Comesaña (2005: 397). Comesaña's example involves a knowledgeable directions-giver who will tell the truth unless she believes she is speaking to Michael, in which case she will lie and give misleading directions. Comesaña then claims that, if you very nearly disguised yourself as Michael, but did not, then upon getting directions from the knowledgeable directions-giver, you have an unsafe true belief. The reason is that, while you in fact received accurate directions, you could easily have dressed up as Michael, and if you had, you would have received misleading directions and believed them. So there is a nearby world where you have a false belief. But this is intuitively a case of knowledge, and so

https://doi.org/10.1017/epi.2016.13 Published online by Cambridge University Press

A second feature of this similarity-relation is worth mentioning here as well. Sometimes the truth of the believed claim features nowhere in the relevant processes. This can happen in many ways. *One* way in which the truth of the believed claim is not involved in the relevant causal processes obtains when evolution operates to produce a belief without regard to its truth. This is one of the premises in the debunkers' arguments, which will receive more attention below. But for now we can simply note that *sometimes* the absence of a truth from a token causal process that produces a belief in that truth does impugn the epistemic status of the belief according to SAFETY. For instance suppose I believe that 154 is the sum of 81 and 73, but only because some neuroscientist who for some odd reason wants me to have this belief has manipulated my neurological state to make me believe it. Does this belief satisfy SAFETY? If the neuroscientist's preference for my believing that 154 is the sum of 81 and 73 has nothing to do with the fact that it is *true* that 154 is the sum of 81 and 73, then my belief does not.[9] She could easily have wanted me to have a different belief about sums in this case, and in these cases the token causal processes that produced my belief about sums would be very similar to the actual one. So the similarity in token causal processes component of SAFETY guarantees that at least some beliefs that are produced by processes that are independent of the truth of those beliefs will not be knowledge.[10]

(Comesaña claims) knowledge need not be safe.

The above reflections on the interaction of SAFETY with the epistemology of testimony suggest that the differences in the similarity-relations between the token causal processes in the Halloween case will in fact be substantial. If a process that involves a testifier who relies on their knowledge is very different from one which does not, then it is also quite natural to think that even if the same testifier is present in both cases, but in one case the testifier speaks knowledgeably, and in the other case lies, then the token causal processes that produce one's beliefs about the relevant directions will also be very dissimilar. (There is no difference between speaking to a reliable testifier with a liar in the vicinity, and speaking to a testifier who has a reason for lying in the vicinity.) So while it is true that one could easily have had a false belief about the directions, since one could easily have dressed up as Michael, one's false belief in that case would have been the product of a very different causal process. One's actual belief in Comesaña's case satisfies SAFETY.

9 I am distinguishing cases like these from a very different kind of case, where the neuroscientist wants you to believe that 154 is the sum of 81 and 73 because it is true (and she knows that it is). Here there is no case that SAFETY is not satisfied. Some may view cases like this as not involving knowledge, but intuitions vary here: it may be plausible to view this belief as knowledge, as it is analogous to a belief formed on the basis of a knowledgeable testifier, though the testimony here produces the belief in a rather unusual way.

10 This aspect of SAFETY is relevant to another class of alleged counterexamples to similar principles that can be found in the literature, for instance in Hiller and Neta (2007: 312) and Clarke-Doane (Forthcoming b: 32). Suppose one is in fake barn country, but has a firm resolve to only look at barns at the top of the tallest hill, and is in an environment where, although there are many fake barns present, none will be at the top of the tallest hill. Hiller and Neta claim in a case like this that the belief that one is looking at a barn satisfies a SAFETY-like principle in this case, although it is not knowledge.

But if one is only brutely concerned with the tallest hill, as the case stipulates, then there is a very natural sense in which it could very easily have turned out that one formed alternative preferences (say, a brute concern with barns on top of medium-sized hills), and thereby had a false similar belief. (Perhaps one couldn't have believed falsely that *that* very thing is a barn, but this complication is handled by the move from SIMPLE SAFETY to SAFETY; other barn-beliefs will be similar enough.) In these nearby worlds, one will believe falsely: some of these worlds contain beliefs to the effect that such-and-such is a barn when in fact it is only a fake barn. The absence of known truths from the token causal processes guarantees that they will be very similar, as brute concern for the tallest hill

https://doi.org/10.1017/epi.2016.13 Published online by Cambridge University Press

There is much more to be said about SAFETY as a formulation of the kind of luck that is incompatible with knowledge. I will not pursue these issues here, but instead will proceed with the following assumptions in hand: SAFETY is a principle that is well-motivated by non-controversial cases where there is an absence of knowledge due to the presence of a kind of epistemic luck. It is not unassailable, but it provides a workable criterion for distinguishing knowledge-destroying luck from other kinds of luck that are irrelevant to epistemology. As the relationship between epistemic luck and evolution is lurking in some debunkers' presentation of their arguments, this relationship deserves to be investigated more carefully by taking a look at the relationship between evolutionary influences on moral belief and SAFETY. It is uncontroversial that, given the debunkers' premises, there is *a* kind of luck present in moral beliefs. But whether a causal history that includes evolution implies the presence of problematic epistemic luck is a question that cannot be answered by appeals to direct intuition. Instead we need to take a closer look at the relationship between evolutionary origins and SAFETY.

As a warm-up, consider a pair of cases. In these cases the risk of error is the same: they do not differ in the probability that one has a false belief, and worlds where one has a false belief are equally nearby in the two cases. But they differ in *where* the risk is situated in the causal history of the relevant belief. Varying cases in only this way changes the epistemic status of the beliefs which are the result of the causal process.

The first case is COIN IN THE HEAD, adapted from a case in White (2010: 599):

> COIN IN THE HEAD: You are forming beliefs about your immediate environment. Inside your head, the various neural pathways involved in your belief-forming process all lead to a box. In the box is a coin that gets flipped around. The outcome of the coin flipping is objectively random, and as it flips your beliefs change in a systematic way: when the coin lands 'heads', you form the beliefs that you are located on Earth, and that you are a human with a body. These beliefs are, however, 'hardwired' into your brain; all that is needed to produce them is your specific cognitive architecture and the coin being heads side up. When the coin lands 'tails', your brain is hardwired to form different beliefs: you form the false beliefs that you are located on Mars, and that you are merely a brain with quasi-perceptual experiences caused by electrode stimulation. As the coin is flipped, you lose memory of your previous beliefs and cannot notice any inconsistency between your quasi-perceptual beliefs over time.

Very plausibly, one can't know that one is on Earth in COIN IN THE HEAD. SAFETY provides a natural and plausible explanation for why this is so, even when the coin is on 'heads' and your beliefs are mostly true. Worlds where the coin is on 'tails' are very close – if an objectively random coin flip had turned out differently, you would have been in one of these worlds. And if one were in such a world, one's beliefs in either outcome are quasi-perceptual, and are sufficiently similar beliefs. They are formed by similar token processes, as in both outcomes one's beliefs are the product of the coin flip plus hardwiring. And in the 'tails' world one's beliefs are false.

---

is not very different from brute concern for medium-sized hills. So SAFETY will count these beliefs as falling short of knowledge. The more general lesson here is that there are some ways of specifying a causal process that produces a belief which operates independently of whether that belief is true, and in some of these cases the specification will make it very easy for similar token processes to produce false beliefs. Thanks to an anonymous referee for raising cases like this.

https://doi.org/10.1017/epi.2016.13 Published online by Cambridge University Press

By varying the causal history of your belief without eliminating the need for luck in arriving at a true belief, we arrive at a very different epistemological verdict. Consider a related case, which we can call God's Coin:

> God's Coin: God created you five minutes ago, and before doing so, God flipped a coin. If the coin landed 'heads', God would create you in the environment you are actually in, and moreover would give you a normal human visual system and cognitive architecture that produces the true beliefs that you are located on Earth and are a human with a body. If the coin landed 'tails', God would not have given you a perceptual system, but rather would have given you credible but misleading testimony, on the basis of which you would form the false beliefs that you are located on Mars and are merely a brain with quasi- perceptual experiences caused by electrode stimulation. The outcome of God's coin flip is objectively random.

Suppose again the coin landed 'heads'. Unlike Coin in the Head, it is quite plausible to say that one knows propositions about one's condition and environment. (After all, one can clearly see that one is on Earth in this case.) But intuitively worlds where one has false beliefs about being an envatted brain on Mars are just as close in God's Coin as they are in Coin in the Head: the outcome of the coin flips are objectively random in both cases, and the coin could easily have landed 'tails' in each. And in each 'tails' scenario, you by and large have the same false beliefs. The crucial difference is these beliefs are plausibly produced by sufficiently dissimilar processes in God's Coin. While the false beliefs produced by a 'tails' outcome are the products of believing on the basis of testimony, the true beliefs in a 'heads' outcome are the products of a human visual system interacting with its environment. Since beliefs produced by similar perceptual processes are true in all nearby worlds, one's true beliefs in the 'heads' outcome satisfy Safety.

There are a number of disanalogies between God's Coin and Coin in the Head, but the central point is that all of these differences derive from the same type of chancy event: a coin flip. What happens downstream from the coin-flip is different in each case, and what this serves to illustrate is that *how* a chancy process operates will also be relevant to the epistemic status of a belief it produces. Let us reserve the term *epistemic luck* for the kind of knowledge-destroying luck that follows from a failure of Safety. Whether a belief is lucky in this kind of way is not only a matter of whether it (or a similar belief) could have been false, but also depends on the nature of the token causal process that produces the belief. Epistemic luck is the kind of luck that is present in Coin in the Head. But there are other kinds of non-epistemic luck that might lead via different causal processes to formation of true beliefs, and these kinds of luck are not necessarily incompatible with knowledge. For instance whatever luck is required for a true belief in God's Coin is not epistemic luck.

We need, then, to keep the distinction between epistemic luck and other kinds of luck in mind when assessing whether the need for luck of some kind in the causal history of a belief implies that it is epistemically suspect. In the next section, I outline the case that appeals to luck in evolutionary debunking arguments have run afoul of these distinctions.

3. EVOLUTIONARY LUCK

3.1 The Appeal to Luck in General form

We outlined in § 1 the intuitive connection between evolutionary debunking and luck. It is easy to see why luck of some kind or other would be necessary for true ethical belief: it is

https://doi.org/10.1017/epi.2016.13 Published online by Cambridge University Press

plausible (and I will assume here) that the evolutionary forces that shape these beliefs are not themselves concerned with producing true beliefs. While evolution might select for creatures who are disposed to form certain ethical beliefs, it is no part of the evolutionary process that these beliefs are selected for because they are true. If one forms the beliefs evolution selects for, and thereby forms true beliefs, it is no stretch of ordinary language to say that one is thereby lucky in some sense.

But this does not get us very far by itself. We concluded the previous section by noting that not every instance of luck in forming a true belief is epistemically relevant, so we should be careful about what kind of luck is present owing to the evolutionary origins of the relevant beliefs. One might be lucky in some sense but still have knowledge, as in God's Coin. The present section carries out this neglected task, and looks at whether evolutionary origins make it at all plausible that ethical beliefs are dependent on epistemic luck.

First, more detail is needed about the contours of the evolutionary origins that feature in the debunker's argument. Debunkers do not typically claim that any belief that has an evolutionary explanation is epistemically suspect. Instead they only wish to target beliefs with a certain *kind* of evolutionary explanation.[11] To get an initial grasp of the kind of explanation that debunkers have in mind, a contrast case might be helpful.

Consider an evolutionary explanation of perceptual beliefs. A system for forming beliefs about one's immediate environment will plausibly confer a selection advantage on an organism only if it produces true beliefs about the organism's environment. Perceptual beliefs are causal inputs into processes that are directly related to survival and reproduction, such as finding food and avoiding predators. And perceptual inputs about the immediate environment will be better suited to producing these behaviors that confer selection advantages if they are true. So it is no evolutionary accident if one finds oneself with a true belief-producing perceptual system: natural selection will favor perceptual systems that generally produce true beliefs about the environment rather than false beliefs.[12]

Arguably analogous claims cannot be made for ethical beliefs. For one illustration of this idea we can take the outline of the evolutionary explanation of why we form ethical beliefs that Joyce (2006) favors: roughly, ethical belief provides a selection advantage because organisms who have the capacity for ethical belief will do better at resisting the temptation to cheat to obtain short-term goals when such cheating is disadvantageous in the long term.[13] According to this proposal, (i) there is a selection advantage being

11 See Vavova (Forthcoming) for more on the details of the debunker's evolutionary premise. Here I will focus on the aspects of evolutionary processes that might plausibly be thought to have consequences relevant to the question of epistemic luck.

12 Wilkins and Griffiths (2013: 139).

13 Here is Joyce:

> [A] person who in addition to being sympathetic judges that cheating is morally wrong will feel very differently if on occasion she succumbs to temptation. She can tell herself that she has done something wrong, that her action was unfair or unjust, that she must make amends, that she not only has risked punishment but deserves it. The emotion of guilt is available to her … The fact that these more robust forms of self-recrimination are available to the moralized thinker when she does cheat strongly suggests that when she is behaving herself her motivation not to cheat is more reliable and resolute than that of her non-moralized counterpart (Joyce 2006: 112–13).

https://doi.org/10.1017/epi.2016.13 Published online by Cambridge University Press

resistant to short-term temptations, and (ii) one way of acquiring this advantage is to acquire the capacity to form ethical beliefs. The crucial feature of this story for our purposes is that this selection advantage is available regardless of whether it is *true* that one ought to avoid succumbing to temptation. (Other evolutionary explanations will have a similar upshot; I focus on Joyce's solely for the purposes of illustration.)

In this respect there is an important contrast with perceptual belief. If our ethical beliefs are true, it is not because natural selection favored an ethical belief-forming module that produces true ethical beliefs. Debunkers have found talk of accidents and luck to be appropriate in light of this difference, and this is the feature of ethical belief that makes the motivations from luck sketched in §1 seem especially plausible. We can grant that the debunker is on to something; there is plausibly some difference between ethical and perceptual belief-forming mechanisms in this respect. We can even use 'luck' to label the feature that the debunker has latched onto. In order to not prejudge matters, let us call it *evolutionary* luck. Ethical beliefs require evolutionary luck in order to be true, while perceptual beliefs do not.

By calling this a species of luck, we haven't thereby shown anything about the epistemic status of these beliefs. Evolutionary luck isn't necessarily a species of epistemic luck. Rather, accusations of a need for luck will be relevant to a successful debunking argument only if an additional premise connecting the presence of evolutionary luck with epistemic luck holds. This is Luck Connection:

Luck Connection If a belief requires evolutionary luck, then it requires epistemic luck.[14]

If Luck Connection is true, then the epistemic consequences that attend to epistemic luck, including its incompatibility with knowledge, are available to the proponent of the debunking argument. Allegations of an objectionable reliance on luck in ethical belief will be entirely apt. In particular, the debunker can claim that since it is a consequence of evolutionary theory that ethical beliefs are subject to evolutionary luck, it follows from Luck Connection that these beliefs are subject to epistemic luck.

In the remainder of this section, I will consider two ways of arguing that Luck Connection is true.

3.2 *Conceptual Possibilities*

At a structural level, epistemic luck can be present in two ways. The first way arises when a belief isn't sufficiently sensitive to the fact it is about. If a true belief is lucky in this way, there are nearby worlds where one has the belief but the facts it is about have changed, and so the belief is false in nearby worlds. One is lucky to not have been in the nearby world with a false belief.

Does the phenomenon of evolutionary luck in ethical belief give rise to epistemic luck? To make use of the first structural feature of epistemic luck, one needs to show that ethical

It should be noted that this idea is not wholly uncontroversial: one can read Enoch (2011: 168 *ff.*) as challenging the claim that every evolutionary explanation for ethical belief will not imply that, or even raise the probability that, those beliefs are true.

14 Terminological note: I will use phrases like 'belief *b* requires evolutionary luck' as shorthand for 'if *b* is true, then *b*'s truth is a matter of evolutionary luck'. Similarly for the phrase '*b* requires epistemic luck'.

https://doi.org/10.1017/epi.2016.13 Published online by Cambridge University Press

belief which fails to be sufficiently sensitive; this involves showing that one retains one's beliefs in worlds where the ethical facts are different. As a textual note, there are some passages in Street's writings which suggest that she has a picture like this in mind when she alleges that ethical beliefs require objectionable luck:

> [A]s a matter of sheer luck, evolutionary pressures affected our evaluative attitudes in such a way that they just happened to land on or near the true normative views among all the conceptually possible ones.[15]

One feature of this passage deserves special attention. Suppose one's basic ethical beliefs are true. What are worlds like where one has those same basic ethical beliefs and thereby has a false belief? These are worlds where the basic ethical facts are different. But, quite plausibly, the basic normative facts – whatever they turn out to be – are *necessary*. So, there are no metaphysically possible worlds where one has the same ethical beliefs and yet they are false. Such worlds are nonetheless *conceptually* possible – and this is a point Street explicitly acknowledges in the passage quoted above, which conspicuously avoids the claim that there are metaphysical possibilities where a different ethical system is true.

One can use these possibilities (viz., counter-possible possibilities where ethical beliefs are held fixed and the basic ethical facts they are about are different) to forge a connection between evolutionary and epistemic luck, and thereby show that Luck Connection is true. First, if a belief requires evolutionary luck, then in the possibilities where the basic ethical facts are different, evolution will not change course and produce correspondingly different ethical beliefs. The same beliefs confer survival advantages as in the actual world; this fact does not depend on the actual ethical facts, and so in worlds where the basic ethical facts are different but evolution operates in the same way, one will have false beliefs. This is a general consequence of the nature of evolutionary luck:

> **1.** If a belief requires evolutionary luck, then in worlds where that belief is produced by a similar (evolutionary) process but the facts are different, one has false beliefs.

This observation only lends support to Luck Connection in conjunction with the following premise:

> **2.** If in worlds where a belief is produced by a similar (evolutionary) process but the facts are different, one has false beliefs, then that belief requires epistemic luck.

1 and **2** together guarantee that if ethical belief requires evolutionary luck, then it also requires epistemic luck. Given that it is also plausible to grant the debunker that ethical belief does require evolutionary luck, we have the makings of an argument for Luck Connection.

The problem for this argument is that **2** is false given some very natural assumptions. First, recall that the worlds which satisfy the antecedent of **2** are worlds that are

15 Street (2008: 208–9). See also Clarke-Doane (2012: 334 *ff.*) and Clarke-Doane (Forthcoming a: §4.4) for more discussion of attempts to turn this insight into a canonical version of the debunking argument.

https://doi.org/10.1017/epi.2016.13 Published online by Cambridge University Press

metaphysically impossible: they are worlds where the basic ethical facts are different, while the processes which produce ethical beliefs remain the same. Second, recall what epistemic luck requires: that there be a *nearby* world where one has a similar but false belief. But third, it is very plausible that metaphysically impossible worlds are not nearby worlds – that is, they are not worlds we could easily have been in; we are not at present at risk of being in such a world. There are some difficult questions for anyone who wants to understand fully what it takes to be a 'nearby' in the sense relevant to epistemic luck. But cases involving metaphysically impossible worlds are not among the difficult cases.[16]

There is more to be said about **2** for the debunker, although it is questionable whether this helps her case that there is a significant role for epistemic luck to play in debunking arguments. Rather, the point might best be construed as a possible diagnosis of why the debunker has been misled into thinking that LUCK CONNECTION is true. Begin with the observation that, according to SAFETY, the distant possibilities where one's ethical beliefs are false (owing to a counter-possible change in the ethical facts) are not irrelevant to what one can know. This is because of the following: while knowing p requires that the belief that p is safe; knowing that one knows p requires safely believing that one knows p. Knowing that one knows p, in other words, requires that the belief in p be *safely safe*: not only could a belief similar to belief that p not easily have been false, but in addition the belief that belief in p is safe itself could not easily have been false.

These do not come to the same thing.[17] One can know p, which means that one's belief in p is safe, and yet not know that one knows p, since one's belief that belief in p is safe isn't safe. Unpacking a bit: suppose that, in any nearby world, beliefs similar to the belief that p are true. One's belief is safe, and can be known. But to know that one knows p, one's belief that one knows p must be safe. Even if it is true that one's belief in p is true in nearby worlds, it might be that in some of these nearby worlds, one's belief that one knows p is false. This happens if, in a world that is near to a nearby world, one's belief in p is false. Worlds that are near to nearby worlds aren't *ipso facto* nearby themselves. And false beliefs in worlds that are near to nearby worlds preclude knowledge that one knows p, without precluding knowledge of p.

The point generalizes to further iterations of knowledge. A false belief in a world nearby to a nearby to a nearby world precludes knowing that one knows that one knows, though not necessary shorter iterations of knowledge. And so on.

The debunker has pointed out that there are guaranteed to be counter-possible worlds where one has false moral beliefs. These are not guaranteed to be nearby worlds, I have pointed out, and so these counter-possible worlds do not show that one's ethical beliefs are lucky. But their being counter-possibilities is no barrier to their being *some* iteration of nearbyness to nearby worlds. So one might argue that the debunker's counter-possible worlds are guaranteed to show that some iteration of knowledge is unavailable because the belief that one has one fewer iterations of knowledge will, while true, be objectionably lucky. One will not be able to know that one knows ethical facts, or will not be able to know that one knows that one knows, or ...

Debunkers, however, are naturally interpreted as claiming that the epistemic import of evolutionary luck is stronger than this: they not only take the presence of evolutionary

16 See Clark and Rabinowitz (2011) for a related point in a different setting.

17 More rigorous discussion of this point can be found in Williamson (2000, 2011). I will stick with informal glosses of the issue here.

https://doi.org/10.1017/epi.2016.13 Published online by Cambridge University Press

luck to show that some iteration of knowledge of ethical claims is unavailable; they take it to show that no knowledge is available. Counter-possibilities where the basic ethical facts are different need to get us from 1 to the claim that the first-order ethical beliefs are epistemically lucky. We could get this result if we collapse iterations of nearbyness: that is, if any world that is nearby to a nearby world is a nearby world simpliciter.[18] But this is a drastic measure. It threatens to allow any possibility where one's belief is false to count as nearby, and hence show that even if one succeeds in forming a true belief, there is a nearby possibility where one believes falsely. Epistemic luck under such an assumption threatens almost any ordinary belief, since there is some possibility of believing falsely on any matter. Debunkers would do well not to rely on such an assumption when claiming that epistemic luck infects ethical belief.

3.3 Belief-similarity

There is a second way in which epistemic luck can manifest itself. Even if the facts targeted by a belief are fairly stable – that is, if the facts are the same across all nearby worlds – one's belief-forming methods might not produce beliefs in a way that matches the factual stability. Epistemic luck in Coin in the Head arises in this way. So there is another way in which ethical beliefs might require epistemic luck, and this does not rely on the relevance of counter-possible conceptual possibilities where the basic ethical facts are different.

To show that ethical beliefs are epistemically lucky in this way, one needs to show that there are nearby worlds where one has similar but false beliefs. This can arise not because the basic ethical facts have changed in these worlds, but rather because ethical beliefs in these worlds fail to accurately represent these facts. This is what happens in Coin in the Head: owing to the coin-flipping hardwiring, one's quasi-perceptual beliefs about the outside world are different (and false) in nearby worlds where the coin lands 'tails'. Nearby changes in one's outside environment (if there are any) are not the source of epistemic luck in this case.

So the question arises: if ethical beliefs require evolutionary luck, is this enough to show that they also require epistemic luck, owing to the nearby possibility of different but false basic ethical beliefs? Evolution might be thought to be working like the box in cases where beliefs are subject to evolutionary luck: if natural selection doesn't favor ethical beliefs because they are true, then it might have produced different beliefs that have the same survival value but a different truth-value. One might think that for similar reasons evolutionary luck gives rise to epistemic luck. This would establish Luck Connection.

There are two premises in this argument. The first is an alleged consequence of evolutionary luck:

> **3.** If a belief requires evolutionary luck, then evolution could easily have produced a similar but false belief.

The second draws a connection between this consequence and epistemic luck:

> **4.** If a belief is such that evolution could easily have produced a similar but false belief, then that belief requires epistemic luck.

18 This would validate the claim Williamson (2000) labels 'KK' and provides compelling counterexamples against.

https://doi.org/10.1017/epi.2016.13 Published online by Cambridge University Press

I have briefly sketched reasons why we might accept 3: roughly, because evolutionary luck by definition implies that a belief was not selected for because it is true. There are a number of questions that remain about 3,[19] but suppose we grant it. Still, 4 is false. The reason is that, with an appropriately nuanced understanding of epistemic luck in hand, its presence does not follow from the nearby possibility of a similar but false belief alone. The nearby false belief needs, in addition, to be the product of a sufficiently similar causal process. Otherwise it is like my belief that John is smiling, formed as I happened to turn my head just as he began to smile. I might easily have had a false belief since I could easily have not turned my head and believed that he is not smiling on the basis of what I know about John's current mood. But the nearbyness of this scenario is not a threat to knowledge, since the false belief is formed by a very different process. There is luck in this situation, but it is GOD'S COIN-type luck, and is not epistemically problematic.

This isn't a mere technical quibble, and is instead very important for a discussion of the import of evolutionary luck. Strictly speaking, evolution doesn't directly explain particular token beliefs such as Joe's believing such-and-such at a certain time and a certain date. Instead, at best, it explains certain cognitive *dispositions* or *propensities* to form beliefs of a particular kind. This fact is acknowledged by some debunkers,[20] but the relationship between evolutionary explanations of cognitive dispositions, and its relationship to luck is rarely explored. The framework outlined here suggests that the issue is highly important: if we *could* have evolved to form different beliefs, but the evolutionary processes that would explain these are processes that involve very different cognitive dispositions, then the possibility of false moral belief would appear prima facie to be irrelevant to the issue of epistemic luck.

At the very least nothing we have set down so far guarantees that moral beliefs will have false nearby counterparts that are also produced by very similar token processes. The possibility that such nearby false beliefs exist, but are the products of very different cognitive processes that would have been deployed if evolution had taken a different route, remains. So we should conclude that the debunker is not entitled to 4 in an argument for LUCK CONNECTION.

To sum up: evolutionary origins of ethical belief do not, by themselves, do much to show that epistemic luck infects ethical belief. Debunkers have been quick to throw around allegations of luck for evolutionarily explained ethical beliefs. But they have not distinguished between kinds of luck that are epistemically problematic, and those that are not. Their allegations of luck are much less compelling when these distinctions are kept clearly in mind.

But we noted that there is simply a gap in one of the ways the debunker might go about making her case: she could claim that, in fact, evolution worked in ways that does predict the existence of nearby false moral beliefs that are *also* the products of token processes that are very similar to the token processes that produced our actual beliefs. Whether this is so is an empirical question. In the closing section, I do not weigh in on the empirical issue, but instead explore in schematic form what the empirical details would have to look like, and some of the resulting issues, for a resuscitation of the second argument for LUCK CONNECTION.

19 See Clarke-Doane (Forthcoming a: 4.5) for criticism.

20 Street (2006: 113 *ff.*)

https://doi.org/10.1017/epi.2016.13 Published online by Cambridge University Press

4. WHEN DO EPISTEMIC AND EVOLUTIONARY LUCK COINCIDE? THE WAY FORWARD

While there is a risk that debunkers conflate the category of epistemic luck with other epistemically benign forms of luck, we haven't shown here that the categories of evolutionary and epistemic luck never overlap. Debunkers might favor a less ambitious claim and try to show that, even if evolutionary luck does not by itself guarantee that ethical beliefs require objectionable epistemic luck, there are other, additional features of the evolutionary explanation for ethical belief that can be brought to bear in securing this conclusion. The prospects for luck to assist the debunker's argument are then not closed, but will have to pay special attention to whether it is empirically plausible that the best evolutionary explanations for ethical belief have the needed features. I will close by outlining some constraints on a project of this kind.

4.1 Belief-forming Processes

The problem for the second route to establishing LUCK CONNECTION is that, for some evolutionary explanations for a belief, the only way for a different belief (and hence a belief with a different truth-value) to arise is for evolution to select for a distinct cognitive mechanism. That is, for some theories, instances of MECHANISM CORRELATION hold where b is a basic ethical belief:

> MECHANISM CORRELATION If the belief b is produced by the cognitive mechanism M, then in any nearby possibilities where one forms the false belief b^* instead of b, b^* is produced by the distinct cognitive mechanism M^*.

If MECHANISM CORRELATION holds, then it is highly plausible that the nearby possibility for forming the false belief b^* does not threaten b with problematic epistemic luck: for the token processes which lead up to the beliefs will go via two distinct cognitive mechanisms and, as a result, will not be very similar. Threats of epistemic luck will not be very credible for empirical theories that accept it.

Of course not every way of working out an evolutionary theory of ethical belief implies MECHANISM CORRELATION. Instead it is empirically possible that the best working out of such a theory posits a single mechanism which is capable of producing both truth and false ethical beliefs. On such a theory, evolution would have selected for an ethical belief-forming mechanism that structurally resembles the box in COIN IN THE HEAD. The same mechanism produces both true and false beliefs in nearby worlds, and hence instances of PROCESS SIMILARITY, where b is a basic ethical belief, are true on such a theory:

> PROCESS SIMILARITY For belief b, there are similar but false beliefs b^* in nearby worlds, where the token process which produces b and the token process which produces b^* are very similar.

Theories which imply that ethical beliefs are subject to epistemic luck must be consistent with PROCESS SIMILARITY for ethical beliefs.[21]

---

21 In order to be consistent with PROCESS SIMILARITY, a theory must entail the denial of MECHANISM CORRELATION. But note that the denial of MECHANISM CORRELATION does not, by itself, entail PROCESS SIMILARITY. Presumably all human visual systems have the same evolutionary explanation, but it is

https://doi.org/10.1017/epi.2016.13 Published online by Cambridge University Press

This provides a template for the claims the debunker needs in order to successfully vindicate allegations of epistemic luck. A plausible case that the template is satisfied will be largely empirical in nature. To illustrate: suppose that the propensity to form beliefs in the supernatural is the product of wishful thinking: roughly, people tend to believe these things because one wants them to be true. This may or may not be empirically plausible. But *if* it is true, it entails that religious beliefs are at best epistemically lucky. For suppose one managed to arrive at true beliefs about the supernatural via wishful thinking: nevertheless, similar but false beliefs could easily have been formed by a similar process of wishful thinking. One could easily have wanted different things (one could want for instance that atheism rather than theism to be true), and in this case one would have formed similar false beliefs by similar processes.[22]

So which theory of the cognitive mechanism that leads to the formation of supernatural beliefs is true matters a great deal when we are assessing whether such beliefs require epistemic luck. If the wishful thinking theory is true, PROCESS SIMILARITY is true for religious belief, and a good case for epistemic luck can be made. But we cannot assume that every alternative explanation for these beliefs will retain the aspects of the wishful thinking explanation that are epistemically relevant. Different hypotheses about the mechanisms involved in belief-formation might make the prospects for the epistemic luck criticism significantly worse.

The epistemic import of evolutionary hypotheses about the formation of ethical beliefs will follow the same contours. Some evolutionary theories will not rule out METHOD CORRELATION, and will not provide a suitable basis for criticisms from epistemic luck. On the other hand theories which imply PROCESS SIMILARITY will be more promising. Whether the debunker has the tools needed to properly raise the issue of epistemic luck is, in large part, an empirical question.

4.2 *Propensities and Environmental Luck*

In closing I will outline one further aspect of the template that needs to be filled out in order for epistemic luck to be relevant to the debunker's purpose. Again it will be an empirical question whether this aspect of the template is satisfied.

---

false that any two token beliefs in the proposition that *War and Peace* is a long book formed by perception are products of relevantly similar causal processes. You might believe this by looking at the table of contents and noticing the page count, and your friend might believe it by looking at the closed book from the side and noticing that it is thick. Suppose both processes are ones that you might have instantiated (you and your friend are in the same room) and you could easily have formed a false belief by looking at the table of contents (since the printer was highly unreliable). If you actually look at the book from the side, it doesn't follow that you could easily have had a false belief by a similar process because you could easily have looked at the table of contents instead. And this is so even though you would be using your visual system in either case. The token causal visual processes producing the belief that *War and Peace* is long are not sufficiently similar.

22 For example even if one believes in a deity because of wishful thinking, one could instead have had the sentiment expressed in Nagel (1997: 130):

> I want atheism to be true and am made uneasy by the fact that some of the most intelligent and well-informed people I know are religious believers. It isn't just that I don't believe in God and, naturally, hope that I'm right in my belief. It's that I hope there is no God! I don't want there to be a God; I don't want the universe to be like that.

https://doi.org/10.1017/epi.2016.13 Published online by Cambridge University Press

It is unlikely that evolution properly explains ethical belief; an empirically plausible evolutionary theory will, at best, explain a widespread propensity to form these beliefs. Propensities constitute only a partial explanation for why a particular individual has a particular belief. To illustrate with the example of belief in the supernatural: there might be a credible evolutionary explanation for belief in the supernatural,[23] but this does not mean that every agent will have such beliefs; plenty of atheists exist, even though their cognitive architectures are the product of the same evolutionary forces. They have an unactualized propensity to form beliefs in the supernatural.

That it is propensities which are the products of evolutionary explanation is doubly relevant to questions of epistemic luck. First, it makes the claim that one could easily have formed similar but false beliefs easier to establish. But, second, it also makes the claim that the false beliefs would have been formed by a similar process more difficult to establish. I take each of these in turn.

First: as an initial approximation, take the relevant propensities to be dispositions that manifest themselves in some but not necessarily all environments. A propensity to form certain ethical beliefs, then, is a disposition to form such beliefs in certain environments. Moreover, this disposition need not manifest itself by producing the same ethical beliefs across a range of different environments. Suppose Joyce's explanation of ethical belief, on which it confers a selection advantage by enabling organisms to resist the temptation to cheat for short-term gains, is correct. The disposition that confers the relevant selection advantage produces different specific ethical beliefs in different environments: one might have this disposition and form utilitarian ethical beliefs in one environment, but might in a different environment have this same disposition and yet believe a rights-based morality. Which specific set of beliefs a token individual ends up with is explained not just by this disposition, but by its interaction with the rest of the individual's cognitive faculties, their cultural setting, and other environmental factors.

If the disposition to form ethical beliefs might have produced different beliefs in different environments, one might naturally reason as follows: one could easily have been in a different environment, and if one were in that environment, one's disposition to form ethical beliefs would have manifested itself by leading to the formation of possibly false ethical beliefs. Taking propensities seriously appears to make one component of epistemic luck easier to establish, since one will be able to find environments where the agents in question have similar but false beliefs.

Now for the second point: the focus on propensities makes another aspect of epistemic luck, which requires Process Similarity, *more* difficult to establish. Let D be the disposition to form ethical beliefs which evolution has selected for. And let $E1$ and $E2$ be different environments broadly construed in which D might be placed, where D produces the belief $b1$ in $E1$ and the belief $b2$ in $E2$. And finally, let us suppose that the beliefs $b1$ and $b2$ are not both true in their respective environments. The question pertinent to epistemic luck is whether the token causal process which involves D in the environment $E1$ and produces $b1$ is sufficiently similar to the token process which involves D in $E2$ and produces $b2$. Even if the dispositions in the two cases are the same, it is still a further question whether the entire token causal processes leading up to the beliefs in the two cases are sufficiently similar.[24]

23 For example, see Bloom (2009).

24 For more on these issues, see Dunaway and Hawthorne (Forthcoming).

https://doi.org/10.1017/epi.2016.13 Published online by Cambridge University Press

If evolution is going to have implications for the presence of epistemic luck, the answer has to be 'yes'.[25] This is a question, however, which cannot be answered in the abstract. Different answers to this question may well be plausible depending on what exactly the disposition *D* is, and how the relevant environments interact with the disposition to produce beliefs. Some dispositions evolution selects for might be realized in very robust cognitive mechanisms where it is quite natural to think that any two beliefs that are products of such a mechanism are formed by similar token processes, even if strictly speaking other inputs to the mechanism are needed for it to produce any beliefs. But there is no guarantee that evolution selects for a disposition that is realized in this way; other hypotheses may deliver different results. Again, crucial questions about the role of epistemic luck in debunking arguments are such that their answers depend on the results of evolutionary biology and cognitive science.

5. CONCLUSIONS

Debunkers have been quick to raise considerations of epistemic luck when motivating their case that there is something epistemically defective about ethical belief owing to its evolutionary origins. But serious reflection on the nature of epistemic luck, and what distinguishes it from benign cases of luck, suggests that the connection between evolution and epistemic luck is tenuous. Those that wish to continue to press a link between evolution and epistemic luck will need to marshal specific (and no doubt controversial) evolutionary explanations of ethical belief which fit the template outlined in §3.

This does not mean that evolutionary debunking arguments can be safely ignored. Debunkers might defend the needed empirical hypotheses that fit aspects of the template outlined in §3. Alternatively they might retract their suggestions that epistemic luck is relevant to the epistemological upshot of evolutionary debunking arguments. There may be something left to these arguments, once the allegations of luck are scaled back.

Nothing I have said here rules out this latter strategy. But it will be worth pointing out how the debate over debunking arguments can be expected to play out, and what obstacles the debunker can expect to face, given such a move. First, the supposed analogies with non-evolutionary cases where epistemic luck explains an intuitive verdict will need to be discarded. It is worth revisiting Joyce's Napoleon Pill case in this connection: since the presence of epistemic luck explains why one's Napoleon belief is epistemically suspect, it will no longer serve as a motivating analogy for the debunker's thought that ethical beliefs explained by evolutionary forces are similarly suspect. We know that the epistemically important features of the Napoleon case are not present in the evolutionary case, and so any intuitive support for the debunker's verdict from the analogy goes missing. Similar points apply to use of terminology like 'coincidence', 'accident', and the like. When we understand the debunker's use of these terms, we need to keep in mind that, if the

25 Obviously the prospects are better here than they are in §2.3: in the present case, we have at least guaranteed that the disposition *D* is a part of the token belief-forming process in both cases, so some similarity is guaranteed. When evaluating premise **4**, we could not even assume this much: the presence of evolutionary luck in a belief guarantees that very little will be held in common between the process producing a true belief and a nearby process producing a similar false belief.

https://doi.org/10.1017/epi.2016.13 Published online by Cambridge University Press

debunker is warranted in using them, they cannot entail that the beliefs in question require epistemic luck.

Second, once these misleading suggestions are discarded, the debunker will need to find another route to motivating the thought that there is anything epistemically worrisome in evolutionary luck. It isn't obvious that there is anything problematic here: even debunkers will acknowledge that plenty of beliefs have a causal history that is epistemically uninteresting. Perhaps there will be an alternative, widely recognized epistemological vice that arises from the evolutionary origins of ethical belief. But luck has been the focus of debunkers for good reason: it is an uncontroversial condition on many epistemically interesting notions, and assessing whether it is (or is not) present is often straightforward. How the debunker might accomplish this task while explicitly eschewing talk of luck is unclear, as epistemology is a contentious area, and there are not many alternatives that will enjoy widespread acceptance. Regardless of the outcome, extended reflection on the role of luck in evolutionary debunking arguments points to some directions which are dead ends for debunkers, sheds light on directions in which a successful argument might be developed, and imposes significant constraints on what would constitute success for the debunker.[26]

REFERENCES

Bloom, P. 2009. 'Religious Belief as an Evolutionary Accident.' In J. Schloss and M. Murray (eds), *The Believing Primate: Scientific, Philosophical, and Theological Reflections on the Origin of Religion*, pp. 118–27. Oxford: Oxford University Press.

Clark, K. J. and Rabinowitz, D. 2011. 'Knowledge and the Objection to Religious Belief from Cognitive Science.' *European Journal for Philosophy of Religion*, 3: 67–81.

Clarke-Doane, J. 2012. 'Morality and Mathematics: The Evolutionary Challenge.' *Ethics*, 112: 313–40.

—— Forthcoming a. 'Justification and Explanation in Mathematics and Morality.' In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics*, vol. 10. Oxford: Oxford University Press.

Clarke-Doane, J. Forthcoming b. 'Debunking and Dispensability.' In N. Sinclair and U. Leibowitz (eds), *Explanation in Ethics and Mathematics*. Oxford: Oxford University Press.

Comesaña, J. 2005. 'Unsafe Knowledge.' *Synthese*, 146: 395–404.

Dancy, J. 1985. *Introduction to Contemporary Epistemology*. Oxford: Basil Blackwell.

Dunaway, B. Manuscript. 'Grounds and Knowledge: Epistemological Motivations for Anti- Realism.'

—— **and Hawthorne, J.** Forthcoming. 'Scepticism.' In W. J. Abraham and F. D. Aquino (eds), *Oxford Handbook of the Epistemology of Theology*. Oxford: Oxford University Press.

Enoch, D. 2011. *Taking Morality Seriously: A Defense of Robust Realism*. Oxford: Oxford University Press.

Gettier, E. L. 1963. 'Is Justified True Belief Knowledge?' *Analysis*, 23: 121–123.

Hawthorne, J. 2004. *Knowledge and Lotteries*. Oxford: Oxford University Press.

Hiller, A. and Neta, R. 2007. 'Safety and Epistemic Luck.' *Synthese*, 158: 303–13.

Joyce, R. 2006. *The Evolution of Morality*. Cambridge, MA: MIT Press.

Kahane, G. 2011. 'Evolutionary Debunking Arguments.' *Noûs*, 45: 103–25.

Manley, D. 2007. 'Safety, Content, Apriority, Self-Knowledge.' *Journal of Philosophy*, 104: 403–23.

Nagel, T. 1997. *The Last Word*. Oxford: Oxford University Press.

26 Special thanks to Matthew Benton, Julien Dutant, Daniel Fogal, John Hawthorne, Brian Hedden, Maria Lasonen-Aarnio, David Manley, Katia Vavova, Dani Rabinowitz, and audiences at University College, London and St. Andrews for many valuable comments and/or discussion of this paper.

https://doi.org/10.1017/epi.2016.13 Published online by Cambridge University Press

Pritchard, D. 2004. *Epistemic Luck*. Oxford: Oxford University Press.
Schafer, K. 2010. 'Evolution and Normative Scepticism.' *Australasian Journal of Philosophy*, 88: 471–88.
Setiya, K. 2012. *Knowing Right from Wrong*. Oxford: Oxford University Press.
Sosa, E. 1999. 'How to Defeat Opposition to Moore.' *Philosophical Perspectives*, 13: 141–53.
Street, S. 2006. 'A Darwinian Dilemma for Realist Theories of Value.' *Philosophical Studies*, 127: 109–66.
—— 2008. 'Reply to Copp: Naturalism, Normativity, and the Varieties of Realism Worth Worrying About.' *Philosophical Perspectives*, 18: 207–27.
Vavova K. Forthcoming. 'Debunking Evolutionary Debunking.' *Oxford Studies in Metaethics*, vol. 9. Oxford: Oxford University Press.
White, R. 2010. 'You Just Believe That Because ...' *Philosophical Perspectives*, 24: 573–615.
Wilkins, J. S. and Griffiths, P. E. 2013. 'Evolutionary Debunking Arguments in Three Domains: Fact, Value and Religion.' In G. W. Dawes and J. Maclaurin (eds), *A New Science of Religion*, pp. 133–46. London: Routledge.
Williamson, T. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.
—— 2011. 'Improbable Knowing.' In T. Dougherty (ed.), *Evidentialism and Its Discontents*, pp. 147–64. Oxford: Oxford University Press.

BILLY DUNAWAY is Assistant Professor of Philosophy at the University of Missouri-St Louis. He received his PhD from the University of Michigan in 2013 and previously was a Postdoctoral Research Fellow at the University of Oxford from 2013–2015. Currently he works on the metaphysics of fundamentality, epistemic luck, and their relevance to traditional problems in metaethics.

https://doi.org/10.1017/epi.2016.13 Published online by Cambridge University Press