Canadian Journal of Philosophy (2026), 56: e47, 1–15
doi:10.1017/can.2026.10127

ARTICLE

Moral Realism and the Argument from Alternate Histories

Daniel Z. Korman

University of California, Santa Barbara, California, USA
Email: dkorman@ucsb.edu

Abstract
Evolutionary moral debunking arguments often emphasize the inevitability of our moral beliefs: given how natural selection operates, we would have had the moral beliefs we do regardless of what the moral facts are. I examine an alternative strategy, which emphasizes the contingency of our moral beliefs: there are things in our evolutionary history that could easily have gone differently and that would have resulted in substantially different moral beliefs. I consider various ways of moving from this observation to a skeptical upshot.

Keywords: meta-ethics; debunking; safety; moral realism; evolutionary arguments; error theory; coincidence

1. Introduction

According to evolutionary moral debunkers, we have the moral beliefs that we do because those beliefs, or at least the underlying attitudes that dispose us to form those beliefs, are adaptations: having such moral attitudes enabled our ancestors to outperform their contemporaries and pass on those attitudes to their descendants. This evolutionary explanation is meant to drive us to skepticism, either about the truth of our moral beliefs or about realist construals of morality.

Standard presentations of evolutionary debunking arguments emphasize the inevitability of our ending up with the beliefs that we did: natural selection was bound to favor the specific moral beliefs we have, irrespective of what the (putative) moral facts are.[1] However, an alternative way of developing these arguments has been gaining prominence, one that instead emphasizes the contingency of our moral beliefs. I will call this style of argument *the argument from alternate histories*, since it turns on the idea that there are things that could easily have gone differently for us or for our ancestors that would have resulted in profoundly different—and, by realist lights, profoundly misguided—moral outlooks.[2]

The shift to an argument from alternate histories is significant and merits special attention for several reasons. First, it aligns with a growing consensus among cognitive scientists and philosophers of biology that highly contingent cultural factors are the primary determinants of the content

[1]See especially Street (2006, Section 4). Here and in what follows, I ignore the complication, alluded to just above, that it arguably is not beliefs themselves but rather the underlying attitudes that are selected for.

[2]See Barkhausen (2016, Sections 2–3), Braddock (2016, Sections 6–7; 2021), Handfield (2016, 68–71), Warren (2017, 1655–1656), Clarke-Doane (2020, 148–149), Koon (2024, 472–474), and McKay (2024, Section 3). Street herself invokes alternate histories in an oft-cited passage (2006, 119–120), though her aim there is not to question the inevitability of moral beliefs but rather to establish that Darwinian forces indirectly influence our beliefs (which are not genetically heritable) by directly influencing underlying evaluative tendencies (which are genetically heritable).

© The Author(s), 2026. Published by Cambridge University Press on behalf of The Canadian Journal of Philosophy, Inc. This is an Open Access article, distributed under the terms of the Creative Commons Attribution-NonCommercial licence (http://creativecommons.org/licenses/by-nc/4.0), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original article is properly cited. The written permission of Cambridge University Press or the rights holder(s) must be obtained prior to any commercial use.

of moral beliefs.[3] Second, it strengthens debunking arguments against a standard realist response, which seeks to explain moral accuracy by appeal to the inevitability of our moral beliefs.[4] This appeal to inevitability is meant to be dialectically unobjectionable, since debunkers themselves rely on it in their arguments. But by grounding the argument in contingency rather than inevitability, the debunker is well positioned to challenge this realist maneuver. Third, it promises to sidestep well-known problems facing debunking arguments from insensitivity and explanatory constraints; more on that below.

Finally, the revelation that one's beliefs are at the mercy of highly contingent factors is exactly the sort of thing that can undermine one's beliefs. Whereas the mere possibility that you are a brain in a vat may well be no threat to your perceptual beliefs, learning of an ongoing wave of recent envattings in your neighborhood would pose a genuine threat. The revelation that our moral outlook is at the mercy of highly contingent features of our ancestral environments seems poised to undermine our moral beliefs in just the same way.

Given its emphasis on ways things could easily have gone differently, it is natural to formulate the argument from alternate histories as targeting the *safety* of our moral beliefs, where S's belief that p is safe iff S could not easily have been mistaken about whether p. I will show that the argument, so construed, fails (Sections 3–6), as does an alternative approach that turns on the idea that the highly contingent processes responsible for our moral beliefs are unlikely to yield accurate moral beliefs (Section 7). Finally, I consider whether the epistemic pressure generated by the contingency of moral evolution is better captured by an "explanationist" construal of the argument from alternate histories (Section 8).

2. The Argument from Alternate Histories

2.a. High contingency

The argument from alternate histories turns on the idea that there are things that could easily have gone differently and that would have resulted in people having profoundly different moral outlooks. What things? The debunker might point to aspects of our personal histories—for instance, where we grew up or whom we had as mentors (more on this in Section 6). But those advancing the argument more commonly point to contingencies in the distant past, for instance variable aspects of the environment of evolutionary adaptation that influence the evolution of moral norms.[5] There are various reasons for thinking that a great many of our moral beliefs are largely the products of cultural evolutionary processes, chief among them the diversity of moral outlooks across human history and the relative speed with which moral outlooks can change. I am going to assume that moral beliefs are indeed highly contingent in this sense: the specific contents of many of our moral beliefs are primarily the products of cultural evolutionary processes that are highly sensitive to features of our ancestral environments that could easily have been different. My interest is not in whether this assumption is correct, but in how it is meant to deliver a skeptical or anti-realist upshot.

For concreteness, it will be useful to work with a specific proposal about the source of this high contingency: Buchanan and Powell's "adaptive plasticity hypothesis."[6] Here is the idea. In the environment of evolutionary adaptation, outsiders can pose a threat (e.g., disease or attack), but also present opportunities (e.g., trade or mate exchange). In the presence of dangerous outsiders, it is advantageous to adopt "exclusivist" (illiberal) norms—for example, norms that prohibit associating with outsiders, permit violent treatment of outsiders, and downplay the extent to which their

[3]See, e.g., Richerson and Boyd (2005), Boehm (2012), Sterelny (2012), Henrich (2016), Buchanan and Powell (2016), and Cofnas (2023).

[4]See, e.g., Clarke-Doane (2020, 109–110).

[5]See especially Barkhausen (2016, 667–671), Braddock (2016, 2021), and Koon (2024, 472–474).

[6]See Buchanan and Powell (2016, 996–1004; 2018, Chapter 6), which builds on a more circumscribed proposal from Fincher and Thornhill (2012). Cf. Sober (1994, Section 2) on phenotypic plasticity.

interests must be taken into account. In the presence of harmless outsiders, it is advantageous to adopt "inclusivist" (liberal) norms that require fair and humane treatment of outsiders. According to the adaptive plasticity hypothesis, we evolved a kind of psychological "toggle" that affects our receptivity to exclusivist norms, and whose setting is determined by the apparent presence or absence of specific environmental conditions—conditions that affect whether it makes better prudential sense to greet outsiders with hostility or with open arms. Those exclusion-conducive conditions include severe competition for resources, the lack of institutions to facilitate cooperation among groups, and high risk of infection (2016, 990). The absence of such conditions increases receptivity to inclusivist norms—manifested in moral intuitions and emotions that align with those norms—and, the idea goes, we owe our own especially high levels of receptivity to inclusivist norms to such things as reduction in crime, the enforcement of property rights, and the emergence of markets (2016, 1000–1001).

For added concreteness, we will focus on one specific moral belief: that it is wrong to enslave prisoners of war (POWs). This choice of example tilts the scales in the debunker's favor, since such beliefs about the acceptable treatment of outsiders are among the most likely to be sensitive to environmental conditions on the adaptive plasticity hypothesis. Other moral beliefs—for example, those concerning incest—likely arise from hard-wired tendencies and are not especially sensitive to environmental conditions. Accordingly, there is a question of just how thoroughgoing a moral skepticism one can obtain from this argument from alternate histories. I mention this potential shortcoming only to set it aside. As we will see, there are problems even with the argument's most promising instances.

2.b. A safety-based approach

There is no universally accepted way of formulating evolutionary moral debunking arguments. However, as indicated above, an argument framed in terms of safety seems well suited for capturing the force of the argument from alternate histories.[7] The safety-based argument from alternate histories, targeting the belief that it is wrong to enslave POWs, can be understood as consisting of two complementary arguments:

The Safety-Based Debunking Argument
(DK1) ∀S∀p: if S should believe that S's belief that p is not safe, then S should not believe that p
(DK2) Informed realists should believe that their belief that it's wrong to enslave POWs is not safe
(DK3) So, informed realists should not believe that it's wrong to enslave POWs

The No Safety Argument
(NS1) There are things that (i) could easily have happened and (ii) would have resulted in you believing that it's not wrong to enslave POWs
(NS2) If so, then you could easily have believed that it's not wrong to enslave POWs
(NS3) If you could easily have believed that it's not wrong to enslave POWs, then you could easily have been mistaken about whether it's wrong to enslave POWs
(NS4) If you could easily have been mistaken about whether it's wrong to enslave POWs, then your belief that it's wrong to enslave POWs is not safe
(NS5) So, your belief that it's wrong to enslave POWs is not safe

As indicated above, I leave it open whether—and to what extent—such an argument would be equally effective against other moral beliefs.[8]

[7]Of the authors listed above as advancing arguments from alternate histories, Barkhausen (2016), Handfield (2016), Warren (2017), and Clarke-Doane (2020) are explicitly focused on safety.

[8]Notably, on more nuanced formulations of safety, one's belief that p can be unsafe even when one could not easily have been mistaken about p itself, so long as one could easily have been mistaken about propositions "sufficiently similar" to p (Williamson, 2000, 147; Manley, 2007, 403–404; Clarke-Doane & Baras, 2021, 164). In that case, moral propositions about which we could not easily have been mistaken may still turn out to be unsafe, so long as they qualify as sufficiently similar to

On premise DK1: There are various roads to the conclusion that the recognized absence of safety is epistemically significant. One is via reflection on Gettier-style cases. Suppose that, after glancing at a clock displaying 2:42, you form the belief that it is 2:42. You then learn that the clock is broken. You should suspend belief. Why? Not implausibly, it is because you recognize that—even if your belief happens to be true—you could easily have ended up mistakenly believing that it is 2:42 if, for example, you had glanced up at 2:41 instead.

On DK2: For our purposes, we can understand a *realist* to be someone who believes that moral truths are, for the most part, independent of our moral beliefs, and an *informed realist* to be a realist who is acquainted with the aforementioned No Safety Argument. The idea behind DK2 is that realists ought to find the No Safety Argument convincing. A realist who wishes to deny DK2 must find some premise of the No Safety Argument to reject.

NS1 is meant to be motivated by the high contingency of the factors responsible for our moral outlooks. Suppose that a prolonged drought or encounter with a belligerent tribe had occurred just at the time that, in the actual world, our ancestors were settling into inclusion-conducive conditions. These are things that could easily have happened. The adaptive plasticity hypothesis predicts that this would have resulted in our ancestors having an increased receptivity to exclusivist moral beliefs, and the cultural transmission of such beliefs to their descendants and ultimately to all of us.

NS2 falls out of the seemingly innocuous assumption that if a certain something could easily have happened that would have resulted in some further thing, then that further thing could easily have happened.

The idea behind NS3 is that the truth about what is right or wrong could not easily have been different, at least by realist lights, and at least when it comes to things like whether it is wrong to enslave POWs. Yet if there are no nearby worlds in which the truths about such matters are different, while there are nearby worlds in which your beliefs about them are different, then there are nearby worlds in which you are making a mistake; you could easily have been mistaken. Here again, my choice of example tilts the scales in the debunker's favor. For when it comes to other beliefs—that senicide is impermissible, for instance—a realist might reasonably contend that the moral facts would themselves have been otherwise in the harsh conditions in which we would have believed otherwise. But find me a realist who thinks it would be okay in sufficiently harsh conditions to keep POWs as slaves.

Finally, premise NS4 is trivial on the definition of safety provided above:

Safety
S's belief that p is safe iff$_{def}$ S could not easily have been mistaken about whether p

To be sure, epistemologists who put safety to serious epistemic work tend to prefer more nuanced definitions. For instance:

Method-Relativized Safety
S's belief that p is safe$_m$ iff$_{def}$ S could not easily have been mistaken about whether p, using the method that S actually used in determining whether p.[9]

To see why the shift to something like safety$_m$ is needed, consider the following example (adapted from Nozick 1981, 193):

Chance Glance
Jesse James is robbing First National Bank, partly as a ruse to frame Butch Cassidy. Jesse arrives masked and puts on an impeccable vocal impression of Butch. The teller, Mabel, is

moral propositions about which we could easily have been mistaken—thereby expanding the potential scope of the safety-based argument from alternate histories.

[9]See, e.g., Pritchard (2005, Chapter 6.2).

initially fooled and believes that Butch is robbing the bank. But she sneaks a peek from her hiding place at just the right moment to see his mask slip, and sees that the robber is Jesse, not Butch.

On our initial definition of safety, Mabel's belief that the robber is Jesse is unsafe. She could easily have gone on mistakenly believing that the robber is Butch, had she not peeked at just the right moment. Because of that, safety-involving epistemic principles like DK1 have the absurd consequence that Mabel should not believe that the robber is Jesse, even after seeing his face.

Mabel's belief is, however, $safe_m$, since the method responsible for her belief that the robber is Jesse is *visual perception*, whereas the mistaken belief she could easily have had in its place would have been formed on the basis of *auditory perception*.[10] Accordingly, one can maintain that Mabel's belief is beyond epistemic reproach, consistently with $safety_m$-involving epistemic principles. For ease of exposition, I will for the most part ignore this complication, but we will revisit the methods clause when it makes a difference (and it will).

2.c. The gold standard

There is a strong case to be made that the Safety-Based Debunking Argument, supplemented by the No Safety Argument, is the gold standard for formulating debunking arguments. As already indicated (in Section 1), by grounding the argument in the high contingency as opposed to the inevitability of moral beliefs, the argument draws significant support from cognitive science and better positions the debunker to resist standard realist vindications of moral beliefs. Additionally, the revelation that our moral outlook is at the mercy of highly contingent features of our ancestral environments does evoke the sort of genealogical anxiety that debunking arguments aim to capture. And the epistemic premise, DK1, pins the subversive effect of genealogical revelations on the absence of what is widely viewed as a critical epistemic virtue: safety.

The merits of a safety-based argument can also be seen by contrasting it with two competing approaches: a sensitivity-based argument and an explanatory argument. Here is a representative sensitivity-based debunking argument.[11]

($DK1_{ss}$) ∀S∀p: if S should not believe that S's belief that p is sensitive, then S should not believe that p [where S's belief that p (formed via method M) is sensitive iff_{def} S wouldn't have believed p (via M) had p been false]

($DK2_{ss}$) Informed realists should not believe that their moral beliefs are sensitive

Both premises are widely viewed as unworkable.

Against $DK2_{ss}$, it is quite difficult to secure a verdict of insensitivity, either for moral contingencies or moral necessities. Had Will Smith not acted wrongly at the Oscars would we have believed that he did not act wrongly? Yes—since had he not acted wrongly, it would have been because he stayed seated instead of slapping Chris Rock—meaning that this moral judgment is sensitive.[12] Had causing gratuitous harm not been wrong, would we have believed it is not wrong? Such "counterpossibles" are arguably not evaluable at all or else trivially true.[13]

[10]On a more coarse-grained individuation of methods, Mabel uses the same method, *perception*, whether she relies on what she heard or on what she saw, yielding the (unwanted) opposite result that her belief is $unsafe_m$. This analogue of the "generality problem" for reliabilism complicates any assessment of what is and is not $safe_m$. It is ultimately the $safety_m$-based debunker's responsibility to provide a way of specifying the method in any given case that is both principled and delivers the result that the moral beliefs in the envisaged alternate histories are formed by the same method.

[11]For sensitivity-based approaches, see Ruse (1986, 254), Ruse and Wilson (1986, 186–187), Joyce (2001, 159 and 170), and Davies (2019, 287).

[12]Cf. Sturgeon (1988) and Clarke-Doane (2015, Section 4.4).

[13]Cf. Clarke-Doane (2015, Section 4.4).

Against $DK1_{ss}$, anti-skeptical beliefs like *that you are not a brain in a vat* and *that you do not mistakenly believe that you are human* are demonstrably insensitive, and yet it is highly implausible that this recognition alone should require you to suspend belief. Indeed, it is precisely such problems that have led many to abandon sensitivity-theoretic epistemologies in favor of safety-theoretic approaches.[14]

On an explanatory approach, the idea would be that revelations about a belief's origins defeat (if and when they do) directly by virtue of revealing that there is no appropriate explanatory connection between the beliefs and the associated facts.[15] An explanationist debunking argument might be framed as follows:

($DK1_{ex}$) ∀S∀p: if S should not believe that there is an appropriate explanatory connection between S's belief that p and the fact that p, then S should not believe that p
($DK2_{ex}$) Informed realists should not believe that there is any appropriate explanatory connection between their moral beliefs and the moral facts

But such arguments have well-known problems. For one, the explanationist must characterize "appropriate explanatory connection" in such a way that it is not trivially easy for the realist to supply such a connection. If, for instance, having a "common cause" qualifies as an appropriate connection between a moral fact and associated belief, then $DK2_{ex}$ is vulnerable to standard third-factor responses to debunking arguments.[16] At the same time, if the constraints on appropriate connections are too demanding, $DK1_{ex}$ will generate sweeping skeptical results. For instance, suppose "appropriate" connections are restricted to those in which the fact that p explains the belief that p. In that case, $DK1_{ex}$ leads straight to inductive skepticism, since our beliefs about future sunrises plainly are not explained by facts about future sunrises.[17] Threading this needle is no small feat.

Additionally, the explanationist faces the question of *why* an explanatory constraint like $DK1_{ex}$ would be true. Why, that is, would the revelation that one's beliefs are not appropriately explained by the associated facts undermine those beliefs? A natural answer is that such explanatory revelations undermine beliefs (when they do) on account of calling the safety of those beliefs into question. After all, as Justin Clarke-Doane (2015, 97) has emphasized, if an explanatory revelation in no way threatens the "modal security" of one's beliefs—if it gives one no reason to think that the beliefs could easily have been false—then it is hard to see how those revelations could undermine the beliefs.[18]

3. Easy Transmission

Now for my objections to the No Safety argument. The first targets NS2. As indicated above, NS2 is an instance of a seemingly innocuous "easy transmission" principle:

Easy Transmission
For any possible events p and q, if p could easily have occurred, and p would have resulted in q, then q could easily have occurred

In other words, the ease with which a counterfactual event could have occurred is transmitted to the consequences it would have had.

[14]See, e.g., Sosa (1999) and Pritchard (2005).
[15]Cf. Lutz (2018, 2020), Faraci (2019), Korman and Locke (2020, 2023a, 2023b), Bogardus and Perrin (2022), Bhogal (2023), and Noonan (2025). See Harman (1977, Chapter 1) for an important precursor.
[16]See Enoch (2010) for a representative third-factor response; though see Noonan (forthcoming).
[17]Cf. White (2010, 582–583).
[18]Though see Korman and Locke (2023b, Section 3).

However, there are clear counterexamples to Easy Transmission. Here is one:

> Almost Spilled
> Lowell wakes up and, like every morning, walks bleary-eyed through the kitchen to make some coffee. The night before, his roommate Sandy nearly spilled dish soap all over the kitchen floor. Had she spilled the dish soap, Lowell would have slipped. But there was no spill and Lowell did not slip.

Here is another:

> Almost Mutated
> Stevie, a mediocre guitarist, enters a regional talent show. During embryonic development, a chance mutation nearly occurred that would have resulted in Stevie becoming a guitar virtuoso. Had that happened, she would have won the talent show. But there was no mutation and she loses.

It is not true (or so it seems to me) that Lowell could easily have slipped this morning, nor that Stevie could easily have won the competition. This is so despite the fact that there are things that could easily have happened that would have resulted in his slip and her victory.[19]

What we learn from these cases of transmission failure is reminiscent of what we learned long ago from Robert Nozick (1981, 197–217) about skeptical possibilities. You believe (a) that you are not a handless brain in a vat and (b) that you have hands. Although (a) and (b) are in some ways intimately connected, there is a stark difference in which scenarios are relevant to assessing counterfactuals about them. When assessing what you would have believed had (a) been false, you are to consider scenarios in which you are an unwitting brain in a vat. In such scenarios, you do believe (a). But when it comes to what you would have believed had (b) been false, far-out scenarios in which you are an unwitting brain in a vat are irrelevant. All that is relevant are closer-by scenarios, for instance ones in which you have had some terrible accident and can see for yourself that your arms end at the wrist. In those scenarios, you do not believe (b).

Our intuitions about Almost Spilled and Almost Mutated suggest that something similar is true for easy possibilities. When assessing whether Sandy could easily have spilled the soap, you are to consider scenarios in which things went differently last night, as Sandy is fumbling with the slippery bottle. When assessing whether Lowell could easily have slipped, by contrast, you do not go all the way back to scenarios in which things go differently last night; all that is relevant are nearby scenarios in which things go differently this morning, as he bounds across a spill-free floor.

Likewise for easy possibilities in the distant past. When assessing whether our distant ancestors could easily have had fundamentally different moral beliefs, we have to look at scenarios in which things go differently thousands (or tens of thousands) of years ago. When assessing whether you yourself could easily have had fundamentally different moral beliefs, those further-back scenarios are irrelevant. All that is relevant to assessing this easy possibility claim are scenarios in which things go differently in your lifetime. Even supposing that NS1 is true, then, there is reason to reject NS2.

4. The Non-Identity Problem

My second objection targets NS1 and begins with the observation that precisely which people exist at any given time is highly sensitive to how things went in the distant and even recent past. As Derek Parfit (1984, 351–355) observes, small changes in the timing of conception, for instance delaying pregnancy by a few months, would have resulted in the birth of a different person. By the same

[19]Cf. Sainsbury (1997, 910), Williamson (2000, 124–125), and Bogardus (2014, 297–299; 2016, 648–649) against easy transmission.

token, small but consequential changes in our distant biological and cultural histories would have led to massive differences in the composition of future generations.[20] The envisaged changes countless generations past would have set off a chain reaction, affecting who couples up, and at what times, almost certainly resulting in a population of humans that includes none of the individuals who actually presently exist. If that is right, then NS1 is almost certainly false. Those things that could easily have gone differently—and that are responsible for your moral outlook—would not have resulted in you having different beliefs. You would not have existed at all. Call this 'the non-identity problem'.[21]

I will consider three strategies for addressing the non-identity problem: modifying the definition of safety (Section 5), shifting attention to things that could easily have gone differently within one's own lifetime (Section 6), and abandoning the safety-theoretic approach altogether in favor of some other argument from the high contingency of moral beliefs to a skeptical upshot (Sections 7–8).

5. Refining Safety

As we saw in Section 2.b, there is reason to be dissatisfied with the simplistic definition of safety with which we began. It is therefore worth exploring whether there is a superior definition, which could underwrite an argument from alternate histories that does not fall victim to the non-identity problem.

To begin, recall the method-relativized definition of safety: S's belief that p is $safe_m$ iff_{def} S could not easily have been mistaken about whether p, using the method that S actually used in determining whether p. Shifting to this construal of safety appears to be well-motivated (by the Chance Glance case), but it takes us no way toward addressing the non-identity problem. For, in order for her belief to be $unsafe_m$, it still must be S herself who could easily have made the relevant mistake.[22]

There is a safety-adjacent notion that could potentially do the job, which Nozick (1981, 186–187) calls *adherence*.[23]

Adherence
S's belief that p is adherent iff_{def} it could not easily have been false that S believes that p

One way for it to be false that S believes that p is for S to believe ~p. But another way is for S not to exist at all. And, as we have seen, there are things that could easily have gone differently in our evolutionary histories that would have resulted in us not existing at all. Bracketing the concerns from Section 3 about easy transmission, there is a fairly direct road from there to the conclusion that our moral beliefs are not adherent. From there, an adherence-based debunking premise—

($DK1_a$) ∀S∀p: if S should believe that S's belief that p is not adherent, then S should not believe that p

—could deliver the sought-after skeptical upshot. The obvious problem, however, is that the argument so modified will generalize to all of our beliefs. After all, by exactly similar reasoning, none of our beliefs are adherent. $DK1_a$ leads straight to global skepticism.[24]

[20]Cf. Lenman (2000) on the "massive causal ramifications" of "identity-affecting actions." The problem of cluelessness advanced in Lenman (2000) suggests an additional challenge to NS1: for all we know, the emergence of exclusion-conducive conditions at relevant points in the distant past would have been a temporary setback, and inclusion-conducive conditions and associated inclusivist norms would have reemerged shortly thereafter. I will not pursue this here.

[21]I am not the first to observe how this non-identity problem arises for safety-based debunking arguments; see Hanson (2017, 517–518) and Braddock (2021, 183–184).

[22]Likewise for definitions of safety requiring that the subject not be mistaken about "sufficiently similar" propositions (see footnote 8) or that she not be mistaken in "normal" worlds (Beddor & Pavese, 2020).

[23]This evidently is how Handfield (2016, 68–69) understands safety in his presentation of the argument from alternate histories.

[24]See Sosa (2005, 273–274), White (2010, 597), Setiya (2012, 91–92), and Schechter (2017, 451) for additional objections to adherence-based principles like $DK1_a$.

Let us consider one last modification:

Variable-Subject Safety
S's belief that p is $safe_v$ iff_{def} it could not easily have been the case that S or anyone sufficiently similar to S is mistaken about whether p

This definition is tailor-made to help with the non-identity problem, since it explicitly makes room for one's own belief to count as $unsafe_v$ on account of someone else's mistake. The opening premises of the No Safety Argument and Safety-Based Debunking Argument could then be revised as follows:

($NS1_v$) There are certain things that (i) could easily have happened and (ii) would have resulted in you or someone sufficiently similar to you believing it isn't wrong to enslave POWs
($DK1_v$) For all p, if S should believe that S's belief that p is not $safe_v$, then S should not believe that p

We can reason from $NS1_v$ to your belief being $unsafe_v$ without supposing that you yourself could easily have believed otherwise, and that would be enough—by $DK1_v$—to deliver the skeptical upshot for realists.

Is the shift to $safety_v$ ad hoc? Probably. But perhaps it could draw independent support from examples like the following.[25] Suppose Briana just now was nearly turned into a brain in a vat. Suppose, implausibly but not utterly implausibly, that on account of how close she came to envatment, Briana does not know that she has hands. Suppose, again implausibly but not utterly implausibly, that, despite the psychological continuity, an envatted person is a different person from the embodied person whose brain it once was. Supposing all of that, we would need something like $safety_v$ in order to account for Briana's lack of knowledge in safety-theoretic terms.

Even supposing that the shift to $safety_v$ is well-motivated, defenders of the $safety_v$-based argument will face the following dilemma. In order for $NS1_v$ to come out true, our reading of 'sufficiently similar' needs to be sufficiently relaxed that counterfactual people with very different psychologies, with very different evolutionary histories, and operating with very different moral principles qualify as sufficiently similar to us. At the same time, if it is too relaxed, then $DK1_v$ will be far too demanding. Presumably even the debunker will agree that your belief that it is wrong to enslave POWs is not undermined by the mere existence of one actual sociopath who thinks it is just fine to do so. (Not to mention that your belief that $8 \times 7 = 56$ had better not be undermined by some third-grader mistakenly believing otherwise.) So counterfactual but not actual slavery-indifferent people need to come out as sufficiently similar to us. The burden would be on the debunker to thread that needle.

6. Alternate Personal Histories

A second strategy for circumventing the non-identity problem involves shifting focus from things that could easily have gone differently in the distant past to things that could easily have gone differently within one's own lifetime. Justin Clarke-Doane suggests something along these lines:

> Had we grown up in a different town, read different books, or had different mentors—scenarios that could have "easily" obtained—our moral outlook might have been very different. Whether knowledge of this is undermining by itself, it at least threatens to be qua evidence that we could have easily had false moral beliefs. (2020, 148)

[25]Thanks to Billy Dunaway for helpful discussion.

In motivating NS1 by appeal to alternate personal histories, the debunker neatly sidesteps the non-identity problem, since trivially it is we ourselves who have the different beliefs in scenarios in which our personal histories unfold differently.

Is there, however, anything that could easily have gone differently that would have resulted in our believing that it is not wrong to enslave POWs? There almost certainly is no town in which any of us could easily have been raised that would have resulted in our believing that it is permissible to enslave POWs. An argument from alternate hometowns might pose a challenge to beliefs about hot-button issues (e.g., abortion or immigration), but it will not underwrite the sort of fundamental challenge to our moral outlook that the debunking argument is meant to pose. It is likewise hard to imagine books we could easily have read or mentors we could easily have had that would have resulted in our believing that it is not wrong to enslave POWs.

With one possible exception. Perhaps the case can be made that there are things we (moral realists) could easily have read or mentors we could easily have had that would have resulted in our rejecting realism in favor of error theory, the thesis that all (atomic, first-order) moral propositions are false. We would in that case have believed that it is false that it is wrong to enslave POWs. If that is right, then NS1 is true of us.

I see two problems with this argument from error-theoretic influences. The first is that it almost certainly is not universally or even generally true of moral realists that their histories could easily have been different in ways that would have led them to embrace moral error theory. NS1 *may* be true of you if, for instance, your first philosophy course was nearly taught by an outspoken error theorist who would have made a deep and lasting impression on your philosophical outlook. On the other hand, it may be false because all the professors at your undergraduate institution and at any others you could easily have attended would have given unsympathetic presentations of error theory. Or perhaps you are not especially impressionable and would have ended up a realist despite early and sympathetic exposure to error theory. Likewise for the suggestion that reading different things could easily have made the difference. Many of us are realists despite having read canonical defenses of error theory at formative stages in our philosophical development. What this suggests is that it will be highly variable from one realist to the next whether there is anything in their personal history that could easily have gone differently and that would have resulted in their becoming an error theorist. Unless you have special reason to think that there are such things in your own personal history, you have no reason to accept NS1.

The second problem: As indicated in Section 2.b, definitions of safety typically incorporate a methods restriction, and for good reason. But if what ultimately matters epistemically is that one's actual methods could not easily have resulted in mistaken beliefs, then an argument from alternate histories must show, not merely that there are nearby worlds in which you believe otherwise, but that there are nearby worlds in which your actual methods lead you to believe otherwise. Even supposing, however, that early exposure to moral error theory could easily have led you to believe that it is not wrong to enslave POWs, it would not have been by the same method. In the actual world, the method by which you form the belief that it is wrong is by consulting your moral intuitions (or calculating the expected utility, or running it through the categorical imperative). In the envisaged nearby worlds, the method by which you form the belief that it is not wrong is that of inferring it from general error-theoretic principles. Accordingly, even if you could easily have believed differently and mistakenly, we cannot get the result that the belief is unsafe_m—that is, that the same methods could easily have led you to a mistaken belief.[26]

[26]Perhaps the debunker will insist that it *is* the same method in both cases, namely "philosophizing." Here we face an instance of the generality problem mentioned in note 10. The burden is on the debunker to identify a principled way of specifying the relevant method in any given case that would yield something as general as *philosophizing* in the case at hand.

7. Arguments from Coincidence

7.a. The argument

A third strategy for addressing the non-identity problem is to abandon the safety-theoretic approach altogether and opt for a different way of moving from the high contingency of the factors responsible for our moral outlook to a skeptical upshot for realists. One such approach involves replacing the charge of unsafe beliefs with a charge of coincidence.[27]

Let $\Phi_1 \ldots \Phi_n$ be a substantial list of moral propositions that, according to the debunker, we believe to be true as a result of highly contingent features of our histories. For instance, $\Phi_1 \ldots \Phi_n$ might consist of propositions that it is adaptive to believe only in the context of inclusion-conducive conditions (that it is wrong to enslave POWs, that it is permissible to associate with members of other ethnicities, etc.). Let Moral Accuracy be the thesis that people have largely accurate beliefs about $\Phi_1 \ldots \Phi_n$. And let us say that A is a *high-probability explanation* of B iff$_{\text{def}}$ A makes B highly likely. The Coincidence Argument can then be formulated as follows:

The Coincidence Argument

(CA1) There is no high-probability explanation of our believing any of $\Phi_1 \ldots \Phi_n$
(CA2) If so, then there is no realist-friendly high-probability explanation of Moral Accuracy
(CA3) If there is no realist-friendly high-probability explanation of Moral Accuracy, then realists are committed to thinking that Moral Accuracy is a mere coincidence
(CA4) If realists are committed to thinking that Moral Accuracy is a mere coincidence, then realists should not accept Moral Accuracy
(CA5) So, realists should not accept Moral Accuracy

CA1 is meant to be motivated by the high contingency of moral beliefs. The actual explanation of our belief in $\Phi_1 \ldots \Phi_n$—namely, that these beliefs are the product of cultural evolutionary processes—does not make it especially likely that people would end up with these particular moral beliefs. Cultural evolutionary processes, the idea goes, were just as likely to result in people believing exactly the opposite, had history unfolded differently in ways that it easily could have.

Here is the idea behind CA2. By realist lights, the explanation of the *truth* of $\Phi_1 \ldots \Phi_n$ is going to cite facts about people that are invariant across exclusion- and inclusion-conducive conditions—for instance, facts about our cognitive sophistication that (by realist lights) ground our rights to autonomy and dignified treatment. Such explanations make the truth of $\Phi_1 \ldots \Phi_n$ highly likely (if not flat-out necessary).[28] Accordingly, any explanation that makes it likely that people are correct about $\Phi_1 \ldots \Phi_n$ would, by realist lights, have to be one that makes it likely that people believe $\Phi_1 \ldots \Phi_n$ to be true. Contraposing, we get CA2.

The idea behind CA3 is that when there is some striking correlation with no adequate explanation—no high-probability explanation—that is just what it is for it to be a coincidence. And Moral Accuracy is quite the striking correlation if realism is true, since it posits a correlation between the truth values of mind-independent moral propositions and beliefs about those truth values. (By contrast, there is nothing especially striking about Moral Accuracy if moral

[27]See Braddock (2021, 186–188). Braddock targets only *secular* realists, acknowledging that theists have a distinctive way of resisting the argument. Below, I follow Braddock in framing the coincidence arguments in terms of high-probability explanations.

[28]Though, as indicated in Section 2.b, a realist might reasonably contend that at least some moral truths (e.g., concerning senicide) are made true by inclusion-conducive conditions and would have been otherwise in the harsh conditions in which we would have believed otherwise. I leave it to proponents of the coincidence argument to deal with this complication, perhaps by placing further restrictions on $\Phi_1 \ldots \Phi_n$. Thanks to Chris Noonan for helpful discussion.

constructivism is true, since it would in that case be a correlation between mind-dependent moral truths and the very beliefs on which their truth values depend.[29])

Finally, the idea behind CA4 is that it is irrational for realists to believe that they simply got lucky. Once one appreciates that it would have taken a massive coincidence for one's beliefs to be correct, the idea goes, one should abandon those beliefs.[30]

This argument sidesteps the non-identity problem: it turns on the idea that cultural evolutionary processes could easily have given rise to different moral outlooks, but at no point requires that any one of *us* could easily have had a different moral outlook.

7.b. The problem

CA1 is false: the realist can supply a high-probability explanation of our having the moral beliefs that we do.[31] What is needed is a more filled-in explanation, one that cites not just the bare fact that our beliefs are the product of cultural evolutionary processes but also the specific contingent factors that account for why those processes steered us toward this particular moral outlook—the relative abundance of resources, for instance, and the lack of perceived threat from outsiders in the relevant ancestral environment. It is highly probable that cultural evolution unfolding in those sorts of conditions would give rise to a moral outlook like ours.

Perhaps the debunker will grant the point, but insist that what is needed to combat a charge of coincidence is not merely a high-probability explanation but rather a "doubly" high-probability explanation, where A is a doubly high-probability explanation of B iff (i) A makes B highly likely and (ii) A itself is highly likely. The idea, then, would be to reframe the opening premises of the argument in terms of doubly high-probability explanations:

(CA1*) There is no **doubly** high-probability explanation of our believing any of $\Phi_1 \ldots \Phi_n$
(CA2*) If so, then there is no realist-friendly **doubly** high-probability explanation of Moral Accuracy
(CA3*) If there is no realist-friendly **doubly** high-probability explanation of Moral Accuracy, then realists are committed to thinking that Moral Accuracy is a mere coincidence

The above objection to CA1 poses no problem for CA1*. For that, it would have to be established that the envisaged contingencies, which account for the cultural evolution of morality taking the course that it did, were themselves highly likely to have occurred. But there is reason to believe that just the opposite is true: the sorts of political and economic conditions that pave the way for a highly inclusivist morality like ours are something of an anomaly in human history.[32] Accordingly, the sort of filled-in explanation invoked above in resisting CA1 will not qualify as a doubly high-probability explanation and does not impugn CA1*.

The shift to doubly high-probability explanations merely moves the bump in the rug, however, since requiring a doubly high-probability explanation sets the bar for being non-coincidental implausibly high. Suppose we have a coin that consistently comes up tails when flipped.[33] We investigate and discover that, as a result of a freak accident, the coin is weighted toward tails. This is a completely satisfying explanation of the string of tails and shows it to be non-coincidental. But it is not a doubly high-probability explanation. After all, although the coin's being weighted makes the string of tails highly likely, its being weighted is not itself highly likely. It was a freak accident.

[29]Though see Dunaway (2018).
[30]Though see Dworkin (1996, 125), Setiya (2012, Chapter 2), and Mogensen (2016, 1810–1812).
[31]The discussion that follows draws heavily from Baras's (2022, 113–117) discussion of a related proposal.
[32]See Buchanan and Powell (2016, Section 2B), who describe the requisite conditions as a "luxury good"; cf. Braddock (2021, 194).
[33]The example is from Baras (2022, 115).

A striking correlation with no doubly high-probability explanation need not be a coincidence; there is no good reason to accept CA3*.

8. An Explanationist Approach

We have witnessed the failure of multiple attempts to move from the high contingency of the processes responsible for our moral outlook to a skeptical upshot for realists, grounded in the easy possibility of error. Yet this was meant to be the "gold standard" formulation of the debunking arguments. Perhaps it is time to conclude that the evolutionary moral debunker's project is hopeless.

Perhaps. My own view, however, is that it would be a mistake to conclude that there is nothing to the argument from alternate histories. The revelation that our moral outlook is at the mercy of highly contingent features of our ancestral environments clearly and, I think, rightly evokes genealogical anxiety. If some particular framing of the argument from alternate histories fails to capture its bite, then we have to keep looking.

In particular, I think it is worth revisiting the explanationist understanding of debunking arguments sketched above in Section 2.c:

($DK1_{ex}$) ∀S∀p: if S should not believe that there is an appropriate explanatory connection between S's belief that p and the fact that p, then S should not believe that p

($DK2_{ex}$) Informed realists should not believe that there is any appropriate explanatory connection between their moral beliefs and the moral facts.

The argument for $DK2_{ex}$ would proceed as an argument from elimination, considering and rejecting candidate explanatory connections. The debunker's evolutionary explanations would then play a role in closing off the idea that moral facts figure in the evolutionary explanation of our moral beliefs. Something like the adaptive plasticity thesis serves as what Street (2006, Section 6) would call an "adaptive link account" of our belief that it is wrong to enslave POWs: this and other such inclusivist beliefs inhibit behavior that would jeopardize one's ability to take advantage of opportunities afforded in inclusion-conducive conditions. Any competing "tracking account" of the adaptive value of the beliefs—according to which their adaptive value is explained by moral facts—is bound to be inferior on multiple measures. First, the tracking account will be less parsimonious insofar as it invokes a domain of moral facts. Second, the tracking account is obscure: it is entirely unclear how the putative fact that it really is wrong to enslave POWs is supposed to explain, even in part, why believing it is wrong bestows a survival advantage.

That an explanatory constraint like $DK1_{ex}$ is doing the heavy lifting would help account for a curious fact about moral debunking arguments, namely that evolutionary explanations of our moral beliefs evoke genealogical anxiety irrespective of whether they point toward the inevitability of moral outlooks like ours (because moral beliefs are the result of hard-wired dispositions) or the high contingency of moral outlooks like ours (because moral beliefs are the result of arbitrary features of our ancestral environments). How can explanations pointing in opposite directions be equally worrisome? A natural answer is that both underscore the idea that there is bound to be a fully adequate naturalistic explanation of our moral outlook, one that leaves no room for the supposed influence of moral facts on our beliefs—which, on an explanationist approach, is precisely why the beliefs are in jeopardy.[34]

That being said, as we saw in Section 2.c, the explanationist faces serious challenges of her own, challenges that must be addressed before we can properly assess whether an explanationist

[34]Cf. Noonan (2025, 334–336).

approach is any more viable than the safety-based approach criticized here. Such a task is well beyond the scope of this paper. But count me optimistic.[35]

Daniel Z. Korman is a professor at the University of California, Santa Barbara. His research focuses on metaphysics and epistemology.

References

Baras, D. (2022). *Calling for explanation*. Oxford University Press.
Barkhausen, M. (2016). Reductionist moral realism and the contingency of moral evolution. *Ethics*, 126(3), 662–689. https://doi.org/10.1086/684708
Beddor, B., & Pavese, C. (2020). Modal virtue epistemology. *Philosophy and Phenomenological Research*, 101(1), 61–79. https://doi.org/10.1111/phpr.12562
Bhogal, H. (2023). Explanationism versus modalism in debunking (and theory choice). *Mind*, 132(528), 1005–1027. https://doi.org/10.1093/mind/fzad029
Boehm, C. (2012). *Moral origins*. Basic Books.
Bogardus, T. (2014). Knowledge Under Threat. *Philosophy and Phenomenological Research*, 88(2), 289–313.
Bogardus, T. (2016). Only All Naturalists Should Worry About Only One Evolutionary Debunking Argument. *Ethics*, 126(3), 636–661.
Bogardus, T., & Perrin, W. (2022). Knowledge is believing something because it's true. *Episteme*, 19(2), 178–196. https://doi.org/10.1017/epi.2020.18
Braddock, M. (2016). Evolutionary debunking: Can moral realists explain the reliability of our moral judgments? *Philosophical Psychology*, 29(6), 844–857. https://doi.org/10.1080/09515089.2016.1163667
Braddock, M. (2021). The contingency of the cultural evolution of morality, debunking, and theism vs. naturalism. In J. De Smedt & H. De Cruz (Eds.), *Empirically engaged evolutionary ethics* (pp. 179–201). Springer International Publishing.
Buchanan, A., & Powell, R. (2016). Toward a naturalistic theory of moral progress. *Ethics*, 126(4), 983–1014. https://doi.org/10.1086/686003
Buchanan, A., & Powell, R. (2018). *The evolution of moral progress: A biocultural theory*. OUP.
Clarke-Doane, J. (2015). Justification and explanation in mathematics and morality. *Oxford Studies in Metaethics*, 10, 80–103. https://doi.org/10.1093/acprof:oso/9780198738695.003.0004.
Clarke-Doane, J. (2020). *Morality and mathematics*. Oxford University Press.
Clarke-Doane, J., & Baras, D. (2021). Modal security. *Philosophy and Phenomenological Research*, 102(1), 162–183. https://doi.org/10.1111/phpr.12643
Cofnas, N. (2023). How gene-culture coevolution can—but probably did not—track mind-independent moral truth. *The Philosophical Quarterly*, 73(2), 414–434. https://doi.org/10.1093/pq/pqac047
Davies, P. S. (2019). Darwinizing debunking arguments. *Ratio*, 32(4), 275–289. https://doi.org/10.1111/rati.12231
Dunaway, B. (2018). Epistemological motivations for anti-realism. *Philosophical Studies*, 175(11), 2763–2789. https://doi.org/10.1007/s11098-017-0981-7
Dworkin, R. (1996). Objectivity and truth: You'd better believe it. *Philosophy and Public Affairs*, 25(2), 87–139. https://doi.org/10.1111/j.1088-4963.1996.tb00036.x
Enoch, D. (2010). The epistemological challenge to metanormative realism: How best to understand it, and how to cope with it. *Philosophical Studies*, 148(3), 413–438. https://doi.org/10.1007/s11098-009-9333-6
Faraci, D. (2019). Groundwork for an explanationist account of epistemic coincidence. *Philosophers' Imprint*, 19, 1–26.
Fincher, C., & Thornhill, R. (2012). Parasite-stress promotes in-group assortative sociality: The cases of strong family ties and heightened religiosity. *Behavioral and Brain Sciences*, 35(2), 61–79. https://doi.org/10.1017/S0140525X11000021
Goldman, A. I. (1967). A causal theory of knowing. *The Journal of Philosophy*, 64(12), 357–372. https://doi.org/10.2307/2024268
Handfield, T. (2016). Genealogical explanations of chance and morals. In U.Leibowitz & N.Sinclair (Eds.), *Explanation in ethics and mathematics* (pp. 58–82). Oxford University Press.
Hanson, L. (2017). The real problem with evolutionary debunking arguments. *The Philosophical Quarterly*, 67, 508–533.
Harman, G. (1977). *The nature of morality*. Oxford University Press.
Henrich, J. (2016). *The secret of our success*. Princeton University Press.
Jenkins, C. S. (2006). Knowledge and explanation. *Canadian Journal of Philosophy*, 36(2), 137–164. https://doi.org/10.1353/cjp.2006.0009
Joyce, R. (2001). *The myth of morality*. Cambridge University Press.

[35]Thanks to Bob Beddor, Matthew Braddock, Chad Carmichael, Billy Dunaway, Al Hájek, R.J. Leland, Dave Liebesman, Sebastian Liu, Dustin Locke, Chris Noonan, Justin Plumridge, and audiences at UC Irvine and the University of Manitoba.

Koon, J. (2024). Recalibrating evolutionary debunking. *Philosophy and Phenomenological Research*, 108(2), 456–478. https://doi.org/10.1111/phpr.12975

Korman, D. Z., & Locke, D. (2020). Against minimalist responses to moral debunking arguments. *Oxford Studies in Metaethics*, 15, 309–332. https://doi.org/10.1093/oso/9780198859512.003.0014.

Korman, D. Z., & Locke, D. (2023a). An explanationist account of genealogical defeat. *Philosophy and Phenomenological Research*, 106(1), 176–195. https://doi.org/10.1111/phpr.12848

Korman, D. Z., & Locke, D. (2023b). Modal security and evolutionary debunking. *Midwest Studies in Philosophy*, 47, 135–156. https://doi.org/10.5840/msp202311243

Lenman, J. (2000). Consequentialism and cluelessness. *Philosophy and Public Affairs*, 29(4), 342–370. https://doi.org/10.1111/j.1088-4963.2000.00342.x

Lutz, M. (2018). What makes evolution a defeater? *Erkenntnis*, 83(6), 1105–1126. https://doi.org/10.1007/s10670-017-9931-1

Lutz, M. (2020). The reliability challenge in moral epistemology. *Oxford Studies in Metaethics*, 15, 284–308. https://doi.org/10.1093/oso/9780198859512.003.0013

Manley, D. (2007). Safety, content, apriority, self-knowledge. *The Journal of Philosophy*, 104(8), 403–423. https://doi.org/10.5840/jphil2007104813

McCain, K. (2014). *Evidentialism and epistemic justification*. Routledge.

McKay, N. (2024). Can rational reflection save moral knowledge from debunking? *Episteme*, 21(4), 1418–1433. https://doi.org/10.1017/epi.2023.37

Mogensen, A. L. (2016). Do evolutionary debunking arguments rest on a mistake about evolutionary explanations? *Philosophical Studies*, 173(7), 1799–1817. https://doi.org/10.1007/s11098-015-0579-x

Noonan, C. (2025). Evolutionary debunking arguments, explanationism and counterexamples to modal security. *Erkenntnis*, 90(1), 325–347. https://doi.org/10.1007/s10670-023-00700-8

Noonan, C. (forthcoming). An explanationist account of what goes wrong with third-factor replies to debunking arguments. *Synthese*. https://doi.org/10.1007/s11229-023-04270-8

Nozick, R. (1981). *Philosophical explanations*. Harvard.

Parfit, D. (1984). *Reasons and persons*. Clarendon.

Pritchard, D. (2005). *Epistemic luck*. Oxford University Press.

Richerson, P. J., & Boyd, R. (2005). *Not by genes alone*. Chicago University Press.

Ruse, M. (1986). *Taking Darwin seriously*. Prometheus.

Ruse, M., & Wilson, E. O. (1986). Moral philosophy as applied science. *Philosophy*, 61(236), 173–192. https://doi.org/10.1017/S0031819100021057

Sainsbury, R. M. (1997). Easy possibilities. *Philosophy and Phenomenological Research*, 57(4), 907–919. https://doi.org/10.2307/2953809

Schechter, J. (2017). Explanatory challenges in metaethics. In T.McPherson & D.Plunkett (Eds.), *Routledge handbook of metaethics* (pp. 443–459). Routledge.

Setiya, K. (2012). *Knowing right from wrong*. Oxford University Press.

Sober, E. (1994). Prospects for an evolutionary ethics. In *his* From a biological point of view (pp. 93–113). Cambridge University Press. https://doi.org/10.1017/CBO9780511624940.006

Sosa, E. (1999). How to defeat opposition to Moore. *Philosophical Perspectives*, 13, 141–153.

Sosa, E. (2005). Tracking, competence, and knowledge. In P. K. Moser (Ed.), *The Oxford handbook of epistemology* (pp. 264–286). Oxford University Press.

Sterelny, K. (2012). *The evolved apprentice*. The MIT Press.

Street, S. (2006). A Darwinian dilemma for realist theories of value. *Philosophical Studies*, 127(1), 109–166. https://doi.org/10.1007/s11098-005-1726-6

Sturgeon, N. (1988). Moral explanations. In G.Sayre-McCord (Ed.), *Essays on moral realism* (pp. 229–255). Cornell University Press.

Warren, J. (2017). Epistemology versus non-causal realism. *Synthese*, 194(5), 1643–1662. https://doi.org/10.1007/s11229-015-1010-z

White, R. (2010). You just believe that because… *Philosophical Perspectives*, 24(1), 573–615. https://doi.org/10.1111/j.1520-8583.2010.00204.x

Williamson, T. (2000). *Knowledge and its limits*. OUP.

Cite this article: Korman, D. Z. 2026. Moral Realism and the Argument from Alternate Histories. *Canadian Journal of Philosophy* 56: e47, 1–15, doi:10.1017/can.2026.10127