

WHAT'S THE COINCIDENCE IN DEBUNKING?

Harjit Bhogal
University of Maryland

Abstract

Many moral debunking arguments are driven by the idea that the correlation between our moral beliefs and the moral truths is a big coincidence, given a robustly realist conception of morality.

One influential response is that the correlation is not a coincidence because there is a common explainer of our moral beliefs and the moral truths. For example, the reason that I believe that I should feed my child is because feeding my child helps them to survive, and natural selection instills in me beliefs and dispositions that help my children survive since that is conducive to my genes continuing through the generations. Similarly, the reason that it's morally good to feed my child is because it helps them to survive, and survival is morally valuable.

But if we look at some cases from scientific practice, and from everyday life, we can see, I argue, why this response fails. A correlation can be coincidental even if there is a common explainer. I give an account of the nature of coincidence that draws upon recent literature on scientific explanation and argue that the correlation between moral belief and moral truth is a coincidence, even given such common explainers. And I use this to defend a certain form of debunking argument.

Consider this argument:

(1) On robustly realist conceptions of morality the correlation between our actual moral beliefs and the moral truths is a big coincidence.

(2) There is reason to reject theories that lead to big coincidences.

So,

(3) There is reason to reject either the correlation between our moral beliefs and the truth, or the robustly realist conception of morality.

Thanks to Martin Abreu Zavaleta, Dan Baras, Michelle Dyke, Dan Waxman, Mike Zhao, the Work in Progress group at the University of Maryland and an anonymous reviewer.

This, of course, is a type of moral debunking argument. It's one that has a very simple form. For premise (1) we need a notion of coincidence on which the correlation between our actual moral beliefs and the moral truths is a big coincidence, given robust realism. For premise (2) we need a coincidence-avoidance norm. And that's it.

This way of understanding debunking is in the spirit of Street's (2006) evolutionary debunking argument against moral realism. (And in the spirit of classic arguments against mathematical realism (Field, 1989).) But this type of formulation has dropped in popularity in the recent literature. Instead there has been a move to formulating such arguments in terms of weighty epistemic concepts like reliability, sensitivity, safety and defeat with the conclusions of such arguments being that we are not justified in our moral beliefs. A vast and incredibly complicated literature has developed around such arguments – one that I certainly can't do justice to here.

My aim in this paper, though, is to make progress on behalf of the debunker by going back to formulations of debunking arguments in terms of coincidence and coincidence-avoidance norms.¹ This way of formulating the argument, I suggest, will help us see the weakness of some common responses to debunking arguments.

In particular, there are a couple of major lines of response to debunking arguments that have been most influential in the modern literature. Firstly, there are third-factor responses — those which say that there is some factor which is a common explainer of both our moral beliefs and the moral truths (e.g. Enoch, 2011; Skarsaune, 2011; Wielenberg, 2010) And, secondly, there are 'modalist' responses, which say that the modal characteristics of our moral beliefs and the moral truths make it the case that the correlation between them is unproblematic, even given a robustly realist conception of morality.²

The reason that simple arguments like (1)-(3) can help the debunker make progress with respect to these responses is that arguments which have the form of (1)-(3) don't just arise when we are considering the reliability of our beliefs. Rather, very similar arguments apply widely across science and everyday life — some theory says that a particular matching of events or facts is coincidental, we should tend to reject (big) coincidences, so we should reject the correlation or reject that theory. The matching in question need not be between our beliefs and the truth. There are important arguments in cosmology, for example, which note that phenomena on opposite sides of the universe match in a striking way, and so we have reason to reject certain theories which imply that this matching is a coincidence. By looking at this broader class of coincidences from scientific practice and from everyday life we can understand the force of arguments like (1)-(3) better and understand how this

¹Berry (2020) makes a similar suggestion with respect to arguments against mathematical realism.

²Justin Clarke-Doane has been most influential in developing this response to debunking arguments in a series of papers. See Clarke-Doane and Baras (2019) and Clarke-Doane (2020) for recent versions of the response.

argument can avoid objections to debunking arguments.

Fully defending (1)-(3) and discussing both third-factor and modalist responses is too much for one paper though. So the focus on this paper will be understanding the sense in which the correlation between our moral belief and the moral truths is a coincidence and showing how that can allow the argument (1)-(3) to avoid third-factor responses. A detailed defense of the coincidence-avoidance norm and a rejection of modalist responses will have to wait.³ (I'm also putting aside some other possible responses to debunking arguments, for example, what Korman (2019) calls 'unstable minimalist' responses (e.g. Dworkin, 1996; Setiya, 2012). These responses accept that the correlation between belief and truth is a lucky coincidence, but assert that this luck is not problematic for robust realism. In the context of the argument (1)-(3) that again amounts to a rejection of the coincidence-avoidance norm, and so discussion of it will have to wait.)

More specifically, looking at the broader class of coincidences in scientific practice and everyday life allows us to see that some correlations, or matchings, between events can be coincidental even when there is a common explainer of the events. This, on its own, raises doubt about third-factor responses since they purport to dispel the coincidence of the correlation between our moral beliefs and the moral truths by pointing to a common explainer. But further, I'll give a more detailed analysis of the notion of coincidence which shows us how the common explainers that third-factor responses appeal to are exactly of the type that do not dispel coincidence.

In section 1 I will make some preliminary clarifications about the notion of coincidence. In section 2 I'll discuss the broader notion of coincidence, as it has been discussed in the modern philosophy of science literature, and how coincidences are consistent with there being common-explainers or 'third-factors'. In section 3 I'll argue that this sheds light on what is wrong with third-factor responses — the matching between our moral belief and the moral truths is still a coincidence even when there are common explainers. In section 4 I'll discuss the connection between my strategy and some prior attempts from debunkers to avoid third-factor responses. In section 5 I'll discuss in more detail the relevant notion of coincidence to help validate my claim about what is wrong with third-factor responses. Sometimes when we point to a common explainer of correlated events that does dispel the coincidence. But, I argue, the common explainers appealed to in third-factor responses are not of the type that dispel coincidence. In section 6 I'll briefly compare the discussion to the modalist strand of responses to debunking arguments. In section 7 I'll conclude.

³Though in section 6 I give a preview of how this will go.

1 PRELIMINARIES ABOUT COINCIDENCE

Let's start with some preliminary points about the notion of coincidence that is going to be at work in the paper.

The literature on coincidence suggests that a coincidence has, roughly speaking, two major parts — (i) some striking or compelling way in which the component events 'match' and (ii) some sense in which those component events are 'disconnected' — particularly some causal or explanatory sense. (See for example, Hart and Honoré (1985, p. 74), Lando (2017), Bhogal (2020), and the wide range of views that Lando (section 2) labels as 'traditional views'.)

For example, if I toss a fair coin and it lands heads 40 times in a row, that seems like a coincidence. There is a striking match between the component events — all the tosses landed heads — but a sense in which the tosses were independent or disconnected. Similarly if you end up on a cruise with your old enemy (Owens, 1992, pp. 8-9) that seems like a coincidence. There is a striking matching between component events — the location of both you and your old enemy — but the events appear to be disconnected and unrelated. (Of course, if your old enemy had in fact sought you out and followed you on to the cruise then it would not have been a coincidence.)

If, on the other hand, you tossed the coin 40 times and it landed HHHTHTTTTHTTTTHT... that would not be a coincidence, since there is no striking match between the tosses. Similarly, you ending up on a cruise, not with your old enemy, but with some random person that has no importance to you is not a coincidence.

It's not easy to give a theory of exactly what makes for a striking match between component events. But that won't matter for us since we will be considering cases where there isn't any controversy about whether there is a striking match. More important will be to make sense of what type of 'disconnection' makes an event coincidental.

Here's a second preliminary point. It's very important to distinguish the notion of coincidence being discussed here, the one that is at work in the argument (1)-(3), from the notion of an *epistemic coincidence*.

There's a very common and natural thought in epistemology that some true beliefs don't constitute knowledge because they are in some sense 'lucky'. It's natural to say that in such cases the matching between your belief and the truth is an epistemic coincidence (see, for example, Faraci (2019, pp. 19-20)).

This concept of epistemic coincidence is deeply tied up with the notions of justification and knowledge, and with the project of solving the Gettier problem. And further, epistemic coincidences always involve matching between belief and truth.

The sense of coincidence I'm working with is rather different. It is not distinctively about belief or truth. There are coincidences where the relevant matching is not between belief and truth, but rather between coin tosses or other component events. But more importantly, even when the relevant matching is between belief and truth, there is a difference between the notion of coincidence that I'm using and epistemic coincidence. There can be certain kinds of matching between belief and the truth which are epistemic coincidences but not coincidences in my sense, and vice versa. (Consequently, an epistemic coincidence, as I'm using the term, is not just a coincidence that is about epistemic issues.)

Consider for example, this case from Faraci (2019, p. 5):

Bad Sheep Mary sees a sheep-façade in a field and judges there to be a sheep in the field. Unbeknownst to Mary, there is a sheep hiding behind the sheep-façade, which is there because it was attracted by the façade.

Plausibly, Mary's true belief is no coincidence (let's assume, as I think Faraci does, that there is a systematic reason why the sheep is attracted to the sheep-façade so there typically is a sheep behind the sheep façade). There is a clear explanation of why Mary has this true belief, even if this explanation is not available to Mary. But, Faraci says that Bad Sheep is intuitively a case of epistemic coincidence — it seems like Mary lacks knowledge that there is a sheep in the field. Perhaps one could quibble with the details of this case but it illustrates the space for coincidence and epistemic coincidence to come apart.

The point of making this distinction is to stress that the notion of coincidence at work in arguments like (1)-(3) is not distinctively epistemic or tied in any special way to concepts like belief, truth, justification, or knowledge, but rather applies much more widely.

2 COINCIDENCE AND CONNECTION

With those preliminaries in mind we can look in more detail at the the concept of coincidence and the sense in which component events in a coincidence are 'disconnected'. As we noted, intuitively coincidences consist in a striking match between component events which are 'disconnected'. In the case relevant for the debunking arguments the matching is between our moral beliefs and the moral facts.

In this section though, I'm going to note that coincidences are consistent with a natural sense in which there is a 'connection' between the component events. This will end up shedding light on both third-factor responses and the sense in which coincidences involve disconnection.

Consider this case:

Big Bang Coins I have a fair coin and so do you. We each toss our respective coins in our respective houses once each day. As it happens, whenever your coin lands heads, so does mine, and vice versa. These coin tosses, however, have a common cause, the Big Bang. And, more generally, the initial conditions of the universe.⁴

It's clear that the matching between my coin and yours is still coincidental, even in light of the fact that the Big Bang is a common cause. We could even assume that the laws are deterministic so the laws and the facts about the Big Bang necessitate that there is such a matching. In this case it's clear that the Big Bang is a common explainer of the coins, But this, again, would not change the fact that the matching is coincidental. So there can, in this sense of 'connection', be a connection between component events even when the matching between those events is a coincidence.

Here is another case, from Lando (2017, p.135):

Pianos A boy is playing with a ball in the courtyard of an apartment complex. He throws the ball too high, and it bounces off of the balcony of one apartment, sails through the air, bounces onto the balcony of another apartment, and finally falls to the ground. On each of the two balconies sits a grand piano. As the ball lands on the first balcony it strikes a note on the first piano, and as the ball lands on the second balcony, it strikes a note on the second piano. On each of the two pianos, the note struck is the high A.

The fact that the ball struck the same note on both pianos is a coincidence. But the two component events — the two strikings — are clearly connected. The ball hitting the high A on one Piano and bouncing off it in a certain way is part of the reason it hit the high A on the second piano. The two strikings are on the same causal and explanatory chain.

We can also give a version of **Pianos** where the two strikings have a common explainer, rather than one striking causing the other.

Common Cause Pianos A boy is playing with two balls in the courtyard of an apartment complex. He throws both balls too high and they collide with each other in mid air and fly off in different direction. One of the balls hits the piano on the first balcony and strikes a note. The other ball hits the piano on the second balcony and strikes a note. On each of the two pianos, the note struck is the high A.

Clearly this is a coincidence, but the two strikings are connected. The first ball collided with the second in precisely the way that would lead the second ball to strike the high A, and this precise collision is also what lead the first ball to hit the high A. The collision is a common explainer of the two strikings of the respective high As.

⁴Owens (1992, pp. 8-9) discusses similar cases.

Further, we can adapt the Pianos case so that the relevant matching is between a belief and the truth — though the case I have here is rather inelegant.

Epistemic Pianos The boy regularly throws the ball too high so that it hits the first piano. My partner and I live in adjacent apartment so we hear the piano when it is hit. This happens so regularly that my partner has taken to testing her perfect pitch by telling me what note on the piano the ball hit. In fact, it's become a bit of a running joke that she keeps telling me what note was hit, even though I don't care. So much so, that when I'm out of the apartment she leaves me messages on slips of paper that she puts on my desk telling me what note was hit — 'it was an A', 'it was a B' etc.. This time, though, when my partner wasn't in the apartment, the boy threw the ball, it hit the high A on the first piano and then bounced in through the window of our apartment and knocked one of the previously written messages onto my desk. So, when I get back to the apartment I find a message on my desk saying 'it was an A' and so I come to believe that that ball hit an A on the piano, even though my partner did not leave this message for me.

It's important to be clear about why I'm discussing this case. It's not because it's a (pointlessly complicated) Gettier case. Whether I, in the case, have knowledge or justification or anything like that is not important. The point is that it is a coincidence, in just the same way as **Big Bang Coins** or **Pianos** are coincidences, even though the component events of the coincidence are connected. But, in this case, the coincidence consists in a matching between a belief and the truth. And this is so even though the belief and the truth are importantly connected — the fact that the ball hit an A on the piano is part of what explains why I believe it did.

We can find cases of coincidences where there is a common cause or common explainer in scientific practice too. Cosmology is an excellent source of such cases — there are multiple so-called 'coincidence problems'.

Here's one such case:

CMB There is a type of radiation — cosmic microwave background radiation (the CMB) — that is ubiquitous throughout the universe. Interestingly, it's temperature is (almost) the same in every region of the universe — specifically it's 2.725 Kelvin. Normally uniformity of temperature is what we expect because of the process of thermal equilibration — if we put a hot thing and a cold thing close together they will tend to equilibrate to the same temperature. But, given a particular type of Big Bang theory — non-inflationary Big Bang theory — this can't be what's going on because there are regions of the universe where the CMB is the same temperature but which are so far apart that they could never have interacted with each other (given that physical processes can't propagate faster than the speed of light). Given this theory then, the uniformity of temperature is a coincidence. And this is true even though there is a common explainer of the temperature of the CMB in different

regions of the universe — the particular physical state of the universe just after the Big Bang (Guth and Steinhardt, 1989; Earman and Mosterin, 1999; Maudlin, 2009).

The scientific details of this case don't really matter. What's important to recognize, though, is that this case is structurally identical to **Big Bang Coins**, and similar to **Common Cause Pianos**. There is a matching between distant regions, and seemingly those regions do not communicate with each other. Consequently, the matching is a coincidence. If we look back on the causal chain we can find a common explainer of the features of these regions. This common explainer does not, however, undermine the existence of the coincidence. Further, the fact that the uniformity of the CMB would be a coincidence given non-inflationary Big Bang theory was taken to be a powerful reason to reject that theory.

Cases of the type we have considered in this section appear in important parts of scientific practice, then — they are not merely toy cases. (And this example isn't an isolated case – there are other similar cases across physics, and science more generally.)

It's clear, then, that there can be coincidences even when the component parts of the coincidences are connected, in the sense of having a common cause or common explainer. Still, though, it's natural to claim that in the cases we have discussed in this section there is some other sense in which the component events are disconnected, or merely very minimally connected. Perhaps there is no connection, or just a very minimal connection, between the two notes being hit in **Common Cause Pianos**, for example. Understanding what this other sense of 'connection' could be will be important to understand the sense of coincidence at work in (1)-(3).

3 THIRD-FACTOR RESPONSES

The previous section has shown us that there can be coincidences even when there is a 'connection' between the component events — in the sense of there being a common cause or common explainer. Seeing this sheds a lot of light on third-factor responses to debunking arguments.

Third-factor responses say that although there is a correlation between our moral beliefs and the moral truths the correlation is not suspicious. The reason for this isn't that our beliefs are explained by the truth, neither is it that the truth is explained by our beliefs. Rather, there is a third factor that explains both. For example, facts about what encourages survival explain both our moral beliefs and the corresponding moral truths (Enoch, 2011). For example, the reason that I believe that I should feed my child is because feeding my child helps them to survive, and natural selection instills in me beliefs and dispositions that help my children survive since that is conducive to my genes continuing through the generations. Similarly, the reason that it's morally good to feed my child is

because it helps them to survive, and survival is morally valuable.⁵ And, the thought goes, if there is a common explainer of our moral beliefs and of the moral facts then the correlation between them is not so suspicious or surprising.

Of course, such a response only applies to a certain subset of our moral beliefs — those moral beliefs that encourage survival of our offspring. But there are suggestions for other third-factors which can apply to a wider range of our moral beliefs. (For example, Copp (2008), Skarsaune (2011), Wielenberg (2010), Brosnan (2011), Lott (2018).)

A common response from debunkers is that there is something epistemically faulty about third-factor responses because they require certain normative premises — for example, the argument we just mentioned involved the premise that part of what explains why it is morally good to feed my children is that it helps them survive. It's not clear whether third-factor responders are entitled to such normative premises since part of what many debunking arguments do is cast doubt on the justification we have for such premises. Street, for example, has developed this response in a variety of places (2008, pp. 214-217) (2011, section 6), (2016).

On many approaches to debunking this question of whether the third-factor responder is entitled to their premises is a very complex question — one which gets at deep issues in epistemology. (See, for example, Fraser (2014) and Dyke (2019), for example, who agree with Street that third-factor responses are dialectically inappropriate. Vavova (2014), Copp (2019), Berker (2014), Lott (2018) and Morton (2019), among others, argue that at least some third-factor responses can be dialectally appropriate.)

But on my way of understanding the issue, where the debunker is making an argument like (1)-(3), things are much simpler. In (1)-(3) the debunker is arguing that the correlation between your moral beliefs and the moral truths is a coincidence. And, in general, if you want to argue that some correlation is not a coincidence then it's perfectly reasonable to assume that the correlation holds in order to make that case.

So, it's legitimate to respond to debunking arguments like (1)-(3) by assuming that the correlation between your moral beliefs and the moral truths holds, and then attempting to explain the correlation. That is, it's perfectly reasonable for the third-factor responder to initially assume that their moral beliefs are true.

However, I'm going to argue that there is a different problem with third-factor responses to (1)-(3).

⁵This example is from Korman and Locke (2020).

3.1 THIRD-FACTORS AND COINCIDENCE

So, what is wrong with third-factor responses to arguments like (1)-(3)? Perhaps the discussion of section 2 has given it away: Even though third-factor responses identify a connection, in the sense of there being a common explainer, between our moral beliefs and the moral truths, the correlation is still a coincidence.

The cases in section 2 showed us that sometimes matchings or correlations between events can be a coincidence, even when there is a common explainer. In **Big Bang Coins**, for example, if you attempted to explain the fact that my coin tosses always matched yours by appealing to the Big Bang as a common explainer that wouldn't help at all. I would still be very suspicious about the coincidental nature of the matching. I'd start to suspect that maybe you were lying to me, and your coin tosses did not always match mine. Or I'd suspect that there was some trickery involved and the coins were not, in fact, fair.

Similarly, simply identifying a common explainer of our moral beliefs and the moral truths does not mean that the coincidence is dispelled and we should no longer be suspicious about the correlation. In fact, I think, it's reasonable to continue to find the correlation between our moral beliefs and the moral truth coincidental and suspicious, even when a third-factor responder tells me that there is a common explainer.

However, a third-factor responder might naturally reply to the cases we discussed in section 2 by agreeing that there is a class of common explainers that do not dispel coincidence but asking what reason we have to think that the particular common explainers that they appeal are part of that class. After all, there are many obvious cases where the existence of a common explainer does dispel coincidence. For example, in Bhogal (2020) I discuss the coincidence of sitting next to the same person on my flight to and from Denver. But imagine I was sitting next to my partner on the flights, and we booked the trip to Denver together. In that case there would be a common explainer for me sitting next to that person on the flight there, and on the flight back, and the appearance of coincidence would be dispelled.

Similarly, it's easy to give such cases where the relevant matching is between belief and truth. For example:

Concert My friend tells me that she intends to go to the concert tomorrow evening and then, tomorrow evening, I form the belief that she is at the concert. She is, in fact, at the concert so there is a matching between my belief and the truth. In such a case my belief and the truth have a common explainer — her intention to go to the concert. This common explainer dispels any appearance of coincidence.

The reader, I take it, can easily come up with other examples. So, we need to distinguish the cases where having a common explainer does dispel coincidence and cases where they do not. And we need reason to think that the common explainers appealed to in third-factor explanations are of the type that do not dispel coincidence. I will take on those tasks in section 5.

4 EXPLANATIONISTS AND THIRD-FACTORS

Before that, though, it will be useful to note how some other debunkers have reacted differently to the existence of third-factor responses so we can see how my approach connects to existing approaches.

One possible reaction that debunkers might have to third-factor responses is that they fail since the responder isn't entitled to the premises that they need. We talked about this response above.

Another reaction, that we mentioned in the introduction, is to move away from formulations of debunking arguments that rest upon any ideas of coincidence or a 'lack of connection' between moral truth and moral belief. Bedke (2014), for example, points to certain modal features of our moral beliefs that he claims act as defeaters for those beliefs and he formulates his debunking argument in this way because of the way third-factor responses establish an explanatory connection between moral belief and moral truth.

Most important for our purposes, though, are views that have stuck with understanding debunking arguments in explanation-theoretic terms.

The 'explanationist' views that have appeared in the literature typically start from the intuition that for a belief of yours to be justified there must be some explanation of the matching between the belief and the truth. This is rather different from (1)-(3), since (1)-(3) involves no premises about justification or knowledge or anything like that. Rather it just involves a claim of coincidence and a coincidence-avoidance norm.

But in another respect these explanationist views are importantly similar to (1)-(3), since their success depends upon ruling out third-factor responses. Third-factor responses appear to establish that there is an explanation of the matching between moral beliefs and moral truths. This is a problem for (1)-(3) because it threatens to establish that the matching is no coincidence. It's also a problem for the explanationists I'm considering because it shows that the premise that for a belief of yours to be justified there must be some explanation of the matching between the belief and the truth is not strong enough to debunk our moral beliefs. The point here is not to survey all explanationist approaches, but rather, to look at how some explanationists deal with this problem in order to help motivate and clarify my approach.

4.1 KORMAN AND LOCKE

Korman and Locke (2020) react to third-factor responses by, in effect, saying that common cause or third-factor explanations of the matching between belief and truth are *never* relevant for epistemic justification. Rather, they appeal to a principle in the spirit of what they call (EC):

(EC) If S believes that her belief that p neither explains nor is explained by the fact that p, then S is thereby rationally committed to withholding belief that p. (p. 324)

Notice that this principle says that the explanations that are relevant to whether S should believe p are explanations which say that the belief is explained by the truth, or those that say that the truth is explained by the belief. If the only explanation available to S of the matching between belief and truth is one which appeals to some common cause or third-factor then (EC) says that S should not believe p.

Building a debunking argument using this premise avoids third-factor responses — the fact that our moral beliefs and the moral truth have a common explainer is not enough to make those beliefs justified. However, Korman and Locke are very clear that (EC) is false. They reject (EC) because it causes problems with justifying inductive inferences.

Here is the counterexample they give⁶:

Smoke ‘You observe the fire in the fireplace and are justified in believing that there is smoke coming out the chimney. Of course, the fact that smoke is coming out of the chimney does not explain (causally or otherwise) the belief that it is. But this realization surely does not undermine the belief.’ (p. 324)

The problem with (EC) goes beyond these cases of induction. Consider again **Concert** – a testimonial case. My friend tells me that she intends to go to the concert tomorrow evening and then, tomorrow evening, I form the belief that she is at the concert. My belief does not explain the truth, and the truth does not explain the belief, rather, there is a common explainer. (EC) would wrongly say that I should not believe that my friend is at the concert. This problem generalizes to many beliefs we have about the future.

After rejecting (EC) Korman and Locke defend a variant principle:

(EC*) If p is about domain D, and S believes that her belief that p is neither explained by nor explains some D-facts, then S is thereby rationally committed to withholding belief that p.

This principle is similar to (EC). In particular, it denies that common explainers are relevant to what we should believe. However, while (EC) identifies the relevant explanations as those where your

⁶Originally from Goldman (1967, pp. 365-6) as part of his discussion of the causal theory of knowledge.

belief that P explains the truth that P, or vice versa, (EC*) says that the relevant explanations are where your belief that P explains truths *in the domain of P*, or vice versa.

This move to domains helps with the smoke counterexample to (EC). ‘Your belief that smoke is coming from the chimney is explained by facts about smoke and chimneys: it is the result of an inductive inference from past observations of smoke coming from chimneys with active fireplaces. So (EC*) doesn’t prescribe withholding belief that smoke is now coming out of the chimney.’ (p.325)

But the move also leads to problems. As Korman and Locke note, it’s sometimes rather unclear what domains a proposition is part of. And answering this question matters — some plausible ways of connecting propositions with domains lead to counterexamples to (EC*) [p.325]. But, perhaps more importantly, this appeal to domains seems rather strange. Why are such domains epistemically relevant? Why does the justification of my belief in a proposition depend upon the explanatory connections between my belief and *other* propositions? Doesn’t this move push us away from the intuition that when beliefs are justified there has to be an explanation of the matching between belief and the truth? And if it does, how do we motivate (EC*)?

Of course, these questions don’t constitute arguments against Korman and Locke. And filling them out into arguments would take much more space than we have available. But the impression is that the move to (EC*) is going down the wrong path. Korman and Locke want to rule out third-factor responses to debunking arguments and they, correctly, I think, recognize that sometimes our beliefs look unjustified even when there is a common explainer of the belief and the truth. (**Epistemic Pianos** is a case of exactly this kind.) But they go too far in saying that common explainers of the belief and the truth are never relevant to epistemic justification. This puts them in a difficult situation with respect to cases like **Smoke** and **Concert** where the explanationist would like to say that we have justification precisely because of a common explainer of our belief and the truth.

The move to domains is an ingenious attempt to recover the intuitive judgment in such cases. But I think the mistake was made earlier. The cases from section 2 suggest that Korman and Locke should, instead of ruling out all common explainers, recognize that some common explainers are good and some are bad. Some, like those in **Big Bang** or **Epistemic Pianos** don’t dispel the appearance of coincidence. What we, and, I think, what Korman and Locke, need is a story about what features make a common cause explainer good and what make them bad. I will discuss this more in section

5.^{7,8}

4.2 LUTZ

Lutz's (2018) reaction to third-factor responses is more radical. He starts by considering the intuitive thought that for a belief to be justified there must be an explanatory connection between the belief and the truth. But he reacts to third-factor responses by moving away from this intuitive thought in favor of a very different type of approach to debunking — though one that is still based on explanatory considerations.

In particular the principle he accepts is:

Explaining Away Defeats (EAD): New evidence, D, defeats the support that E provides for S's belief that P if: D is evidence in favor of a complete explanation, A, of E, such that S may not infer P from A and S's independent background information. (p. 1110)⁹

⁷In a newer paper (Korman and Locke, *fort*) Korman and Locke reject their prior appeal to domains. They give a new account of when beliefs can be defeated but it's not clear, at least to me, whether it rules out third-factor responses. Their final proposal is that 'If S is not entitled to believe that the facts she treats as reasons to believe that p support* her belief that p, then S's belief that p is defeated'. But imagine I treat the fact that feeding my child leads to his survival as a reason to believe that I should feed him. Am I not entitled to believe that that fact supports my belief that I should feed him? That's not clear. Or, at least, it's not clear prior to determining whether third-factor responses work or not. Though perhaps their intent is to rule out third-factor responses in a different way, for example, by claiming that it's not dialectically appropriate for the responder to appeal to the premises that they need — in line with the discussion at the start of 3. Or perhaps they don't intend to rule out third-factor responses at all — they don't discuss them explicitly in Korman and Locke (*fort*).

⁸There are some explanationist approaches to epistemology more generally which share some similarity in structure to Korman and Locke's approach. A particularly interesting one is that of Jenkins (2008) since it's designed to apply to an *a priori* domain — that of arithmetic. (Thanks to a reviewer for suggesting this.)

Her view is that A knows P just in case P is a good explanation of A's belief that P for someone not acquainted with the particular details of A's situation. This is, in one sense, similar to explanationists like Korman and Locke, because it requires there be an explanatory connection between belief and truth for your belief to be in epistemically good standing.

But there are some very notable differences too. Firstly, Jenkins is using a conception of 'explanation' where an explanation is a communicative act. And so reasoning about knowledge, on her view, involves reasoning about what is apt communication, what is misleading given certain background knowledge, and so on. This view then, should be categorized very differently from explanationists who think of explanation as a mind-independent relation between facts.

More importantly for our purposes, though, is that it's hard to see what this account would say about third-factor responses, or about debunking more generally. (And Jenkins doesn't use her account to investigate these issues.) For example, consider a particular third factor — the fact that feeding my child leads to survival is part of the explanation of why I believe I should feed my child and why I actually should feed my child. In light of this, would it be reasonable to say to an outsider that the reason I believe I should feed my child is because I actually should feed my child? Or would that be too misleading? That's the question that determines whether my moral beliefs are in good standing in this case. But I don't think Jenkins gives us enough of a theory of what counts as a good explanation in the relevant sense for us have grip on the answer to this question. So I think Jenkins' approach is ultimately not particularly helpful for evaluating third-factor explanations.

⁹My focus here is on Lutz (2018). Lutz (2020) is a little different since he identifies a version of the intuition that for a belief to be justified there must be an explanatory connection between the belief and the truth and says that EAD is consistent with this version of the intuition. Regardless, the objection I am about to raise to EAD applies to both papers.

The thought is that we have evidence in favor of Darwinian explanations of our moral beliefs and this evidence defeats our justification in our moral beliefs, because we can't infer the truth of any moral proposition from those Darwinian explanations.

EAD however, faces problems from cases where there are good explanations at multiple 'levels'. Imagine that I see someone getting a candy bar from a vending machine. They put in a five dollar bill and get the candy bar and three dollars back. I come to believe that the candy bar cost two dollars. So my observation is evidence, E, for my belief that P, the candy cost two dollars. Now I gain some new evidence D. Specifically, a Laplacian demon comes along and tells me the complete microphysical state of the person/vending machine system just before the transaction, and the deterministic microphysical laws. This new evidence is evidence in favor of an explanation, A, of E — my observation was necessitated by the prior microphysical state of the system and the deterministic laws. Presumably this explanation, A, counts as 'complete' if any explanation does. But, from this explanation, A, and my independent background information, I cannot infer P. That is, I cannot use this microphysical explanation to infer that the candy bar cost two dollars, given the difficulty of bridging the gap between this microphysical description of the world, and the economic description.^{10,11}

Consequently, we have new evidence, D, in favor of a complete explanation, A, of old evidence, E. And I may not infer proposition P from A and my independent background information. EAD tells us, then, that this evidence in favor of the microphysical Laplacian explanation should defeat the support I have for believing the candy bar cost two dollars. But this is clearly the wrong result. I should continue to believe that the candy bar cost two dollars.

There's a general point here that doesn't depend upon the details of this case. EAD is based on the idea that one explanation of my evidence crowds out another — when I get a new, complete, explanation of my evidence then that undermines other explanations of the evidence that I cannot infer from the new explanation. But EAD faces problems in cases where two different explanations of the same phenomenon can be compatible, rather than competitors. Explanations at different 'levels' are paradigm cases of this.

The appropriate reaction to third-factor responses is not to move to EAD. As we've been arguing, it's to say that third-factor responses don't dispel coincidence.

¹⁰This case is inspired by Potochnik (2010, pp. 62-3).

¹¹We can say something stronger here, though this is not required for the counterexample to EAD. Notice that the microphysical explanation, A, doesn't even necessitate the relevant economic facts. This is because properties like 'being a five dollar bill' depend on much more than the local microphysical facts about the person-vending machine system. Part of what makes the collection of atoms that was fed into the vending machine constitute a five dollar bill is facts about the broader monetary system in the country — for example, past decisions made by the Federal Reserve. So we can't infer merely from the local microphysical facts about the person-vending machine system that there is five dollar bill present (Potochnik, 2010, p. 63).

4.3 FARACI

Faraci's (2019) explanationist approach is closest to the approach that I am developing. He discusses the notion of epistemic coincidence, and although he doesn't explicitly develop a debunking argument based on epistemic coincidence he expresses optimism about such an argument (section 8.3). And his notion of epistemic coincidence is based on the intuition that non-coincidences have an explanatory connection between belief and truth. This is in the spirit of my approach.

But, Faraci's approach to coincidence faces problems — particularly if we want it to be embedded in a debunking argument that avoids third-factor responses. Building upon Lange's (2010) notion of a mathematical coincidence Faraci provides a sufficient condition on coincidence: It is a coincidence that the members of Γ are all true if (a) there is no unified explanation for the members of Γ and (b) no explanatory chain links all members of Γ .

What it is to be a 'unified' explanation is as follows: E is a unified explanation for the members of Γ if and only if no isolable part of E explains some members of Γ at least as well as E but fails to explain other members of Γ at least as well as E.

Notice though, that there are coincidences where (b) fails to hold. In **Pianos** there are explanatory links connecting each of the two strikings of the high A but it is still a coincidence. Further, there are coincidences where condition (a) fails to hold — that is where there is a unified explanation of the component events. We discussed such cases at length in section 2. In **Common Cause Pianos**, for example, there is a unified explanation of the two balls hitting the high-A on the two pianos — it's due to the precise way in which those balls collided in mid-air. Similarly, in **Big Bang Coins** the unified explanation of my fair coin and your fair coin always landing the same way was the initial conditions of the universe and the deterministic laws. And, of course, this kind of unified explanation could be given for any matching between events.

To be clear, Faraci gives a *sufficient* condition on coincidence, so these cases are not counterexamples. But seeing that there are many coincidences with unified explanations suggests that the notion of coincidence has little to do with lack of unified explanations. Consequently, the notion of coincidence is left rather mysterious.¹²

Further, it's not clear how we could successfully embed this notion of coincidence into a debunking

¹²Similarly, see, for example, Linnebo (2006, section 2) and Berry (2020, section 5b) making (in the context of the debate about mathematical realism) the related point that separately explaining, or merely 'stapling together' explanations of, the component events is not satisfactory. I agree with this — merely stapling together unrelated explanations is unsatisfactory — and this point can help the debunker avoid certain responses to their arguments. But notice that cases like **Pianos** and **Big Bang Coins** suggest that we should go further, since, as we just discussed, it seems, at least on the face of it, like there are unified explanations in those cases — ones which don't just staple together separate explanations but rather identify a common explanatory factor. (Thanks to a reviewer for discussion here.)

argument since it's not clear how third-factor responses would be ruled out. Third-factor responses seem to provide a unified explanation of the matching between my moral beliefs and the moral truth.

In footnote 24 Faraci suggests that his response, developed in a manuscript with Aaron Elliott, is to argue that third-factor responses do not, in fact, provide a unified explanation of the matching between my moral beliefs and the moral truth. It's not clear how such a case would be made — much more would have to be said about the notion of a unified explanation (and in particular the notion of an 'isolable part' of an explanation). But, more importantly, as we noted, cases like **Common Cause Pianos** suggest that coincidences have little to do with the absence of a unified explanation.

There is a lot in common between our approaches — importantly, though, I appeal to a broader notion of coincidence that is unrelated to solving the Gettier problem, and I am doubtful of his account of coincidence so the nature of my argument against third-factor responses will be rather different.

5 WHAT ARE COINCIDENCES?

I've been arguing that the problem with third-factor responses to (1)-(3) is that the matching between moral truth and moral belief is still a coincidence, even in light of the common explainer. The key to finishing this argument is to distinguish cases where having a common explainer dispels coincidence from cases where they do not. Why does the type of explanation given by third-factor responders not dispel coincidence?

As we noted in section 1 there are two parts to the notion of coincidence — firstly, there being some striking match between component events, and secondly there being some sort of explanatory disconnection between those events. In the cases we are considering there is clearly a striking match. So claiming that matching is a coincidence involves claiming that there is explanatory disconnection.

But in what sense is there explanatory disconnection between the component events in cases like **Common Cause Pianos** — cases of coincidences with common explainers? Similarly, the third-factor responder says that there is a common explainer of our moral beliefs and the moral truth — in what sense is there explanatory disconnection in this case?

To answer this, I'm going to look at some recent work on the notion of coincidence from the literature on scientific explanation and adapt those ideas to our context. In particular, Lando (2017) and my own work (Bhagal, 2020) have both addressed the question of how coincidences can have common explainers.

An important point, which both Lando and I stress, is that in a putative coincidence there are multiple propositions that we can use to characterize the coincidence. The way I put it in Bhagal (2020,

section 1) is that in **Pianos** there is a *matching proposition* — the fact that the two balls hit the same note on both pianos. This matching proposition describes the property that the component events share. Further, we can describe the *particular proposition* in virtue of which the matching holds — it's that the first ball hit the high A on the first piano and the second ball hit the high A on the second piano. Similarly, Lando says that the two balls hitting the same note on both pianos is a 'relational fact' while the particular proposition is not (p. 144). (Note that there isn't one unique particular proposition — we can describe the situation in virtue of which the matching holds in more or less detail (Bhogal, 2020, footnote 1). But that doesn't really matter in what's to come — the distinction between particular propositions and the matching proposition is what's important here. So I will continue to talk, loosely, as if there is one particular proposition.)

Notice that the particular proposition has to be more specific than the matching proposition. The fact that the first ball hit the high A on the first piano and the second ball hit the high A on the second piano is more specific than the fact that the same note was hit on both pianos.

We can also identify the particular and matching propositions in non-coincidences too, as long as there is a striking match between component events. Consider, for example, about the much-discussed case of Arbuthnot's regularity. Eighteenth century doctor John Arbuthnot found that, in each of the last 82 years in London, more males had been born than females. The particular proposition, in this case 'describes the numbers of male and female births each year — there was 10,341 male births last year, and 10,113 female; there was 10,238 males births two years ago and 9,987 female and so on' (Bhogal, 2020, p. 679). Ultimately this isn't a coincidence — more on this later — but still we can identify the particular proposition and the matching proposition.

Again, the particular proposition is more specific than the matching proposition. The matching proposition is that in each of the last 82 years in London, more males had been born than females. The particular proposition isn't merely that last year in London there were more males born than females, the year before that there were more males more than females, and so on back 82 years. That would be a mere restatement of the matching proposition. The particular proposition is more specific, describing the way in which the matching came about.

Importantly, the explanations we want of the particular and matching propositions are very different. Lando (2017) and Bhogal (2020) both stress that properly explaining why the first ball hit the high A on the first piano and the second ball hit the high A on the second piano is a very different project from properly explaining why both balls hit the same note.

The same is true when we consider the matching between our belief and truth. In this putative coincidence the matching proposition is that my moral beliefs match the moral truths. The particular proposition is one which lists all my moral beliefs and all the moral facts. For example, 'I believe that

it's good to feed your child and I believe that killing is wrong...and it's morally good to feed your child and killing is morally wrong...'. It is in virtue of the particular proposition that the matching between moral belief and moral truth holds.

In this case, too, explaining the matching proposition and explaining the particular proposition are different projects. Explaining the particular moral facts which hold, and the particular moral beliefs I have, is a different project from explaining the matching between my beliefs and the truth.

Why is explaining the particular proposition a different project to explaining the matching proposition? It's because explaining the particular proposition involves explaining the component events. Explaining the particular proposition in **Pianos** and **Common Cause Pianos** — that the first ball landed on the high A on the first piano and the second ball landed on the high A on the second piano — involves going into detail about how each of the balls were caused to strike the pianos in the exact way that they did. Similarly, to explain the particular proposition in the case of our moral beliefs is to explain why I hold each particular moral fact that I in fact hold, and to explain why each of the actual moral facts hold.

But explaining the matching proposition is very different — it not about explaining why the component events are as they are, but rather it's about explaining why the component events match. Trying to explain why first ball landed on the high A on the first piano and the second ball landed on the high A on the second piano is different from trying to explain why they hit the same note.

In light of this one might claim, as Lando does (p. 144), that even if we give a fully adequate explanation of the particular proposition that does not constitute an explanation of the matching proposition (even though the matching proposition holds in virtue of the particular proposition).

This suggestion could, perhaps, be made plausible, but this is *not* what I am arguing. Rather, my argument is based on the idea that, in cases that are intuitively coincidences, the best way for us to explain the matching proposition is simply to repeat an explanation of the particular proposition.

Consider **Pianos**. The way we could explain the particular proposition is by giving the *Trajectory Explanation*:

We would note where the pianos were placed and then note how the ball was thrown — how hard and at what angle, with what spin and so on. And we would show how this leads it to hit the high A on the first piano, and how that leads to the ball bouncing to hit the high A on the second piano (Bhagal, 2020, p. 628).

And if we then try to explain the matching proposition — that the ball hit the same note on both Pianos — the best we can do is just to repeat this explanation. There is no additional story we can appeal to. The same is true of our variant case **Common Cause Pianos**.

And the same is also true of **CMB**. In **CMB** the particular proposition is that in all the different regions of the universe the temperature of the microwave background radiation is 2.725 Kelvin. The matching proposition is that the temperature in all the regions is the same. Given non-inflationary Big Bang theory our best explanation of this matching proposition just involves repeating the explanation of the particular proposition — that the precise physical state of the universe just after the Big Bang leads to each region having this temperature.

This is very different from cases which are intuitively non-coincidental. Imagine, for example, a variant of **Big Bang Coins** were it wasn't a coincidence that my coin tosses and your coin matched. How would this be? Well, perhaps you and I are both extremely skilled at coin tossing, and we are testing our skills by both trying to make our coin tosses match a pre-agreed sequence. We both succeed in this and so our coin tosses match. If this was the situation then we have an explanation of the matching proposition that doesn't just consist in explaining the particular way in which my coin tosses land and the particular way in which your coin tosses land. We would be able to explain the matching proposition by noting that we are very skilled at coin tossing and are both trying to replicate the same sequence.

Or consider again Arbuthnot's regularity — the fact about more males being born than females in each of the last 82 years in London. Explaining the particular proposition — that there was 10,341 male births last year, and 10,113 female; there was 10,238 males births two years ago and 9,987 female and so on — involves going into detail about the factors that lead to each particular birth in each year being a male or a female one. But to explain the matching proposition — that is, to explain why more males were born than females in each of the past 82 years we can do more than repeat this explanation of the matching proposition — we can give a completely different explanation.

In particular, a classic argument in evolutionary biology establishes that there is evolutionary pressure towards a one to one sex ratio at the time of sexual maturity (this is known as Fisher's principle). But, in humans, infant mortality occurs at a higher-rate among males than females. So, there is evolutionary pressure for more males to be born than females so that the ratio becomes equal at the time of sexual maturity.

In (Bhagal, 2020, p. 684) I argue, following Kitcher (2001),¹³ that it is precisely the existence of this evolutionary argument of the matching proposition and the way that it is different from the explanation of the particular proposition that makes Arbuthnot's regularity *non-coincidental*.

This is the key point then: In coincidences our best explanation of the matching proposition involves just repeating an explanation of the particular proposition while this is not so with non-coincidences.¹⁴

¹³And in the spirit of claims made by Roberts (2011), Frisch (2011) and Weslake (2014).

¹⁴In (Bhagal, 2020) I develop this idea by giving an account of explanatory goodness – an account, that is, of when

And we can take this a step further. Explaining the particular proposition is about explaining the precise way in which the component events occurred — such explanations do not exhibit a substantial explanatory relationship between the parts (as Lando p.145 puts it, they only exhibit a ‘thin’ explanatory relationship). Consequently, if the best explanation we have of the matching proposition is just to repeat the particular proposition then we haven’t exhibited a substantial explanatory connection between the component events.

In fact, I claim, when the best way of explaining the matching proposition is just to repeat an explanation of the particular proposition then this is *just what it is* for there to be no substantial explanatory connection between the component events. So, there can be no substantial connection between the component events in a coincidence, even when there is a common cause or common explainer.

Explanations that dispel coincidence are of a different kind. In those cases the best way for us to explain the matching is distinct from the best way of explaining the particular way the matching occurred. And this is so whether or not there is a common explainer.

But now we can see why third-factor explanations are of the type that do not dispel coincidence — why we should still be suspicious of the correlation between moral belief and moral truth even in light of such explanations. With respect to the matching between our moral beliefs and the moral truths to explain the particular proposition is to explain the particular moral beliefs we have and the particular moral truths. To explain the matching proposition is to explain why our moral beliefs match the moral truths. On robustly realist conceptions of morality, though, our best explanation of the matching proposition just consists in repeating an explanation of the particular proposition. In fact, this is precisely what third-factor explanations do — they explain the matching between moral beliefs and moral truths, but they do so merely by being an explanation of our particular beliefs and the particular moral truths.

Consider, again, the explanation of your moral belief that you should feed your child — feeding your child leads to survival and so natural selection has inculcated in you the belief that you should do that — and the explanation of the moral fact that you should feed your child — feeding your child leads to survival and survival is morally good. This is (part of) an explanation of the particular proposition — it explains your particular moral belief and the actual moral truths. The third-factor response involves appealing to this same explanation as (part) of an explanation of the matching proposition. In general this is what third-factor responses do — they explain the matching by explaining the particular moral truths that hold and the particular moral beliefs we have.

explanations are better or worse – that implies that, given this account of coincidence, in coincidences the matching proposition is explained less well than the particular proposition. My focus in that paper was on contingent domains, and not necessary domains, like that of morality. That account of explanatory goodness works much less smoothly with respect to domains like morality and mathematics, which are the focus of the debunking arguments addressed in this paper.

But in doing so, such responses fail to exhibit a substantial explanatory connection between the moral beliefs and the moral truths. Consequently, the matching between belief and truth is still a coincidence, even though a third-factor has been identified.

In cases like **Big Bang Coins** and **Common Cause Pianos** the existence of a common factor in the explanation of the component events doesn't stop the matching between those events being a coincidence. Similarly, third-factor responses can point to a common explainer of moral beliefs and moral truths, but not in a way that exhibits a substantial explanatory connection between them.

The thought here is somewhat related to an intuition that Street (2016, p.322) expresses. She says that 'One may explain each side of the coincidence in as much depth as one likes—going into wonderful normative depth about why family and friendship are valuable, and wonderful scientific depth about why we were selected to think this. But all this goes nowhere toward explaining the thing that really needs to be explained, namely the coincidence itself.' This intuition seems compelling, but the problem is that the third-factor responder has claimed that they have explained the coincidence — after all, they have identified a common explainer. They haven't merely stapled together unrelated explanations. But in this section we have seen how even when the third-factor responder can point to a common explainer between our moral beliefs and the moral truth, the particular way in which they appeal to a common explainer leaves the matching between moral belief and moral truth still coincidental.

6 MODALIST RESPONSES

The focus in this paper has been on third-factor responses to debunking arguments. But, as I noted in the introduction, this isn't a full defense of debunking arguments like (1)-(3) – there are other responses that I haven't argued against in this paper.

In particular, some of the most influential recent work responding to debunking arguments has been Justin Clarke-Doane's *modalist* approach. So it would be useful to say a little about the relation between this approach and the third-factor views that I focus on in this paper, even though full details will have to be left for another paper.¹⁵

The reason I call Clarke-Doane's work 'modalist' is because I think his approach is best interpreted as being importantly different from the approaches that focus on *explanation* that we have been considering in this paper.

The key point of third-factor responses is that there is an explanatory connection between moral belief and moral truth – a common explainer – and that's what makes the correlation unproblematic.

¹⁵Thanks to a reviewer for discussion here.

Clarke-Doane's strategy, on the other hand, is to argue that we don't need an explanation of the correlation. As long as we can argue that our beliefs have certain appropriate modal characteristics with respect to truth – in particular, *safety* and *sensitivity* – then our beliefs are immune from undermining (see, e.g. Clarke-Doane and Baras (2019) and Clarke-Doane (2020)).

(The relation between this view and explanatory considerations is complicated slightly by the fact that Clarke-Doane sometimes says that there can be some senses of 'explain' in which we can 'explain the reliability' of our beliefs merely by showing them to be safe and sensitive (Clarke-Doane, 2016). But this involves Clarke-Doane explaining how our beliefs can have certain modal features which plausibly constitute their reliability – in particular, safety and sensitivity – *on the assumption that the beliefs are true*. This is a very interesting point, but it doesn't constitute an explanation of the correlation between belief and truth.)

Again, Clarke-Doane's overriding point is that these judgments about explanation don't matter as long as we have safety and sensitivity. If we hold fixed the account of coincidence that we have given in this paper then this strategy is, in effect, a rejection of premise (2) in our original argument.

So, in order to respond we would need a completely different line of argument to the one developed here. We would need to either argue that our moral beliefs are not safe and sensitive or argue that the safety and sensitivity of our moral beliefs are not enough to render their correlation with moral truth unproblematic.

In other work¹⁶ I make an argument of this second kind. The idea there is that we can look to scientific practice to see other cases where there is a puzzling correlation, like the correlation between moral belief and moral truth. Studying the nature of theory choice in such cases suggests that it is the explanatory features of the correlation that matter more than the modal ones. In particular, scientific cases suggest that we have reason to reject theories with substantial unexplained correlations, even if those correlations have the modal features that the modalist values. But that full story is for another time.

7 CONCLUSION

The debunker can, I've argued, make progress by understanding debunking arguments not as a distinctive issue about the reliability of our moral beliefs, or when we they are justified, or anything like that, but as an instance of a much broader issue about what to do when we are faced with a striking correlation between facts. The debunker can make progress, that is, if they shift the debate away from terrain that is normally thought to be part of the territory of epistemology, to terrain that

¹⁶See my manuscript 'Explanationism vs Modalism in Debunking (and Theory Choice).

is mostly explored under the banner of philosophy of science. In particular, the debunking argument I gave is driven by considerations of theory choice and how that interacts with explanation.

Looking at other cases where there are striking correlations allowed us to see that correlations, like the correlation between moral belief and moral truth, can be coincidental even when there is a common explainer. We have reason to be suspicious about the correlation in cases like **Big Bang Coins** and **Common Cause Pianos**, for example, even when a common explainer is cited. This raised doubts about the efficacy of third-factor responses – perhaps the correlation is still a coincidence in light of these third-factors.

And these doubts were validated by seeing the sense of coincidence at work in arguments like (1)-(3). It's a concept of coincidence that makes sense of cases from everyday life and from scientific practice. And it captures the intuition that coincidences are about a lack of explanatory connection, even in cases of coincidence with common explainers. Given this sense of coincidence we can see that the type of common explainers that third-factor responses appeal to are not enough to make the matching between moral belief and moral truth non-coincidental.

References

- Bedke, M. (2014). No coincidence? *Oxford studies in metaethics*.
- Berker, S. (2014). Does evolutionary psychology show that normativity is Mind-Dependent?
- Berry, S. E. (2020, August). Coincidence avoidance and formulating the access problem. *Canadian journal of philosophy* 50(6), 687–701.
- Bhagal, H. (2020). Coincidences and the grain of explanation. *Philosophy and phenomenological research* 100(3), 677–694.
- Brosnan, K. (2011, January). Do the evolutionary origins of our moral beliefs undermine moral knowledge? *Biology & philosophy* 26(1), 51–64.
- Clarke-Doane, J. (2016). What is the benacerraf problem? In *Logic, Epistemology, and the Unity of Science*, pp. 17–43. Cham: Springer International Publishing.
- Clarke-Doane, J. (2020, March). *Morality and Mathematics*. Oxford University Press.
- Clarke-Doane, J. and D. Baras (2019, September). Modal security. *Philosophy and phenomenological research* 19, 1.

-
- Copp, D. (2008). Darwinian skepticism about moral realism. *Philosophical Issues. A Supplement to Nous* 18, 186–206.
- Copp, D. (2019, December). How to avoid begging the question against evolutionary debunking arguments. *Ratio* 32(4), 231–245.
- Dworkin, R. (1996, April). Objectivity and truth: You'd better believe it. *Philosophy & public affairs* 25(2), 87–139.
- Dyke, M. M. (2019). Bad bootstrapping: the problem with third-factor replies to the darwinian dilemma for moral realism. *Philosophical studies*, 1–14.
- Earman, J. and J. Mosterin (1999, January). A critical look at inflationary cosmology. *Philosophy of science* 66(1), 1–49.
- Enoch, D. (2011, July). *Taking Morality Seriously: A Defense of Robust Realism*. OUP Oxford.
- Faraci, D. (2019). Groundwork for an explanationist account of epistemic coincidence. *Philosophers Imprint*.
- Field, H. (1989, January). *Realism, Mathematics & Modality*. Realism, Mathematics & Modality. Oxford: Basil Blackwell.
- Fraser, B. J. (2014, March). Evolutionary debunking arguments and the reliability of moral cognition. *Philosophical studies* 168(2), 457–473.
- Frisch, M. (2011, December). From arbutnot to boltzmann: The past hypothesis, the best system, and the special sciences. *Philosophy of science* 78(5), 1001–1011.
- Goldman, A. I. (1967). A causal theory of knowing. *The journal of philosophy* 64(12), 357–372.
- Guth, A. and P. Steinhardt (1989, January). The inflationary universe. In P. Davies (Ed.), *The New Physics*.
- Hart, H. L. A. and T. Honoré (1985, May). *Causation in the Law*. OUP Oxford.
- Jenkins, C. S. (2008, August). *Grounding Concepts: An Empirical Basis for Arithmetical Knowledge*. OUP Oxford.
- Kitcher, P. (2001, January). *Science, Truth, and Democracy*. Oxford University Press.
- Korman, D. Z. (2019, December). Debunking arguments. *Philosophy Compass* 14(12), e12638.

-
- Korman, D. Z. and D. Locke (2020). Against minimalist responses to moral debunking arguments. *Oxford Studies in Metaethics* 15.
- Korman, D. Z. and D. Locke (fort.). An explanationist account of genealogical defeat. *Philosophy and phenomenological research*.
- Lando, T. (2017, March). Coincidence and common cause. *Noûs* 51(1), 132–151.
- Lange, M. (2010, July). What are mathematical coincidences (and why does it matter)? *Mind; a quarterly review of psychology and philosophy* 119(474), 307–340.
- Linnebo, Ø. (2006, June). Epistemological challenges to mathematical platonism. *Philosophical studies* 129(3), 545–574.
- Lott, M. (2018, January). Must realists be skeptics? an aristotelian reply to a darwinian dilemma. *Philosophical studies* 175(1), 71–96.
- Lutz, M. (2018). What makes evolution a defeater? *Erkenntnis. An International Journal of Analytic Philosophy* 83(6), 1105–1126.
- Lutz, M. (2020). The reliability challenge in moral epistemology. *Oxford Studies in Metaethics* 15.
- Maudlin, T. (2009). *The Metaphysics Within Physics*. The Metaphysics Within Physics. New York: Oxford University Press.
- Morton, J. (2019, April). When do replies to the evolutionary debunking argument against moral realism beg the question? *Australasian journal of philosophy* 97(2), 265–280.
- Owens, D. (1992, January). *Causes and Coincidences*. Cambridge University Press.
- Potochnik, A. (2010, January). Levels of explanation reconceived. *Philosophy of science* 77(1), 59–72.
- Roberts, J. T. (2011). Extra-Physical structure in a physical world? or, is the study of life provincial? *The Monist* 94(2), 221–243.
- Setiya, K. (2012, November). *Knowing Right From Wrong*. OUP Oxford.
- Skarsaune, K. O. (2011). Darwin and moral realism: survival of the fittest. *Philosophical studies* 152(2), 229–243.
- Street, S. (2006). A darwinian dilemma for realist theories of value. *Philosophical studies* 127(1), 109–166.

- Street, S. (2008). Reply to copp: Naturalism, normativity, and the varieties of realism worth worrying about. *Philosophical Issues. A Supplement to Nous* 18, 207–228.
- Street, S. (2011). Mind-independence without the mystery: Why quasi-realists can't have it both ways. *Oxford studies in metaethics* 6(1), 1–32.
- Street, S. (2016). Objectivity and truth: You'd better rethink it. *Oxford studies in metaethics* 11(1), 293–334.
- Vavova, K. (2014). Debunking evolutionary debunking. *Oxford studies in metaethics* 9(7), 6–101.
- Weslake, B. (2014). Statistical mechanical imperialism. In A. Wilson (Ed.), *Chance and Temporal Asymmetry*. Oxford University Press.
- Wielenberg, E. J. (2010, April). On the evolutionary debunking of morality. *Ethics* 120(3), 441–464.