

From Moral Realism to Axiarchism: A Metaphysical Response to the Debunking Challenge

Brian Cutter

Abstract

Moral realism faces a well known genealogical debunking challenge. I argue that the moral realist's best response may involve abandoning metaphysical naturalism in favor of some form of axiarchism—the view, very roughly, that the natural world is “ordered to the good.” Axiarchism comes in both theistic and non-theistic forms, but all forms agree that the natural world exists and has certain basic features because it is *good* for it to exist and have those features. I argue that theistic and non-theistic forms of axiarchism are better positioned than metaphysical naturalism to avoid two commitments that a moral realist should seek to avoid: that the correctness of our moral beliefs is a major coincidence, and that there is a complete explanation of our moral beliefs that does not mention any moral truths.

1. Introduction

Moral realism, the view that there are moral truths that hold independently of anyone's attitudes, faces a well known genealogical debunking challenge. Evolutionary forces have played a major role in shaping our moral beliefs. But these forces seem to operate independently of any moral truths there may be. It's a safe bet, for example, that the evolutionary explanation for why we are predisposed to believe that parents have a moral duty to provide for their children will not mention any moral truths. The explanation may invoke the hypothesis that the offspring of parents who were not so disposed were less likely to survive long enough to reproduce. But presumably it will not invoke the hypothesis that parents *in fact* have a moral duty to provide for

their children. The same goes for evolutionary explanations of other moral beliefs toward which we plausibly have some innate predisposition, such as the belief that free-riders are blameworthy, that one ought to reciprocate benefits, or that incest is wrong. Similar points may extend to the cultural and historical forces that shape our moral beliefs. For example, it may be possible to give a historical and cultural explanation for why belief in universal human rights gained widespread acceptance without relying on the assumption that this belief is true.

These observations put the moral realist in an uncomfortable position. One concern is that the moral realist is committed to a massive coincidence. Since our moral beliefs are produced by forces that operate independently of the moral truths, it would be a major coincidence if our moral beliefs turned out to be largely correct, given moral realism. We can avoid a problematic coincidence between our attitudes and the moral truths either by denying that there are moral truths or by maintaining that all moral truths depend on our attitudes. But neither strategy is compatible with moral realism. Another concern is that the genealogy of our moral beliefs provides an undercutting defeater for those beliefs. Even if, by a lucky fluke, our moral beliefs turned out to be largely correct, once we realize that our moral beliefs can be fully explained without mentioning any moral truths, our justification for holding those beliefs is undermined.¹

I have felt pulled in two directions by genealogical debunking challenges to moral realism. On the one hand, there seems to be a genuine difficulty for moral realism here. Tricky issues arise when we get into the details (for example, it can be difficult to frame the objection without relying on principles that are self-defeating, or which would lead to far-reaching skepticism beyond the moral domain). But I find it hard to believe that the genealogical

¹ For moral debunking arguments along these lines, see Street (2006), Joyce (2006), and Harman (1977: ch. 1).

debunking objection is somehow fundamentally confused. It's clear that there is a real problem here, even if it's unclear how best to formulate it.

On the other hand, moral realism is *immensely plausible*—so plausible that (for me, anyway) rejecting it isn't really a live option. It is about as obvious as anything that torturing cats for fun is wrong. It is only slightly less obvious that the wrongness of cat torture has some kind of objectivity or attitude-independence. It is obvious, for example, that torturing cats for fun would be wrong even if the perpetrator approved of the activity. It would be wrong even if his whole society approved of it. It would be wrong even if everyone in his society would continue to approve of it under conditions of full information, or if they would desire to desire cat torture under idealized reflection that brings their values into reflective equilibrium, structural coherence, and so forth. And not only would members of this society act wrongly in torturing cats, but they would be in error when they say things like “cat torture is not wrong.”²

Given the immense plausibility of moral realism, it is tempting to respond to the evolutionary debunking challenges with a “Moorean shift.”³ According to this response, it is more reasonable to conclude that a major coincidence has occurred than to deny the truth or objectivity of claims like “torturing cats for fun is wrong.” But ideally we can do better. There is a glaring internal tension in any combination of views that includes moral realism, commonsense first-order moral beliefs, and the claim that these beliefs could only be correct by a major coincidence (or the claim that the complete explanation for these beliefs makes no mention of the

² This last claim is added to rule out “rigidifying” strategies, whereby “wrong” functions to rigidly pick out whatever property tends to evoke (say) disapproval in the speaker (or her community, or whatever) (cf. Street 2006: 138-9). On this rigidified view, since we disapprove of cat-torture, we can truly assert a counterfactual like: “if our society were to approve of cat torture, cat torture would still be wrong.” (Compare: If XYZ had filled the lakes and oceans, XYZ still wouldn't be water.) However, a proponent of the rigidified view *can't* accept the following: if our society had approved of cat torture, we would be speaking falsely when we say “cat torture is not wrong.” But the latter counterfactual seems just as plausible as the former. Cf. Horgan and Timmons (1991).

³ Cf. Crummett and Swenson (2020: 44).

moral truths). If one finds moral realism immensely plausible, one should be prepared to make significant adjustments elsewhere in one's web of beliefs to resolve such glaring tensions. This paper explores the idea that the best way to do so might be to abandon metaphysical naturalism in favor of some non-naturalistic metaphysical outlook, such as theism or other views on which natural world is somehow "ordered to the good." I will argue that these non-naturalistic views are better positioned than naturalism to avoid both the claim that the correctness of our moral beliefs is a major coincidence and the claim that the moral truths play no role in explaining our moral beliefs. For those who are unwilling to abandon moral realism, this provides some reason to prefer some such non-naturalistic view over metaphysical naturalism.

Related arguments from moral knowledge or moral reliability to theism or some other non-naturalistic metaphysical view have been developed by Swinburne (2004: ch. 9), Ritchie (2012), Crummett and Swenson (2020), Cutter (2019), Nagel (2012), Linville (2009), Bogardus (2016), Jeffrey (2019: ch. 3), and McKay (2023). Others have suggested an inverted "modus tollens" form of the argument. For example, Bedke (2009: 190) objects to one form of moral realism on the grounds that our reliability in moral matters would require "something like a god rigging the ethical facts and the causal order to ensure their serendipitous coincidence."⁴ The current paper adds to these discussions in a few ways. One is by considering a wider range of models for how theism could help with moral debunking challenges, with more development of theistic voluntarist models, as well as more attention to non-theistic forms of axiarchism. Another is by more clearly distinguishing different aspects of the evolutionary debunking challenge (e.g., the challenge of avoiding a coincidence vs. that of securing an explanatory connection between moral truths and moral beliefs). A third is by more thoroughly exploring

⁴ See also Setiya (2012: 114-5).

strategies for showing that metaphysical naturalism and moral realism lead to a problematic coincidence.

The arguments below could be cast in the form of an inference to the best explanation.⁵ We might suppose that our “data” include the truth of moral realism, the fact that we are more-or-less reliable about the moral domain, and the fact that we know many moral truths. From here, we might argue that these data are best explained by some non-naturalistic metaphysical hypothesis. But it seems artificial to draw a data/theory distinction in a way that puts claims like moral realism and the existence of moral knowledge/reliability on the “data” side, with metaphysical claims like theism or naturalism on the “theory” side. I find it more natural to think of the current project as an application of wide reflective equilibrium—an attempt to bring some of our most obstinate moral beliefs, whether mundane or theoretical, into coherence with our metaphysical beliefs. Of course, the outcome of the process of reflective equilibrium is sensitive to starting points. The arguments below are not going to move one whose commitment to metaphysical naturalism is as firmly entrenched as her belief that cat torture is wrong. But they may move someone who sees metaphysical naturalism as no more than a respectable conjecture, while seeing moral realism as closer to non-negotiable. I suspect that many philosophers are in this position.

Some working definitions: By “theism” I mean the view that the natural world is the intentional creation of a transcendent being of perfect goodness, power, and knowledge. By “metaphysical naturalism” (hereafter, “naturalism”) I shall understand the conjunction of two claims. First, there is nothing (concrete) beyond the natural world. Second, the natural world lacks anything like “fundamental purposiveness;” it is not teleologically ordered toward the

⁵ Ritchie (2012) and Crummett and Swenson (2020) often frame their theistic arguments from moral knowledge in this way.

realization of value, whether by an external designer, as on theism and other design hypotheses, or by its own intrinsic striving toward the good, as on Nagel's (2012) view that the universe is teleologically ordered toward the realization of valuable outcomes by fundamental teleological laws. (Note that *metaphysical* naturalism is entirely compatible with *metaethical* non-naturalism, the view, roughly, that moral facts/properties aren't identical to any facts/properties that figure in scientific descriptions of the world.⁶)

I will be specifically concerned with alternatives to naturalism that reject (at least) its second conjunct. The non-naturalistic views of most interest to the current project fall under the broad banner of "axiarchism." Roughly speaking, axiarchism is the view that the natural world is "ordered to the good." The axiarchist holds that, in certain basic and general respects, the natural world is the way it is because it is *good* for it to be that way, or the natural world is that way *in order to realize some good*. For example, an axiarchist might hold that there are conscious organisms, rational subjects, elegant laws, or agents with reliable moral faculties because it is good for there to be conscious life, rational agents, or elegant laws, or agents with reliable moral faculties. And she might say that certain other facts about the natural world obtain at least partly *in order* that such goods would be realized. She might say, for example, that such-and-such physical parameters fall within a narrow life-permitting range in order that there would be the good of conscious (or rational) life, or that evolution took this course rather than that *in order* that there would be intelligent agents with reliable moral faculties.⁷

⁶ For defenses of the latter, see Enoch (2011) and Huemer (2005).

⁷ The term "axiarchism" comes from John Leslie, who defends a version of the view (e.g., in his 1989). However, it's worth flagging that, unlike Leslie, I use "axiarchism" to designate a thesis specifically about the *natural world*---that it exists, and has certain basic features, because it is good for it to exist and have those features---not a thesis about concrete reality in general. So, as I use the term, axiarchism does not entail that there is a value-based explanation for why there is something (concrete) rather than nothing. For example, a theist could be an axiarchist in my sense without holding that God's existence has an axiarchic explanation (i.e., without holding that God exists because it is good for him to exist).

One view is that these axiarchic or purposive explanations are rooted in the purposes or intentions of some intelligent being. This would be a natural line for theists to take. Another view is that these axiarchic explanations hold true without the backing of designer intentions. This is the position of many non-theistic forms of axiarchism. Parfit (1998) and Leslie (1989) suggest that it could just be a bedrock explanatory truth that things are such-and-such way because it is good for them to be that way, without any further rider like “and that’s why God chose to make things that way.” Nagel (2012) similarly suggests a non-theistic view on which the universe is governed by fundamental teleological laws by which the universe “strives” toward certain valuable outcomes.⁸

One *could* accept fundamental teleology without axiarchism. It might be coherent to maintain that the natural world exhibits fundamental teleology, but is not teleologically ordered toward anything of *value*.⁹ However, non-axiarchic teleology does not seem well suited to account for strikingly *valuable* or “suspiciously fortunate” outcomes like moral reliability (or fine-tuning, or irreducible consciousness, or any of the other data that axiarchic hypotheses are commonly invoked to explain). The generic hypothesis that the universe is teleologically ordered toward *some outcome or other* would not seem to predict strikingly valuable outcomes. The same goes for non-axiarchic design hypotheses. The generic design hypothesis that the universe was designed by an agent *for some purpose or other* does little to account for a strikingly valuable outcome. To the extent that theism is supported by such outcomes, it is not just because theism invokes a (sufficiently powerful) designer, but because it invokes a perfectly *good* designer, a designer who acts for the sake of the good.

⁸ See Goff (2023) for an exploration of this and several other non-theistic alternatives to naturalism.

⁹ This possibility is suggested by Hawthorne and Nolan’s (2006) account of fundamental teleology.

I focus on axiarchic views, theistic and otherwise, for two reasons. First, there are considerations of intrinsic plausibility. The hypothesis that there is a designer who acts for the sake of *the good* seems to enjoy greater intrinsic plausibility than, say, the hypothesis that there is a designer who just aims at the realization of spaghetti. Likewise for the hypothesis that the universe is intrinsically teleologically ordered toward the good, or Parfit's suggestion that "the universe is the way it is because it is good for it to be that way" is a bedrock explanatory principle. These seem more intrinsically plausible than the hypothesis that the universe is intrinsically ordered toward the realization of spaghetti, or that "the universe is the way it is because this way involves spaghetti" is a bedrock explanatory principle. Whence the difference in intrinsic plausibility? A partial answer, I think, is a thought that goes back at least to Plato: goodness just seems more primitive, more fundamental, simpler, less arbitrary, than something like spaghetti, in a way that makes goodness more apt to belong to a foundational metaphysical theory.

Second, there are considerations of explanatory power. Axiarchic views have the potential to give a *unified explanation* of a range of striking and valuable facts about the natural world—facts that seem *prima facie* to call for explanation, and to which the naturalist either gives no explanation or piecemeal explanations. In addition to the facts about moral reliability/knowledge that are the subject of this paper, these striking facts include: cosmological fine-tuning for life (Leslie 1989, Swinburne 2004: ch. 8), the existence of irreducible conscious states (Chalmers 1996), the existence of "psychophysical harmony" (Cutter and Crummett forthcoming, Pautz 2020, Goff 2023) and "nomological harmony" (Cutter and Saad forthcoming), the elegance, simplicity, and intelligibility of the laws of nature (Parfit 1998, Plantinga 2011: ch. 9), and the existence of other kinds of knowledge, like mathematical or

metaphysical knowledge, that seem puzzling on naturalism for reasons closely related to the moral case (Field 1980, Rea 2002, Korman 2015, Pittard 2021). A metaphysical outlook that has at least a modest degree of intrinsic plausibility, plus the potential to subsume a wide range of striking facts under a unified explanatory principle, deserves a seat at the table in discussions of foundational metaphysics.

2. Two faces of the genealogical debunking challenge

We can distinguish two related worries about moral realism in standard presentations of the genealogical debunking argument. The first is a worry about coincidence, which might be put as follows:

Coincidental Correspondence: Given the extent to which evolutionary (and perhaps other) forces have shaped our moral beliefs, and given moral realism, it would be a major coincidence—a happy accident, a lucky fluke, a suspiciously convenient stroke of luck—if our moral beliefs turned out to be largely correct.

The second worry has to do with explanatory independence. It rests on something like the following claim:

Explanatory Disconnect: There is a complete explanation of our moral beliefs that does not mention any moral truths.

The two challenges are related, but distinct. There may be strategies for avoiding Explanatory Disconnect that don't avoid Coincidental Correspondence. Suppose that, if a creature endowed with reason has affective responses that align with the moral facts, then these affective responses facilitate a direct “intellectual perception” of the relevant moral facts. For example, suppose that when a rational creature has an affective tendency to feel indignation in response to free-riding, this affective tendency enables the creature to intellectually perceive the

moral fact that free-riding is (pro tanto) wrong. Suppose further that, when one believes that *p* on the basis of an intellectual perception of the fact that *p*, the fact that *p* thereby figures in the explanation of one's belief (perhaps because it is a constituent of the perception that caused the belief).¹⁰ But even if this proposal can avoid Explanatory Disconnect, it does not address the coincidence worry. Given evolution's role in shaping our moral affective sensibilities, one might still worry that it would be a major coincidence that our natural affective sensibilities broadly align with the moral facts, and thus a major coincidence that our beliefs so align. By analogy: if you correctly guess the combination to the vault lock, your beliefs about the secret documents would be explained by the secret documents themselves. But it would be an incredible coincidence for you to correctly guess the combination.

Conversely, there are strategies for avoiding Coincidental Correspondence that would not avoid Explanatory Disconnect. One might argue (i) that the correspondence between our moral beliefs and the moral facts is an objectionable coincidence only if evolution could easily have produced creatures whose moral beliefs fail to correspond to the moral truths (to the extent that ours do), and (ii) there is no such "easy possibility" of non-correspondence. In support of the second claim, consider the belief that parents have a duty to provide for their children, a belief that (the realist will probably suppose) corresponds to the moral fact that parents *do* (typically) have a duty to provide for their children. There is arguably a robust tendency in the evolutionary process to promote some such belief about parental duties among creatures intelligent enough to have moral beliefs. In that case, it looks like evolution couldn't easily have produced intelligent creatures that get things wrong, at least when it comes to whether parents have a duty to provide

¹⁰ Cf. Bogardus and Perrin (2022) and Climenhaga (forthcoming).

for their children. But even if this strategy defuses the coincidence challenge (later I will argue that it doesn't), it wouldn't touch the explanatory challenge.

My view is that a fully satisfactory response to the genealogical debunking challenge will find a way to resist both Coincidental Correspondence and Explanatory Disconnect. The desirability of avoiding Coincidental Correspondence doesn't need much defense. In general, we should try to avoid views that imply major coincidences. But the desirability of avoiding Explanatory Disconnect is less obvious. One view is that Explanatory Disconnect is only a threat to moral realism to the extent that it suggests that a major coincidence has occurred, or that our moral beliefs could easily have been mistaken. I think this is incorrect. I am convinced by the arguments of Korman and Locke (2020) that a lack of explanatory connection can directly undermine one's beliefs, even where the lack of explanatory connection doesn't suggest that the beliefs could easily have been false, or that they fail to satisfy some other modal condition relevant to knowledge. They support this conclusion with several cases. Here is a case from Dustin Crummett that is similar to, and inspired by, a case from Korman and Locke:

Goblins: All goblins have a very strong intuition that there exists a necessarily existent deity whom they call "the Judge"—a deity who judges everyone after death, rewarding the good and punishing the bad, although he is otherwise not active in the world and did not create the world. The goblins believe in the Judge on the basis of these intuitions. The goblins are aware that the Judge is not responsible for their Judge intuitions (because he is not active in the world). In fact, they will tell you, "the reason why we have Judge intuitions is that goblins are naturally mischievous, and cannot be made to cooperate without belief in post-mortem judgment. This is why evolution selected for these intuitions among our ancestors." (Crummett and Antony 2022)

By the goblins' own lights, their belief in the Judge couldn't easily have been false, nor could evolutionary forces easily have given them false beliefs about post-mortem judgment. Nevertheless, it seems that it would be irrational for the goblins to maintain their belief in the Judge while accepting these claims about its genealogy. Moreover, even if their Judge beliefs turn out to be correct, it seems that those beliefs would not amount to knowledge (at least if the goblins are correct about the genealogy of their Judge beliefs). The goblins' epistemic position seems to be relevantly similar to that of a moral realist who believes that moral truths play no role in the explanation of her moral beliefs.¹¹

On the basis of cases like these, Korman and Locke conclude that the moral realist "must embrace some account on which the moral facts explain our moral beliefs" (310). I agree, and will hereafter assume that any fully satisfactory response to the evolutionary debunking challenge must find a way of avoiding both Coincidental Correspondence and Explanatory Disconnect. I will argue that theism, and perhaps other axiarchic views (with some qualifications) are better positioned than naturalism to avoid Coincidental Correspondence and Explanatory Disconnect. For simplicity, I will focus primarily on theism. But I will note the points in the argument where the differences between theism and other forms of axiarchism matter, as well as where they don't.

3. Theism avoids coincidence and secures an explanatory connection

On theism, we are not left with the sense that the reliability of our moral faculties would be a lucky coincidence. God can arrange things so that we end up with more-or-less reliable moral

¹¹ These verdicts do not presuppose the simple principle that, if (one justifiably believes that) one's belief that p is not explained by the fact that p, then one's belief is thereby undermined, or one fails to know that p. This simple principle has plausible counterexamples, such as my belief that the sun will rise tomorrow (White 2010: 582-3). See Korman and Locke (2023) for an attempt to formulate a general explanatory-defeat principle that avoids these counterexamples while delivering intuitively correct verdicts about Goblins and similar cases.

intuitions or moral belief-forming faculties. And it would not be especially surprising, on theism, for God to arrange things in this way, since there is great value in agents having the capacity to rationally form correct moral beliefs and guide their behavior in accordance with those beliefs. A capacity for correct moral beliefs may also be a precondition on morally responsible agency, the value of which is emphasized in standard theodicies.¹² Setiya (2012: 114) suggests that a theistic story along these lines is “the only hope for ethical knowledge if the facts are constitutively independent of us.”

The hypothesis that God arranged for our moral faculties to be reliable is compatible with an evolutionary account of our origins. There are several available models here. On one model, God guides the evolutionary process by intervening at key junctures in our evolutionary history, putting his thumb on the scales to ensure that certain favored mutations prevail against others that might have won in a “fair fight.” On this model, God acts like a government that occasionally intervenes in a market economy to favor firms with socially valuable practices that might not have prevailed under conditions of laissez faire competition. On another model, God sets up the initial conditions and laws so that the desired result occurs without any such intervention. There are also more metaphysically inflationary models. God might bestow a rational soul at a certain late stage of human evolution which confers a capacity for rational insight into basic moral truths. (Even if the specific capacity to apprehend moral truths has a neutral or negative impact on one’s prospects of survival and reproduction, it might prevail under Darwinian competition if it is an aspect of a more general capacity for rational insight that confers compensating advantages.) Or God might institute fundamental teleological laws of the kind hypothesized by Nagel (2012), which direct the world toward such valuable outcomes as

¹² E.g., Swinburne (2004: ch. 11), van Inwagen (2006).

moral reliability. (But it's not clear that this proposal is preferable to Nagel's view that fundamental teleological laws hold without God or any other supernatural designer.) No doubt there are many other models.

Just as God can guide the process of biological evolution toward moral reliability, he can steer processes of cultural evolution toward the same end. In this connection, it is interesting to consider a common talking point among Christian apologists: that many central ideas within modern western moral thought, such as universal human rights, the equal dignity of all persons, and a special concern for the poor and oppressed, come to us largely through the cultural influence of Christianity.¹³ I am concerned here with theism in general rather than Christian theism specifically. But to the extent that these historical assumptions are true, there may be an argument that Christian theism is especially well positioned to avoid a coincidental correspondence between our moral beliefs and the moral facts (assuming "our" moral beliefs are significantly aligned with modern western moral sensibilities).¹⁴ The force of this argument would depend on various issues that I won't try to adjudicate here, such as the extent of Christianity's influence on modern moral sensibilities, the likelihood that we would have converged on similar ideas without Christianity's influence, and whether Christianity's memetic success could itself be explained by the truth of many of its moral teachings, even if its metaphysical teachings are false.

I've suggested that theism is well positioned to avoid the result that the reliability of our moral faculties is a major coincidence (and, to a large extent, the same can be said of non-theistic

¹³ For a lengthy defense of this historical claim, see Holland (2021).

¹⁴ This coincidence-avoidance argument should be distinguished from the argument, defended by Linville (2009) and Cutter (2019), that first-order moral claims about (say) the dignity of all human beings only make sense given, or are best accommodated by, the metaphysical assumptions of something like Christian theism. (See Rachels [1990] for an inverted "modus tollens" version of the argument.)

forms of axiarchism). But how does theism help with the threat of Explanatory Disconnect? Here are two theistic models for how the moral truths might partly explain our moral beliefs.

First model: Indirect explanation via divine intellectual perception. Suppose W.D. Ross (1930) is right that there is a basic prima facie duty to keep one's promises. And suppose that, in order that we would believe this truth, God designs our cognitive faculties (by guiding evolution or whatever) so that, in relevant circumstances, we are disposed to have an intuition or "intellectual seeming" with content in the ballpark of <one has a prima facie duty to keep one's promises>. We then believe the content of this intuition, and thereby have a true moral belief.

This story can be filled out in such a way that the moral truth—that there is a prima facie duty to keep one's promises, which we'll call "M"—explains our belief. First, note that the story above evidently presupposes that God believes M (an unproblematic assumption, since God is supposed to know every truth), and that we come to believe M in part because God believes M. So, if M in turn explains the fact that God believes M, then M will figure indirectly in the explanation for why *we* believe M. How could M explain why God believes M? Here is a proposal I find attractive. First, God believes M because he directly intellectually perceives M. (I would say, more strongly, that God believes (and knows) M *in virtue of* intellectually perceiving M, in the grounding sense of "in virtue of." But I don't think this stronger claim is needed.) Second, when one directly intellectually perceives that p, then the fact that p is a constituent of one's perception in such a way that a belief based on, or grounded in, this perception is thereby explained (at least partly) by the fact that p. A similar model might be available to non-theistic forms of axiarchism that agree with theism that there is something Mind-Like (but without all the

omni-properties) at the foundation of nature.¹⁵ But not all forms of axiarchism accept this assumption.

Second model: direct explanation via human intellectual perception. On a second model, God arranges things so that we have powers of intellectual perception, by which we can directly perceive basic moral facts like M. This claim goes naturally with the more general view that our intuitive knowledge of abstract necessary truths is based on intellectual perception.¹⁶ These abstract necessary truths might include elementary mathematical truths (e.g., $2 + 1 = 3$), simple metaphysical truths (e.g., the part isn't greater than the whole), basic moral truths (e.g., M), principles of logic (e.g., the validity of modus ponens), and so forth. Again, it is plausible that if one's belief that p is based on an intellectual perception of the fact that p, then the fact that p thereby partly explains one's belief (perhaps because it is a constituent of a perception that causes the belief).

This proposal should be contrasted with what is probably a more common intuitionist account. According to the latter, one's belief that there is a prima facie duty to keep one's promises is merely based on an "intellectual seeming," where intellectual seemings are non-factive representational states (unlike intellectual perceptions, which would be factive and success-entailing).¹⁷ On the intellectual-seeming picture, it is not clear how M could have any explanatory role with respect to one's belief. When one believes on the basis of an intellectual seeming, the explanatory role of the seeming is presumably the same whether the seeming has a true or false content. In the bad case, as when one intuits and believes a falsehood like "necessarily, if y is composed of pairwise disjoint xs, then y's volume equals the sum of the

¹⁵ Goff (2023).

¹⁶ See Bengson (2015) for a defense of this view.

¹⁷ See, e.g., Huemer (2005: ch. 5).

volumes of the *xs*,” no abstract fact plays a role in explaining one’s belief.¹⁸ Thus, in the good case—where the seeming has a true content—the seeming will evidently explain the belief in a way that leaves no direct explanatory work to be done by the abstract fact. The latter’s explanatory relevance appears to be “screened off” by the seeming.¹⁹

If one prefers the intellectual-seeming view (as I do, though very tentatively), perhaps the best hope for giving moral truths an explanatory role *vis-a-vis* our moral beliefs is to adopt the first model involving indirect explanation via divine intellectual perception. In that case, of course, *God* will need to have intellectual perceptions, not just intellectual seemings. But I don’t think there is anything awkward about the view that God has intellectual perceptions while we have mere intellectual seemings. The main reason for doubting that *we* have intellectual perceptions is that our putative intellectual perceptions plausibly share a common kind with our false intellectual intuitions. Divine infallibility precludes a similar motivation for denying intellectual perceptions to God.

In our second model, God’s only role is to arrange things so that we have powers of intellectual perception that put us in direct epistemic contact with abstract facts. But of course, a non-theist can say that we have these powers without God’s help. This might be an attractive view for non-theist axiarchists. Indeed, even a naturalist can say that we have these powers. At least, this view is compatible with our definition of naturalism. But it won’t be attractive to naturalists in practice. For one thing, naturalists tend to be physicalists about the mind, and it’s not at all clear how to make sense of intellectual perception along physicalist lines.²⁰ Granted,

¹⁸ If you suffer from this intuition, consider a spherical region with a volume of 1 cubic meter composed of continuum-many points.

¹⁹ This is a variant on the screening-off objection to naive realist theories of sense perception. For discussion, see Martin (2004) and Pautz (2021: ch. 5).

²⁰ Cf. Bogardus and Perrin (2022: 10, fn. 19).

naturalism doesn't have to be combined with physicalism about the mind. David Chalmers (1996) defends "naturalistic dualism," which combines metaphysical naturalism with the view that there are irreducible qualia linked to physical states of the brain by fundamental psychophysical laws. One could extend Chalmers' naturalistic dualism to include psychophysical laws to the effect that, when a brain reaches a certain level of physical complexity, it comes to have (not just sensory and affective qualia, but also) powers of intellectual perception that put one in contact with the third realm. Any naturalist who accepts this view can, I think, avoid the threat of Explanatory Disconnect. But even naturalists with dualist leanings, like Chalmers, do not tend to carry their dualism this far. It would certainly be odd for a universe that lacks anything like fundamental purposiveness—a universe that is not fundamentally ordered towards goods like knowledge—to have fundamental laws of this kind.²¹

4. A theistic voluntarist model for avoiding coincidence

In this section, I will outline another theistic strategy for explaining the non-coincidental truth of our moral beliefs. The strategy has something in common with Street's (2006) anti-realist explanation of moral reliability. For Street, our moral beliefs are largely reliable because the moral truths depend on our attitudes. Like Street's account, the model below gives psychological truths a role in explaining the moral truths. But unlike Street's account, my model adopts a thoroughly realist metaethics. It also avoids the most serious objections to Street's anti-realism. For example, Street's view has the problematic consequence that an "ideally coherent Caligula"—roughly, someone who desires to torture innocents, and would desire this (and desire to desire it, etc.) under conditions of full descriptive information, ideal structural rationality,

²¹ Bogardus (2016) similarly argues that metaphysical naturalists can respond to moral debunking arguments by embracing intellectual perception ("rationalism" is his term), but notes that this response would be "incongruous" with naturalism.

etc.—lacks any moral reason to refrain from torturing innocent people.²² The account below has no such consequence.

As background, let's start with the idea, familiar from standard evolutionary debunking arguments, that evolution has endowed us with certain innate “proto-evaluative dispositions.” That is, we are innately predisposed to find certain actions right/good/attractive and others wrong/bad/unattractive. Thus, evolutionary debunking arguments standardly assume *something like* the following:

Moral Nativism: There are certain innate cognitive structures, which are more-or-less universal among humans, that predispose us to judge that certain actions are right and others are wrong, and which predispose us toward the corresponding moral emotions (e.g., guilt at the recognition of having done certain things and resentment or disapprobation towards others who do them).

Given Moral Nativism, a theist might offer the following conjecture:

Divine Etiology: God intentionally brought it about (directly or indirectly) that we have these innate cognitive structures, and did so *in such a way* that the cognitive structures themselves constitute/realize *commands*—laws “written on the heart” in St. Paul’s phrase.²³ Because they are commands from a legitimate authority, they have prescriptive force. That is, we are morally required to follow them, as we are in general morally required to follow the commands of legitimate authorities (with many qualifications, of course).

In effect, we can think of Moral Nativism as the hypothesis that there is something loosely akin to a rudimentary code of precepts inscribed in the innate structure of the human

²² For a defense of this consequence, see Street (2009).

²³ Romans 2:15.

mind, precepts that influence our judgments, feelings, and actions. Divine etiology adds that these precepts have an etiology that invests them with normative force. Given *naturalism* and moral realism, the fact that we are innately predisposed toward finding an action wrong or right is a mere psychological fact, without direct normative significance. That is, our psychological predisposition toward finding a given action right or wrong does not make it the case that the action is right or wrong. This is part of what gives rise to the sense that the correctness of our moral beliefs would be a lucky coincidence. On theism, however, the fact that we are innately disposed to find certain actions right or wrong *could* have normative significance. This would be true, I suggest, if Divine Etiology is true. This, then, is another way in which theism has resources to avoid a major coincidence in the correspondence between the moral truths and the moral beliefs toward which we are naturally predisposed. (However, it is doubtful whether non-theistic axiarchic views can offer an analogous account. Could a Nagelian teleological law or a Parfitian “selector” play the normative role of a *legitimate authority*, a source of normatively binding commands?)

Now for a litany of qualifications: First, in my view Divine Etiology should *not* be combined with a divine-command metaethics. The theist can and should say that there are many normative truths that hold independently of God’s commands. For example, one does not need to invoke divine commands to explain why it is wrong to torture innocent people. There are sufficient grounds for this and many other moral truths that make no mention of divine commands. But among the normative truths that hold independently of God’s commands is the truth that one has a (pro tanto) duty to obey the commands of legitimate authorities (subject to many qualifications that I won’t attempt to spell out here).²⁴ I do not have a theory of what it

²⁴ Cf. Swinburne (2004: ch. 9).

takes for someone to be a legitimate authority over someone else. But if God exists, it's plausible that he meets the conditions of legitimacy with respect to his creatures. Thus, without adopting anything like divine-command metaethics, we should allow that, if God exists, he can issue commands, and in so doing, he can create obligations that wouldn't otherwise have existed. The claim that there is a presumptive moral duty to obey God's commands no more commits one to divine command metaethics than the claim that citizens have a presumptive moral duty to follow the laws of a legitimate government commits one to a government-command metaethics. Divine Etiology is fully compatible with a realist metaethics on which the most basic moral truths hold independently of anyone's attitudes, including God's.

A second qualification is that I do not assume that every moral belief that comes naturally to us corresponds to a divine command. More plausibly, only a proper subset do. There could be moral beliefs that we are innately disposed to form that God does not intend for us to follow. I assume that whatever role God plays in the etiology of these beliefs, or the underlying cognitive mechanisms that give rise to them, it is not the sort of etiology that makes for a command. (I won't try to spell out what kind of etiology is needed. Perhaps the key factor is just that God intends that we act in accordance with the relevant belief.)

A third qualification: what St. Paul calls the "law [...] written on [our] hearts" has historically been called the "natural law," which obviously figures prominently in the tradition of "natural law ethics." The model suggested here is committed to some core claims of the natural law tradition, such as that we have a natural capacity to know the basic precepts of the natural law. But it remains neutral with respect to the most controversial claims in the natural law tradition. For example, there is no suggestion here that the natural law has any direct connection to natural teleology, or that the natural law forbids acts that are contrary to the ends of our

natural faculties. Indeed, there is no suggestion here that the biological functions of our organs or faculties have any moral authority whatsoever. Thus, many of the usual objections that other natural law theories struggle to answer (why isn't it wrong to chew gum, or walk on one's hands, or wear earplugs, given that these acts seem to be using our natural powers/parts in ways contrary to their biological functions²⁵) are not even *prima facie* threats to the voluntarist natural law theory suggested here. (Relatedly, the current proposal does not encourage the attempt to derive a system of sexual ethics from the natural functions of our sexual organs.) This proposal also remains neutral with respect to other contentious commitments of traditional natural law theory, such as that the good is prior to the right, or that following the dictates of the natural law will promote the agent's own happiness.

Among the actions God has forbidden via the law on the heart, some may have been wrong independently (e.g., murder). But as with human law, others may be wrong only because they are forbidden by God. We've already noted a general reason why God might command X-ing even if X-ing wouldn't otherwise have been obligatory—*viz.*, the good outcomes that result from widespread practice of X-ing. There are many more specific or closely related reasons: X-ing might be one among several equally good ways to solve a coordination problem. Compare human laws requiring driving on the right side of the road. Or the benefits of X-ing might only result if sufficiently many people do X, which is unlikely to happen without a command. Compare human laws limiting emissions or the practice of conscription in a just defensive war. Or it could be that the command to do X achieves the optimal tradeoff of in-practice followability and getting people to follow their independent obligations. This could be the case if X-ing is independently wrong except in a range of relatively rare exception cases. A rule like

²⁵ See, e.g., Feser (2015: ch. 16).

“don’t X, except in cases A, B, C, ...” might not be practically followable, either because it’s too complex or because it’s too easy to deceive oneself about whether one is in an exception case. In that case, a simple rule against X-ing might be preferable to a more complex rule that more closely aligns with the independent moral truths. Probably most human laws are motivated by some such tradeoff.

This possibility—that the natural law forbids some actions that, in the absence of divine command, wouldn’t have been wrong—has an interesting connection to some moral debunking arguments. I have in mind “selective” debunking arguments, which aim to debunk specific moral beliefs without debunking our moral beliefs in general. For example, de Lazari-Radek and Singer (2014) give a selective evolutionary debunking argument against moral views that favor partiality toward one’s near and dear. Huemer (2008) defends the use of selective debunking arguments to undermine beliefs about the wrongness of incest, bestiality, and many other categories of actions where our beliefs may result from evolutionary programming or other biasing factors.

Selective debunking arguments commonly focus on actions that we naturally find morally repugnant, but where it’s very difficult to explain why they are wrong. Good examples of such actions are found in the empirical literature on “moral dumbfounding.”²⁶ One famous case involves consensual incest, with details filled in to block obvious candidate explanations for the act’s wrongness (they use two forms of birth control, they keep it a secret, it doesn’t harm the siblings’ relationship, etc.)²⁷ Another involves necrobestiality, with details filled in so that no

²⁶ See Haidt (2001).

²⁷ Haidt, Bjorklund, & Murphy (2000).

person or animal is hurt by the act.²⁸ Subjects tend to judge that these acts are wrong, but are “dumbfounded” in attempting to explain what makes them wrong.

An attractive feature of the voluntarist model under consideration is that it has resources to explain what could make such actions wrong. It could thereby help to resist selective debunking arguments focused on actions that we naturally find wrong but whose wrongness is otherwise hard to explain. On the current model, the fact that we naturally find (say) consensual incest or bestiality morally repugnant could itself be morally significant. If this psychological fact has the right etiology, it could be the realization of a divine command (say, a general command against sexual relations with siblings or animals). In this way, our natural reactions to such actions might be a part of the explanation for what makes them wrong.

This model might thereby vindicate a degree of *self-trust* concerning some natural and stubborn moral intuitions and emotions that would otherwise be hard to justify. In an (in)famous essay entitled “The Wisdom of Repugnance,” Leon Kass (1997: 20) writes,

[D]isgust is the emotional expression of a deep wisdom that is beyond the capacity of reason to articulate. Can anyone give a fully adequate argument about the atrocity that is incest (even consensual incest) between father and daughter, or about the internship of sex with animals, or about mutilating a corpse, or about eating human flesh [...]? Could anyone's inability to give a fully rational justification for such practices make ethical disgust suspect? Not at all.

Apart from something like the voluntarist model suggested here, I do not see an easy way to credit moral disgust with the “deep wisdom” Kass claims for it.

5. From Naturalism to Coincidental Correspondence and Explanatory Disconnect

²⁸ Haidt, Koller, and Dias (1993).

5.1 Explanatory Independence and Coincidence

Given metaphysical naturalism, it is very plausible that there will be a complete explanation of our moral beliefs that does not mention any moral truths. In other words, it looks plausible that naturalism leads to Explanatory Disconnect. As I've said, the Darwinian explanation for why we are disposed to believe that parents have a moral duty to provide for their children will presumably not invoke the hypothesis that parents *in fact* have a moral duty to provide for their children. Cultural factors may play a role in addition to biological factors in explaining this belief. But on naturalism, we would expect the cultural explanation to invoke descriptive factors to do with social expectations, legal requirements, practices of social shaming, and so forth, not normative hypotheses about the objective moral duties of parents. Thus, it seems plausible on its face that Explanatory Disconnect is true if metaphysical naturalism is true. This point could be disputed, but I shall focus here on what I take to be the more subtle and interesting question of whether naturalism leads to Coincidental Correspondence.

Recall that the latter says:

Coincidental Correspondence: Given the extent to which evolutionary (and perhaps other) forces have shaped our moral beliefs, and given moral realism, it would be a major coincidence—a happy accident, a lucky fluke, a suspiciously convenient stroke of luck—if our moral beliefs turned out to be largely correct.

A common motivation for Coincidental Correspondence is that evolution only “aims at” survival/reproduction, and only “cares about” giving us true beliefs insofar as truth contributes to survival/reproduction. From here, it is argued that, for *moral* beliefs specifically, their contribution to a creature's survival/reproduction is independent of their truth. Even if the beliefs *are* true, their truth does not help to explain their contribution to survival/reproduction. The fact

that it is my duty to provide for my children plays no role in explaining why this belief causes my children to survive. Thus, it can seem that it could only be a coincidence if the moral beliefs instilled in us by evolution correspond to an independent moral truth. In this respect, moral beliefs are very different from, say, beliefs about the locations of berry bushes, predators, or mates. If Og's belief that a berry bush is over yonder helps Og to survive, this is probably because there *really is* a berry bush over yonder. The truth of the belief is part of the explanation for why the belief helps him survive. Thus, while evolutionary forces influence our beliefs about berry bushes no less than our moral beliefs, the evolutionary story in the former case does not tempt us to conclude that it could only be a coincidence if our berry-bush beliefs turned out to be true.

More generally, it seems that, given naturalism and moral realism, basic moral truths (such as that there is a *prima facie* duty to avoid harming others, or that pain has final disvalue) are explanatorily independent of our moral beliefs in a very strong sense. That is, it is neither the case that moral truths explain our moral beliefs, nor that our moral beliefs explain the moral truths, nor does there appear to be any common cause or explanation that accounts for both.²⁹ It is arguably this strong explanatory independence that makes the correspondence of moral beliefs and basic moral truths seem like a striking coincidence.

In my view, this explanatory independence is enough to support the verdict that naturalism and moral realism lead to an objectionable coincidence. In the remainder of this

²⁹ There could be a "common-cause" or "third-factor" explanation in the case of certain *non-basic* moral truths and beliefs. For example, if survival is good, then the fact that eating leads to survival both explains why eating is good (instrumentally, by contributing to the good of survival) and explains (via Darwinian pressures) why we believe that eating is good. (See Enoch (2011: ch. 7) for a "third-factor" explanation of moral reliability along these lines.) But no such common cause/explanation will hold for the most basic moral truths. For example, if it is a basic moral truth that we have a *prima facie* duty to avoid harming others, then this truth is not explained by the instrumental contribution of harm-avoidance to survival or anything else.

section, I consider two further (mutually compatible) strategies for supporting the same verdict. The first emphasizes the wide range of (nomologically) possible moral beliefs that evolution might easily have produced, had the evolutionary process gone somewhat differently. The second emphasizes the wide range of conceptually possible moral truths, most of which do not correspond to any belief that evolution has a tendency to instill.³⁰

5.2 Alternative possible courses of evolution

One way to argue that naturalism and moral realism yield an objectionable coincidence is to argue for something like the following:

Evolutionary Luck: Given naturalism and moral realism, the process of (biological) evolution could very easily have produced intelligent creatures whose moral beliefs fail to correspond to the moral truth to the degree that ours do.

In support of Evolutionary Luck, we can argue that evolution could easily have produced intelligent species who are disposed to form moral beliefs very different from ours. This is suggested by the fact that evolution *has* produced species whose “proto-evaluative dispositions” are different from ours. Crummett and Swenson (2020) argue in this way that naturalism faces a coincidence problem, citing Darwin’s (1902: 137) remark:

If . . . men were reared under precisely the same conditions as hive-bees . . . our unmarried females would, like the worker-bees, think it a sacred duty to kill their brothers, and mothers would strive to kill their fertile daughters, and no one would think of interfering.

Similarly, Linville (2009: 397) writes,

³⁰ My reading of Street (2006) is that she favors both strategies, though without disentangling them. (“[Moral reliability would] require a fluke of luck that’s [...] extremely unlikely, in view of the huge universe of logically possible evaluative judgements and truths” (122).

Wolves in a pack know their place in the social hierarchy. A lower-ranked wolf feels compelled to give way to the alpha male. Were he endowed with [human-level] intellectual powers, then, presumably, his “moral sense” would tell him that obeisance is his moral duty. He would regard it as a moral fact that alpha interests trump beta or omega interests

One might likewise argue that if evolution produced a rational non-social species, they probably would not recognize moral duties of reciprocity. Or consider that different species vary dramatically with respect to parental investment in their offspring. Rational creatures evolved from a lineage with minimal parental investment might have a much weaker sense of the moral importance of caring for offspring.

I find this way of arguing for the coincidence verdict unpersuasive. First, in some of these cases the moral truths could very well be different. For an intelligent species in which the young require little parental investment to survive and flourish, it’s plausible that parental duties would be much weaker than they are for us. The same goes for intelligent species whose litters are large enough that significant parental investment in each child is infeasible. (The claim that truths about parental moral duties are contingent on these biological factors is, of course, compatible with the standard view that *the most basic* moral truths are necessary. What is the underlying necessary truth here? Perhaps something like: there is a Rossian prima facie duty, stronger than our duties of impartial beneficence, to meet the basic needs of (and beyond this baseline, a weaker duty to promote the well-being of) any welfare subject that one has a direct hand in creating.)

But this response has limits. It is implausible for Darwin’s hive bees, for example. Even if intelligent hive bees were strongly disposed to believe that fratricide is morally permissible or

praiseworthy, and to act accordingly, it would presumably still be morally wrong (contra Doane-Clarke 2012: 320). (We're imagining that these hive bees are intelligent enough to make moral judgments, so they presumably qualify as moral persons.) Given naturalism and moral realism, it therefore looks somewhat lucky that, on the issue of fratricide anyway, we have proto-moral dispositions that roughly align with the moral truth. At least they are better aligned than those of Darwin's hive bees.

However, it's not clear that our proto-moral dispositions are better aligned with the moral truth *in general*. Still less is it clear that human proto-moral dispositions are generally better aligned with the moral truth than those of most other animal species. For example, we are arguably innately disposed to find genocide morally unproblematic. Peter van Inwagen (2010: 15-6) suggests that it was natural for the ancient authors of Deuteronomy to represent God as commanding the Israelites to commit what we call genocide because

they were typical human beings, and typical human beings see nothing wrong with genocide [...] Most people have taken it for granted that a tribe, when it moves into a new territory, will kill those of the previous inhabitants that it does not enslave. That's what people do — the Old Common Morality says — and they'd be crazy to do otherwise.

Perhaps we couldn't have ever come to the right verdict on the morality of genocide if our evolutionary endowment were corrupt to the core. But it's not clear that we needed an evolutionary endowment that is especially good in comparison with most other species.

If one wants to support the coincidence verdict by appeal to alternative possible courses of evolution, a more promising strategy might be to consider alternative possible courses of *cultural* evolution. (The context of van Inwagen's remark is his argument that we now recognize genocide as wrong due to the cultural influence of Christianity. "The morality to which critics of

the moral character of the God of the Bible appeal is a gift to the world from Israel and the Church and is by no means self-evident.”) That is, instead of Evolutionary Luck, one could argue for:

Cultural Evolutionary Luck: Given naturalism and moral realism, the process of *cultural* evolution could very easily have turned out people whose moral beliefs failed to correspond to the moral truth nearly as well as ours do.

To what group does “ours” correspond? Perhaps we should say something like “all humans alive today.” Of course, the claim would be more plausible if we focus on a narrow group, like you and your like-minded friends (surely cultural evolution *has* turned out groups with less enlightened moral beliefs than this group). However, the narrow construal doesn’t seem to give us a problematic coincidence of the kind that would cast doubt on naturalism. There might be *some* problematic coincidence here. It would be awfully lucky if cultural forces have given you and your friends more accurate moral beliefs than just about any other group on the planet. This would be like winning the epistemic lottery. However, as with winning the actual lottery, it’s hard to see how this striking fact could cast doubt on naturalism, because it’s hard to see how this striking fact would be more likely on any remotely plausible non-naturalistic view.

Although the argument for coincidence from Cultural Evolutionary Luck is worth pursuing, I will not do so here.

5.3 Alternative conceptually possible moral truths

Instead of emphasizing alternative possible courses of evolution, we can argue for the coincidence verdict by considering the vast range of alternative conceptually possible moral truths.³¹ The basic moral truths are often assumed to be metaphysically necessary. But they seem

³¹ Cf. Clarke-Doane (2012: sect. II).

to be logically or conceptually contingent. It is not analytic that pain is bad, or that paperclips aren't the sole bearers of final value, although these claims may be necessary and knowable *a priori*. The range of conceptually possible moral truths is vast. For most conceptually possible moral truths, evolution has no tendency to instill the corresponding belief. Consider all the conceptually possible moral truths pertaining to how we should treat our offspring. One possibility is that we should care for them. Another is that we should kill them. Another is that we should live between 7 and 9 miles away from them. Another is that we should surround them in a circle of leaves. Another is that we should turn them into paperclips (a belief that might be held with great conviction by Bostrom's (2014) paperclip maximizer). Only a very small proportion of these conceptual possibilities does evolution have any tendency to get us to believe. But as luck would have it, *these are the ones that turn out to be true*. How convenient!

Is this a good argument that, given naturalism and moral realism, a major coincidence has occurred? I'm not sure. It is hard to say whether the mere conceptual possibility of alternative moral truths gives us a robust enough kind of contingency to support a coincidence verdict. But there are reasons to think that metaphysical contingency is not required for a problematic coincidence (*pace* Wielenberg 2014: 167). Consider the alleged fact that certain basic physical parameters fall within a very narrow life-permitting range. Many have suggested that it would be an unbelievable coincidence if these parameters just happened to fall within the life-permitting range without some further story. One response is that the laws and constants are brutally metaphysically necessary. But this doesn't seem to remove the appearance of an incredible coincidence, absent some further story as to why the necessary parameter values should fall within the narrow life-permitting range. Or suppose one accepts the Spinozist view that all truths are metaphysically necessary. This shouldn't prevent one from finding some hypotheses

implausible on the grounds that they imply an incredible coincidence. If two students submit matching term papers, even a Spinozist should doubt that they worked independently, since it would be an incredible coincidence for two students to independently produce identical papers.³²

The idea that conceptual (rather than metaphysical or nomological) contingency is what matters for an objectionable coincidence fits naturally with a Bayesian framework. In this framework, we distribute probabilities over a space of epistemic/conceptual possibilities, which needn't correspond to metaphysical possibilities. Although it's metaphysically necessary that Hesperus is Phosphorus, the ancient Babylonians could rationally assign substantial probability to conceptually possible scenarios in which Hesperus isn't Phosphorus.

That said, it's not entirely clear that there is a legitimate Bayesian argument for a problematic coincidence in the moral case. In a Bayesian setting, objectionable coincidences typically involve a correspondence between two or more variables that is antecedently improbable, absent some further story of a kind that would render the correspondence non-coincidental. For example, it is highly improbable that two students would produce matching term papers conditional on the assumption that they worked honestly and independently. This is part of the reason why we should doubt the latter assumption when we receive matching term papers. (The other part is that there are alternative hypotheses, such as cheating, that both predict the data and whose prior probability, even for seemingly honest students, is many orders of magnitude higher than the probability that the papers would match by sheer coincidence.)

Can we similarly say, in light of the vast space of conceptually possible moral truths, that there is an extremely low probability, on naturalism, that the moral truths would turn out to

³² The term paper example is from White (2005). For a defense of the claim that there can be epistemically problematic coincidences without metaphysical contingency, see Faraci (2019). See Baras (2020) for a related discussion of when a metaphysically necessary truth "calls for explanation."

match the beliefs that evolution has a tendency to produce? For example, in view of the vastly many conceptual possibilities for how we ought to treat our offspring, can we say that there was an extremely low prior probability that there would turn out to be a duty to *care* for our offspring? We might imagine that there is some ur-prior probability function P (or a relatively limited set of such functions) that serves as a constraint on rational credences, such that any ideally rational agent with total evidence E assigns to any given proposition X a credence equal to $P(X|E)$. And we might suppose that the rational ur-priors treat it as a wide open question whether we ought to care for our offspring, surround them with leaves, turn them into paperclips, or whatever, distributing probabilities in some natural (uniform?) way over this vast set of alternative conceptual possibilities. The “duty to care” hypothesis, being just one of vastly many alternatives, thus receives a miniscule prior probability. (Compare: from the perspective of the rational ur-priors, it is a wide open question what the exact diameter is, to the nearest inch, of the first life-supporting planet. No specific answer should receive more than a sliver of probability.)

One feels that this is a brazen abuse of the Bayesian apparatus. I suspect that this issue does not lend itself to a Bayesian treatment. I have no objection to the “objective Bayesian” idea that there is some privileged “ur-prior” probability function P , or a limited set of such functions, that constrain rational credences in the manner suggested above. But even on this (already controversial) assumption, it is entirely unclear how to think about ur-priors for basic moral propositions. (One of many complications: if ur-priors correspond to credences that are justified *a priori*, or independently of any empirical evidence, then it’s hard to maintain that the ur-prior for “agents have a prima facie duty not to harm others” is roughly on a par with bizarre moral hypotheses like “agents have a *prima facie* duty not to benefit others.” I would have thought that the former is knowable *a priori*, via something like rational intuition. But perhaps one could

think of ur-priors as corresponding to the credences that are rationally justified independently of any empirical data, *or data from rational intuition*. But in that case, it's unclear why we should expect ur-priors to satisfy the usual probabilistic axioms. For example, it is unclear how we could be justified in assigning credence 1 to logical truths prior to any "data from rational intuition.")

I am inclined to think that this issue is not fruitfully viewed through a Bayesian lens, at least not before sorting out some difficult and foundational questions in Bayesian epistemology. This is not to dismiss the general strategy of appealing to alternative conceptually possible moral truths to argue that naturalism and moral realism lead to a problematic coincidence. But I am pessimistic about a Bayesian implementation of this strategy.

To sum up: the main claim under consideration in this section is that if naturalism is true, then Coincidental Correspondence is true. I have considered three strategies for supporting this claim. The first emphasizes the strong explanatory independence of our moral beliefs and the moral truths. On naturalism, it seems that neither explains the other, nor is there a common explanatory source of both. This is arguably enough to generate a problematic coincidence. The second emphasizes the alternative moral beliefs that might easily have been produced by alternative courses of (biological or cultural) evolution. The third emphasizes the vast range of alternative conceptual possibilities concerning the moral domain. My own view is that the first strategy is successful. I am less sure about the second and third (and somewhat more optimistic about the third than the second).

6. Conclusion

Moral realism sits uneasily with naturalism. I've argued that theism and other forms of axiarchism are better positioned than naturalism to avoid two conclusions that the moral realist

should seek to avoid—that the correctness of our moral beliefs is a major coincidence, and that our moral beliefs can be fully explained without mentioning any moral truths. Many philosophers, myself included, find moral realism so plausible that rejecting it isn't a live option. If my conclusions are correct, they have reason to prefer theism or some other form of axiarchism over naturalism.

References

- Baras, Dan (2020). How can necessary facts call for explanation. *Synthese* 198 (12):11607-11624.
- Bedke, Matthew (2009). Intuitive non-naturalism meets cosmic coincidence. *Pacific Philosophical Quarterly* 90 (2):188-209.
- Bengson, John (2015). The Intellectual Given. *Mind* 124 (495):707-760.
- Bogardus, Tomas (2016). Only All Naturalists Should Worry About Only One Evolutionary Debunking Argument. *Ethics* 126 (3):636-661.
- Bogardus, Tomas & Perrin, Will (2022). Knowledge is Believing Something Because It's True. *Episteme* 19 (2):178-196.
- Bostrom, Nick (2014). *Superintelligence: paths, dangers, strategies*. Oxford University Press.
- Chalmers, David (1996). *The Conscious Mind: In Search of a Fundamental Theory*. Oxford University Press.
- Clarke-Doane, Justin (2012). Morality and Mathematics: The Evolutionary Challenge. *Ethics* 122 (2):313-340.
- Climenhaga, Nevin (forthcoming). How infallibilists can have it all. *The Monist*.

Crummett, Dustin & Antony Louise (2022). Does Moral Knowledge Prove God?

https://www.youtube.com/watch?v=5QJ_e15a81w

Crummett, Dustin & Swenson, Philip (2020). God and Moral Knowledge. In Kevin Vallier & Joshua Rasmussen (eds.), *A New Theist Response to the New Atheists*. New York, USA: Routledge. pp. 33-46.

Cutter, Brian (2019). Down the Labyrinthine Ways. In Brian Besong & Jonathan Fuqua (eds.), *Faith and Reason: Philosophers Explain Their Turn to Catholicism*. San Francisco: Ignatius Press. pp. 79-96.

Cutter, Brian & Crummett, Dustin (forthcoming). *Psychophysical Harmony: A New Argument for Theism*. Oxford Studies in Philosophy of Religion.

Cutter, Brian & Saad, Bradford (forthcoming). *The Problem of Nomological Harmony*. *Noûs*.

Darwin, Charles (1902). *The Descent of Man*. 2nd ed. New York: American Home Library Company.

de Lazari-Radek, Katarzyna & Singer, Peter (2014). *The Point of View of the Universe: Sidgwick and Contemporary Ethics*. New York: Oxford University Press.

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

Faraci, David (2019). *Groundwork for an Explanationist Account of Epistemic Coincidence*. *Philosophers' Imprint* 19.

Feser, Edward (2015). *Neo-scholastic essays*. South Bend, Indiana: St. Augustine's Press.

Field, Hartry (1980). *Science Without Numbers: A Defence of Nominalism*. Princeton, NJ: Princeton University Press.

Harman, Gilbert (1977). *The nature of morality: an introduction to ethics*. New York: Oxford University Press.

Haidt, J. (2001). The emotional dog and its rational tail: a social intuitionist approach to moral judgment. *Psychological review*, 108(4), 814.

Haidt, J., Koller, S., & Dias, M. (1993). Affect, culture, and morality, or is it wrong to eat your dog? *Journal of Personality and Social Psychology*, 65, 613-628

Haidt, J., Bjorklund, F., & Murphy, S. (2000). Moral dumbfounding: When intuition finds no reason. Unpublished manuscript, University of Virginia.

Hawthorne, John & Nolan, Daniel (2006). *What Would Teleological Causation Be? In Metaphysical Essays*. Oxford University Press.

Holland, Tom (2021). *Dominion. How the Christian Revolution Remade the World*. Basic Books.

Horgan, Terence & Timmons, Mark (1991). New Wave Moral Realism Meets Moral Twin Earth. *Journal of Philosophical Research* 16:447-465.

Huemer, Michael (2005). *Ethical Intuitionism*. New York: Palgrave Macmillan.

Huemer, Michael (2008). Revisionary intuitionism. *Social Philosophy and Policy* 25 (1):368-392.

Jeffrey, Anne (2019). *God and Morality*. Cambridge: Cambridge University Press.

Joyce, Richard (2006). *The evolution of morality*. Cambridge, MA: MIT Press.

Kass, Leon (1997). Why we should ban human cloning: the wisdom of repugnance. *The New Republic*, 17-26.

Korman, Daniel (2015). *Objects: Nothing Out of the Ordinary*. New York: Oxford University Press.

- Korman, Daniel Z. & Locke, Dustin (2020). Against Minimalist Responses to Moral Debunking Arguments. *Oxford Studies in Metaethics* 15:309-332.
- Korman, Daniel Z. & Locke, Dustin (2023). An Explanationist Account of Genealogical Defeat. *Philosophy and Phenomenological Research* 106 (1):176-195.
- Leslie, John (1989). *Universes*. London: Routledge.
- Linville, Mark D. (2009). The Moral Argument. In William Lane Craig & J.P. Moreland (eds.), *The Blackwell Companion to Natural Theology*. Oxford, UK: Wiley-Blackwell. pp. 391–448.
- Martin, Michael (2004). The limits of self-awareness. *Philosophical Studies* 120 (1-3):37-89.
- Nagel, Thomas (2012). *Mind and Cosmos: Why the Materialist Neo-Darwinian Conception of Nature is Almost Certainly False*. New York: Oxford University Press.
- Parfit, Derek (1998). Why Anything? Why This? In Tim Crane & Katalin Farkas (eds.), *Metaphysics: A Guide and Anthology*. Oxford University Press.
- Pautz, Adam (2020). Consciousness and Coincidence. *Journal of Consciousness Studies* 1 (5-6):143-155.
- Pautz, Adam (2021). *Perception*. New York, NY: Routledge.
- Pittard, John (2021). Deceptive Worlds, Skepticism, and Axiarchism. *Inquiry: An Interdisciplinary Journal of Philosophy*:1-36.
- Plantinga, Alvin (2011). *Where the Conflict Really Lies: Science, Religion, and Naturalism*. New York. Oxford University Press.
- Goff, Philip (2023). *Why? The Purpose of the Universe*. Oxford University Press.
- Rachels, James (1990). *Created from animals: the moral implications of Darwinism*. New York: Oxford University Press.

- Rea, Michael (2002). *World without design: the ontological consequences of naturalism*. New York: Oxford University Press.
- Ritchie, Angus (2012). *From morality to metaphysics: the theistic implications of our ethical commitments*. Oxford: Oxford University Press.
- Ross, David (1930). *The Right and the Good*. Oxford: Oxford University Press.
- Setiya, Kieran (2012). *Knowing Right From Wrong*. Oxford: Oxford University Press.
- Street, Sharon (2006). A Darwinian dilemma for realist theories of value. *Philosophical Studies* 127 (1):109-166.
- Street, Sharon (2009). In defense of future tuesday indifference: Ideally coherent eccentrics and the contingency of what matters. *Philosophical Issues* 19 (1):273-298.
- Swinburne, Richard (2004). *The Existence of God*. Oxford: Oxford University Press.
- van Inwagen, Peter (2006). *The problem of evil*. New York: Oxford University Press.
- van Inwagen, Peter (2010). The Inspiration of the Hebrew Bible and the Morality of God's Commands. *The Harvard Ichthus*: 13-18.
- White, Roger (2005). Explanation as a guide to induction. *Philosophers' Imprint* 5:1-29.
- White, Roger (2010). You just believe that because.... *Philosophical Perspectives* 24 (1):573-615.
- Wielenberg, Erik (2014). *Robust Ethics: The Metaphysics and Epistemology of Godless Normative Realism*. Oxford: Oxford University Press.