

THE LIMITS OF RATIONAL BELIEF REVISION: A DILEMMA FOR THE DARWINIAN DEBUNKER

Forthcoming in Noûs; please cite final version.

Katia Vavova
Mount Holyoke College

Abstract. We are fallible creatures, prone to making all sorts of mistakes. So, we should be open to evidence of error. But what constitutes such evidence? And what is it to rationally accommodate it? I approach these questions by considering an evolutionary debunking argument according to which (a) we have good, scientific, reason to think our moral beliefs are mistaken, and (b) rationally accommodating this requires revising our confidence in, or altogether abandoning the suspect beliefs. I present a dilemma for such debunkers, which shows that either we have no reason to worry about our moral beliefs, or we do but we can self-correct. Either way, moral skepticism doesn't follow. That the evolutionary debunking argument fails is important; also important, however, is what its failure reveals about rational belief revision. Specifically, it suggests that getting evidence of error is a non-trivial endeavor and that we cannot learn that we are likely to be mistaken about some matter from a neutral stance on that matter.

1. Rationality, belief revision, and evidence of error.

We are fallible creatures, prone to making all sorts of mistakes. So, we should be open to evidence of our own error. But what constitutes such evidence? And what is it to rationally accommodate it?

I approach these questions by looking at an argument according to which we have good reason to think that our moral beliefs are mistaken and rationally accommodating this evidence requires revising our confidence in, or altogether abandoning, the suspect beliefs. This argument aims at moral skepticism via the less traditional route of evidence of error. The evidence here is empirical, consisting of facts about our evolutionary history. The skeptical conclusion is meant to follow from the realization that, in some sense, we believe what we do because we evolved to.¹

I begin by stating the *evolutionary debunking argument*: what it is and how it is supposed to work. I then show that the debate is at an impasse over the rules of the game. Specifically, the disagreement is about whether responses to the debunker can legitimately rely on the very moral claims under attack. This is the most difficult question in this debate.² Rather than answer it, I use it to formulate a dilemma that, if it works, shows that we cannot get from evolution to moral skepticism. It also reveals something important about the nature of rational belief revision; namely, that getting evidence of error is a non-trivial endeavor and that we cannot learn that we are likely to be mistaken about some matter from a neutral stance on that matter.

¹ For an introduction to these arguments see Vavova (2015). For versions of the debunking argument itself see Bedke (2009), Greene (2008), Joyce (2006), Kitcher (2007), Ruse and Wilson (1986), Ruse (1998), and Street (2006). Vavova (2018) and White (2010) discuss a more general version of this challenge.

² Shafer-Landau agrees; see his (2017) section II.

The dilemma, briefly, is this. Either moral assumptions are legitimate in response to the debunker or they are not. If they are, then learning about the evolutionary origins of our moral beliefs may give us good reason to think that our moral beliefs are mistaken. But if moral assumptions really *are* legitimate, then there are moral claims we can take for granted and therefore use to self-correct. In this way we may stagger, but needn't fall from the debunker's hit. On this horn, we can self-correct and thus alleviate the debunker's worries.

If, instead, moral assumptions are not legitimate, then we have no moral claims we can use to self-correct. But then, learning about the evolutionary origins of our moral beliefs cannot give us reason to think we are mistaken about morality. This is because we cannot get evidence to think that we are mistaken about something that we can make no assumptions about. Evidence is evidence only against a background of beliefs we take for granted, and if no moral assumptions are allowed, then we do not have enough such background. So, we cannot get evidence of our error.

Either way, then, the evolutionary debunker's argument needn't worry us.

If this is right, then we have achieved one sort of victory—we've shown that we don't have reason to think our moral beliefs are mistaken. But notice, we lack something stronger: reason to think that our moral beliefs are *not* mistaken. That is, we haven't *vindicated* our moral beliefs, and thus established that we are largely accurate or reliable about the moral realm. We might hope for this sort of vindication, especially in response to more general skeptical worries. But that is work for another day. Disarming the evolutionary debunker doesn't require us to demonstrate that our beliefs are not mistaken. We need only show that we have no reason to think they are mistaken.

If my dilemma succeeds, it shows more than that the debunker fails. It also points to a bigger epistemological lesson. It suggests that there are limits to rational belief revision and that getting evidence of error is a non-trivial matter. Understanding the evolutionary debunker's argument, and the way that it fails, thus sheds light not only on the debate about debunking, but also on bigger epistemological questions about rationality and belief revision.³

2. The argument.

The evolutionary debunker's argument is empirical, targeted, and epistemological: it aims to undermine the epistemic status of our beliefs with a scientific claim about their origins. The idea

³ This dilemma is similar to that which appears on p. 15 of my (2014). This paper can be read as offering a full, in depth treatment of that argument, though both my interpretation of the debunker's argument and the dilemma differ in important ways from those in my (2014). The reconstruction of the debunker's argument that I present here makes fewer controversial assumptions. In particular, it does not assume (1) the Independence Requirement or (2) moral realism, which most authors on both sides accept for *reductio*. This improved reconstruction helps clarify what is really at stake in the debate. We will see that it is not all about how spooky realism is, and that something much more mundane and widely accepted is causing trouble. The dilemma I present here, for this improved version of the debunker's argument, also allows us to draw more general lessons, not just about the nature and potential of debunking arguments, but also about the way rational belief revision works more generally.

is this: just as evolutionary forces shaped our eyes and ears, so they shaped our moral beliefs—at least early normative tendencies, or proto-versions of those beliefs. But evolutionary forces aim at survival and fitness—not moral truth. Left to evolution alone, we would only ever be inclined to believe what was true about morality if doing so were in our genetic interest. So, if it were in our interest to believe only falsehoods, then that is what evolution would, in its slow-moving, subtle way, select for us to believe. Insofar as evolution has influenced our beliefs, then, this influence is epistemically bad. Or, so the debunker argues.

The science is, of course, subtler than this, but we can omit the details and grant the debunker whatever she needs to support her empirical claims. What I am really interested in here is not whether evolution in fact influenced our beliefs in such and such a way, but whether *if* it did, we'd have reason to worry.⁴ Answering this question in the negative makes the evolutionary story irrelevant—even if that story goes exactly as debunkers need it to. Answering it in the positive teaches us something more general about irrelevant influences—and this is important whether or not evolution really works this way.

Let's start, then, with the first premise of the debunker's argument: her empirical claim.

Influence. Evolutionary forces influenced our moral beliefs.

This claim doesn't settle what sort of influence evolution has been: good or bad, weak or strong. Influence does not, by itself, entail anything epistemological, much less skeptical, about our beliefs' origins.

Some think that the skeptical conclusion comes when we combine Influence with:

Off-track. Evolutionary forces aim at fitness, not moral truth.

To get something like:

Off-track Influence. A process that aims at fitness, not moral truth, influenced our moral beliefs.⁵

But an off-track influence is one merely concerned with something other than the truth. Evolution is also off-track in this sense with respect to our beliefs about ordinary mid-sized objects like tables and chairs, cliffs, tigers and bears. It doesn't select for their truth either—it selects for whatever would be most adaptive. Of course, the true beliefs about ordinary, mid-sized objects are generally adaptive. If we tended toward false beliefs about such matters, we would bump into things, get eaten, etc. Thus, an influence can be off-track *and* reliable. So, we cannot conclude anything from *Off-track Influence* alone about whether our beliefs are epistemically suspect.

This shows that we need a further premise to conclude that this off-track influence is a bad influence. Specifically, we need:

⁴ Hanson (2017) and Mogensen (2016), in contrast, engage directly with the debunker on the empirical claim.

⁵ This seems to be Kahane's reading in his (2011) p. 106.

Gap. The true moral beliefs and the adaptive moral beliefs come apart.

Gap is plausible. It captures the thought that whatever morality is about, it is not *only* about maximizing fitness—that there’s more to doing what we ought, and treating one another as we should, than doing whatever would be evolutionarily advantageous.

Notice that, as stated, Gap doesn’t say how much these two sets come apart: how many true moral beliefs are not also adaptive. We might thus accept different versions of Gap. All of them, however, commit us to holding that morality is about more than maximizing fitness.

If we reject Gap, then, we are committed to holding that whenever morality requires us to *phi*, it requires us to do something that is evolutionarily advantageous—that maximizes fitness. So, rejecting Gap commits us to thinking that morality is *only* about maximizing fitness.

To see why we need a premise like Gap, notice that the aforementioned empirical version of Gap is false: the true beliefs about ordinary mid-sized objects and the adaptive beliefs about such objects decidedly do *not* come apart. In fact, they coincide.⁶ This makes all the difference, and explains why no one is trying to debunk our beliefs about tables and chairs.

Gap thus serves two important functions in the debunker’s argument: it captures the intuitive thought that morality is about more than what is adaptive and it distinguishes the good case from the bad. It is also, importantly, *not* an epistemological claim—nor does it have any immediate epistemological upshots. It is merely a claim about the overlap, or lack thereof, between two sets. This is important because the epistemological moves in the argument come later. If there’s some epistemological upshot hiding inside a claim about the nature of moral truth, then this argument might not be the new and exciting empirical argument we thought it was. So, Gap shouldn’t be doing any covert debunking work here.

Thus, Gap, or something very much like it, is an important premise in the debunker’s argument. It should also, I think be relatively uncontroversial. Whatever our ethical or metaethical commitments, I suspect most of us believe that morality is about more than what is in our evolutionary interest. So, even if it enhances fitness, say, to have as many babies as possible, we are not morally required to embark on such a project. (Indeed, we might think that we are morally prohibited from doing so.)⁷ Sometimes, of course, moral goodness and genetic fitness coincide: we ought to take care of our children and doing so is in our evolutionary interest. But this is a mere coincidence: caring for children is not morally good *because* it is evolutionarily good.

⁶ Not all empirical versions of Gap are this straightforward. Many empirical beliefs, such as abstract scientific ones, don’t clearly have an evolutionary advantage. Evolutionary psychologists and philosophers of science are actively debating how exactly to tell the story in the empirical case. But, I suspect, most think such a story can be told. I set these important complications aside here.

⁷ Those who *do* think that morality’s dictates and our evolutionary tendencies coincide, tend to understand the latter so that they don’t entail moral absurdities like requirements to reproduce as much as possible. See, e.g., Foot (2001).

If all this is right, then we can see the trouble. We don't think that the true moral beliefs substantially coincide with the fitness maximizing moral beliefs (Gap). Yet if the debunker is right, a fitness maximizing maniac has shaped our moral beliefs (Off-track Influence). From this we can infer:

Bad Influence. Our moral beliefs reflect the influence of an epistemically bad process.⁸

In other words, Gap turns Off-track Influence into Bad Influence. It is thus what's missing in the easy empirical case, where beliefs about our manifest surroundings are generally true *and* adaptive.

We are now a short step to the conclusion. It makes sense that an epistemically bad influence—one that isn't interested in either the truth or anything that correlates with the truth—would push us toward error. We can capture this idea thus:

Plausible Principle. If a belief reflects the influence of an epistemically bad process, then that belief is likely to be mistaken.

Combine it with Bad Influence and we get:

Mistaken. Our moral beliefs are likely to be mistaken.

The 'likely'—in both the principle and the conclusion—is that of objective probability. The thought is that an epistemically bad influence makes our beliefs objectively more likely to be mistaken. 'Mistaken' here is an umbrella term for the many ways we might be in error: evidence that I'm irrational, overconfident, underconfident, or incoherent counts as evidence that I am mistaken, just as evidence that my beliefs are false or held for the wrong reasons does. Evidence that my belief is *mistaken*, then, needn't be evidence that my belief is false—it could be higher-order evidence of irrationality or some other sort of error.

Although not the final stop for all debunking arguments, Mistaken is an important premise on the way to establishing moral skepticism, anti-realism, nihilism, or the like. Focusing on it also underscores that this argument is *epistemological*: whatever its final conclusion, the introduction of evidence of error is a crucial step. The evidence, of course, is the evolutionary story; the *mistake* is having the wrong doxastic attitudes about the moral realm.

This, then, is the debunker's argument:

1. *Influence.* Evolutionary forces have influenced our moral beliefs.
2. *Off-track.* Evolutionary forces aim at fitness, not moral truth.

⁸ The inference actually has a few more steps. Off-Track Influence and Gap entail something like, Non-Truthful Influence: a non-truthful process influenced our beliefs. This, combined with a premise that says non-truthful influences are bad influences, gets us to bad influence. I think, actually, that it is very important how exactly these steps go—in particular, what the right, precise version of Non-Truthful Influence is. Fortunately, it doesn't make a difference for the purposes of this paper, so I set it aside.

3. *Off-track Influence*. A process that aims at fitness, not moral truth, influenced our moral beliefs. (1,2)
4. *Gap*. The true moral beliefs and the adaptive moral beliefs come apart.
5. *Bad Influence*. Our moral beliefs reflect the influence of an epistemically bad process. (3,4)
6. *Plausible Principle*. If a belief reflects the influence of an epistemically bad process, then that belief is likely to be mistaken.
7. *Mistaken*. Our moral beliefs are likely to be mistaken. (5,6)

Notice that this formulation doesn't assume moral realism: the view that the moral facts are attitude-independent.⁹ This is unusual and important. It's unusual because evolutionary debunking arguments are often run as *reductios* of realism. It's important because it shows that one needn't be a moral realist to be vulnerable to this attack. More metaethical views are within the debunker's sights than previously thought. This changes the standard way of understanding the evolutionary debunking debate. It challenges the thought that metaphysically spooky, epistemologically suspect moral realism is, yet again, in *more* trouble than alternative views. It also has implications for what we can do with this argument: if we needn't assume realism to get Mistaken, then we can't straightforwardly get a *reductio* of realism only. Nor can we easily avoid the skeptical conclusion by becoming moral anti-realists. The dialectic is thus more complicated.

This way of running the argument showcases these complications. At the same time, it also clarifies what's really at stake. The problematic commitment—the one that makes us vulnerable to attack—isn't a metaethical commitment to realism. Rather, it is the commonsense thought that morality is about more than spreading our seed. So long as we are committed to the existence of some true moral beliefs that are not adaptive, the debunker has us in her sights.

3. The impasse.

The evolutionary debunking debate is stalled at the question of whether we may rely on moral beliefs when responding to the debunker's challenge. Can we, for example, assume that pain is bad or survival good? If we can, then maybe things aren't so dire.

“Third-factor” responses think this move is legitimate.¹⁰ They acknowledge that evolutionary forces do not select for true moral beliefs, but point out that there is, nevertheless, an important link between what evolution selects for and the moral truth. After all, evolution aims at survival, and survival is good! Evolution also, plausibly, selects for creatures who are good at avoiding pain, since doing so increases the chances of survival. But pain is bad! Objectively morally bad!

⁹ The definition of realism is both controversial and standard (see Shafer-Landau (2005) and Street (2006)). However, it is less important here that we get realism precisely right, since this argument doesn't assume realism, for *reductio* or otherwise. This is one way my interpretation of the debunker's argument differs from most others, including my own earlier formulation.

¹⁰ Brosnan (2011), Copp (2008), Enoch (2010) and (2011), Sarksaune (2011), and Wielenberg (2010) all employ this strategy. They vary primarily in what they select as the third-factor that correlates with the moral truth. Enoch argues that survival is both morally good and what evolution selects for; Sarksaune focuses on the claim that pain is bad; Wielenberg starts with the assumption that creatures like us have rights.

So, it seems, evolution indirectly and accidentally (but quite reliably!) selects for at least some of the right moral tendencies and true moral beliefs. If this is right, then we're not totally hopeless about morality. Evolution has influenced our beliefs, yes, but it has done so in a way that is, for an important class of beliefs, reliable.

Third-factor responses hit the debunker's argument at Gap, but not by outright rejecting it. Indeed, they grant that the moral beliefs and the adaptive beliefs come apart—they must, or else morality would be only about what is adaptive. But, they argue, what's good and what's adaptive does overlap, and, crucially, we know where that overlap is. Furthermore, that overlap is sufficiently substantial so that we have a stock of moral beliefs we can take for granted. If this is right, then we can conclude both that we are not *massively* mistaken about morality, and that, where we are mistaken, we can self-correct.

Of course, debunkers and opponents both agree that evolution selects for creatures who are good at survival and pain avoidance. And, I suspect, most of us think that the fact that something causes pain is, generally, a reason to avoid it and the fact that something increases our chances of survival is, generally, a reason to go for it.

But can we just assume this, in light of the challenge laid before us? The debunker claims to have given us undermining evidence for exactly these beliefs. Assuming these beliefs are true seems as question-begging as the following response to an optometrist who says that your color vision might be deceiving you: (*points to a chair*) "But that *is* blue, and it looks blue to me. So I must not be colorblind."¹¹

Responding thus to your optometrist won't do. But third-factor responses seem similarly problematic. The debunker says: your moral beliefs are likely to be mistaken. We respond: But look, they are true! So we are not mistaken. What, exactly, is our evidence that they are true, however? The very moral assumptions that are likely mistaken if the debunker's right.

This is the impasse. Both sides agree that if we are allowed substantive assumptions like these, then we have a way out of the debunker's challenge. They disagree, however, about whether those assumptions are legitimate. On the one hand, they seem question-begging. On the other, they seem obvious—almost trivially true.¹² There is much more to say here, but rather than explore this impasse further, I will use it to formulate a dilemma.

4. The dilemma.

The crux is this: in evaluating the debunker's challenge, either we are allowed to make moral assumptions or we are not. If we are, then we can answer the challenge; if we are not, then the challenge doesn't arise.

¹¹ Cf. Vavova (2015) p. 106.

¹² How can it be illegitimate to assume that pain is generally bad for creatures like us? If we bracket that, are we even talking about morality anymore? Cuneo and Shafer-Landau (2014)'s moral fixed points view can be used to argue that we aren't. More on this shortly.

We've already seen how the first horn of the dilemma should go. If we can legitimately assume the truth of specific moral claims, then we can tell a plausible story about why these off-track influences aren't bad—about why evolution has not made us hopeless at recognizing or reasoning our way to the moral truth. With these assumptions in hand, we can reason like this:

Pain is bad, survival is good, and so forth. But this is exactly what evolution inclined us toward believing! So, we don't have to worry about these moral beliefs. Others? Maybe. But that's okay. We have enough now so that we can self-correct.

On this horn, remember, we still accept a version of Gap, but not a problematic one. Although the true moral beliefs and the adaptive moral beliefs come apart, there is some overlap and we can identify it.

This move works, if it does, with just the assumption that pain is bad—that the fact that something causes pain is, generally, a reason not to do it. If *any* moral assumptions are legitimate in the context of such a debate, surely one this weak is.

Of course, we know what the debunker would say: however uncontroversial or obvious such assumptions may seem, to employ them is to beg the question. This brings us to the second horn of the dilemma, on which moral assumptions are not legitimate.

If moral assumptions aren't allowed, then we cannot presuppose anything about morality. We cannot, for example, assume that morality is about what's good for us, or that being in pain and falling off cliffs just aren't those sorts of things. Thus, absent moral assumptions, we must be agnostic about morality. We cannot know its interests or dictates.

This, initially at least, seems to be exactly what the debunker wants. After all, she says, morality could be about anything! It is conceptually possible, she argues, that morality is about throwing ourselves off cliffs and causing ourselves pain.¹³ If it had been, evolution would have still inclined us to think it was about survival, pain avoidance, and caring for our children. So, we cannot trust our judgment that morality is about these things.

That's how the debunker's argument is supposed to go. But this suggestion, that morality could be about anything, hurts rather than helps the debunker. Briefly, the trouble is this. If morality could be about anything, then we have no idea what morality is about. And if we have no idea what morality is about, then we cannot get good reason to think we are mistaken about morality. But this is exactly what the debunker needs to do: provide evidence of our error, or reason to think that our moral beliefs are likely to be mistaken. This reason must be a good reason by our own lights—a reason that shows us the error of our ways. But we cannot determine that we are likely to be mistaken about morality if we can make no assumptions at all about what morality is like.

It's true that, once we set all of our moral beliefs aside, we lack reason to think that we're *not* mistaken—that is, we lack reason to think our moral beliefs are *true*. But this doesn't help the debunker. The demand that we demonstrate we have such reason was never her challenge. We

¹³ As Street (2006) puts it, “the universe of logically possible evaluative judgements is huge...” (133).

aren't here required to *vindicate* our moral beliefs, as more traditional skeptics demand. The debunker presents a different kind of challenge, which aims to *undermine* our beliefs. The burden is therefore on her to show us we're mistaken—not on us to show her that we're not.

We can now see what's wrong with the debunker's argument. She wants to demonstrate that we have good reason to accept:

Mistaken. Our moral beliefs are likely to be mistaken.

To do that, she must provide evidence of error. This is the aim of her evolutionary story. But for this story to worry us, we must think that the off-track influence affecting our beliefs is a *bad* influence. And to think that we must also accept:

Gap. The true moral beliefs and the adaptive moral beliefs come apart.

Remember why we need this. If, contrary to Gap, the set of true moral beliefs and the set of adaptive moral beliefs overlap, we have reason to think that evolution would not have pushed us completely off track. If we know *where* they overlap, then we have at least some moral beliefs we can safely use to self-correct. This was horn one. On horn two, the debunker bans even the most uncontroversial moral assumptions, thereby making it impossible to identify an overlap between the two sets. She thus blocks our hope for self-correction. However, by doing this—by insisting that morality could be about anything—she makes it impossible for us to accept Gap. Why? Because if morality could be about anything, and we can make no assumptions about what morality is about, then we cannot hold that the true moral beliefs and the adaptive moral beliefs come apart. After all, to hold that is just to hold Gap—it is to say: morality is about more than what is adaptive.

The debunker could object, here, that this line of reasoning misses the point. Yes, it's true that if morality could be about anything, then there could be overlap between the true moral beliefs and the adaptive moral beliefs—overlap that dissolves the debunker's challenge. But we have no reason to think there is such an overlap. In the absence of such reason, our moral beliefs are in trouble.¹⁴

However, this response conflates the debunker's challenge with an importantly different one. To see this, notice what is implicit in that final sentence: our beliefs are in trouble absent reason to think they are right (true, reliably formed, etc.). This demand—for a non-question-begging vindication of our beliefs—is familiar from more standard skeptical arguments. These arguments assume that our beliefs are guilty until proven innocent—and that proof cannot assume anything that is in question. If this reasoning is right, then our moral beliefs are in trouble—but so are all of our beliefs, and the trouble has nothing to do with how we evolved. Thus, this skeptical challenge is neither empirical nor targeted, as the debunker's is meant to be. She, recall, aims to give you evidence of error. This skeptical argument does no such thing. At best, it makes salient the possibility of error. So, we should not conflate it with the debunker's.

¹⁴ There are other, perhaps better, ways for the debunker to respond, and I'll consider some shortly. I focus on this one here because responding to it allows me to make an important distinction that illuminates the both debunker's argument and the greater dialectic.

The debunker's challenge, in contrast, imposes no guilty until proven innocent requirement on our beliefs. It is compatible with the view that our beliefs are innocent until proven guilty—it's just that the debunker claims to have supplied proof of their guilt.

I have argued, however, that we cannot get such proof in a vacuum. We need to know something about morality in order to judge that we're likely to be mistaken about it. In other words, to accept Gap, and thus believe that moral truths do not coincide with the adaptive judgments, we must assume *something* about what those moral truths are, or what they are like. If we make no such assumptions, then morality really could be about anything. And if morality could be about anything, then we have no idea what morality is about. So we have no reason to think that the true moral beliefs and the adaptive ones don't overlap. But without this, we have no sense of what the chances are that we are mistaken. So, we can't get to the conclusion that they are.

The upshot is that if substantive moral assumptions are illegitimate, then I cannot accept Gap. But without Gap, I can't get evidence of my own error. However, we have to be careful here about what we mean by "substantive" moral assumptions. There are some non-trivial moral assumptions that don't seem question-begging in this context. For example, we can make formal claims about the nature of morality or moral truths: we can assume that morality is about right and wrong, good and bad; that 'right' and 'wrong' range over acts. We can assume, if we want, that morality is or is not attitude-independent. These assumptions are, in a sense, substantive. And they are clearly moral. However, none of them can get us to Gap. To believe that the moral truths and the adaptive truths come apart, I need to know more about the nature of the moral truths than their metaphysical or formal properties. I need something more specific, such as that morality is about what's good for us, and unnecessary pain or sacrificial cliff jumping just aren't those sorts of things. This is the sense in which these moral assumptions are not substantive: they're not enough to get us to Gap.

I can't prove there are no merely formal assumptions that can get Gap, because I can't prove a negative existential. But it makes sense that there wouldn't be: anything sufficiently non-committal compatible with the claim that morality could be about anything will not be sufficiently substantive to get us that the moral and adaptive truths come apart. To accept that premise, I need, at least, to accept that morality is about more than what is adaptive. But, the debunker argues, that's an illegitimate view in this context. (If it weren't, I could use it to argue my way out via the first horn.)

To sum up, then, this is the dilemma. Either we are allowed to rely on moral assumptions in responding to the debunker's challenge, or we are not. If we are, then we can tell a story about why evolutionary influences have not made us completely hopeless at recognizing and reasoning our way to the normative truths. If, instead, we aren't allowed to rely on any moral assumptions, we cannot tell such a story. But nor can we tell anything about whether evolution would have led us toward or away from the moral truths. We thus cannot conclude that we are likely to be mistaken about morality.¹⁵

¹⁵ There is at least one other way of defending the assumptions of the dilemma as well as the more general claim, which I'll defend later, about what assumptions must be in place for us to get evidence of error. Cuneo and Shafer-

5. Two worries dismissed.

I've argued that the debunker aims to provide evidence of error.¹⁶ I've also argued that we can't get evidence that we're mistaken about morality in the absence of substantive moral assumptions. If that's right, and if my interpretation of the debunker's argument is right, then the second horn of the dilemma goes through. I'll come back to the claim about evidence of error in section 6. Here, I'll defend my interpretation of the debunker's argument by considering two objections. Both attack horn two of the dilemma by arguing that the right interpretation of the debunker's challenge can get us to evidence of error without substantive moral assumptions. The first fails because its analogy fails; the second has more problems, the worst of which is that it collapses the debunker's argument into a more general skeptical one.

Implicit in these responses are just two of the many interpretations of the debunker's argument. I focus on them because they are familiar from the literature, and because they fail in telling and important ways. But more on this later; for now, consider the first objection.

5.1. First objection: the point is precisely that evolution is random!

The first objection begins with an example along the following lines:

Hat. I decide to investigate who the 19th US president is by randomly pulling a name out of a hat filled with the names of all 45 US presidents. I know nothing of US politics—historical or contemporary, so I can't eliminate even the most recent presidents. There is thus only a 1/45 (roughly 2%) chance that I'll come to the true belief about the 19th US president by this means.

It seems like I don't need to know anything about the US presidents to know that this method of belief formation will lead me astray. This, the objection goes, shows that I can get evidence of error in the absence of assumptions about the relevant matter. And, it continues, this is exactly how the debunker's argument is supposed to work: evolution is a random process, like picking names out of a hat.¹⁷

Landau (2014) argue for the existence of *moral fixed points*—substantive moral propositions that are also conceptual truths. These truths, they argue, “constitute any reasonably comprehensive moral system for beings such as ours in a world such as ours,” and “fix the boundaries of moral thought: one could not engage in competent moral thinking while rejecting them” (400-1). Elsewhere they say something stronger: that competence implies “a disposition to sincerely affirm the moral fixed points” (423). See also their comment about error theorists (438). The sorts of assumptions we need for the first horn of the dilemma are strong candidates for moral fixed points. If that's right, then we can defend the assumptions we need in the first horn thus: how could it be question-begging to assume that without which we wouldn't even be talking about morality? We could also defend the second horn thus: if these assumptions are conceptual truths, then we cannot be agnostic about them while maintaining our concept of morality. The moral fixed points view can thus help explain why learning that we are mistaken about morality requires taking for granted specific assumptions about morality's dictates.

¹⁶ Cf. Shafer-Landau (2017) and Vavova (2015).

¹⁷ This is my example, but debunkers have provided similar examples in the literature. Cf. Street's discussion of hypnosis on p. 686 of her (2015) and Joyce's belief pill example on p. 179 of his (2006).

The objector is right that, in Hat, I should think that I am likely to be mistaken about the 19th president. This is precisely because I know that the process by which I came to that belief is random, and, by my own lights, a random process is not a reliable one. The objector is also right that this is how the debunker wants us to view the evolutionary process: as random, or unreliable, with respect to the moral truth. (And, this is in fact how we view it. But only, I'll argue shortly, because we're not agnostic about morality.)

Where the objection goes wrong is in claiming that we are as agnostic about the presidential truths in Hat, as the debunker demands we be about the moral truths in horn two. Cases like Hat are worrying because, and only because, we already think that there is *no connection* between who the 19th US president is and which piece of paper I pull. We don't think, for example, that there is some magnetic force drawing my hand to the paper with the right answer. Nor do we think that the presidential truth is somehow determined by what happens inside my hat (if only!). We think I'm equally likely to pull *any* president's name. This is why the objector is wrong: we cannot get reason to think that a given process is random or unreliable with respect to some matter from a standpoint that is agnostic about that matter. Even in Hat, we know quite a lot about the nature of the presidential truths.

To make the cases analogous, we'd have to assume that there is *no connection* between the true moral beliefs and the adaptive ones—that evolution is as likely as chance to pick out true moral beliefs as false ones. So, it'd be an accident, a *coincidence*, if it did. The debunker wants us to view the evolutionary process as random with respect to the moral truth, so, *no connection* is exactly what she needs. But the debunker can't simply posit it—that would be question begging. And to defend it, she needs to make substantive moral assumptions.

Consider this second point first. To argue for *no connection* we'd have to say: "*this*... is the deal with the true moral beliefs, *that*... is the deal with the adaptive moral beliefs. Notice: *this* and *that* are totally different!" In other words, we need to argue that the true moral beliefs and the adaptive ones come apart. This was Gap. But accepting Gap requires substantive moral assumptions. So, the debunker must argue for *no connection*, but she cannot do so without moral assumptions. Contrary to the objector, then, absent such assumptions, I can't get reason to think that the evolutionary process is random in the same way that pulling a name out of a hat is random.

Back to the first point: that it'd be question-begging for the debunker to simply posit *no connection*, and, therefore, that evolution is a random process. To see this, recall the dialectic. The debunker doesn't merely challenge us to provide a moral epistemology: a story about how we could possibly know, or even access, these weird moral facts. Many before her have made *this* demand. The debunker, instead, aims to give us good, empirical evidence that we couldn't know or access these things. To accomplish this, she has to show that the process that caused our beliefs is, at best, random with respect to the truth. She cannot just posit this: it is a crucial premise of her argument, for which she claims to be providing evidence. Furthermore, to establish this premise, as I've argued, the debunker needs to also grant some substantive moral assumptions. And, as we saw in the first horn, once we have these assumptions, we can use them to self-correct.

To sum up: the debunker *does* need the evolutionary case to be like Hat, so that the process that led to our beliefs is random or unreliable. But she needs to argue for this, and doing so requires making substantive moral assumptions of the sort that are either illegitimate or exactly what we need in order to be able to self-correct.

5.2. Second objection: the point is that the chances of being right are so low.

This objection interprets the debunker's argument somewhat differently—although the focus is still on what a coincidence it would be if our moral beliefs were true. On this view, the debunker's argument is best understood as follows:

We acquire moral beliefs through either causes or reasoning (with beliefs that were caused). There is no third way of getting moral beliefs. Ideally, the moral facts would be doing the causing, but they are causally inert, so that isn't possible. Now, whatever the actual causes of our moral beliefs are—be they evolutionary, social, perceptual, or whatever—they need to get us to the moral truth. There's no reason to think these causes would be reliable. And given how many ways my moral beliefs might go, it seems, the chance of our getting at the true ones is infinitely low. So, it'd be a massive coincidence if our beliefs are true. So, we're likely to be mistaken about morality.¹⁸

On the face of it, this argument seems to give us evidence of our own error without making any substantive moral assumptions. I'll now argue that this is not the argument that the debunker should go for—because of its scope, because it's not empirical, and because of the epistemic principle it presupposes. In short: this is not an evolutionary argument, nor does it provide evidence of error.

The first thing to note about this argument is that it assumes that the only way we acquire moral beliefs is through causal processes and reasoning from beliefs acquired by these processes. If we have some other way of getting at the moral truths, then the conclusion doesn't follow. Simply positing this, we might worry, seems question-begging. Presumably, those who hold this view think we have some way of getting true beliefs about the causally inert facts. True, we might want such a story—a moral epistemology that explains how we could come to believe these weird moral facts. But the opponent can't simply assume there is no such story. To give us reason to think we are mistaken, she must do more than this.

Second, not everyone agrees that the moral facts are causally inert. Even among moral realists, only non-naturalists accept this. So, this version of the argument is more limited than the debunker intended. She wanted to take out all the moral realists, and perhaps some others. This argument can't help her do that.

¹⁸ Versions of this argument in the literature include Bedke (2014) and Street (2016) 305. There are differences between their arguments, and the one I present here is more general than either. But all three contain the same moving parts, namely: a claim about a causal influence, a claim about the moral truths or facts, a claim about a low probability or coincidence, and a skeptical conclusion.

Third, the argument is more far reaching than the original: it targets not only our moral beliefs but also our beliefs about any other causally inert facts. So, if we think that mathematical facts or facts about epistemic normativity or reasons are causally inert, and this argument works, then those beliefs are in trouble too. In this way, the argument is much more ambitious than the debunker wanted.¹⁹ It resembles, both in its scope and assumptions, other, familiar arguments against our moral and mathematical beliefs.²⁰ These arguments demand an epistemology (moral or mathematical). They don't claim to provide empirical evidence of error, as the debunker does. If her argument is just an instance of this familiar argumentative strategy, then it may well have skeptical bite, but not for the novel reasons we thought. The debunker should thus resist this way of formulating the argument.

The fourth point is related: this is not really an evolutionary argument. One story we might tell about how our moral beliefs are caused is an evolutionary one. But it needn't be. The argument works, if it does, without a supplemental story. It is supposed to be enough, to get the worry going, that the moral beliefs are caused and the facts are inert. It doesn't matter what the causes are. (Crucially, the causes could be even seemingly legitimate stuff like visual inputs.) On the one hand, that's unsurprising. We know there are other ways of raising worries about our moral beliefs—ways that don't make empirical or evolutionary assumptions. On the other, raising those other worries was not what the debunker was purportedly doing. So, again, she should avoid this version of the argument.

Finally, let's think about what gets us to the conclusion that the chance of getting at the right moral beliefs is low. Notice, it's not that there's been an undue influence on our moral beliefs. This argument doesn't posit such an influence and cannot simply stipulate that causal influences are bad influences. (That'd be question-begging.²¹) Maybe then we get to the conclusion that the chance of getting at our moral beliefs is low because we don't know how these causal processes that give us our beliefs work—we don't know whether they are truth-conducive or not. But this, again, is to pose the wrong kind of challenge. Arguments whose aim is to show us that we're likely to be mistaken don't work by challenging us to prove, in a non-question-begging way, that we're likely to be right. That's what more standard skeptical arguments do. Rather, these more sensible skeptics are, as I argued, like concerned optometrists who aim to give you good reason to think that you are likely to be mistaken.

The difference here is between two sorts of epistemological principles, which disagree about what it takes for our beliefs to be in trouble. The first says that our beliefs are in trouble just in case we *lack* good reason to think they're *not mistaken*. The only way to maintain belief, if this is the right principle, is to show that we have good reason for thinking that our beliefs are correct, reliably formed, or etc.²² This is the sort of principle the standard skeptic uses when arguing that

¹⁹ Some debunkers might be okay with this (e.g., Street (2009)), but many won't be.

²⁰ See e.g., Benacerraf (1973) and Field (1989). As I put it elsewhere, "The parallels are worth exploring (Cf. Clarke-Doane (2012)). But that there are parallels at all should raise suspicions" (Vavova, 2015, fn. 22)

²¹ Many realists grant that the influence is causal, but they still think it indirectly gets us to the moral truths. E.g., Enoch (2010) and (2011) and Wielenberg (2010).

²² In my (2014) I call this principle "No Good". It is similar to Christensen's principle B in his (2011). Street (2015) endorses a principle like this but does not discuss its skeptical implications.

since we have no non-question-begging way to prove that, e.g., we're not dreaming, we cannot believe that we're not.²³ This is not the principle that the debunker should use.

The correct epistemic principle for the debunker says that our beliefs are in trouble when we *have* good reason to think they *are mistaken*. One needn't *have* reason to think one is likely to be *right*, to maintain belief according to this principle. It is enough to *lack* reason to think one is *wrong*. We can see that this is the right principle for the debunker because it reflects the fact that the burden isn't on us to show we're likely to be right, but on the debunker to show that we are likely to be wrong.²⁴

Perhaps the problem arises, according to the debunker, because of what she says at the end of her argument: that there are infinitely many possible coherent belief systems and only one right one. That does seem worrying. This move relies on nothing but probabilities. It is, in that sense, not an empirical or targeted argument.²⁵ It is also a variation of the standard skeptical move (which I briefly canvassed in section 4): show me you're not massively mistaken about this and do it without assuming anything I'm calling into question. Again, this relies on the wrong sort of principle, and leads the debunker's argument to overgeneralize beyond the moral realm.²⁶

At best, then, this is a non-evolutionary argument against our moral beliefs that attacks only a subset of realist views. At worst, it's the same old skeptical argument in disguise.

5.3. The upshot.

Intuitively, it seems that if evolution shaped our moral beliefs, then those beliefs are probably wrong. And if our beliefs are probably wrong, then we shouldn't believe them. This, in a brief, compelling nutshell, is the debunker's argument.

The trouble starts when we try to understand those conditionals. Why does evolution shaping these beliefs mean that they are probably wrong? And why does something being probably wrong mean we shouldn't believe it?

In section two, I presented a valid reconstruction of the debunker's argument from which we can glean clear, explicit answers to these questions. I think it is a formidable version of the argument. Some of its assumptions are controversial, but they are all plausible. It also meets some important constraints: it is a targeted, epistemological, and empirical argument.

Of course, I think the argument fails—or so I have argued.

²³ Cf. Elga (2008) p. 7 and Christensen (2011) p.15 on this point.

²⁴ To say that this is the right principle for the debunker isn't to take a stand on the truth or plausibility of the other principle. It is just to say that this principle captures the more ordinary sort of error the debunker thinks we're likely to succumb to, whereas the other captures the more extraordinary sort of error the skeptic is pushing.

²⁵ It assumes something like the plausible but controversial Principle of Indifference. Even setting aside the controversy, an argument that employs this principle doesn't need anything empirical to get going, so it won't be the right kind of argument for the debunker's purposes.

²⁶ In defense of this last claim see Berker (2015) section 7, Shafer-Landau (2012) 13-14, and Vavova (2015) section 3.2.

It is tempting to respond that there is some other, better way to understand the debunker's argument. Indeed, this is the response I've most often received—and, along with it, a plethora of suggestions for how the argument should go. The details of these various suggestions differ, but there is one common thread: all of them try to establish moral skepticism via a claim about how improbable, random, coincidental, or the like, it would be if our moral beliefs were true. This is why I have focused on the above two suggestions.

Dialectically, these responses start by trying to show that we don't need substantive moral assumptions to get evidence of error. In doing so, they present two ways of running the debunker's argument via a claim about improbability, randomness, or coincidence. The first can't establish this claim without the sorts of moral assumptions that would pave our escape route, at least. The second establishes the claim, but is of no use to the debunker, since it isn't not empirical and does not give us reason to think our moral beliefs are mistaken.

Suppose, however, that I have been too quick to dismiss these suggestions. Suppose that the debunker *can* establish the low probability or coincidence claims. This doesn't mean the game is over. The debunker cannot win merely by showing that it would be a coincidence if our beliefs were true. To establish that something is a coincidence or improbable is not yet to establish anything epistemologically problematic. The debunker would then need to either explain what the relevant, epistemologically problematic sense of coincidence is, or explain how we get from an epistemologically innocent sense of coincidence to moral skepticism. I think that to do the first, she would have to use substantive moral assumptions; to do the second, she would need to rely on the wrong kind of epistemic principle. So, we would be back where we were when we started this section.

My suspicion is that most ways of running the debunker's argument succumb to one of these two problems. I cannot prove that here, so I must leave it for the reader to investigate. At any rate, the gauntlet has been thrown. If there is a better way to run the argument, it is time for the debunker to step up and give us a clear, valid reconstruction of it.

To sum up, I've reviewed some of the many ways that we might run the debunker's argument. There might be some other way of running the argument that both doesn't rely on moral assumptions and concludes that we are likely to be mistaken. But I am not optimistic. There seems to be a fundamental problem here: spinning this argument so that it doesn't need substantive moral assumptions seems to make it either not empirical, not evolutionary, or insufficiently targeted. But those are supposed to be the strengths of the debunker's argument, which set it apart from other familiar arguments for moral skepticism. If it is to be a novel challenge, and if it is to have these strengths, then the debunker's argument is stuck with my dilemma.

6. The lesson of the dilemma.

Recall the dilemma: either moral assumptions are legitimate, or they are not. If they are legitimate, then we don't have reason to think we're massively mistaken and where we are

mistaken we can self-correct. If they are illegitimate, then we have no reason to think we're mistaken: we cannot rationally get any.

The upshot of the dilemma is, of course, that the debunker's argument fails. But how, exactly, does it fail? It's not invalid, and we haven't identified a false premise. So, what have we done? We have shown that we cannot rationally get from the debunker's premises to a problematic version of her conclusion. Thus, what the dilemma shows is less about the details of the debunker's argument and more about the limits of rational belief revision. This is why the dilemma doesn't identify a false premise or show that the conclusion doesn't follow. Rather, it shows that we cannot *both* accept that morality could be about anything and accept that we're likely to be mistaken about morality. That is, we cannot accept Mistaken while remaining agnostic about Gap. Accepting Gap, along with the other premises, *can* get us to Mistaken, but not to a problematic version of it. After all, if we accept Gap, then we're not clueless about morality, and we can use what we know to self-correct. There is thus no stable position from which we can judge that our epistemic situation is hopeless with respect to morality. And that, ultimately, is what the debunker needs.

The debunker's failure suggests a more general conclusion: that we cannot get evidence of our own error about some domain in the absence of substantive assumptions about the truths in that domain: what they (at least some of them) are, or what they are like. Why should we think this is true? Consider an analogy.

You present me with two sealed boxes and ask, "What are the chances that these boxes contain the same sorts of things?" In other words, what are the chances that both boxes contain, say, puppies, staplers, or ice cubes. The less I know about the sorts of things the boxes contain, the harder it will be for me to answer this. It'd help me to come to a promising answer, if you'd agree to play Twenty Questions, so I could ask, "Do both boxes contain fluffy things? Sharp things? Cold things?" That would help me determine whether these boxes have properties in common. If they don't, their contents must differ. If they have some properties in common, there is some chance the contents are the same; if they have many, the chance is greater. Suppose, however, that you refuse to play the game. In fact, you refuse to tell me anything at all about the boxes' contents. In that case, I have no idea what sorts of things each box contains, and I have no idea what the chances of overlap are. I am rightly at a loss.

The lesson is this. To judge that two sets do not have the same contents, I need to know *something* about the contents of those sets—what they are or what they are like. This is why I cannot accept Gap if I am not allowed any assumptions about what the moral truths are or what they are like.

We can support this idea further with some general considerations about evidence. The epistemic significance of a piece of evidence—how, if at all, it should lead you to alter your belief—cannot be determined in a vacuum. Consider this bit of evidence: the sound of water drops on my office window. This could be evidence of rain. Or, it could be evidence that the gardener is trying to

flood my office again. The correct way to interpret the evidence depends on my background: on what the rest of my evidence supports.²⁷

Undermining evidence, or evidence of error of the sort the debunker aims to provide, works this way too. A bit of evidence cannot be undermining without the right sort of background in place. In this case, evidence of a belief influence is only undermining if you think that influence was a bad influence. One reason you might have for thinking this is that you know that the influence pushed you away from the truth—but to know *that* you have to have some conception of what the truth is. This underscores why, on horn two, we can receive no evidence of error without moral assumptions.

Again, none of this shows that we have good reason to think our beliefs are *not* mistaken—we haven't vindicated our moral beliefs. Rather, we've blocked one kind of attack. (Why ignore the other? First, one thing at a time. Second, remember, it is a different challenge, a more general skepticism.) It does, however, show the way in which debunking arguments like this are inherently limited. They cannot debunk most or all of what we believe. This is because they aim to give us evidence of error—and we cannot rationally get such evidence except against the background of substantive epistemological and, in this case, moral assumptions. This teaches us something more general about the possibility of getting evidence of error and thus, about rational undermining. I cannot get such evidence without knowing what would count as error in this context. And knowing that requires knowing much more than we might have antecedently expected.

7. Conclusion.²⁸

Debunkers and opponents are at an impasse: they cannot agree on the rules of the game. I have argued that whatever these rules, the evolutionary debunker's attempt to undermine our moral beliefs fails. It fails either because we have hope of self-correction, or because we get no evidence of error.

This is an important result for the evolutionary debunking debate. But it is also important for the greater battle between science and morality. This time, at least, the two are not in conflict. Or, at least, we have no reason to think that they are.

²⁷ This example comes from my (2014) p. 97. There, I used insights from this example and the Independence Principle to argue for the Inverse Rule of Debunking, which says that the potential strength of a debunking argument is inversely proportional to its ambition. That is, the more of our beliefs the debunker tries to debunk, the more likely she is to fail. What I have argued here can be used to vindicate this point, but it is a more general lesson that isn't just about debunking and which does not invoke the controversial independence requirement.

²⁸ For helpful comments and conversations, thanks to David J. Barnett, John Bengson, Sarah Buss, David Copp, Terence Cuneo, Catherine Elgin, Nina Emery, Branden Fitelson, Alexander George, Elizabeth Harman, James Harold, Sophie Horowitz, Don Loeb, Sarah McGrath, Sam Mitchell, Luis Oliveira, Alejandro Pérez Carballo, Gupreet Rattan, Simon Rippon, Said Saillant, Nishi Shah, Matthew Silverstein, Paulina Sliwa, Mark van Roojen, Kenny Walden, Mason Westfall, Joeri Witteveen, and a reviewer for this journal. Thanks also to audiences at Brandeis University, Cornell University, the Institute Jean Nicod, the Massachusetts Institute of Technology, Princeton University, the University of Michigan, UNC Chapel Hill, the University of Texas-Austin, the University of Toronto, Upsalla University, Vrije Universiteit Amsterdam, and Washington University in Saint Louis.

Finally, this result is important for what it teaches us about evidence of error and rational belief revision. Awareness of our fallible nature should make us humble and open to evidence of our own error. But there are limits to how much we can give up, even from a maximally humble stance. The failure of the debunker's argument teaches us that getting reason to think that we are mistaken is not trivial and we cannot get such evidence about all of our beliefs at once. There are, therefore, limits to rational belief revision.

These limits constrain the debunker. Her noble aim—noble when compared to the standard skeptic—is to provide us with reasons that are good reasons *by our own lights*. She doesn't trade in mere possibilities of error. She aims to provide us with evidence of error that we can recognize as such. This is a difficult task she's set for herself. Unfortunately, she has failed.

References.

- Bedke, M. (2014). No Coincidence. *Oxford Studies in Metaethics*, 9, 102-125. doi: 10.1093/acprof:oso/9780198709299.003.0005
- (2009). Intuitive Non-Naturalism Meets Cosmic Coincidence, *Pacific Philosophical Quarterly*, 90, 188-209. doi: 10.1111/j.1468-0114.2009.01336.x
- Berker, S. (2015). Does Evolutionary Psychology Show that Normativity is Mind-Dependent? In J. D'Arms & D. Jacobson (Eds.), *Moral Psychology and Human Agency: Essays on the Science of Ethics*, Oxford University Press: 215-252. doi: 10.1093/acprof:oso/9780198717812.003.0010
- Benacerraf, P. (1973). Mathematical Truth. *Journal of Philosophy*, 70, 661-679. doi: 10.2307/2025075
- Brosnan, K. (2011). Do the Evolutionary Origins of our Moral Beliefs Undermine Moral Knowledge? *Biology and Philosophy*, 26, 51-64. doi: 10.1007/s10539-010-9235-1
- Christensen, D. (2011). Disagreement, question-begging and epistemic self-criticism. *Philosopher's Imprint*, 1(6), 1-22.
- Clarke-Doane, J. (2012). Morality and Mathematics: The Evolutionary Challenge. *Ethics*, 122(2), 313-340. doi: 10.1086/663231
- Copp, D. (2008). Darwinian Skepticism about Moral Realism. *Philosophical Issues*, 18, 186-206. doi: 10.1111/j.1533-6077.2008.00144.x
- Cuneo, T. and Shafer-Landau, R (2014). The Moral Fixed Points: New Directions for Moral Nonnaturalism. *Philosophical Studies*, 171, 399-443. doi: 10.1007/s11098-013-0277-5
- Enoch, D. (2011). *Taking Morality Seriously*. Oxford University Press.
- (2010). The Epistemological Challenge to Metanormative Realism: How Best to Understand It, and How to Cope with It. *Philosophical Studies*, 148(3), 413-38. doi: 10.1007/s11098-009-9333-6
- Field, H. (1988). Realism, mathematics and modality. *Philosophical Topics*, 16(1): 57-107. doi: 10.5840/philtopics19881613
- Foot, P. (2001). *Natural Goodness*. Oxford University Press. doi: 10.1093/mind/112.445.130
- Greene, J. (2008). The Secret Joke of Kant's Soul. W. Sinnott-Armstrong (Ed.), *Moral Psychology 3*: 35-79.

- Hanson, L. (2017). The Real Problem with Evolutionary Debunking Arguments. *Philosophical Quarterly*, 67(268), 508-33. doi: 10.1093/pq/pqw075
- Joyce, R. (2006). *The Evolution of Morality*. Cambridge, MA: MIT Press. doi: 10.1111/j.1467-9973.2008.00572.x
- Kahane, G. (2011). Evolutionary Debunking Arguments. *Nous*, 14(1), 103-25. doi: /10.1111/j.1468-0068.2010.00770.x
- Kitcher, P. (2007). Biology and Ethics. D. Copp (Ed.), *The Oxford Handbook of Ethical Theory*, Oxford: Oxford University Press: 163-185. doi: 10.1093/0195147790.003.0007
- Mogensen, A. (2016). Do Evolutionary Debunking Arguments Rest on a Mistake about Evolutionary Explanations. *Philosophical Studies*, 173(7), 1799-1817. doi: 10.1007/s11098-015-0579-x
- Ruse, M. (1998). *Taking Darwin Seriously*. Oxford: Blackwell Publishing.
- Ruse, M. and E. O. Wilson (1986). Moral Philosophy as Applied Science. *Philosophy*, 61, 173-92. doi: 10.1017/S0031819100021057
- Sarksaune, K. O. (2011). Darwin and Moral Realism: Survival of the Iffiest. *Philosophical Studies*, 152, 229-43. doi: 10.1007/s11098-009-9473-8
- Shafer-Landau, R. (2017). Moral Realism and Evolutionary Debunking Arguments. M. Ruse and R. Richards (Eds.), *The Cambridge Companion to Evolution and Ethics* (Cambridge Handbooks in Philosophy, pp. 175-187). Cambridge: Cambridge University Press. doi: 10.1017/9781316459409.013
- (2012). Evolutionary Debunking, Moral Realism and Moral Knowledge”. *Journal of Ethics and Social Philosophy*, 7(1), 1-37. doi: 10.26556/jesp.v7i1.68
- (2005). *Moral Realism: A Defence*. Oxford University Press. doi: 10.1215/00318108-2006-017
- Street, S. (2016). Objectivity and Truth: You’d Better Rethink It. R. Shafer-Landau (Ed.), *Oxford Studies in Metaethics*, 11, 293-333. doi: 10.1093/acprof:oso/9780198784647.003.0012
- (2015). Does Anything Really Matter or Did We Just Evolve to Think So? A. Byrne, J. Cohen, G. Rosen and S. Shiffrin (Eds.), *The Norton Introduction to Philosophy*, W. W. Norton, New York: 685-695.
- (2009). Evolution and the Normativity of Epistemic Reasons. *Canadian Journal of Philosophy*, 39, supplement 1, 213-48. doi: 10.1080/00455091.2009.10717649

- (2006). A Darwinian Dilemma for Realist Theories of Value. *Philosophical Studies* 127(1), 109-66. doi: 10.1007/s11098-005-1726-6
- Vavova, K. (2018). Irrelevant Influences. *Philosophy and Phenomenological Research*, 96(1), 134-152. doi: 10.1111/phpr.12297
- (2015). Evolutionary Debunking of Moral Realism, *Philosophy Compass*, 10(2), 104-116. doi: 10.1111/phc3.12194
- (2014). Debunking Evolutionary Debunking. *Oxford Studies in Metaethics*, R. Shafer-Landau (Ed.) 9, 76-101. doi: 10.1093/acprof:oso/9780198709299.003.0004
- White, R. (2010). You just Believe that Because.... *Philosophical Perspectives*, 24(1), 573-615. doi: 10.1111/j.1520-8583.2010.00204.x
- Wielenberg, E. (2010). On the evolutionary debunking of morality. *Ethics*, 120(3), 441-464. doi: 10.1086/652292