

Debunking Evolutionary Debunking

Katia Vavova
Mount Holyoke College

1. The evolutionary challenge.

Worries about the compatibility of evolution and morality are not new—even Darwin had them. A number of recent arguments revive these concerns. These *evolutionary debunking arguments* take the following form: you just believe what you do because you evolved to, therefore you're not justified in believing what you do. They typically target evaluative realism: the view that evaluative facts are attitude-independent—that what is valuable is valuable whether or not we happen to value it.¹

The worry is that just as evolutionary forces shaped our eyes and ears, so they shaped our evaluative attitudes. But, the debunker argues, we have no reason to think that these forces would track the attitude-independent evaluative truths that the realist posits.² Worse yet, we seem to have a good reason to think that they wouldn't: evolution selects for characteristics that increase genetic fitness—not ones that correlate with the evaluative truth. Plausibly, the attitudes and judgments that increase a creature's fitness come apart from the true evaluative beliefs. If this is so, then it seems that evolutionary forces have had a distorting effect on our evaluative attitudes. The debunker concludes, insofar as we are realists and insofar as the evolutionary facts are thus-and-so, we are not justified in our evaluative beliefs.

Evolutionary debunking arguments are sometimes meant to establish just this: evaluative skepticism. Other times the skeptical conclusion is in the service of the greater goal of undermining evaluative realism. In either case, the debunker must first establish that learning about the evolutionary origin of our evaluative beliefs gives us, qua realists, good reason to worry about our evaluative beliefs. I will argue that the considerations she puts forth cannot give us such reason. I will conclude that there is little hope for distinctly evolutionary debunking arguments. This is bad news for the debunker who hoped that the cold, hard scientific facts about our origins would undermine our evaluative beliefs.

2. The Debunker's Argument.

¹ This understanding of realism follows the evolutionary debunking literature. Similar definitions can be found in metaethics more generally (see Shafer-Landau [2005] 15 on 'stance-independence'). For present purposes, evaluative propositions are of the form: that X is a normative reason to Y, that one should or ought to X, that X is good, valuable or worthwhile, that X is morally right or wrong, and so on. Evaluative attitudes include (conscious or unconscious) beliefs in evaluative propositions, as well "as desires, attitudes of approval and disapproval, unreflective [...] tendencies such as the tendency to experience X as counting in favor of or demanding Y," etc. (Street [2006] 110).

² From here on I'll drop the 'attitude-independent' qualifier on evaluative attitudes or truths.

“[T]here can hardly be a doubt,” Darwin speculated, that if we had evolved under the same conditions as hive-bees, “our unmarried females would [...] 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” ([1871] 73). If instead we had evolved as lions did, Street argues, males would have “a strong unreflective evaluative tendency to experience the killing of [other’s] offspring [...] as ‘demanded by the circumstances.’” Not only would females lack an “unreflective tendency to ‘hold it against’ a male when he killed her offspring,” but would tend to become “receptive to his advances soon afterwards” (121).

These observations are meant to support this counterfactual: if we had evolved differently, we would have believed differently—our evaluative beliefs, in particular, would have been different. In turn, this counterfactual is meant to support the claim that the content of human evaluative judgments has been “tremendously influenced [...] by the forces of natural selection” (Street [2006] 121).

The debunker hopes to use this story to undermine our evaluative beliefs. We cannot rationally maintain our opinions about good and bad, right and wrong, reasons and values, she argues, once we realize from where they came. The debunker thus aims to somehow get from

INFLUENCE. Evolutionary forces have influenced our evaluative beliefs.

to

REVISION. We cannot rationally maintain our evaluative beliefs.³

To be sure, INFLUENCE is not equally worrying for everyone. Antirealists take the evaluative truths to be attitude-dependent—somehow a function of our (actual, ideally rational, etc.) beliefs and desires. Since antirealists hold that our values determine what is valuable, they needn’t worry from where those values came. Realists are more vulnerable. Since they take the evaluative truths to be independent of our beliefs and desires, they are committed to the possibility of evaluative error: what we value and what is valuable can come apart.

Some varieties of realism are importantly different and may be better placed to dismiss the debunker. I won’t explore that here. First, understanding the debunker’s challenge doesn’t require digging into the details of realism. The evolutionary story is at least initially worrying for anyone who holds that the true evaluative beliefs come apart from the adaptive evaluative beliefs.⁴

³ Street doesn’t say that we cannot rationally maintain belief, but rather that we “should suspend belief” ([forthcoming] 2). I think we mean the same thing here. Presumably Street’s ‘should’ is an epistemic one, but that doesn’t narrow it down much. What we ‘should’ believe could depend on what we actually believe, what our evidence supports, etc. Here it won’t matter exactly how we understand this ‘should’ or the relevant notion of ‘rational’ because Street’s argument proceeds by first trying to establish a lemma that I will argue she cannot.

⁴ This is in contrast with the claim that the challenge is best understood as aimed at non-naturalist or non-reductive realists (e.g., Bedke [ms.] 1). The challenge may be more formidable for this particular variety of realism, but a more minimal commitment suffices to get it going.

Second, since I will present structural problems with the debunker's challenge, my strategy should be one that realists of any stripe may deploy in self-defense.

Let us grant then that some form of evaluative realism is the target, and assume for the sake of argument that the true evaluative beliefs come apart from the adaptive evaluative beliefs. Given this much, the question is how to get from INFLUENCE to REVISION. To seal this gap, we need to know what is the epistemic significance of the evolutionary story for our evaluative beliefs.

In the next sections, I will consider two ways of filling in the debunker's story.⁵ I will extract valid arguments to REVISION from both. The first, which Street suggests, is compelling, but too strong for the debunker's purposes. It collapses her challenge into a more general skeptical challenge. The second is more promising and the right way to understand distinctly empirical debunking arguments. But consider the first suggestion first.

3. Do we have good reason to think we're right?

The evolutionary debunker claims that in some sense of 'evolved' and in some sense of 'belief', we evolved to hold our evaluative beliefs. The thought is that just as "creatures inveterately wrong in their inductions have a pathetic but praiseworthy tendency to die before reproducing their kind" (Quine [1969] 126), so creatures with deep-rooted inclinations to kill themselves and their offspring tend to have quite short evolutionary histories. Given that different evaluative tendencies can have "extremely different effects on a creature's chances of survival and reproduction," we should expect "over the course of our evolutionary history, relentless selective pressure on the content of our evaluative judgments" (Street [2006] 114).

This is the evolutionary story. The debunker doesn't suggest, implausibly, that evolution directly shaped our more sophisticated evaluative beliefs. The evolutionary story is meant to undermine directly only more basic and less controversial beliefs,⁶ like the belief that the fact that something would promote one's survival is a reason in favor of it, or that we have greater obligations to help our own children than complete strangers. But the evolutionary story is also meant to undermine indirectly the rest of our evaluative beliefs, including our much more sophisticated judgments. If our belief that we have reason to avoid inflicting unnecessary suffering goes, so does the moral theory that rests, partly, on it. Hence, the debunker concludes: "our system of evaluative judgments is thoroughly saturated with evolutionary influence" (Street [2006] 114).

This is the empirical claim. No one, not even the debunker, thinks it conclusive.⁷ So, why take it seriously? Because the philosophically interesting question is not whether some empirical claim is true, but what follows about the rationality of our beliefs if something like it were true. This

⁵ There is textual evidence for both readings, though I do not know of others who distinguish them. For the first see Street's [ms.] and [forthcoming]; for the second see Street's talk of distorting influences in her [2006].

⁶ Or some sort of proto-belief states or tendencies (Street [2006] 115).

⁷ Cf. Street [2006] §3. For reasons to think that the evolutionary story is "a long way from even beginning to fill out the empirical details needed to fully secure" these premises see the just quoted Kahane ([2011] 111), Sliwa [ms.], and FitzPatrick [forthcoming].

question has implications for our epistemology and our metaethics, but it is also of practical interest. Even if the evolutionary debunker fails, some of our other beliefs might reflect some other suspect influence. We need to know how to respond to such evidence if, or when, we do get it.

Grant the evolutionary story for argument's sake. Why should it worry us? Because if it is true, the debunker argues, then the best explanation for why we hold the evaluative judgments we do is that they are adaptive.⁸ And this explanation is epistemically unflattering: that we evolved to hold a judgment is no reason to think that it is true.

The debunker then asks: knowing just about the evolutionary origin of our evaluative beliefs nothing else, do we have reason to think that those beliefs are true? We know that, by hypothesis, evolution selects for adaptive beliefs regardless of their truth. So it may be that the evaluative beliefs we should hold are such-and-such, but that the ones we do hold are this-and-that, because the latter are adaptive and the former aren't. Our evaluative beliefs may, then, be massively mistaken and our origin story gives us no reason to think that they are not.⁹

This is Street's suggestion. Since we evolved to hold our evaluative beliefs, we have no reason to think they are true. Rationality requires we have good reasons for thinking our beliefs are true. So we cannot rationally maintain our evaluative beliefs. Skepticism follows.

This version of the debunker's story relies a principle like this:

NO GOOD. If you have no good reason to think that your belief is true, then you cannot rationally maintain it.

Street explicitly endorses this kind of principle. She argues that it captures the difference between being hypnotized to believe that Hayes was the twentieth US president and learning it in school ([forthcoming] 2). In the former case you have no reason to think that the process by which you gained your belief would have lead you to form true beliefs. We don't typically think that magicians use their powers of hypnosis for good—e.g., to implant in their victims true beliefs about US history. Competent high school teachers, on the other hand, are concerned with just this task. The explanation of your historical beliefs in terms of hypnosis is thus undermining; the one in terms of education is vindicating.

⁸ Cf. Street [2006] on the adaptive link account.

⁹ There are two relevant ways of understanding 'mistake' here. On the first, a belief is mistaken just in case it is false. On the second, a belief is mistaken just in case it is not supported by the believer's evidence. What sort of mistake does the debunker point to? That's for her to say. I will follow much of the literature and focus on the first. This mostly won't matter for my purposes, but I will make a note when it does.

Street argues that evolution is more akin to a careless hypnotist than a teacher.¹⁰ We have no good reason to think that selective pressures would push us toward the truth. Learning about the influence of evolutionary forces on our evaluative beliefs should thus undermine those beliefs.

Many have found this puzzling, insisting that we have plenty of good reasons to think our evaluative beliefs are true. Even if evolution caused us to believe that “pain and injury are bad, and that we have strong reasons to promote the survival and well-being of ourselves and our children,” Parfit writes, “these beliefs are not badly mistaken, but correspond to some of the independent normative truths. Pain is bad, and we do have strong reasons to promote the survival and well-being of ourselves and our children” ([2011] 533). Discussing an analogous case, Dworkin wonders what the fuss is about. Why shouldn’t we, he writes, “count it as a piece of luck—a special example of what Bernard Williams has called moral luck [that our adaptive beliefs and the true ones] here coincide? ([1996] 125)

Other defenses of realism begin with similarly substantive moral assumptions: that pain is bad, that survival is good, that we have rights, and so on.¹¹ Street argues, however, that such assumptions are illegitimate in this context. To presuppose the truth of particular evaluative judgments is to presuppose exactly what the evolutionary story is meant bring under scrutiny. This is “trivially question-begging”, Street argues. Our reasons for thinking that our judgments are true cannot simply assume “the very thing called into question” the truth of those judgments (Street [Ms.] 15-6).

Whatever we think of the best version of this response, we should grant that there is something *prima facie* fishy about it. This is most evident in Dworkin. He begins by granting that evolution has been a suspicious, epistemically no-good influence on our evaluative beliefs. He then insists that we happened to have gotten things right. After all, we believe we have reason to take care of our kids, and we are right in so believing. How lucky that the adaptive beliefs and the evaluative truth here coincide!

If the onus is on us to demonstrate that we are not mistaken, we cannot simply insist that our beliefs are true and count ourselves lucky. We would be like the dogmatist who reasons that since he knows that *p*, any evidence he gets against *p* must be misleading, so he can ignore it.¹² We cannot safeguard our beliefs from defeating evidence like this. Nor can we dismiss the debunker’s challenge so easily.

¹⁰ I agree, though I’ll soon argue that this principle doesn’t capture these differences.

¹¹ Wielenberg’s [2010] response assumes that we have rights. Enoch’s [2010] assumes that “survival or reproductive success (or whatever else evolution ‘aims’ at) is at least somewhat good” ([2010] 18). Dworkin repeatedly insists that we can just count ourselves lucky ([1996], [2011]). Parfit earlier claims that moral beliefs can be justified by their intrinsic credibility (see his [2011] 490). I won’t say more about these here. I take them up in my [ms.b].

¹² Cf. Harman [1973] 148 and Kripke [2011] 49.

We can now see what the debunker thinks we need if we are to avoid her challenge: a reason to think that we are not mistaken in our evaluative beliefs that doesn't simply presuppose the truth of those beliefs. This reason is, in some sense, independent of what is called into question.¹³

This explains why the debunker asks us to bracket our evaluative beliefs—even those that we know or rationally believe—and to focus only on the origin story. If we do not do this, we stack the deck in our own favor. The danger, of course, is that if we do, then we may well lack reason to think our beliefs aren't mistaken.

3.1 Why NO GOOD is no good.

The debunker thus needs a 'good' reason to be an appropriately 'independent' reason. This stringent understanding allows the debunker to dismiss Parfit et. al. and claim that we have no good reason to think our evaluative beliefs are right. But if we understand 'good' reason this way here, we must understand it in the same way in NO GOOD. This, I will now argue, entails a skepticism far more pervasive than the debunker ever intended.

Start with an explicit statement of this version of the argument.

1. INFLUENCE. Evolutionary forces have influenced our evaluative beliefs.
2. We have no good reason to think that our evaluative beliefs are true. [1]
3. NO GOOD. If you have no good reason to think that your belief is true, then you cannot rationally maintain it.
4. REVISION. We cannot rationally maintain our evaluative beliefs. [2, 3]

Every premise in this argument is controversial. I granted the first, and I will grant for argument's sake that it somehow entails the second. Do not worry that this concedes too much to the debunker. Such generosity will not give the game away. Focus instead on the third premise. NO GOOD seems compelling because it raises a familiar sort of skeptical challenge. But it also collapses the debunker's challenge into that more ambitious one for which no empirical premise is necessary and which undermines much more than evaluative realism. To see this, consider:

Perception. We come to hold beliefs about our manifest surroundings on the basis of signals that hit our sensory organs.

¹³ This independence requirement is crucial to the debunker's argument, and yet has no defense in the debunking literature. Elga [2007], Christensen [2007], and others explicitly endorse similar independence requirements for disagreement. White questions them in his ([2010] 588-9). More must be said about what counts as independent, how to set aside what is not, and how to characterize this 'setting aside' formally. These questions have been little addressed in the literature and I won't be remedying that here. Though rough, the above characterization suffices. I think independence requirements are plausible and Street is therefore right to insist that a proper response to her challenge cannot just assume the very thing in question. I won't argue for this here, though see my [ms.a] and [ms.b]).

Unless we are skeptics, we should grant that sensory perception is a perfectly good belief forming method. *Ceteris paribus*, if you perceive that p , you are rational in concluding that p . Do we have good reason to think that perception would lead us to true beliefs about our surroundings? Not if ‘good’ reason is understood as an appropriately independent reason: for if we set aside all that is in question, we must set aside all beliefs gained by perception. This includes all scientific beliefs, like the belief that evolutionary theory is true. Without those, we cannot evaluate the rationality of beliefs formed by perception. We can test the reliability of a particular sense modality by granting the reliability of others. We can test our eyes against our ears, and so on. But if we cannot rely on any of our senses, we have nothing with which to evaluate reliability. We have set aside too much.

This might just be what the skeptic aims to demonstrate: that our justifications eventually run out and our beliefs ultimately rest on nothing. This, however, was never the debunker’s point. She aimed to undermine a particular, limited set of our beliefs using good scientific evidence that they are mistaken. NO GOOD commits her to much more. If this argument works, it undermines all that we believe and the evolutionary premise drops out. Worse yet, if we aren’t justified in believing *anything*, then it is the end of the world and everything is awful and there is no special problem for the evaluative realist.

Some have argued that the evolutionary story is not essential to the argument. This is only true in an uninteresting sense: any suspect influence could do the job. It needn’t be evolution. But an empirical claim of some sort is essential—this is the novel, distinctive feature of such arguments.¹⁴

This isn’t always clear in presentations of the argument. Elsewhere Street begins by pointing to the phenomenally low “odds that among all the possible coherent normative systems, one’s own is the right one” ([ms.] 21). Since there are infinitely many possible coherent normative systems, she argues, it would be a “striking coincidence” if one’s own normative system happened to be the correct one (ibid.).¹⁵ Given that “one has no non-trivially-question-begging evidence that one’s own system is the right one,” it is unreasonable to conclude that it is (ibid.). Street thus concludes that we have no good reason to think that our evaluative beliefs are roughly on-track, for we have no reason that does not assume the very thing called into question: the truth of those beliefs.

On this presentation, the debunker’s challenge brings nothing new to the table. It raises a worry and it demands we explain why we wouldn’t be massively mistaken about morality. This more general explanatory demand may be legitimate and important, but it isn’t the debunker’s.¹⁶ It is just an instance of a general skeptical worry, which suspiciously resembles this one:

Possibility of Error. Some possible states of belief are coherent and stable—they look fine “from the inside”—and yet are mistaken. There are infinitely many of these and just one

¹⁴ Cf. Bedke [ms.] 3 and Street [2006] 155.

¹⁵ Bedke presents the challenge this way: as that of explaining this striking coincidence. He does think an empirical premise is necessary, however, so it isn’t obvious which way he goes.

¹⁶ I argue for this in my [ms.b], first presenting the explanatory demand and then distinguishing it from the debunker’s.

that is right. Furthermore, we have no good reason to think we're not in such a state. So it would be unreasonable for us to be confident that we're not in such a state.¹⁷

This challenge doesn't and needn't rely on empirical claims. You are asked to justify your entire body of belief—and, on the relevant understanding of 'good reason', you must do it without presupposing the truth of any of the beliefs that have been called into question. But all of your beliefs have been called into question, so the skeptic asks you to put them all aside. She then asks: have you one good reason to think that your beliefs are true? You do not, of course. And it isn't because you have some reasons, but they aren't any good. The problem is that once you put aside all that you believe, you don't have any reasons left.¹⁸ You do not even have beliefs, so how could you have reasons?¹⁹

This challenge can be raised against any subject matter. It isn't peculiar to the evaluative, it isn't uniquely a problem for realism, and it can be raised without empirical premises. If the debunker accepts NO GOOD, she commits herself to the legitimacy of this reasoning. She thus ends up with the conclusion that we should all—regardless of our metaethics—suspend judgment about everything. But that was never her goal.

Focusing on the many coherent evaluative states that we might be in is thus misleading. That there are many such states, and that we have no good reason to think we aren't in one of these bad ones may be a problem, but it isn't the debunker's problem. What the debunker is trying to show, I will now argue, is that we have good reason to think that we are in one of the bad states. Her principle should reflect this.

4. Why GOOD is good.

What is the epistemic significance of the evolutionary story for our evaluative beliefs? I argued that it couldn't be that it leaves us with no good reason to think we are not massively mistaken about the evaluative. If we understand a 'good reason' as we must, to avoid begging any questions, then we certainly lack such reason. But we lack it for our entire body of beliefs. While that may be a problem, it isn't the debunker's problem. So her point cannot be that we lack good reason to think we're right.

What is her point? It has something to do with the epistemically unflattering picture the evolutionary story provides. What is epistemically unflattering, however, isn't that we cannot independently establish that these beliefs are right. Rather, it is that in learning this story about the origin of our evaluative beliefs, we get good reason to think that our beliefs are wrong. Since

¹⁷ Elga [ms.] 7.

¹⁸ Do you have anything left with which to even comprehend the skeptic's question? That is another difficulty. There is a more general anti-skeptical strategy in this spirit, most commonly attributed to Wittgenstein [1969]. Wright [2004] develops a view in the same spirit. My goal is not so ambitious. It is just to distinguish skeptics from debunkers.

¹⁹ Of course, there is a sense of 'reason' on which I can have one even if I do not or cannot believe I have one. For the record, here and throughout, I will use 'having a reason' and 'believing you have a reason' interchangeably.

evolutionary forces select for adaptive beliefs—and not true ones—evolution is a bad, potentially distorting influence on our evaluative beliefs. On this alternative line of thought, the problem is not that we cannot dismiss the possibility of error—it is that good scientific evidence makes this possibility more probable.

This version of the debunker’s argument is distinct from traditional skeptical arguments since it rests on an empirical claim. It is more selective than traditional skeptical arguments because it targets all and only the suspiciously influenced beliefs. The epistemic principle it relies on is:

GOOD. If you have good reason to think that your belief is mistaken, then you cannot rationally maintain it.²⁰

The difference between GOOD and NO GOOD is subtle but crucial. Roughly, it is the difference between taking our beliefs to be innocent until proven guilty and taking them to be guilty until proven innocent. NO GOOD requires you to launch a defense on behalf of your belief; GOOD requires you to hear out the prosecution. Both of these principles can be used to formulate a valid debunking argument, but the debunker should accept GOOD only.

The debunker’s point is that evidence of evolutionary influence is evidence of error. When we get such evidence, we must accommodate it with appropriate revision. This is exactly what GOOD expresses. It rightly shifts the burden to the debunker. It isn’t up to us to show her that we aren’t mistaken. It is the debunker’s job to show us that we are mistaken. GOOD reflects this dialectic and provides a plausible link between the discovery that a belief reflects the influence of a suspect process and the conclusion that we cannot rationally maintain that belief.

Earlier we granted, for the sake of argument, that we have no good (independent) reason to think our evaluative beliefs are not mistaken. With NO GOOD, this entailed that we could not rationally maintain our evaluative beliefs. If we accept GOOD only, the debunker must do more. Our lack of good (independent) reason to think our evaluative beliefs are right leads nowhere without something like NO GOOD. The onus is now on the debunker to show that the evolutionary story supports something stronger. She must do more than merely demand an explanation and watch

²⁰ The caveat from fn. 9 is relevant here. I use ‘mistaken’ to mean ‘false’, but these principles could be formulated in terms of rationality, justification, or evidential support. E.g.,

GOOD*. If you have good reason to think that your belief is not supported by your evidence, then you cannot rationally maintain it.

This is more controversial. Christensen [2011], Elga [2007], and Vavova [ms.a] defend principles along these lines. Kelly [2005] and Weatherson [ms.] reject them. They argue that “higher-order” evidence about p—evidence about your evidence for p—should not affect your “first-order” attitude about p. There might be nothing wrong, on their views, in believing both that p and that your evidence does not support p. So they would reject GOOD*. They could still accept GOOD, however, for that commits them to something weaker and more plausible: that you cannot rationally believe both that p and that p is false.

us squirm. She must show us that we have good reason to think that our evaluative beliefs are mistaken.²¹

A good reason is here, as before, an appropriately independent one. Your evaluation of whether you have good reason to think that you are mistaken about p should not rely on p or on the evidence or arguments on which p is based. This is for the same reason as before: to block a certain kind of question-begging response. If I can take for granted that pain is bad and survival is good, then I have a quick and easy explanation for why evolution is concerned with exactly the attitude-independent moral truths.

The independence requirement is also important here for another reason. Since the onus is now, rightly, on the debunker to give us evidence of error, this evidence should be good evidence we can recognize as such. It should follow from our other beliefs about reasons and evidence. But notice how strange it would be for her to rely on the beliefs she does not allow us to rely on—the ones we are supposedly mistaken about. Her argument would something like this one: p is probably false, but it entails q , so you should believe q . The debunker cannot simply rely on the beliefs that are supposed to be mistaken—the very same ones she won't let us take for granted. She must build her case upon solid, independent grounds. She thinks she can, but I will argue to the contrary.²²

Consider first this revised version of the argument:

5. INFLUENCE. Evolutionary forces have influenced our evaluative beliefs.
6. MISTAKEN. We have good reason to think that our evaluative beliefs are mistaken. [1]
7. GOOD. If you have good reason to think that your belief is mistaken, then you cannot rationally maintain it.
8. REVISION. We cannot rationally maintain our evaluative beliefs. [2,3]

Every premise of this argument is also controversial, but GOOD is weaker and more plausible than NO GOOD. It provides a framework within which the debunker can pose an appropriately selective and distinctive challenge. It is at least possible to construct the right kind of debunking argument.

The action is now with the second premise: have we, realists, been given good reason to think that our evaluative beliefs are mistaken? I will examine three evolutionary debunking arguments, which aim at a different set of our evaluative beliefs. I will argue that in all three, the debunker fails to give us good reason to think we are mistaken. Since we can reject second premise, we aren't pushed into evaluative skepticism.

²¹ There is some evidence for this reading (cf. fn. 5). See especially Street [2006] where she often talks of the 'distorting' Darwinian forces having led us off-track, or "having pushed us in evaluative directions that have nothing whatsoever to do with the evaluative truth" (121). I do not think anyone is consistent on which way to understand the evolutionary debunker's challenge: like this or as a more general skeptical challenge. My [ms.b] more thoroughly defends this interpretation of the dialectic.

²² Cf. Street [ms.] where she argues that the particular normative assumptions in question are not needed for either raising or responding to the challenge.

5. Debunking evaluative realism.

The most familiar evolutionary debunking argument targets moral realism, and aims to undermine our beliefs about what we have reason to do. I will start with a more ambitious argument, which aims to undermine evaluative realism wholesale: not just our beliefs about what we have reason to do, but also our beliefs about what we have reason to believe. This debunker thus targets realism about both practical and epistemic reasons.²³

To see how the trouble is supposed to arise, consider our belief that frequency facts like

[TIGERS] the fact that all previously encountered tigers were carnivorous,

give us reason to believe inductive claims like

[NEXT TIGER] the next tiger we encounter will also be carnivorous.

It is clear why we evolved a tendency to form beliefs like [NEXT TIGER] on the basis of frequency facts like [TIGERS]: if we hadn't, tigers would have eaten us. But why did we evolve to take frequency facts like [TIGERS] as reasons to believe facts like [NEXT TIGER]?²⁴ Is it because grasping this attitude-independent normative truth was itself adaptive? Unlikely, Street argues: natural selection favored a tendency to take considerations of truth to bear on what to believe “not because it constituted a perception of an independent fact about reasons, but rather simply because it guided the formation of creatures' beliefs in ways that turned out to be advantageous for the purposes of survival and reproduction—in particular, because it got them to believe things that turned out to be true, or at least roughly true, about tigers and much else” ([forthcoming] 17).

In other words, we wouldn't believe that [TIGERS] is a reason for believing [NEXT TIGER] if concluding [NEXT TIGER] on the basis of [TIGERS] weren't to our evolutionary benefit. Since evolution has no interest in the attitude-independent epistemic truth, the beliefs it influences are likely to be mistaken. Insofar as we are realist, the debunker argues, and continue to maintain that what is epistemically valuable is valuable whether or not we value it, we seem pushed to skepticism.

This argument rests on the claim that the same kinds of considerations meant to undermine beliefs like we have reason to take care of our children would also undermine beliefs such as we have reason to believe this rather than that on this evidence. Even as she launches a formidable defense of this claim, arguing both that evolutionary forces influenced our beliefs and that this should worry us, Street admits that this case is much harder to make.

²³ Cf. Street [2009].

²⁴ I assume here a view on which taking [TIGERS] to be a reason to believe [NEXT TIGER] is something more than merely having the disposition to infer one from the other.

Grant her the first bit again (namely, INFLUENCE) and ask: if evolution had shaped our beliefs about epistemic reasons, would this give us a good epistemic reason to worry about those beliefs? I will argue that it does not and it cannot, for there is a deep structural problem with an argument this ambitious.

The debunker aims to give us good reason to believe that we cannot trust our beliefs about reasons for belief. But this itself—what the debunker wants to give us—is a reason for belief. So we cannot trust it. We are therefore not permitted to take for granted the very thing we need to call our evaluative beliefs into question. This is because, recall, the debunker must give us good independent reason that is, by our own lights, reason to think we are mistaken. But on this version, what we are supposed to be mistaken about includes, crucially, epistemic principles about how to revise our beliefs in light of evidence. We need to take for granted the truth of GOOD and MISTAKEN. Both of these claims, however, are about what we have reason to believe, which is exactly what we're supposed to be mistaken about.

The debunker thus faces a dilemma. She may relax her standards for what counts as a 'good' reason, or she may maintain them. If she maintains them, then she cannot give us good reason to think we are mistaken about the evaluative. In short, this is because to evaluate we must rely on the evaluative. But in aiming to debunk all of our evaluative beliefs, the debunker leaves us with nothing with which to evaluate whether those beliefs have been debunked.

If instead the debunker relaxes her understanding of 'good reason', then GOOD is back. But so are our other beliefs about epistemic reasons, like the belief that [TIGERS] really does give us reason to believe [NEXT TIGER], and so on. And if we are allowed these assumptions, then the question-begging response Street blocked is open again.

There is a natural response available to the debunker here. She could reply that her point is dialectical, not skeptical. Though some debunkers are skeptics or nihilists, others, like Street, are not. They do not really aim to debunk our evaluative beliefs—they think those are true. Instead, they aim to debunk realism. The skeptical conclusion is only for the purposes of *reductio*, for these debunkers. It follows from realism and science, they argue, and it is absurd. We cannot give up science, so realism must go.

Unfortunately, this response won't do. Even if the debunker does not ultimately endorse the skeptical conclusion, she must still show that it follows from realism and the evolutionary story. It is only if she can demonstrate this that she has what she needs for her *reductio*. To do so, the debunker must give us realist's good reason to think we are mistaken. I have argued that the debunker is in principle incapable of providing evidence of such global error. The *reductio* thus cannot go through. MISTAKEN is false. We do not have good reason to think we are mistaken. The evolutionary story, at least, hasn't given us any.

Such is the fate of the debunker who attacks evaluative realism wholesale. Perhaps it isn't surprising that this most ambitious debunker failed in just this way. There are well-known puzzles about whether we can revise, or even be anything short of certain of, our most fundamental

principles of belief revision.²⁵ But perhaps the debunker can sidestep these difficulties and avoid such a fate, if she can narrow target.

6. Debunking moral realism.

There is more hope for the debunker who aims only at moral realism. Since she does not target our beliefs about epistemic reasons, both GOOD and MISTAKEN are potentially in play. The question is whether she can actually establish the latter—whether she can use her evolutionary story to give us good reason to think we are mistaken about morality alone. There are two impediments in her way.

The first is that the debunker must show that evolution causes trouble for our moral beliefs only—that there is some disanalogy between this argument and the previous one. But the two arguments are presented as exactly analogous (Street [2009]). If the debunker cannot narrow down her target in a principled way this less ambitious argument collapses into the previous, thereby sharing its fate.

The second is that even an appropriately narrowed challenge calls too much into question. Since it targets all of our moral beliefs, we are left knowing nothing about morality. But how can we tell if we are likely to be mistaken about morality, if we know nothing about it? This concern will occupy the rest of this section. To see it more clearly we need to zoom in to the first inference of the argument.²⁶

So far, we have either granted or glossed over the move from INFLUENCE to MISTAKEN. But MISTAKEN simply doesn't follow without, at least, reason to be suspicious of the purported influence. As Street puts it:

...genealogical information by itself implies nothing one way or another about whether we should continue to hold a given belief. Rather, in order validly to draw any conclusions about whether or how to adjust one's belief that *p*, one must assess the rational significance of the genealogical information, locating it in the context of a larger set of premises about what counts as a good reason for the belief that *p*. ([forthcoming] 2)

Kahane [2011] suggests, as a possible supplementary premise, that evolution is an 'off-track' process since, by hypothesis, it doesn't track the attitude-independent evaluative truths.²⁷ So long as we think that the adaptive beliefs come apart from true beliefs, we can accept this premise. Expanding the argument thus we get:

²⁵ See Field [ms.a], [ms.b], and Lewis [1971].

²⁶ In fact, the previous debunker faces an exactly analogous problem: if we know nothing about the evaluative, how can we tell we are likely to be mistaken about it.

²⁷ Bedke rightly warns that a process being off-track, "is ambiguous between the claim that the process has been shown to be unreliable and the claim that explanations for the process do not aver to the target facts" ([ms.] 4-5). I think the debunker should be claiming something more like the former. The latter claim is more akin to the aforementioned explanatory demand, which I take up and distinguish from the debunker's in my [ms.b].

1. Evolutionary forces select for creatures with characteristics that increase fitness.
2. The true evaluative beliefs and the adaptive evaluative beliefs come apart.
3. Evolutionary forces are off-track: they do not track the evaluative truth. [1, 2]
4. INFLUENCE. Evolutionary forces have influenced our evaluative beliefs.
5. OFF-TRACK. Off-track forces have influenced our evaluative beliefs. [3, 4]

If the debunker can establish OFF-TRACK, she is a short step from MISTAKEN. After all, an off-track influence pushes your beliefs in directions having “nothing whatsoever” to do with the truth. Reason to think your belief reflects the influence of an off-track process thus looks like good reason to worry about the truth of that belief. If the above argument gives us good reason to think that our evaluative beliefs reflect an off-track influence, then it seems that we have good reason to think that those beliefs are mistaken. GOOD then takes the debunker home:

6. MISTAKEN. We have good reason to think that our evaluative beliefs are mistaken. [5]
7. GOOD. If you have good reason to think that your belief is mistaken, then you cannot rationally maintain it.
8. REVISION. We cannot rationally maintain our evaluative beliefs. [6, 7]

We’ve granted INFLUENCE and GOOD. We could resist the inference from OFF-TRACK to MISTAKEN, but we shouldn’t. It isn’t so controversial: it doesn’t say that learning about an off-track influence should all-things-considered worry you; just that it gives you a reason to worry.

Focus instead on OFF-TRACK. To get there, the debunker needs P2: the claim that the evaluative truths and the adaptive beliefs come apart—that there isn’t any helpful overlap between these two sets. Why should the realist accept this? Can’t she point to an apparently obvious overlap? Pain is bad, survival is good, and these are exactly the things evolution tracks! It may not track the evaluative truth directly, but evolution tracks it indirectly, by selecting for features with which it correlates (cf. Parfit et. al.).

Street hoped to block this move. Our beliefs that pain is bad and survival is good are exactly the sorts of beliefs we would expect evolution to have led us to, whether or not they were true. A legitimate response to the debunker’s challenge, Street argued, cannot just assume the very things called into question. We must set aside the suspect beliefs and independently evaluate whether we have good reason to think we are mistaken.

The problem here is that our entire body of moral beliefs is suspect. It follows that we must set all of our moral beliefs aside, if we are to block such question-begging responses. We cannot, then, simply assume that we have reason to avoid pain—that morality is about what is good for us, and that needlessly throwing ourselves off of cliffs just isn’t that sort of thing. These assumptions aren’t appropriately independent. Taking them for granted threatens to stack the deck against the debunker. I will now argue, however, that taking these assumptions off the table threatens to undermine the debunker’s argument.

Recall that we are meant to be getting good reason to think that we are mistaken about morality. But we cannot determine if we are likely to be mistaken about morality if we can make no assumptions at all about what morality is like. I argued that the debunker's challenge threatens anyone who holds that the attitude-independent moral truths do not, in any helpful way, coincide with the evolutionarily advantageous beliefs—anyone who accepts P2. But even to make this crucial judgment, that these two sets do not have the same contents, we need to know something about the contents of those sets—what they are or what they are like.

Compare: I cannot demonstrate that I am not hopeless at interacting with external objects in my manifest surroundings without knowing something about what those objects and surroundings are like. Likewise, I cannot show that I am not hopeless at understanding right and wrong without being allowed to make some assumptions about what is right and wrong.

If we can make no moral assumptions, then we cannot get P2: the claim that the true evaluative beliefs and the adaptive evaluative beliefs come apart. Now, I think P2 is plausible, and probably you do too. Certainly any realist should believe it. However, we find P2 plausible in part because of our substantive moral beliefs. For example, we believe it is wrong to discriminate against someone on the basis of race. At the same time, there are evolutionary explanations of racism, on which it is adaptive to be suspicious of those who do not look like you. In this case, then, the adaptive belief and the true moral belief come apart. Thus, to believe P2, one must also believe that the evaluative beliefs are such-and-such, while the evaluative truths are this-and-that. But if we cannot take for granted any of our beliefs about the evaluative truths, then we cannot infer that they come apart from the adaptive beliefs.

Again the debunker faces a dilemma. She may relax her standards for what counts as a 'good' reason, or she may maintain them. If she relaxes them, she cannot give us good reason to think we are mistaken. Worse yet, if we are permitted to assume that pain is bad, etc., then we can give her good reason to think we are not mistaken and her purportedly undermining story vindicates our evaluative beliefs.

If, instead, the debunker maintains her standards, she blocks such responses. But she also blocks herself. If we cannot make any moral assumptions—not even that pain is bad—then morality could be about anything.²⁸ To hold that the moral truths do not coincide with the adaptive judgments, we must assume something about what those moral truths are, or are like. If we may assume nothing about morality, then morality 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 attitude-independent truths and the adaptive beliefs don't overlap. But without that, we have no sense of what the chances are that we are mistaken. Therefore, we cannot get to the conclusion that we probably are mistaken.²⁹ Not, at least, via an evolutionary story.

²⁸ You might worry here that we are even talking about morality any more. The debunker assumes that morality really could be about anything—it is conceptually possible that morality is about throwing ourselves off of cliffs and causing each other pain. I'm not so sure about this. Cuneo and Shafer-Landau [ms.] argue that some of the very basic moral claims (like that pain is bad) are conceptual truths: if we don't have them we don't have our concept of 'morality'. This seems right to me, but I won't explore it further here.

²⁹ I expand on this discussion in my [ms.b].

6. Debunking Deontology.

The third debunking argument aims to undermine neither realism nor our entire body of moral beliefs. It targets a restricted class of those beliefs: those based on deontological intuitions.³⁰ This should be the most promising argument yet. Leaving intact most of our belief system gives this debunker an abundance of resources with which to construct her challenge.

Unfortunately, this debunker's evolutionary story is either idle or too strong. On the first point: worries about the targeted intuitions arise independently and are not worsened when supplemented with an origin story. On the second point: even if we lack other reason to worry, we should be reluctant to rely on an evolutionary story. It just isn't selective enough.

But first, the argument. It begins with a sociological observation: most think it permissible to divert a trolley away from five people toward one, but impermissible to push one in front of a trolley to save five. Why the discrepancy? We are killing one person in both cases, after all. The answer, of course, is evolutionary. Pushing the one, rather than diverting the trolley onto the one involves "up close and personal" violence of the sort that, unlike button pushing or lever pulling, has been around for a long time (Greene [2008] 43). Evolution selects for negative responses to this direct way of killing; it doesn't select for similarly negative responses to more indirect ways of killing. But the fact that "I have killed someone in a way that was possible a million years ago, rather than in a way that became possible only two hundred years ago" is morally irrelevant (Singer [2005] 348). If our deontological intuitions have this suspect origin, then we should worry about the beliefs we rest upon them. They are likely to be mistaken. The debunker concludes that we can only trust our utilitarian intuitions, which come from our uncontaminated "rational intuition" (Singer [2005] 350-1).

Two questions arise for this debunker. First, did we need an evolutionary story to make us worry about these particular intuitions? Second, why should we think that our consequentialist intuitions are less suspect?

On the first point. It is true that we feel a greater pull to help the nearby needy than the distant needy. Greene says "the only reason that faraway children fail to push our emotional buttons is that we evolved in an environment in which it was impossible to interact with faraway individuals" ([2008] 76). This should make us uncomfortable, he argues, if we think we are justified in ignoring the distant needy. For it was just an accident of evolution that we are "emotionally insensitive to their plight" (ibid.).

Recognizing that we are emotionally responsive to only nearby suffering should worry us, but for more familiar reasons.³¹ What, after all, is the moral difference between the drowning child in

³⁰ Here I follow Greene: deontological judgments are those "in favor of characteristically deontological conclusions (e.g., 'It's wrong despite the benefits')," and *mutatis mutandis* for consequentialist judgments ([2008] 39).

³¹ Singer [1972].

front of you and the starving child across the world? Our intuitive judgment that we may be selectively altruistic in these ways is already under pressure in the same way our judgments about trolley cases are under pressure. Try as we might, we can't seem to find satisfying reasons for these diverging judgments. But this problem is a distinctly moral one. We can recognize it from the armchair—no empirical origin story is necessary.

Likewise with intuition that it is impermissible to push one person off of a bridge to save five. Many of us feel this quite strongly. Many of us also believe that there is no morally significant difference between killing by lever-pulling and killing by person-pushing. Recognizing that we are making a distinction without a difference should already make us quite confident that we are making a mistake. Furthermore, we are rightly more confident in this judgment than we are in any origin story. Learning, then, that we evolved to make this distinction without a difference shouldn't further increase our confidence that we are making a mistake: it is already maxed out. The evolutionary story is thus, at best, an idle premise in this argument.³²

Why doesn't this just mean that this debunker is lucky—that the undermining of our deontological beliefs is over determined? This is where the second point comes in. Suppose we lack this other reason to worry, so the full weight of the conclusion falls on the evolutionary story. Wouldn't we then have good, evolution given, reason to worry about our deontological beliefs? Only if the evolutionary story could be employed selectively against just our deontological intuitions and it cannot.

First, there isn't an uncontroversial evolutionary explanation of our altruistic tendencies.³³ Without one we cannot accept this debunker's influence claim. Second, if we had such an explanation, it would debunk more than our belief in selective altruism: "if a disposition to partial altruism was itself selected by evolution, then the epistemic status of its reasoned extension [impartial altruism of the sort utilitarians promote] should also be suspect" (Kahane [2011] 119). This echoes a claim Street [2006] makes for a different purpose: that rational reflection cannot correct for evolution's unsavory influence. If our most basic moral intuitions are infected, and they are the starting points for our moral reasoning, then any result of that reasoning will also be infected. Even Green recognizes this point, succinctly putting it thus: "garbage in, garbage out" ([2008] 116). Why, then, isn't he worried? Surely he should be. For if these considerations are right, then Greene's argument targets utilitarian intuitions too. It thus collapses into the previous, more ambitious argument. This is bad for two reasons. First, it is no longer an argument against our deontological beliefs only, as it was intended to be. Second, as an argument against all of our moral beliefs, it suffers from the same trouble as the previous.

This debunker thus faces a different dilemma. She either relies on an evolutionary claim or she doesn't. If she doesn't, she accomplishes nothing new philosophically. She merely reiterates one side of the same old debates about drowning children and runaway trolleys. If the debunker instead decides to rely on an evolutionary claim, she presents a new argument, but it isn't a good one. The considerations she cites undermine more than she intended, and the argument collapses

³² These considerations reinforce similar points made by Berker [2009].

³³ Cf. Okasha [2009].

into the more ambitious and less promising one. This is why the evolutionary bit is, at best, idle, and at worst, too powerful, for this debunker's purposes.

8. Some general lessons.

I argued that we have reason to worry about each of the available evolutionary debunking arguments. This doesn't show that evidence of an off-track influence could never give us good reason to think we are mistaken. On the contrary, I think we often can.³⁴ What the foregoing shows is limits of such arguments, and so, more generally, the limits of our ability to get evidence of our own error. These limits are not of our cognitive architecture. Our ability to acquire evidence of our own error is not limited because we are, say, bad at recognizing such evidence. The limits arise out of the way that such evidence works. Let me explain by extracting two lessons from the above: one about debunking arguments and one about undermining evidence more generally.

8.1 Is evolution the problem?

I argued that the most modest of the available debunking arguments fails: that insofar as it poses an epistemic threat it does so without an evolutionary story and adding one threatens to weaken it. I now want to suggest that this situation will be quite typical of evolutionary debunking arguments more generally.

Consider a debunking argument that should work, if the empirical facts are as they need to be. Studies show that on the basis of only CV perusal people are much more likely to favor job candidates with stereotypically white names over candidates with stereotypically black names.³⁵ This is so even though the CVs are otherwise identical: same content, different names. If these studies are right, then we have reason to worry about our judgments in such situations.

We could try to give an evolutionary explanation to diagnose the worry. Whether it works depends on the empirical details. Is there an evolutionary story about why preferring people who look like us is evolutionarily advantageous, and whether this story is sufficiently selective? Notice, however that we needn't answer these questions. We can establish a problematic pattern of CV selection much more easily than we can establish the influence of evolutionary forces on these particular beliefs. The empirical details are far less clear and the debunking is far less promising in the evolutionary case. Given that we do not take a candidate's name or race to be a relevant hiring consideration, this pattern of selection is sufficient cause for concern. Given everything else we believe, we have good reason to think that we are making a distinction without a difference here. Insofar as we can control our confidence in the quality of this candidate over that one, we should revise it when it fits this troublesome pattern. If we cannot revise it, then we should judge our confidence that Emily is better suited for this job than Lakisha to be less than fully rational and

³⁴ See my [ms.a].

³⁵ Bertrand & Mullainathan [2004].

take whatever correcting measures we can. We have good independent reason to think we are mistaken.

The upshot here is that both of these more modest debunkers are better off dropping the evolutionary bit of their arguments. This will be typical for evolutionary debunkers. Their stories just cannot provide an appropriately selective argument that targets all and only the intended beliefs. This might not be surprising. It can seem as if an evolutionary story can be told about any of our beliefs. Advances in the relatively young field of evolutionary psychology might change this. Currently, however, things do not look promising for the evolutionary debunker. She is better off dropping the evolutionary story altogether.

8.2 On undermining.

I have argued that all three debunking arguments fail to give us reason to worry about our beliefs. This assumes, of course, that the burden is on the debunker to give us such a reason to revise our beliefs. It also assumes that rational belief revision works a certain way. I should make this background picture explicit, though it not controversial. Two minimal assumptions guide the foregoing thoughts:

A1. A reason is a reason, and evidence is evidence, only against a backdrop of beliefs we take for granted.

Consider: the sound of water drops on my office window is typically evidence that it is raining. Suppose, however, I believe that the college gardener is out to get me, so he regularly aims the garden hose at my window in the hope of flooding my office. Water drops on my window, relative to these background beliefs, is not evidence of rain. It is evidence that the gardener is at it again.

The second assumption is this:

A2. The undermining power of a reason or a piece of evidence is not all-or-nothing.

Hearing a trusted colleague say that the gardener is in the shed is a good reason to think that the gardener is in the shed. But seeing him there, with my own eyes, putting the hose away, may be a better reason to think so.

Just as you can get various strengths of reasons for thinking that the gardener is out in the shed, you can get varying strengths of reasons for thinking that you are mistaken about some p . The stronger your reason for thinking that you are mistaken, the more substantial revision you will probably have to make. The strength of this reason will depend on what you have to go on. This ties in with the first assumption: the more substantial the body of beliefs you can take for granted, the more potential you'll have for getting a good reason. If, for example, you cannot take for granted the trustworthiness of testimony, you won't be able to get testimonial evidence that the gardener is out in the shed. Things would be worse yet if you couldn't take for granted the trustworthiness of your own eyes and ears, or even your powers of reasoning.

These observations about how evidence works apply to the evidence the debunker presents in the following way. Evidence of error is a piece of evidence like any other: the better ground you have from which to evaluate the evidence, the more potential you'll have for getting good reason to revise. This is important because, recall, the debunker requires you to set aside the targeted beliefs when evaluating her challenge. You must not take those for granted if you are to avoid stacking the deck in your own favor. The more substantial the body of beliefs that the debunker calls into question then, the less substantial your independent ground will be. And the less substantial the independent ground is, the worse the resources for both presenting and evaluating evidence of error.

Notice, however, that having many uncontested beliefs does not guarantee having 'substantial' independent ground, since these beliefs might be quite superficial or otherwise irrelevant. Perhaps I have memorized the phonebook, so that I have a large number of true beliefs. These are useless for determining if I am likely to be mistaken about evaluative matters. Nor is this independent ground of help to the debunker: she cannot make her case on phone numbers alone. What determines whether we have good ground from which to evaluate the debunker's challenge is thus not how many beliefs are appropriately independent, but whether the right sorts of beliefs are appropriately independent. These include, at least, beliefs about rationality, evidence, and belief revision.

These assumptions about evidential support combined with the GOOD principle of belief revision suggest something like the following rule of thumb.

The Inverse Rule of Debunking. The potential strength of a debunking argument is inversely proportional to its ambition.

The 'strength' of the debunking argument has to do with how extensively we must revise. The 'ambition' of the debunker's argument refers to how much it targets.

Again, what matters in determining the strength of a debunking argument, or the undermining evidence it provides is not how many of your beliefs it calls into question but whether it leaves you enough of the right sorts of beliefs with which to evaluate the evidence that has been put before you. What matters for the ambition of the argument is also thus not how many beliefs it calls into question, but whether it calls the right sorts of beliefs into question, namely: those necessary for evaluating the relevant evidence.

This rule of thumb thus issues in the following prediction. The most ambitious debunker, she who aims to undermine all we believe, has the lowest chance of success. The evidence she aims to give has no undermining potential. This is because, by calling all of our beliefs into question, she leaves us nothing with which to question.

A moderately ambitious debunking argument may be able to cast some doubt on the targeted beliefs. The extent to which it succeeds depends on what exactly is called into question, and thus how substantive are the relevant independent grounds. The debunker who aims to debunk your most fundamental beliefs, those on which everything else you believe rests and with which you

judge what to believe—she is out of luck. Her challenge just is the ambitious one in disguise. If she aims at a fairly superficial set of your beliefs, she has a decent chance of undermining them.

The most modest debunker thus has the best prospects. She aims low, but she may score high.

This makes sense. A comparison might help. Consider three disagreements. You disagree with Anne about the permissibility of abortion. You agree on other moral and political matters. You disagree with Beth about the permissibility of abortion, but also about a myriad of other moral and political matters. You disagree with Clarisse about the permissibility of abortion. But you also disagree about every other moral matter. Clarisse is a psychopath.

Your disagreement with Anne has the most undermining potential; your disagreement with Clarisse the least. Generally, the more common ground you share with someone, the more significant their disagreement may be. This is because the more common ground you share, the more independent ground you have from which to get evidence of your error. You have much independent ground on which to evaluate your disagreement with Anne; you have none with Clarisse.

Since evidence of a suspect belief influence is also evidence of error, we should expect the same pattern: evidence of a more pervasive belief influence should be less worrying. This may seem counterintuitive, but it is a good result. It makes sense that the most modest, targeted sort of debunking argument should be the most effective, if it works. That sort of argument provides me with good reason to think I am wrong about some p. This good reason is good by my own lights: it follows from my other beliefs about reasons and evidence. This is the kind of reason the debunker must provide. It is also, I have argued, the kind that the distinctly evolutionary debunker cannot provide. This is good news for the realist. Whatever her epistemic troubles, this scientifically grounded one is not of them.³⁶

³⁶ Ancestors of this paper were given at the First Biennial Mentoring Workshop for Women in Philosophy, the University of Sydney, the Australian National University, the Australian Association of Philosophy 2012 meeting, the Ninth Annual Wisconsin Metaethics Workshop, Union College, and the University of Edinburgh. Thanks to the organizers for the opportunity to share my work and to the audiences for their helpful comments. Many thanks also to Lee Bowie, Sarah Buss, David Christensen, Terence Cuneo, Tom Dougherty, Adam Elga, Catherine Elgin, Alexander George, Jyl Gentzler, Daniel Greco, Toby Handfield, Caspar Hare, James Harold, Sophie Horowitz, Justin Horn, Hilary Kornblith, Dustin Locke, Joseph Moore, Elisa Mai, Sam Mitchell, Alejandro Pérez Carballo, David Plunkett, Paolo Santorio, Nishi Shah, Paulina Sliwa, Judith Jarvis Thomson, Jonathan Vogel, Kenneth Walden, Angela Ruohan Wang, Tom Wartenberg, Roger White, two anonymous referees from Oxford University Press, and students in my Mount Holyoke College seminar on moral epistemology. Thanks especially to Nishi Shah for many conversations on this topic, including those from which this paper emerged.

References

- Bedke, M. [ms.] "No Coincidence?" draft for Wisconsin Metaethics Workshop 2012.
- Berker, S. [2009] "The Normative Insignificance of Neuroscience" in *Philosophy & Public Affairs* 37(4): 293-329.
- Bertrand, M. & Mullainathan, S. [2004] "Are Emily And Greg More Employable Than Lakisha And Jamal? A Field Experiment On Labor Market Discrimination" in *American Economic Review* 94(4): 991-1013.
- Christensen, D. [2011] "Disagreement, Question-Begging and Epistemic Self-Criticism" in *Philosophers Imprint* 6: 1-22.
- Christensen, D. [2007] "Epistemology of Disagreement: The Good News" in *Philosophical Review* 116(2): 187-217
- Cuneo, T. & Shafer-Landau, R. [ms.] "The moral fixed points".
- Darwin, C. [1871] *The Works of Charles Darwin*, vol. 21, NYU Press, 2010.
- Dworkin, R. [2011] *Justice for Hedgehogs*, Harvard University Press.
- Dworkin, R. [1996] "Objectivity and Truth: You'd Better Believe it" in *Philosophy and Public Affairs* 25(2): 87-139.
- Elga, A. [ms.] "Lucky to be Rational" draft as of June 6, 2008, URL = <<http://www.princeton.edu/~adame/papers/bellingham-lucky.pdf>>
- Elga, A. [2007] "Reflection and Disagreement" in *Nous* 41(3): 478–502.
- Enoch, D. [2010] "The Epistemological Challenge to Metanormative Realism: How Best to Understand It, and How to Cope with It" in *Philosophical Studies* 148(3): 413–38.
- Field, H. [ms.a] "The revision of logical laws".
- Field, H. [ms.b] "Logic, Normativity, and Rational Revisability" from his 2008 Locke Lectures.

- FitzPatrick, W. [forthcoming] "Debunking Evolutionary Debunking of Ethical Realism," in *Philosophical Studies*.
- Greene, J. [2008] "The Secret Joke of Kant's Soul" in W. Sinnott-Armstrong, ed., *Moral Psychology* 3: 35–79.
- Harman, G. [1973] *Thought*, Princeton University Press.
- Kahane, G. [2011] "Evolutionary Debunking Arguments" in *Nous* 14(1): 103-125.
- Kelly, T. [2005] "The Epistemic Significance of Disagreement" in J. Hawthorne and T. Gendler Szabo (eds.) *Oxford Studies in Epistemology* 1: 167-196.
- Kripke, S. [2011] "Two Paradoxes of Knowledge" in *Philosophical Troubles: Collected Papers*, Vol. 1, Oxford University Press.
- Lewis, D. [1971] "Immodest Inductive Methods" in *Philosophy of Science* 38(1): 54-63.
- Okasha, S. [2009] "Biological Altruism" in *The Stanford Encyclopedia of Philosophy*, URL = <http://plato.stanford.edu/archives/win2009/entries/altruism-biological/>.
- Parfit, D. [2011] *On what matters*, Vol. 2., Oxford University Press.
- Quine, W. [1969] "Natural Kinds" in *Ontological Relativity and Other Essays*, Columbia University Press: 114-138.
- Shafer-Landau, R. [2005] *Moral Realism: A Defence*, Oxford University Press.
- Singer, P. [2005] "Ethics and Intuitions" in *The Journal of Ethics* 9(3-4): 331–352.
- Singer, P. [1972] "Famine, Affluence, and Morality" in *Philosophy and Public Affairs* 1(3): 229-243.
- Sliwa, P. [ms.] "Darwin and the Moral Realist".
- Street, S. [ms.] "Objectivity and Truth: You'd Better Rethink it".
- Street, S. [forthcoming] "Does Anything Really Matter or Did We Just Evolve to Think So?" in A. Byrne, J. Cohen, G. Rosen, and S. Shiffrin (eds.) *The Norton Introduction to Philosophy*, W. W. Norton.

- Street, S. [2009] “Evolution and the Normativity of Epistemic Reasons” in the *Canadian Journal of Philosophy* 39, supplement 1: 213-248.
- Street, S. [2006] “A Darwinian Dilemma for Realist Theories of Value” in *Philosophical Studies* 127: 109-166.
- Vavova, K. [ms.a] “Irrelevant Influences”.
- Vavova, K. [ms.b] “A Dilemma for the Darwinian Debunker”.
- Weatherson, B. [ms.] “Do Judgments Screen Evidence?”
- Wielenberg, E. [2010] “On the evolutionary debunking of morality” in *Ethics* 120(3): 44-464.
- White, R. [2010] “You just believe that because...” in *Philosophical Perspectives*, 24(1): 573–615.
- Wittgenstein, L. [1969] *On Certainty*, G.E.M. Anscombe and G.H. von Wright, eds., Basil Blackwell.
- Wright, C. [2004] “Wittgensteinian Certainties” in D. McManus, ed., *Wittgenstein and Scepticism*, Routledge.