



**EVOLUTIONARY DEBUNKING, MORAL REALISM
AND MORAL KNOWLEDGE**

BY RUSS SHAFER-LANDAU

JOURNAL OF ETHICS & SOCIAL PHILOSOPHY

VOL. 7, No. 1 | DECEMBER 2012

URL: WWW.JESP.ORG

COPYRIGHT © RUSS SHAFER-LANDAU 2012

Evolutionary Debunking, Moral Realism and Moral Knowledge

Russ Shafer-Landau

THERE HAVE RECENTLY BEEN A NUMBER OF sharp criticisms directed against moral realism – the tripartite view that (i) sincere moral judgments express beliefs, rather than conative attitudes; (ii) some of these beliefs are true; and (iii) such beliefs, when true, are not true by virtue of being the object of, or being implied by, the attitudes of (even idealized) agents. The arguments I am going to focus on all have the same form: Given moral realism, and given what we know about the evolutionary origins of our “moral faculties” – roughly, whatever interconnected, internal mechanisms and processes lead us to have the moral beliefs we do – there is no reason to think that these faculties are reliable, and plenty of reason to think that they are not.¹

The evolutionary debunkers rarely lay out a master argument in premise-conclusion form, and I think that their antirealist criticisms can be interpreted in a variety of ways. While these critics disagree amongst themselves on some important matters, I intend to elide most of these differences and concentrate on what I take to be the most serious common challenge that emerges from their writings. The debunkers claim that if moral realism is true, and if selective pressures have heavily influenced the development of our moral faculties, then we can have no moral knowledge. This by itself does not refute moral realism, but it leaves realists in the deeply unappealing position of being saddled with a thoroughgoing moral skepticism – a logically coherent position that contains about zero appeal.

1. Genealogical Critiques

Evolutionary arguments against moral realism are instances of a general kind of epistemological objection – that of the genealogical critique. Such critiques always begin with an empirical hypothesis about the origins of our beliefs in a

¹ See Joshua Greene (2008) “The Secret Joke of Kant’s Soul,” in Walter Sinnott-Armstrong, ed., *Moral Psychology, Vol. 3: The Neuroscience of Morality, Emotions, Brain Disorders and Development* Cambridge, MA: MIT Press, pp. 35-80; Justin Horn “Evolution and Moral Realism,” unpublished manuscript; Richard Joyce (2007) *The Evolution of Morality*, Cambridge, MA: MIT Press, ch. 6; Philip Kitcher (2007) “Biology and Ethics,” in David Copp, ed., *The Oxford Handbook of Ethical Theory*, Oxford: Oxford University Press; Michael Ruse and E. O. Wilson (1986) “Moral Philosophy as Applied Science,” *Philosophy* 61: 173-92; Michael Ruse (1998) *Taking Darwin Seriously*, Oxford: Blackwell Publishing, ch. 6; and Sharon Street (2006) “A Darwinian Dilemma for Realist Theories of Value,” *Philosophical Studies* 127: 109-66. (Street actually targets forms of evaluative and normative realism that are broader than moral realism, but I will focus almost exclusively on her criticisms just as they apply to morality.)

given domain, and proceed to raise enough doubts about the reliability of those origins that the ensuing beliefs are claimed to be unwarranted.²

Here is one way to raise such doubts – call this the *knowledge-based* genealogical critique. We can reveal the likely falsity of a given belief by showing that its origins are of a kind that we know to be typically distorting – either across the board (as when making pronouncements when dead drunk, or after an astrological consult) or, more narrowly, when it comes to specific topics. Knowledge of such distorting tendencies presupposes knowledge of many relevant truths on the matter at hand. For instance, we know that it is a bad idea to form color judgments while wearing tinted glasses. We know this because we know the colors of things, and what effect tinted lenses have, and this knowledge enables us to assess such a source as a distorting one in this context. As a result, when we encounter new beliefs that are formed in these ways, we can defeasibly discredit them even before investigating their truth.

Evolutionary critics of morality are *not* leveling this kind of critique. Debunkers do not seek to reveal the unreliability of evolutionary origins by showing that they have generated moral beliefs we know to be false. Indeed, this sort of critique would be self-defeating, for if we were able to sort true moral beliefs from false ones, then we could rely on that knowledge to correct for any epistemically pernicious evolutionary influence on the formation of our moral beliefs.

A different sort of argument – what I will call an *agnostic* genealogical critique – begins without any assumption of where truth in a given domain lies. Those who advance such critiques argue that even in our ignorance, we are well positioned to determine that beliefs caused in specific ways are unlikely to be true. The key element to such a critique is the claim that when it comes to certain topics, there is no discernible connection between the relevant doxastic origins and the truth (whatever it may be). Given our background beliefs about the nature of the doxastic influence in question, and of the general status (though not the content) of the target truths, we rightly believe that it would be at least a minor miracle were the influence a reliable one.

Suppose, for instance, that someone told you that there were exactly 5,422,000,000,000 fish in the world's oceans. You have no idea whether this number is even close. Then you discover that he landed on this figure by assuming that it was identical to the U.S. trade deficit for 2011. You now have all you need in order to discredit his belief. Even though you are in no position to verify the correct number, and even if the methods that went into figuring out the deficit numbers are in other contexts perfectly reliable, the

² I am centrally concerned with whether genealogical critiques undermine our chances for moral knowledge. Throughout the paper, I will focus on the threat to epistemic warrant, understanding such warrant as a necessary condition for knowledge. I do so purely for ease of exposition and expression, as I want to remain officially neutral on whether warrant or justification plays this role. I intend my remarks concerning warrant to apply equally well to justification.

doxastic origins in this case bear no discernible connection to whatever the truth might be.

The recent genealogical critiques offered by moral antirealists are of this second sort. These critics argue that if the moral facts are as realists depict them, then there is no discernible connection between evolutionary origins and stance-independent moral truths that could sustain confidence in the reliability of our moral faculties.

This is compatible with our moral judgments actually being more or less correct. After all, we might just as well luckily land on the truth about the number of fish in the sea. But precisely because luck would be essential to forming a true belief in such matters, the resulting beliefs are ones that we should not trust. And that undermines whatever presumptive warrant such beliefs might enjoy.

2. The Darwinian Argument

Suppose that evolutionary pressures have greatly influenced our moral faculties. Given the nature of such pressures, these causal origins selected for adaptive mechanisms of moral belief formation. Adaptive doxastic mechanisms will tend to yield beliefs that foster survival and reproductive success. But as we know from many other cases, the pragmatic gains of holding a belief are one thing – its truth, another. It would be a miracle were there anything like a close correspondence between the deliverances of a faculty shaped by evolutionary pressures, and a set of moral truths whose contents (if realists are correct) are fixed independently of the outputs of this faculty. If you have to rely on a miracle to sustain your position, then it is time to look for a different position. At the very least, the views generated from within the vulnerable position are subject to defeat – after all, it would be a massive coincidence were they true, and once we come to appreciate that, then we have decisive reason to abandon those beliefs.

That is the essence of the new genealogical critique of moral realism. Here is the line of reasoning presented a bit more carefully, reconstructed in what I will call the *Darwinian Argument*. None of the debunkers actually lays things out as explicitly as this, but this version seems to me to capture what is most challenging about the new debunking efforts:

D1. If moral realism is true, and if evolutionary forces have thoroughly shaped our moral faculties in ways that are *doxastically discriminating* (i.e., such that they dispose our moral faculties to generate beliefs with certain propositional contents rather than others)³, then it would be a massive coincidence were our moral faculties reliable.

D2. If it would be a massive coincidence were our moral faculties reliable, then absent any independent confirmation of their reliability, there is decisive reason to believe that our moral faculties are unreliable.

D3. There is no independent confirmation of the reliability of our moral faculties – their reliability can be confirmed only by showing that they have generated moral beliefs in which we have a high degree of warranted confidence.

Therefore,

D4. If moral realism is true, and evolutionary forces have thoroughly shaped our moral faculties in doxastically discriminating ways, then there is decisive reason to believe that our moral faculties are unreliable.

D5. Evolutionary forces have thoroughly shaped our moral faculties in doxastically discriminating ways.⁴

Therefore,

D6. If moral realism is true, then there is decisive reason to believe that our moral faculties are unreliable.

For those who prefer all loose ends tied up, we can continue the argument by means of highly plausible premises:

D7. If there is decisive reason to believe that our moral faculties are unreliable, and if those reasons are accessible to us, then this defeats any warrant we might have for our moral beliefs.

D8. Those reasons (in the form of D1-D6, and the considerations that support them) are accessible to us.

Therefore,

D9. If moral realism is true, then our moral beliefs are unwarranted.

D10. If our moral beliefs are unwarranted, then we have no moral knowledge.

Therefore,

D11. If moral realism is true, then we have no moral knowledge.

³ A brief word about the phenomenon that I am calling *doxastic discrimination*. There is little question that evolutionary forces have shaped our cognitive capacities in a wide variety of ways. But that by itself casts no doubt on the reliability of our doxastic faculties. Our ability to discern the truths of (e.g.) chemistry, or metaphysics, depends in a way on our brains having evolved as they did. But these evolutionary origins do not undermine the beliefs of chemists or those doing mereology. That is precisely because evolutionary forces are not in these cases doxastically discriminating. They have not pushed us to hold some beliefs rather than others on these topics.

⁴ Justin Clarke-Doane ((2012) “Morality and Mathematics: The Evolutionary Challenge,” *Ethics* 122: 313-40) argues that debunkers misunderstand the nature of their own critiques, and do not need to assume that selective pressures have actually had any influence on our moral faculties. According to Clarke-Doane, all debunkers need to do is to establish the truth of a counterfactual: Had the (basic) moral truths been very different, our moral beliefs would have been the same. It seems to me, though, that what would put us in a position to favorably assess this counterfactual is possession of excellent evidence that selective pressures actually work in the way postulated by D5. That, combined with fidelity to the debunkers’ intentions, makes me think it acceptable to incorporate D5 into their master argument. I offer a further reason for incorporating D5 below, in n. 21. I discuss Clarke-Doane’s views in section 5b.

I will not challenge any of the steps from D7-D11. Both D7 and D10 strike me as extremely plausible, and I will assume their truth in what follows. D8 is also plausible, at least for anyone reading this essay, and for others familiar with the literature that has grown up around the debunking efforts. Further, it would be discomfiting (to say the least) if the only way to preserve the presumptive warrant of moral beliefs while endorsing moral realism is to have closed oneself off to the genealogical critiques that have recently been advanced. So I will proceed on the assumption that D8 is correct.

The upshot: Once debunkers have secured D6, the path to D11 is clear. Henceforth, then, I will confine myself to what I consider the Darwinian Argument proper, i.e., D1 through D6.

D2 is plausible. Its qualifying clause – that there are no independent means of confirming the reliability of our moral faculties – is crucial. For we are sometimes warranted in our beliefs, even if it would be massively coincidental were the faculties that generated them reliable ones. Such warrant comes in the form of independent confirmation, i.e., confirmation that relies on warranted beliefs generated by faculties other than those whose reliability is in question. Suppose, for instance, that we know that someone's beliefs on a given topic have resulted from brainwashing. Still, she (or we) may rightly regard those beliefs as plausible if they are also supported by evidence that is free of brainwashing and independently credible.

D3 is plausible as well. It represents an autonomy thesis about the moral, and is essential to the success of the debunkers' efforts. For if it were false, then even if we have to initially suspend judgment about all of our moral beliefs, there might be routes to vindicating our moral faculties that do not rely on such beliefs to do the job. I will raise a potentially serious concern about this premise toward the end of this essay. For now, though, we can proceed on the assumption of its truth, and see whether we might locate a vulnerability elsewhere in the argument.

This leaves two premises, D1 and D5. Before examining D1, we should pause to consider perhaps the most natural strategy for replying to evolutionary debunkers. This strategy, as we will see, raises some concerns about the truth of D5.

3. The Natural Reply

Genealogical critiques will be plausible only to the extent that they imply a means of distinguishing the real origins of our moral faculties from mere pretenders. Such means, however, would also provide us with a way of determining the *limits* of a given type of doxastic influence. Focusing here just on evolutionary origins, realists should thus be able to point to specific moral beliefs, doxastic dispositions or operations of moral faculties that are immune to such influence. These can serve as untainted benchmarks for assessing the probity of those beliefs that originate in evolutionary pressures.

Applying these thoughts to the moral case, the realist can resist genealogical critiques if she can locate moral beliefs that meet three conditions:

- (i) they are immune from evolutionary influence;
- (ii) we can know which beliefs those are, and
- (iii) such beliefs are highly presumptively warranted.

We could deploy moral beliefs that satisfy these conditions to assess the merits of those with evolutionary origins.

We can formulate this strategy in what I will call the *Natural Reply*:

N1. If some of our highly presumptively warranted moral beliefs are not the product of evolutionary influence, and we can know which ones are thus immune, then we can utilize such beliefs to ascertain the plausibility of those that are the product of evolutionary pressures.

N2. Some of our highly presumptively warranted moral beliefs are not the product of evolutionary influence.

N3. We can know which beliefs those are.

Therefore,

N4. We can utilize such beliefs to ascertain the plausibility of those that are the product of evolutionary influences.

The truth of N2 casts doubt on D5. Further, once armed with a set of warranted moral beliefs that are immune from evolutionary influence, we can resuscitate the warrant of at least some of those that are not. Depending on the outcome of this effort, we could thereby confirm the general reliability of our moral faculties by showing that they had in many cases generated moral beliefs in which we can have warranted confidence.

Though Michael Huemer never explicitly refers to the line of reasoning that I am calling the Natural Reply, it is clear that his response to evolutionary debunkers exemplifies its workings. The moral beliefs that Huemer cites as epistemically pure are wholly formal ones. Huemer provides six examples of the sort of beliefs he has in mind.⁵ A couple of characteristic ones:

- If x is better than y and y is better than z , then x is better than z
- If it is wrong to do x , and it is wrong to do y , then it is wrong to do x and y .

The problem with Huemer's line of argument is that even if he is right in his choice of influence-free beliefs, such beliefs do not seem, either by themselves or in conjunction with nonmoral claims, sufficient to warrant any specific, substantive moral beliefs over others. Here we are up against a standing problem for classical foundationalism: that of having too few privileged beliefs from which to derive the wide variety of beliefs we intuitively regard as highly justified.

⁵ Michael Huemer (2008) "Revisionary Intuitionism," *Social Philosophy and Policy* 25: 368-92. All of the examples appear on p. 386.

Can we move beyond such formal beliefs to substantive moral commitments? Perhaps. Here is how we would do that. The primary evidence for the influence of evolutionary pressures on our moral faculties is the fact that many of our paradigm moral beliefs are highly adaptive.⁶ So if there are warranted moral beliefs that are unlikely to be adaptive, then, by the standards that obtain among those seeking to posit such evolutionary origins, we have excellent reason to think that we have located moral beliefs that have their origins elsewhere.

There is bound to be controversy about how adaptive a given type of belief is likely to be. But there do seem to be some that appear anything but adaptive. Among many possibilities: those that counsel impartial benevolence, compassion for vulnerable strangers, kindness to small animals, concern for distant peoples and future generations, and speaking truth to power. We can use such beliefs and dispositions to do the familiar work of subjecting our other moral beliefs to critical scrutiny.⁷

Evolutionary debunkers should reply by trying to show that even these beliefs are the products of evolutionary influences. Any success on this front will undermine N2.

⁶ The strength of this evidence is unclear, however, since we do not know that people hold those beliefs *because* they are adaptive. (A point made by Derek Parfit (2011) *On What Matters*, Vol. 2, Oxford: Oxford University Press, p. 535; Roger White makes a similar point in his fine, wide-ranging paper (2010) “You Just Believe That Because...,” *Philosophical Perspectives* 24: 573-615, at 586.) The evolutionary account is in many ways a just-so story: We have the data here, in the form of our moral beliefs and dispositions, and are trying to infer a literally prehistoric causal story on this basis. The claim of widespread evolutionary influence (of a doxastically discriminating kind) is not, on its face, implausible. But we should not pretend that it is free of speculation, either.

⁷ This is a strategy utilized by Greene (2008) and Peter Singer ((2005) “Ethics and Intuitions,” *The Journal of Ethics* 9: 331-52), both of whom share Huemer’s suspicion that evolutionary doxastic influence amounts to taint. Singer and Greene argue that characteristically deontological intuitions are heavily subject to evolutionary pressures, and are to be discounted on that basis, while characteristically consequentialist ones are largely free of such influence and originate in rational reflection.

There is some reason to doubt that the Natural Reply can be developed along these lines. In arguing for the insulation of consequentialist intuitions from evolutionary influence, Singer tells us, for instance, that the application of consequentialist reasoning often relies on the intuition that “it is a bad thing that a person is killed.” And he says of this intuition that it “does not seem to be one that is the outcome of our evolutionary past” (2005: 350). But by the lax standards that typically prevail in moral discussions that postulate evolutionary influence, we seem to have extremely strong reason to assign such influence here, since it is surely adaptive to be very averse to death, and adaptiveness is (for Singer and others) about all that is needed to supply an evolutionary causal explanation.

The arguments of Singer and Greene are subject to a variety of objections, many of which are ably noted by Selim Berker (2009) “The Normative Insignificance of Neuroscience,” *Philosophy and Public Affairs* 37: 293-329. Guy Kahane ((2010) “Evolutionary Debunking Arguments,” *Nous* 45: 103-25) is quite good about how difficult it is to limit the scope of evolutionary-debunking arguments within normative ethics. His criticisms of Singer and Greene are especially apt on this point. See Kahane 2010: 113, 119-20.

N2 may indeed be less plausible than it initially appears to be. To see this, begin with a concession: Not all presumptively warranted moral beliefs are adaptive. Hence, to the best of our knowledge, not all such beliefs have been selected for. So it is implausible to suppose that all such moral beliefs are the direct product of evolutionary forces.

But they might all be the *indirect* product of such forces. Even maladaptive beliefs may be the products of evolutionary pressures, if such pressures explain the origins of the doxastic dispositions that in turn produced such beliefs. For instance, debunkers might argue that the allegedly immune moral beliefs mentioned above (*viz.*, those counseling impartial benevolence, kindness to vulnerable strangers, etc.) in fact originate in dispositions of (say) sympathy and reciprocity, which were themselves selected for because they were so adaptive.

Such indirect explanations may be plausible. But in order to entirely fend off the Natural Reply, debunkers must provide an indirect explanation of the origins of *every* moral belief that is both highly presumptively warranted and apparently maladaptive. Debunkers are *very* far from having done this. In the absence of specific and substantiated indirect explanations, we lack reason to doubt the credibility of these apparently maladaptive moral beliefs, and we can develop the Natural Reply by using those beliefs as secure points from which to rebuild our stock of credible moral commitments.

If, in order to resist this move, debunkers nevertheless insist that there *must* be an evolutionary genealogy that can best explain the origins of apparently maladaptive (though highly presumptively warranted) moral beliefs, then their view threatens to become untestable. After all, the basis on which they confidently assign evolutionary origins to so many of our moral beliefs is precisely the supposition that such beliefs would enhance the fitness of those who held them. The absence of any evidence of this sort for particular moral beliefs implies a limitation on the scope of justified debunking efforts. And that opens the door to the development of the Natural Reply.

That said, N2 might nonetheless be false, and D5 might be true. But absent a great deal of further empirical and philosophical work, this remains an unsubstantiated possibility that cannot rightly cast doubt on N2 or vindicate D5. That is a much greater problem for the debunker than for the realist. After all, the burden is on the evolutionary debunker to substantiate every one of the premises of the Darwinian Argument. That argument is meant to serve as a defeater of the presumptive warrant that even debunkers grant to many moral beliefs. Such warrant is not defeated until we have warranted confidence in all of the premises of the Darwinian Argument. At this point, we do not. D5 is far from well confirmed.

4. The Argument from Uncertainty

Still, this leaves realists in a precarious position, hostage to the fortunes of future empirical research. Relying on optimism that the empirical outcomes will fall the realist's way is an awfully tenuous way of proceeding.

Indeed, debunkers can offer a modified evolutionary argument that begins with our puzzlement about the etiology of our moral beliefs and dispositions. Precisely because the actual origins of our moral beliefs are so hard to discern, we cannot be sure which ones result from evolutionary pressures and which do not.

This amounts to the claim that N3 is false. Here is how to turn that claim into a different debunking argument – call it the *Uncertainty Argument*:

U1. If, for all we know, there is a reasonable chance that evolutionary forces have thoroughly shaped our moral faculties in doxastically discriminating ways, and if moral realism is true, then, for all we know, there is a reasonable chance that our moral faculties are unreliable.

U2. For all we know, there is a reasonable chance that evolutionary forces have thoroughly shaped our moral faculties in doxastically discriminating ways.

Therefore,

U3. If moral realism is true, then, for all we know, there is a reasonable chance that our moral faculties are unreliable.

U4. If, for all we know, there is a reasonable chance that our moral faculties are unreliable, then we must suspend judgment about all of our moral beliefs.

Therefore,

U5. If moral realism is true, then we must suspend judgment about all of our moral beliefs.⁸

U2 represents the relevant weakening of the empirical thesis, D5, that is at the heart of the Darwinian Argument. U2 is plausible, and I will assume its truth in what follows. Though U4 is contestable, I will not argue against it here. Let us turn instead to U1.

U1 incorporates two kinds of uncertainty. The first is uncertainty about the extent of evolutionary influence on our moral faculties. The second is uncertainty about how deleterious such influence would be, were moral realism true. The first sort of uncertainty is affirmed in U2, which we are accepting. But what of the second sort – why think that the combination of realism and evolutionary influence yields a reasonable chance that our moral faculties are unreliable?

The answer, I think, must refer again to the idea of massive coincidence at the heart of the Darwinian Argument. Moral realism and evolutionary influence together create a reasonable chance of moral unreliability just because that combination would make the reliability of our moral faculties a massive coincidence. In other words, the best explanation of why anyone would find

⁸ Something like this line of argument is given an airing (without being endorsed) by Kahane 2010: 118.

U1 plausible is because she already finds D1 appealing. What we need to do, then, is to tackle D1 directly.

5. Against D1: Diagnosing Massive Coincidence

Premise D1 says that:

If moral realism is true, and if evolutionary forces have thoroughly shaped our moral faculties in doxastically discriminating ways, then it would be a massive coincidence were our moral faculties reliable.

Debunkers have thus far not been clear enough about the precise nature of the alleged coincidence. In this largest section of the paper, I explore the five most promising interpretations of massive coincidence within the context of evolutionary critiques, and try to show that none of these interpretations gives us adequate reason to accept D1.

a. The odds

To speak of massive coincidence is to speak of great improbabilities. Perhaps the simplest understanding of the relevant improbability here is in terms of numbers. There is an indefinitely large set of possible moral commitments, and evolutionary pressures have led us to adopt just a small number of these. If we suspend judgment about which outlook is actually correct – as we must do, according to debunkers – then the chance that selective pressures have led us to the realistically correct set of moral commitments is extremely small. Thus it would be an extraordinary coincidence were our core moral beliefs on target, given their evolutionary origins.⁹

Assessments of coincidence must be rendered from some epistemic standpoint or other. For those with no knowledge of magnetism, for instance, it will appear miraculous that two pieces of metal repeatedly attract or repel one another. Once we impart background assumptions about magnetism, however, we eliminate any appearance of coincidence.

The idea, then, is that we are to assess the probability that our moral faculties are reliable *without* the aid of any background moral beliefs. For if we were allowed to introduce such beliefs in assessing the reliability of our moral faculties, then we would have stacked the deck in favor of the realist. We would be in a position to certify the reliability of our moral faculties simply by showing that they have produced a set of privileged (background) moral beliefs whose truth we are taking as a fixed point in our investigations. That would make things too easy for the realist.

⁹ This line of reasoning is suggested at various points in Street (2006), and can be found explicitly in her (2008) “Reply to Copp: Naturalism, Normativity, and the Varieties of Realism Worth Worrying About,” *Philosophical Issues* 18: 207-28, at 208ff.

So we must suspend judgment about the truth of our moral beliefs. Once we do that, however, we have no feasible way of determining the credibility of those beliefs – on the plausible assumption that (trivial, aberrant cases aside¹⁰) nonmoral claims alone are incapable of certifying any substantive moral claims. From the perspective of moral ignorance, which we are required to adopt in order to avoid begging questions, it would be hugely coincidental were our moral beliefs largely on target. That is because of the numbers – our actual moral beliefs represent only a small portion of all possible moral beliefs.

We can summarize this line of reasoning in the *Probability Argument*:

P1. If moral realism is true, and if evolutionary forces have thoroughly shaped our moral faculties in doxastically discriminating ways, then the odds are very low that our actual moral beliefs, among all psychologically possible moral beliefs, are correct.

P2. If the odds are very low that our actual moral beliefs, among all psychologically possible moral beliefs, are correct, then it would be a massive coincidence were our moral faculties reliable.

Therefore,

D1. If moral realism is true, and if evolutionary forces have thoroughly shaped our moral faculties in doxastically discriminating ways, then it would be a massive coincidence were our moral faculties reliable.

I think that P2 is plausible. But P1 is problematic, for three reasons.

(1) There may be certain conceptual constraints on what can qualify as a moral view. Street claims that realists are in no position to exclude, *ex ante*, the possibility that the inventory of stance-independent moral truths might include the badness of survival, the worthlessness of our children's lives, the goodness of returning a kindness with aggression.¹¹ If every possible moral proposition is on the table as a candidate for stance-independent truth, and we, as a result of selective pressures, believe just a small fraction from this array of possibilities, then the odds really are stacked against us, and (we might think) the chances of our having been influenced in reliable ways are quite slim.

But realists might deny that every possible moral proposition is a viable candidate for stance-independent truth. They can argue instead for a moral semantics according to which certain moral propositions are reference-fixing, such that denial of (enough of) these propositions shows that an agent is no longer talking about *morality* at all. On this view, semantic competence with moral concepts is determined in part by whether a speaker is prepared to sin-

¹⁰ I am thinking here of the sort of examples that Arthur Prior advanced in (1960) "The Autonomy of Ethics," *Australasian Journal of Philosophy* 38: 199-206. One such: *grass is blue*; therefore, *either grass is blue or smoking is immoral*. If this disjunction is a moral claim, then we have deduced a moral claim from a nonmoral one. If the disjunction is not a moral claim, then add to it the nonmoral premise that *grass is not blue*, and so deduce from two nonmoral premises the moral conclusion that *smoking is immoral*.

¹¹ These examples appear in Street (2008: 208).

cerely affirm certain moral propositions and deny others. For instance, a person who understands and rejects the propositions that there are wrong-making features of genocide, rape, betrayal and sadism is not really speaking our language. She may use the words “wrongness” or “immorality,” but she is utilizing a different concept from the one we rely on when we condemn the practices just mentioned. A willingness to sincerely affirm certain platitudes about morality is a precondition of semantic competence in the area.

If this picture is correct, then the realist’s strategy of invoking those reference-fixing moral claims when assessing the reliability of our moral faculties will *not* be question-begging. If it is a precondition of our talking about morality that we affirm certain paradigmatically moral propositions, then these propositions are common ground that all competent speakers are entitled (indeed, required) to take for granted. The realist’s invocation of such claims will simply ensure that she is relying on the same concepts as her antirealist interlocutors. Affirming such claims does not presuppose the reliability of our moral faculties; rather, it presupposes minimal semantic competence with moral concepts.¹² If this is so, then the realist can use these conceptually privileged propositions as a secure basis from which to develop the Natural Reply, and to insulate her position from this charge of coincidence.

I know that some will find this appeal to semantic competence less than persuasive. For the doubters, realists can offer two further replies to the “low-odds” understanding of massive coincidence.

(2) The line of reasoning that supports P1 would prove too much. There is an indefinitely large number of psychologically possible *perceptual* beliefs. Our perceptual faculties have led us to endorse only a small subset of them. Suppose that, in order to avoid begging questions, we are not allowed to take any such judgments for granted in assessing the reliability of those faculties. If the original line of argument is sound, then this should lead us to conclude that we have no perceptual knowledge if perceptual truths are construed realistically. No debunker believes this. Almost no one else does, either.

Of course we can resist this line of reasoning by assuming that our perceptual faculties are largely reliable. But by parity of reasoning, we should then be licensed to make the same assumption about our moral faculties. And once we do, we can easily reinstate our confidence in them.

Or we can resist this line of reasoning by assuming the truth of at least some nearly indubitable perceptual judgments, and then using these to assess other perceptual judgments and the reliability of our perceptual faculties. But

¹² The strategy of taking certain non-tautologous moral claims as partly definitive of morality is of course a controversial one, and I do not want to do anything here other than register some sympathy for the view. The view sketched here is one that Terence Cuneo and I are developing in a paper in progress titled “Moral Fixed Points.” The view is also suggested by William FitzPatrick in his valuable unpublished essay, “Debunking Evolutionary Debunking of Ethical Realism.” The *locus classicus* of such a view is Philippa Foot’s work of the late 1950s. See, e.g., “Moral Beliefs” and “Moral Arguments,” both reprinted in her (1978) *Virtues and Vices*, Berkeley: University of California Press.

again, once we extend such license, we must do so for our moral faculties as well, thus enabling realists to reject the charge of massive coincidence understood in this way.

The important point here is that if we grant that evolutionary pressures have greatly narrowed the scope of doxastic possibilities in the moral realm, this is not *by itself* any evidence that they have done so in a way that has led us away from moral truths, realistically construed. The genealogical critic needs to point to some further feature of the etiology of moral beliefs that creates a special problem for the reliability of our moral faculties. We will soon consider four such features.

(3) The low-odds conception of the relevant coincidence is subject to a *tu quoque* argument that will beset every non-error-theoretic proponent of an antirealist genealogical critique. No matter one's metaethic, it is true that our moral beliefs represent only a tiny fraction of those that are psychologically possible. If we are required to suspend judgment about which moral views are true and which false, then from that agnostic perspective, it would be a miracle were the moral beliefs we actually hold correct. The probability that ours is the (roughly) correct moral outlook will be vanishingly small.

Constructivists have a reply here, but any success it might enjoy can supply the basis for an adequate realist counter to debunking efforts. Constructivists will say that even if we suspend judgment about all substantive moral claims, the odds may be quite high that our moral beliefs are correct. Elevating the odds requires, on their view, that we grasp the correct *constructive function*, which provides the metaphysical grounding for moral truth. This constructive function takes a set of attitudes as inputs, and yields substantive moral standards as outputs. Armed with knowledge of the correct constructive function, and the ability to apply the function to the relevant attitudes in a reliable way, we can come to warranted first-order moral knowledge without presupposing the truth of any substantive moral beliefs. We will effectively have deduced the truth of such beliefs from a statement of the correct constructive function, an accurate report of the content of the attitudes that represent the relevant inputs, and a successful instance of reasoning that evinces an understanding of how to apply the constructive function. No substantive moral beliefs will play a role in this effort. So even if we suspend judgment about all such beliefs, the odds that we have lighted on the moral truth may be quite high. Moral knowledge is within our grasp, if constructivism is true.

This is the path that Sharon Street takes when arguing that constructivism can withstand the charge of massive coincidence.¹³ She believes that the attitudes that serve as the input of her constructivism have been influenced by evolutionary forces. But the origins of the truth-fixing attitudes are, she thinks, irrelevant to constructivist efforts to vindicate the credibility of our moral beliefs. However we came by our attitudes, we can gain moral

¹³ Street (2006: 152-54).

knowledge by reasoning *from* these attitudes, via the appreciation of some constructive function, to the moral truth. So long as we can do this in a reliable way, it will be no coincidence that we are able to distinguish right from wrong.

Suppose we set aside any reservations we might have about Street's specific version of constructivism, the details of which need not concern us here.¹⁴ Still, any success that constructivists may have in obtaining moral knowledge can, at least in principle, be replicated by the realist. For constructivists, the constructive function plays the role of supplying a reference-fixing standard for moral concepts. That, if anything, is what would enable constructivists to move from an appreciation of the constructive function to a determination of first-order moral principles without reliance on substantive moral beliefs. But there is no compelling reason to suppose that realists are unable to supply reference-fixing standards of their own for moral concepts. (Certainly the task is not easy. But it is no easier for constructivists, as is clear from the abiding controversies surrounding the plausibility of every candidate constructive function.)

It may, of course, be dubious to suppose that we are able to reason from reference-fixing standards to first-order moral truths in any event. We might also wonder whether such reasoning will really be immune to the influence of evolutionary pressures. Finally, we might ask whether the reasoning that results in the identification of reference-fixing standards for moral concepts is, as Street seems to think, free of doxastically discriminating evolutionary influences. (She does not entertain the thought that her preference for constructivism, much less her preference for the specific constructive function she endorses, is itself a product of possibly distorting genealogical pressures.)

All such suspicions, to the extent that they have merit, apply equally to constructivism and to realism. If the feats of reasoning described in the previous paragraph are possible, and resistant to the worries just aired, then realists can avail themselves of such reasoning to quell genealogical doubts. If, instead, these kinds of reasoning fall prey to the doubts just described, then constructivism would be prey to them as well. In either case, there is a parity between realism and constructivism that supports the *tu quoque* response.

True, that response does nothing to deflect the worries of moral error theorists. But I take it that the second reply is sufficient to show that even error theorists who are not skeptics about perceptual knowledge should opt for a different understanding of the massive coincidence that allegedly undermines realist hopes for moral knowledge. And the first reply (the one that depends on semantic competence), though at present insufficiently developed, may hold promise for responding to this worry as well.

¹⁴ Her positive view is most carefully worked out in "Constructivism about Reasons" (2008) *Oxford Studies in Metaethics*, Vol. 3, Oxford: Oxford University Press, pp. 207-45, and (2012) "Coming to Terms with Contingency: Humean Constructivism about Practical Reason," in James Lenman and Yonatan Shemmer, eds., *Constructivism in Practical Philosophy*, Oxford: Oxford University Press, pp. 40-59.

b. Insensitivity

If evolutionary forces are shaping our moral faculties in doxastically discriminating ways, then we might worry that these faculties are insensitive to the truth. After all, such forces will be pushing us to doxastic practices that are adaptive, regardless of whether those practices issue in true beliefs. It appears that under the influence of selective pressures, we would continue to have the moral beliefs we do, even if they were false. So whatever accuracy our moral beliefs might possess would be massively coincidental.¹⁵

We can reconstruct this line of reasoning as the *Insensitivity Argument*:

I1. If moral realism is true, and if evolutionary forces have thoroughly shaped our moral faculties in doxastically discriminating ways, then we would hold the moral beliefs we do, even if they were false.

I2. If we would hold the moral beliefs we do, even if they were false, then it would be a massive coincidence were our moral faculties reliable.

Therefore,

D1. If moral realism is true, and if evolutionary forces have thoroughly shaped our moral faculties in doxastically discriminating ways, then it would be a massive coincidence were our moral faculties reliable.

Here is a quick reply to the Insensitivity Argument. Our core moral beliefs express metaphysically necessary truths. Therefore, it is not the case that we would continue to hold them *even if they were false*. They cannot be false; the scenario sketched in the consequent of premise I1 is impossible. So there is no truth about whether we would continue to hold them under those impossible circumstances. And therefore no basis for thinking that I1 is true. So this worry is groundless.¹⁶

There are two problems with this reply. First, most of our moral beliefs do not have necessarily true propositions as their contents, so at best this reply would still leave the vast majority of our moral beliefs prey to insensitivity concerns. Second, we cannot establish an agent's epistemic sensitivity simply by showing that the contents of her beliefs are metaphysically necessary truths. For agents can believe such truths for the wrong reasons, and when they do, this reveals a kind of insensitivity that precludes hopes for knowledge.

¹⁵ Claims that support an insensitivity charge can be found in Joyce (2007: 183) and Street (2006: 132). Ruse and Wilson (1986: 186-87) also raise this worry, though they conflate metaphysical and epistemological issues in doing so: "The evolutionary explanation makes objective morality redundant, for even if external ethical premises did not exist, we would go on thinking about right and wrong in the way we do. And surely, redundancy is the last predicate that an objective morality can possess."

¹⁶ See the fine complementary discussion by Erik Wielenberg (2010) "On the Evolutionary Debunking of Morality," *Ethics* 120: 441-64, at 454-45, for further argument defending the realist against a slightly different insensitivity charge leveled by Ruse (1998: 254).

Justin Clarke-Doane, though not himself a debunker, argues that the realist cannot rest easy just by noting the metaphysical impossibility of there being different basic moral truths. He claims that we ought to consider the relevant counterfactual in the light of conceptual possibility: “[H]ad – for all we can intelligibly imagine – the [basic] moral truths been very different, our moral beliefs would have been the same.”¹⁷ Though Clarke-Doane does not speak of insensitivity per se, it is clear that he is concerned about precisely that. For if this counterfactual is true, then it seems that we are failing to form our moral beliefs in ways that are adequately responsive to the moral truth. If that were so, then it would be a massive coincidence were our moral faculties reliable.

A potential problem for this line of argument is that the counterfactual may fail to describe an intelligible possibility. (See the first reply offered in the previous section.) It may be that we cannot coherently imagine the fundamental moral truths being very different from what they actually are. We might be tempted to say, for instance, that a set of rules that celebrated the intrinsic value of misery and cruelty, that incorporated nonderivative requirements to wantonly kill, rape, torture and betray innocents, could not qualify as a moral system. Those who might sincerely advertise such a system as a moral one would be making a conceptual error.¹⁸

Another possible reply takes a cue from reliabilist efforts to deal with demon worlds. We ought to construe reliability as a relation across normal worlds. That a belief-forming process lands us in inaccuracies in worlds very different from the actual one does not impugn the kind of reliability sufficient for conferring warrant here in our world. If that is so, then even if Clarke-Doane is describing an intelligible scenario, the unreliability of our actual belief-forming processes in that distant possible world would not affect the epistemic status of those paradigm moral beliefs that are necessarily true.

I do not want to take a stand here on the plausibility of either of these replies. For even if they work, there is a residual worry that has not been discharged. The worry is clarified if we distinguish sensitivity from reliability. Reliability is a matter of generating mostly true beliefs. Sensitivity is a matter of non-accidentally believing propositions when they are true, and non-accidentally refraining from believing them when they are false. Insensitive belief-forming methods or processes can be reliable. For instance, we might accrue a much better than average track record for our weather predictions

¹⁷ Clarke-Doane (2012: 321).

¹⁸ Clarke-Doane (ibid., 335) argues that the existence of actual moral disagreement about basic moral principles provides defeasible evidence for the relevant conceptual possibility. Perhaps. Such disagreements need to include interlocutors who meet two conditions: (1) they have sincerely affirmed as moral principles *radically* revisionary rules, of the sort I just described; (2) they suffer from no conceptual deficiencies when issuing those sincere affirmations. We cannot know that the second condition is met just by pointing to examples in which the first is satisfied.

when basing them on a reading of the daily astrology column. But this method is still insensitive, for any success one's predictions enjoy is accidental.

It is not easy to explicate the notion of accidentality here. But I think this will suffice for our purposes: Agents accidentally come to true beliefs when they fail to base their beliefs on adequate insight, evidence, experience or reasons. In the evolutionary context, the idea is that any accuracy enjoyed by our moral beliefs is really accidental, a matter largely or entirely of luck, since we have come to hold those beliefs not on the basis of appropriate experiences, or attentive consideration of the evidence or reasons that support our beliefs, but rather because, to put it crudely, evolution has "programmed" us to have those beliefs. Since the program was not designed to generate beliefs based on their truth, but only based on their adaptiveness, it would appear to be the sheerest luck were our moral beliefs accurate. So: Even if our moral faculties are reliable, they are insensitive, and this undercuts the chances that they generate beliefs that amount to knowledge.

As I read him, this captures Richard Joyce's central debunking argument, one that begins with a nice thought experiment.¹⁹ Suppose that we know of the existence of a pill that causes people who take it to believe that Napoleon was defeated at Waterloo. An immediate side effect of the pill is amnesia about having taken it. Once you know of the existence of the pill and its effects, you lose any warrant you once had for your belief that Napoleon was defeated at Waterloo.

This is not because we know such pills to be unreliable. Indeed, for all we know, they might invariably cause us to have true beliefs. Rather, once we believe that our belief that *p* might (for all we know) have been caused by a pill, we must suspend judgment about *p* because we are unable to discount a doxastic etiology that would reveal our belief to have been formed in a way that is entirely insensitive to the truth. While we cannot be sure that our beliefs originate from such a pill, we can be very confident that if they did, then it would be a massive coincidence were they true. Once you pop that pill, you will think that Napoleon was defeated – even if he wasn't.

Joyce thinks that belief pills are analogous to evolutionary forces in the moral case. We cannot be sure that our moral beliefs have emerged from selective pressures. But if they have, then they are the products of causal processes that are alethically insensitive. That is because selective pressures exert their influence with an eye, so to speak, exclusively to inculcating adaptive doxastic dispositions, rather than truth-tracking ones. If we are unable to exclude the possibility that our moral beliefs have emerged from such insensitive processes, then we must suspend judgment about all of our moral beliefs.

Joyce's critique is sound only if two conditions are true: (1) to the extent that our moral beliefs are the product of evolutionary forces, such beliefs are formed in ways that are insensitive to whatever moral truth there may be; and

¹⁹ Joyce (2007: 181-82).

(2) our inability to exclude the possibility of insensitive doxastic origins for a set of beliefs *S* mandates suspension of judgment regarding all beliefs within that set.

Even if we were to grant condition (1), this general line of reasoning is deeply problematic, since the truth of (2) would lead to wholesale skepticism. For almost every possible belief, we are unable to exclude the possibility that it is the product of evil demons or manipulative neuroscientists. If that inability requires us to suspend judgment, then we must be skeptics about nearly everything. Assuming that such skepticism is erroneous – and also assuming that debunkers are trying to identify something especially worrisome about moral knowledge, rather than developing an argument for global skepticism – this line of reasoning will fail.

I think that Joyce’s argument really does rest on (2), as he is appropriately modest about the nature and extent of evolutionary influences on our moral faculties. It is our inability to exclude the possibility of insensitive doxastic origins that undermines warrant for moral beliefs, rather than the debunkers’ ability to assign a reasonably high likelihood that our moral beliefs have insensitive origins. But debunkers might try for more than Joyce does, and claim that it is *likely*, perhaps even *highly likely*, that all of our presumptively warranted moral beliefs have such origins, and that this likelihood is what justifies moral skepticism.

But if this is the relevant claim, then it is unsubstantiated. The basic reason is simple. As noted above in the discussion of the Natural Reply, debunkers have fallen far short of defending the claim that all of our presumptively warranted moral beliefs are products (even indirectly) of selective pressures. So even if evolutionary doxastic origins are insensitive to moral truth, realistically construed, it has not yet been shown that it is highly likely that all or even most presumptively warranted moral beliefs have such origins. And so we as yet lack reason to think that all or most presumptively warranted beliefs are formed in ways that are alethically insensitive.²⁰

In general, we need to think a bit more carefully about what it is for a belief to be “the product” of evolutionary influences. On the model required by the belief-pill analogy, a belief is the product of such influences just when it is wholly caused by them. The doxastic influence of a belief pill is 100 percent. Its operation is direct and guaranteed. You take the pill, you have the programmed belief, no matter whether the belief is true or not. But so far as we know, evolutionary influences are at best only indirectly efficacious in causing us to have the specific moral beliefs we do. Such influences may affect only a subset of our moral beliefs. For those they do affect, what influ-

²⁰ As Kevin Brosnan argues, there is an even deeper reason to resist any such likelihood argument. The basic concern is that any likelihood argument must make assumptions about the prior probability that certain moral claims are true. But the debunker is in no position to do that, and without such an assumption, any likelihood argument will collapse. See Brosnan (2011) “Do the Evolutionary Origins of Our Moral Beliefs Undermine Moral Knowledge?” *Biology and Philosophy* 26: 51-64, at 54-57.

ence they wield is not a matter of wholesale causation. Instead, such pressures supply us with affective and cognitive dispositions that in many cases only incline us in certain doxastic directions, without entirely fixing the content of our beliefs.²¹

Now this might seem a fairly lame reply. After all, if the belief pill works only 70 percent of the time, rather than every time, it still seems to represent a real skeptical threat. But I think, in fact, that the reply is not at all inadequate. To see why, let us give debunkers all they want, and assume that evolutionary forces are indeed distorting. Yet so long as such forces are not the sole cause of our presumptively warranted moral beliefs, then *even if* evolutionary pressures are distorting ones, there is the possibility that other doxastic operations can compensate for the distortions and lead us to the moral truth in ways that are perfectly familiar from other instances in which we have corrected for distorting doxastic influences.²²

Think, for instance, of non-evolutionary genealogical critiques that cite historical or cultural pressures as untoward influences on our moral faculties, and seek to impugn moral beliefs as distorted artifacts of those causes. Yet Southern antebellum abolitionists were able to transcend the influences of their time and place; so too were advocates of female suffrage in societies suffused with ideologies of male dominance. The precise means by which we are able to overcome such pernicious influences and gain moral insight are of course highly controversial. But we do recognize cases in which agents have been able to correct for very substantial distorting pressures on their moral faculties. It is possible, of course, that while we can detect and resist pernicious historical and cultural influences on our moral beliefs, we are quite helpless at doing so when it comes to indirect evolutionary pressures. But debunkers have said nothing on behalf of this special status for evolutionary forces, nothing to show that evolutionary forces possess an immunity to correction that is unique among all possibly distorting doxastic influences.

The lesson here is this: Once the debunker admits that the influence of possibly distorting doxastic causes is less than complete, the door is open for the realist to argue that there is room for corrective efforts that enable us to appreciate the distortions for what they are, and to lead us to the path of

²¹ Joyce accepts this. He does not believe that evolutionary pressures have directly caused us to have specific moral beliefs. Rather, he argues that evolutionary forces have indirectly influenced all of our moral beliefs, because such forces are primarily responsible for our having moral concepts in the first place. But for selective pressures, we would not have developed moral concepts at all. This of course is a highly speculative hypothesis. Suppose we grant it. Still, it is not clear how damaging it is. After all, we might say the same thing about our mathematical or perceptual concepts – but for the influence of selective pressures, we would not have such concepts at all. That by itself casts no doubt on the plausibility of any of our mathematical or perceptual beliefs. (This explains why I have taken D5, with its focus on doxastic discrimination, to be essential to the strongest antirealist evolutionary critique.)

²² This is perhaps the most important of the many significant points made by FitzPatrick (unpublished manuscript).

wisdom. The contamination of our moral thinking would be local rather than global.

The key worry presented by Joyce's belief pill is that such a pill would undermine an appropriate basing relation needed to confer knowledge. Knowledge must be based, in the right ways, on adequate insight, evidence, reasons or experience. The nature of the basing relation is a subject of great debate, and I want to entirely sidestep that here. Still, we can be confident that beliefs about Napoleon that stem from taking a pill cannot amount to knowledge. That is because we would not be basing our beliefs on any appreciation of the relevant evidence, but rather simply by having had a belief implanted in us by a cause that, so far as we know, is wholly unrelated to any recognizable source of historical knowledge.

If selective pressures operated like a belief pill – effectively short-circuiting an agent's powers of deliberation, and just implanting a belief without allowing an agent's rational insight or assessment of evidence to ground her belief – then we would have real cause for worry on this score. But selective pressures, operating as indirect doxastic influences, still allow us to reflect critically on the nature of morality, assess evidence for moral claims, fully deliberate about the truth of such claims and base our beliefs on the grounds we find most compelling. Beliefs that emerge from such a process can be well grounded.

True, if the beliefs that we are relying on to do the corrective work are themselves implanted in us by possibly distorting evolutionary pressures, then the realist is no further along in replying to the insensitivity charge. Our moral beliefs in that case might be grounded in adequate reflection on available evidence, but if the evidence itself is tainted, as a result of emerging from alethically insensitive doxastic pressures, then no amount of careful reflection is going to save moral knowledge for the realist. But debunkers have failed to show that evolutionary pressures have had a complete, thoroughgoing influence on our moral faculties. They have failed to show that we are incapable of critically stepping back from whatever doxastic inclinations have been implanted in us by selective pressures.

Of course, it is *possible* that selective pressures have thoroughly saturated our moral faculties in doxastically discriminating ways, to the point that any exercise in moral thinking is bound to produce tainted moral beliefs. But this possibility alone does not justify skepticism. (Or, if it does, then the warrant for most nonmoral beliefs will be similarly undermined.) What the debunker needs to show is that selective pressures have actually, as a matter of fact, exercised such a thoroughgoing influence. As the discussion of D5 indicated, debunkers are very far from having done that.

c. The confirmed doxastic effects of selective pressures

Here is another way of understanding the charge of massive coincidence that is at the heart of the Darwinian Argument. Focus now on what we can be sure of when it comes to evolutionary influences. The confirmed doxastic effect of selective pressures is the inculcation of adaptive doxastic dispositions. It is *not* the inculcation of truth-tracking doxastic dispositions. It is possible, of course, that such pressures have the collateral effect of pushing us toward the truth. But when it comes to morality, this possibility is not yet confirmed. Nor can it ever be. Absent such confirmation, it would be a miracle were our moral faculties reliable.²³

We can confirm that selective pressures have had the indirect effect of leading us to moral truth only if we (1) assume the truth of some substantive moral or evaluative claims, and then (2) reveal how it would be adaptive to hold beliefs with that content. (2) probably would not be that difficult, if we could secure (1). But we can't, since assuming the truth of such substantive claims is precisely what agnostic genealogical critiques forbid.

And this prohibition appears to be well motivated. For if we were not required to suspend judgment about moral or evaluative matters in a wholesale way, but were instead allowed to rely on some substantive moral claims in assessing the reliability of our moral faculties, then we would be effectively immunizing a class of moral beliefs from genealogical doubts, thus enabling us to rely on this privileged set of moral beliefs to assess the merits of those that arise from selective pressures. Given that so many of our paradigm moral beliefs also appear to be highly adaptive ones, it would come as no surprise were such a strategy to result in the vindication of many of our moral beliefs and the general reliability of our moral faculties.²⁴

As a result, the debunker should insist that we suspend judgment about all substantive moral claims. And once we do that, then (it is thought) we will have no way to show that adaptive moral dispositions are also ones that reli-

²³ Thus Street: "allowing our evaluative judgements to be shaped by evolutionary influences is analogous to setting out for Bermuda and letting the course of your boat be determined by the wind and tides: just as the push of the wind and tides on your boat has nothing to do with where you want to go, so the historical push of natural selection on the content of our evaluative judgements has nothing to do with evaluative truth. Of course every now and then, the wind and tides might happen to deposit someone's boat on the shores of Bermuda. Similarly, every now and then, Darwinian pressures might have happened to push us toward accepting an evaluative judgement that accords with one of the realist's independent evaluative truths. But this would be purely a matter of chance, since by hypothesis there is no relation between the forces at work and the 'destination' in question, namely evaluative truth" (2006: 121-22).

²⁴ This paragraph summarizes a worry that anyone ought to have about the plausibility of so-called "third-factor" accounts, of the sort developed on behalf of moral realism by David Enoch (2011) *Taking Morality Seriously*, Oxford: Oxford University Press; Knut Skarsaune (2010) "Darwin and Moral Realism: Survival of the Iffiest," *Philosophical Studies* 152: 229-43; Brosnan (2011: 60-62), and (on a certain interpretation) Wielenberg (2010). For discussion of such accounts, see below, section 5e.

ably track the truth. Given that the only confirmed doxastic effect of selective pressures on our moral faculties is to inculcate adaptive doxastic dispositions, we have no basis for thinking that those pressures have also directed us toward the truth.²⁵ In such an epistemic situation, it would be a massive coincidence were our moral faculties reliable.

What we can call the *Doxastic Effect Argument* codifies this line of reasoning:

DE1. If moral realism is true, and if evolutionary forces have thoroughly shaped our moral faculties in doxastically discriminating ways, then the confirmed doxastic effect of selective pressures on our moral faculties is something other than believing the truth.

DE2. If the confirmed doxastic effect of selective pressures on our moral faculties is something other than believing the truth, then, absent an independent confirmation of the reliability of our moral faculties, it would be a massive coincidence were they reliable.

Therefore,

DE3. If moral realism is true, and if evolutionary forces have thoroughly shaped our moral faculties in doxastically discriminating ways, then, absent some independent confirmation of the reliability of our moral faculties, it would be a massive coincidence were they reliable.

DE4. There is no independent confirmation of the reliability of our moral faculties.

Therefore,

D1. If moral realism is true, and if evolutionary forces have thoroughly shaped our moral faculties in doxastically discriminating ways, then it would be a massive coincidence were our moral faculties reliable.

Here are two potential replies that realists might offer. The more ambitious one will try to show that DE4 is false, because we are in a position to offer independent confirmation – i.e., confirmation that does not rely on moral beliefs – of the reliability of our moral faculties. In section 6, I tentatively suggest a strategy designed to do just that. Any success that strategy might enjoy would undermine DE4, and cast doubt on D3 of the Darwinian Argument as well.

A second reply to this argument is to show that it would prove too much. We can replace “moral” with “perceptual” in the premises and conclusion, *salve veritate*. As with morality, evolutionary pressures in the perceptual domain have as their confirmed doxastic effect the cultivation of adaptive dispositions. But no one thinks that it would be a massive coincidence were our perceptual faculties reliable – even if realism about the objects of perception and their properties is true.

²⁵ Thus Joyce (2007: 211): “We have an empirically confirmed theory about where our moral judgments come from (we are supposing). This theory doesn’t state or imply that they are true, [and] it doesn’t have as a background assumption that they are true ... This amounts to the discovery that our moral beliefs are products of a process that is entirely independent of their truth, which forces the recognition that we have no grounds one way or the other for maintaining these beliefs.”

The moral antirealist might argue that DE1 has no parallel in the perceptual case. Suppose that realism about the objects of perception is true, and suppose as well that our perceptual faculties have been thoroughly shaped by selective pressures. Finally, let us grant that the confirmed doxastic effect of such influences is to incline us to adaptive perceptual beliefs. But this effect is *not* a distorting one, since adaptive perceptual beliefs will, at least for the most part, be true. After all, dispositions to hold false perceptual beliefs are not likely to be adaptive. Failure to accurately track the contours of the perceptual world could prove hazardous to your health. So adaptive pressures will likely incline us to accurate perceptual beliefs. By contrast, we lack any reason to think that dispositions to hold false moral beliefs would be any the less adaptive for that. So there is something special, and inferior, about moral faculties that reveals why the Doxastic Effect Argument does not overgeneralize.

I believe that this tempting line of reasoning spells trouble for the debunker. We can know that adaptive perceptual practices are also reliable ones only if we already have a sense of which perceptual judgments are true and which are false. We can tell that dispositions to hold false perceptual beliefs are likely to be maladaptive only if we can identify some false perceptual beliefs, show that they tend to undermine fitness and make inferences from those cases. If we were in no position to know which perceptual beliefs are false, then we could not know that false perceptual beliefs tend to be maladaptive.

By parity of reasoning, we are in a position to determine how (mal)adaptive false moral beliefs are only if we can identify at least some substantive moral beliefs as false, and determine how harmful their possession is or is likely to be. But that, of course, would also position us to identify some substantive moral beliefs as true. And once we are possessed of such knowledge, we could deploy it to plausibly assess the merits of other moral claims that we suspect of having been influenced by selective pressures. We could, in other words, develop the Natural Reply and answer the debunkers.

To put the point another way: If we are required to suspend judgment about all perceptual beliefs – as we must, if required to do so in the moral case – then we will most likely *not* be in a position to confirm the reliability of our perceptual faculties. We must presuppose the truth of at least some central, widely uncontroversial perceptual beliefs in order to get the confirmation of our perceptual faculties off the ground. But if we are allowed such liberties in the perceptual realm, we should be given similar license for morality. And then the debunking game is up.

The Doxastic Effect Argument will overgeneralize to all domains of inquiry that are construed realistically and are relevantly autonomous, i.e., such that claims within that domain can be ratified only by assuming the truth of at least some other claims within it. It is of course controversial just which domains these are. But so long as there are any, and we are unwilling to em-

brace a wholesale skepticism in such domains, we have reason to resist the Doxastic Effect Argument.

d. The Empiricist Argument

Here is another understanding of the alleged miracle. We can offer empirical confirmation of the origins of our moral beliefs, and that confirmation makes no mention of the moral truth, or any aptitude for sensitivity to such truth. When we compare the scientific merits of two stories – a wholly naturalistic one that omits any mention of realistic moral facts and epistemic powers of moral appreciation, and another that incorporates such facts and such powers – it is easy to see which one is scientifically better confirmed. So, in the presence of a scientifically well-confirmed account of the causal origins of our moral faculties that does not imply the truth of our moral beliefs, realistically construed, and the absence of any scientifically confirmed account that does have such an implication, it would be a miracle were our moral faculties reliable.

This is a line of reasoning that Sharon Street (2006, *passim*) pursues on behalf of her allegation of massive coincidence. There are, in fact, two distinct forms of argument on offer here, one that I call the *Empiricist Argument*, and the other, the *Causal Argument*. Perhaps Street has both in mind; I am not sure. But it pays to distinguish them and assess each on its own merits.

The Empiricist Argument is basically a reworking of a traditional line of antirealist critique:

E1. If moral realism is true, and if evolutionary forces have thoroughly shaped our moral faculties in doxastically discriminating ways, then the best-confirmed empirical account of the origins of our moral faculties does not imply their reliability.

E2. If the best-confirmed empirical account of the origins of a doxastic faculty does not imply its reliability, then it would be a massive coincidence were that faculty reliable.

Therefore,

D1. If moral realism is true, and if evolutionary forces have thoroughly shaped our moral faculties in doxastically discriminating ways, then it would be a massive coincidence were our moral faculties reliable.

I think that the first premise of this argument is probably true. But its second is probably false.

The first worry about E2 is this: An empirical account of certain doxastic origins may be the best-confirmed among all competitors and still fail to be evidentially compelling. It may be the best of an epistemically dubious lot. Even if we suppose, for instance, that an evolutionary account of the origins of our moral faculties is the best one we have, it is still in many ways a just-so

story, as previously noted.²⁶ Evolutionary accounts are meant to identify the distal, not proximate, origins of our moral faculties; those accounts concern themselves with literally prehistoric events for which we have relatively thin evidence.

Another reason to doubt E2 is that its truth would prove too much. The best-confirmed empirical account of the origins of our *philosophical* faculties does not imply their reliability. Central philosophical premises are not empirically confirmable. So if E2 is true, then our philosophical faculties are unlikely to be reliable. If that is so, then the philosophical arguments we advance are unlikely to be sound. But a wholesale philosophical skepticism is implausible, and is in any event self-undermining for debunkers, since they can achieve their aims only by advancing philosophical arguments.

Similarly, the beliefs generated by our moral (or other normative) faculties are not the sort of beliefs that are likely to be subject to empirical confirmation. Indeed, so long as there is any plausible form of an is-ought gap, empirical claims *cannot* imply normative truths. If we assume the existence of some such inference barrier, then unless one is willing to give up on the truth and credibility of all normative claims, we will have to accept that empirical inquiry alone will fail to reveal the truth and the epistemic status of our normative beliefs. Thus the fact that empirical investigation does not imply the reliability of normative faculties fails by itself to indict the reliability of those faculties.

Lastly, another *tu quoque*. I am not sure just which doxastic faculties are responsible for generating our metaethical beliefs. But it would be surprising, given their content, to think that the reliability of those faculties is subject to empirical confirmation. Street does not believe, for instance, that constructivism is superior to realism because constructivism has been empirically confirmed and realism has not. It is difficult indeed to see how constructivism (or any other fundamental metaethical theory) could be empirically vindicated. So endorsing E2 would mean discrediting her brand of constructivism.

The very hardy sort of empiricism that underwrites E2 may be true for all that. But we will need a very great deal of argument to show it so. Until we have been convinced by that sort of extensive argumentation, we are right to demur.

e. The Causal Argument

Another argument that might be animating Street's charge of coincidence begins with the assumption that there is a wholly nonmoral causal story of the origins of our moral faculties and moral beliefs. Such an etiology implies that moral facts have played no causal role in the formation of our moral be-

²⁶ See n. 6, above. That there is a non-negligible element of speculation in such accounts is clear even in what is perhaps the best of them, that recently offered by Philip Kitcher (2011) *The Ethical Project*, Cambridge, MA: Harvard University Press.

liefs. And if that is so, then it would be a miracle were our moral faculties reliable.

The wholly nonmoral set of causes that explains the origins of our moral faculties cannot plausibly be exhausted by evolutionary forces. After all, those forces have been fairly uniform across our species, yet there is plenty of variation among people in terms of their moral dispositions and beliefs.²⁷ This variability calls for explanation. We are very far from having the whole of it.

Yet unless the realist can offer a well-confirmed etiology that incorporates essential reference to moral facts, this lacuna is no cause for celebration. For even if it turns out that evolutionary forces have exerted only a little influence on our moral faculties, there are plenty of other nonmoral candidates for filling out the causal story. These include cultural, parental, historical, genetic and economic forces. Even if no single nonmoral source tells the entire explanatory tale, it may be that these combined sources, perhaps in conjunction with other nonmoral causes, entirely explain why we have the moral beliefs we do. And if that is so, then, says the critic, any hopes for moral knowledge will have been defeated.

This line of reasoning can be formulated as the *Causal Argument*:

C1. If moral realism is true, and if evolutionary (as well as genetic, parental, cultural, historical, etc.) forces have thoroughly shaped our moral faculties in doxastically discriminating ways, then our moral beliefs have a complete set of natural causes.

C2. If our moral beliefs have a complete set of natural causes, then our moral beliefs are exclusively caused by things other than moral facts.

C3. If our moral beliefs are exclusively caused by things other than moral facts, then it would be a massive coincidence were our moral faculties reliable.

Therefore,

D1. If moral realism is true, and if evolutionary (or other nonmoral) forces have thoroughly shaped our moral faculties in doxastically discriminating ways, then it would be a massive coincidence were our moral faculties reliable.

Moral realists who are also ethical naturalists will reject C2. Ethical naturalists regard moral properties as natural properties, and moral facts as natural ones. So even if moral beliefs have a complete set of natural causes, this does not

²⁷ Note that the Darwinian Argument is perfectly suited to round out the more traditional versions of genealogical, antirealist critique, ones that rely on the diversity of (e.g.) historical, genetic and psychological influences to attempt to best explain why people have the moral beliefs they do. These familiar doxastic influences are standardly deployed in a complementary a posteriori argument from disagreement. This argument takes the form of an inference to the best explanation; the explanandum is the fact of broad and deep moral disagreement. The argument seeks to offer the most parsimonious explanation of this datum solely in terms of contingent nonmoral facts about the causes of our moral outlooks. In a move that dovetails beautifully with this familiar antirealist critique, the Darwinian Argument posits precisely the opposite explanandum – the widespread *agreement* in belief across a range of core moral claims. Realists have long used this broad consensus as a partial basis for rebutting the abductive antirealist argument from disagreement. But along comes the Darwinian critic to rob realists of even this shred of hope!

exclude the possibility that these beliefs also result from moral causes. On this line, some natural facts are also moral facts, and some of these may well be causally responsible for our holding the moral beliefs we do. If that were so, then C2 would be false.

This may be a successful strategy for handling the Causal Argument. But what of nonnaturalists (like myself) who deny that moral properties and facts are natural ones? Nonnaturalists will have to take aim at either C1 or C3. Consider C1 first.

First, and a familiar point by now: Evolutionary (and psychological, historical, etc.) forces may heavily influence our moral faculties without providing a *complete* causal explanation of their origins and workings. The antecedent of C1 does not entail its consequent. Despite a variety of nonmoral causes impinging on our moral faculties, debunkers have failed to provide anything like a *complete* nonmoral causal story of the workings of our moral faculties.

Second, even if we allow that the nonmoral forces, taken together, are sufficient to determine what moral beliefs we hold, this will still fail to support C1's consequent. In other words, causal sufficiency does not entail causal completeness. That is because of the possibility of causation by supervenient facts.²⁸ Suppose, for instance, that our actions are fully causally determined by some array of neurological facts; still, we might think that beliefs, intentions and desires really do cause actions, and that such mental events are neither type- nor token-identical to their underlying neurological ones.²⁹ If this is plausible, then moral facts might causally explain our moral beliefs, even if there is full causal determination by natural forces of our moral beliefs, and even if, as nonnaturalists claim, moral facts are neither type-identical nor token-identical to natural ones.

Now this sort of reply will not be helpful to those nonnaturalists (almost all) who reject the causal efficacy of moral facts. Perhaps nonnaturalists should rethink this stand. Unless natural facts are just identified as those that are causally efficacious – not a very attractive way of drawing up the taxonomy, I think³⁰ – then nonnaturalists are not per se barred from endorsing the causal efficacy of moral facts.

I am not going to defend the causal efficacy of moral facts here; rather, I want to articulate what seems to be the three most common worries about

²⁸ Here and below I will speak of facts as having causal powers; if that offends, feel free to substitute references to your favorite candidates for causal relata.

²⁹ That they are not type-identical is largely taken for granted nowadays. That they are not token-identical is more controversial, but it seems plausible that a given belief, desire or intention would retain its identity even if one of its subvening neurons went missing. The loss of one such neuron would not alter the content or functional role of these mental-state tokens, and these properties, rather than precise neuronal composition, are the ones that are essential to individuating mental states.

³⁰ I discuss various ways to draw the natural/nonnatural divide in (2003) *Moral Realism: A Defence*, Oxford: Oxford University Press, pp. 58-65, and (2005) "Ethics as Philosophy: A Defense of Ethical Nonnaturalism," in Terence Horgan and Mark Timmons, eds., *Metaethics after Moore*, Oxford: Oxford University Press, pp. 209-32, at 210-12.

this sort of causation, and very briefly say why these worries might be resistible. If I am successful on this front, then the door is open to moral causation. That would not establish it, of course, but it might raise enough doubts about C1 as to undermine the force of the Causal Argument.

First worry: If moral facts are nonnatural, then how could they cause natural facts, given that this would amount to causation across different metaphysical kinds?

Answer: We have familiar examples of cross-kind causal interactions (e.g., physical events causing mental events, and vice versa) and do not regard this as an insuperable barrier to assigning causal efficacy in such cases.

Second worry: Moral facts would be causally superfluous, as all the causal work we might assign to them is in fact done by the subvening natural facts.

Answer: We assign causal power to sets of supervening facts all the time, even while recognizing the causal potency of their subvening facts. It is not clear why we should treat moral facts differently in this regard. Economic, chemical, geological and biological facts cause things – or if, by virtue of their supervenient status, they are thereby deemed to be causally inert, then the Causal Argument would prove too much, by forcing us either to antirealism about every supervenient domain, or to skepticism about all of the claims in such domains.

Third worry: The only thing that we might need moral facts to causally explain are our beliefs about them. But this would make the causal powers of moral facts *sui generis*, and so moral explanations would be objectionably ad hoc, as we would be allowing moral facts a more limited causal role than that assigned to any other kind of fact. All other causally efficacious facts cause things in addition to beliefs about themselves.

Answer: We can point to companions in innocence here. Modal facts, numerical facts and philosophical facts are all such that they do not *directly* cause anything other than beliefs with such facts as their contents. (Of course they may distally cause non-doxastic effects, but this would be via the mediation of beliefs or other psychological states.)

Consider these propositions: (a) necessarily, any two worlds identical in all nonmoral respects are also morally identical; (b) $113 + 223 = 336$; (c) Ayer's principle of verifiability is false. These propositions are true. We might say that we believe them because the facts they represent have caused us (in conjunction, of course, with our capacities for reflection and understanding) to believe them – even though these facts do not do any non-doxastic causal work. *If* that is plausible, then we should say the same thing about moral facts. We believe them because they are true – where this is shorthand for the idea that the moral facts themselves have caused us (in combination, of course, with our appropriate doxastic receptivity) to have beliefs with the propositions representing those facts as contents.

On this way of looking at things, C1 is either false or it begs a question against realists. It is false if moral facts can indeed cause our moral beliefs; it

begs the question if, without further argument, it supposes that moral facts are incapable of causing us to have the moral beliefs we do. On this line of reply, the nonmoral influences on our moral beliefs do not supply a complete causal explanation of those beliefs, because the moral facts can impress themselves upon those with appropriate moral sensibilities in such a way as to enable us to appreciate those facts for what they are.

For those who are suspicious of moral causation, let us now consider the worst-case scenario for the realist, and assume the causal inefficacy of moral facts. If moral facts cause nothing, then, *a fortiori*, they do not cause our moral beliefs. In the absence of moral causation, it seems that we must accept C1 – there would be some complete nonmoral causal explanation of our moral beliefs, even if genealogical debunkers have yet to specify the complete causal story. In that case, moral realists would have to reject C3. How could they do that with a straight face?

Here is a two-part answer that offers the beginnings of a reply. First part: companions in innocence. If moral facts are causally inert, then so too are modal, numerical and philosophical facts. Still, we are rightly confident that we have some knowledge in these domains. And so there is noncausal knowledge. Consider modal conditionals. However we understand their nature, the things that such conditionals express or represent do not cause anything (other than, perhaps, modal beliefs). And yet we can sometimes know that they are true. This example is especially important, because on realist assumptions (which debunkers are granting in order to generate skeptical implications), fundamental moral principles are conditionals that express necessary truths.

Second, even on the assumption that our modal conditional beliefs are not caused by conditional facts, our knowledge of their truth need not be coincidental. There are many stories one can tell about how we might gain access to such truths, and we need not opt for just one here. The essential point is that we have lots of examples in which we have warranted confidence in our conditional beliefs, even though (we are supposing) the facts represented by such beliefs have not caused us to hold them.

Sound fishy? It shouldn't. To see why, suppose that we take the Causal Argument to its limit, as Matt Bedke does.³¹ Assume the causal closure of the physical, and the causal impotence of moral facts. Add the further assumption that the formation of a moral belief or intuition is also a physical event. Bedke argues that this is all we need to successfully mount a charge of massive coincidence against the nonnaturalist realist.³² But that is not so. For if it were, we would be likewise forced to such a verdict in every case of apprehending modal, numerical and philosophical truths. Indeed, if Bedke's as-

³¹ Bedke (2009) "Intuitive Non-Naturalism Meets Cosmic Coincidence," *Pacific Philosophical Quarterly* 90: 188-209.

³² To be fair, Bedke's argument (2009: 190) has eight steps, not three. But the three claims mentioned above form the essence of his critique, and together seem sufficient to raise the core worry that he is pressing.

sumptions really did generate a verdict of massive coincidence, then every instance in which we grasped an a priori truth would be miraculous. But that is highly implausible.

With a few possible exceptions that are irrelevant for present purposes, a priori knowable propositions are metaphysically necessary truths. Such truths, or the facts they represent, do not cause anything. Further, instances of a priori apprehension are themselves physical events. (Or, if they are not, then we have excellent reason to reject the causal closure of the physical, and so the completeness of physical explanations of our a priori beliefs.) But it is no coincidence that those who adequately understand a priori propositions and form their beliefs on the basis of that understanding are reliably in touch with the truth.

It is certainly plausible that our capacities of a priori apprehension and rational insight have themselves been subject to a variety of nonmoral forces – perhaps even to physical determination by physical laws. But this should not undermine our confidence in their reliability, even on the assumption that the facts they discern are causally impotent and best construed realistically. That is because we are rightly confident in a variety of beliefs in a priori propositions – indeed, synthetic a priori propositions³³ – even if there are wholly physicalistic etiologies of the episodes in which we gain the relevant understanding. Debunkers have offered no good reason to suppose that these epistemic capacities, reliable in many other contexts, are broken when their target propositions have moral content.

Of course I have not argued for the existence of a priori moral truths. And there may be special reasons to deny the existence of a priori moral propositions while allowing the existence of nonmoral ones, or even, more radically, for denying the possibility of a priori knowledge altogether. But evolutionary debunkers have done nothing to challenge the existence of a priori truths generally. Nor have they targeted a priori moral truths in particular. And this is not surprising – it is difficult indeed to see how genealogical considerations could do the needed work on that score.

It is true that not all moral knowledge is going to be a priori knowledge. So nothing contained in the short story I've just told can provide anything like a complete account of nonnaturalistic moral knowledge. But that is not the realist's current burden. All she needs to show is that if moral facts are causally inert, then noncausal moral knowledge would not amount to a mas-

³³ Some familiar examples: The moral globally supervenes on the nonmoral; justified true belief is insufficient for knowledge; nothing can be red all over and green all over at the same time; the tallness relation is transitive; closed trilateral figures have interior angles that sum to 180 degrees. These propositions are true. They are not analytic – or if they are, then apparently substantive moral propositions may be as well, thus strengthening the case for their a priority. We know these propositions. And we do not need to rely on empirical evidence for this knowledge. Further, we can ask: What is the epistemic status of the proposition that there are no synthetic a priori truths? It does not seem an analytic truth. Nor does it seem empirically confirmable. If true at all, if knowable at all, it would be synthetic a priori, and so self-undermining.

sive coincidence. So long as a priori knowledge generally is something within our grasp, then, for all debunkers have shown, a priori moral knowledge is also within our grasp. A priori knowledge is not miraculous – or if it is, then given our confidence in the existence of such knowledge, a charge of massive coincidence should no longer suffice to undermine prospects for knowledge. An appeal to the as-yet-undefeated possibility of a priori moral knowledge strikes me as the best way for realists to take issue with C3.³⁴

In sum, there are several ways that moral realists might respond to the Causal Argument. They might first opt for a naturalistic version of realism, in which case they can invoke all of the supporting arguments for their position as a means of criticizing C2. The path is not as straightforward for nonnaturalists. Those who hold out hope for moral causation can challenge C1. Those of a less sanguine bent will be forced to challenge C3. Perhaps the best way to do this is to first identify companions in innocence – kinds of facts, such as modal, arithmetic and philosophical ones, that lack causal powers, but are nevertheless knowable – and then to explore the possibility of a priori moral knowledge.

Since a priori moral truths will not be to everyone's liking, we should conclude our examination of the Causal Argument by exploring an alternative criticism of C3. This option relies on the existence of *indirect tracking hypotheses*.³⁵ Such hypotheses claim that our moral faculties track nonmoral features of the world that are well correlated with moral features, thus establishing the reliability of those faculties. The basic picture looks like this: Evolutionary pressures have shaped us so that our moral faculties are largely reliable, not because these faculties are designed to identify moral facts, but rather because these faculties are designed to be sensitive to nonmoral facts that, happily for us, are reliably correlated with moral ones.

This account has an important implication: While evolutionary pressures on our doxastic faculties have not shaped them in such a way that they are directly sensitive to the moral truth, neither are such pressures thereby distorting influences.³⁶ Evolutionary pressures have influenced our doxastic faculties to track those features of the world that it is adaptive to track. But many of these features bear moral significance. For instance, it is adaptive to

³⁴ Erik Wielenberg defends a specific instance of this general approach within the context of what he calls *The Basic Evolutionary Debunking Argument* (2010: 453). He claims that possession of certain cognitive capacities entails possession of moral rights. Though he does not highlight the a priori of this claim as relevant for replying to the debunkers, he takes the claim to be a conceptual truth, and uses it as his primary example designed to establish the reliability of our moral faculties.

³⁵ I take the term *indirect tracking hypothesis* from Justin Horn (unpublished manuscript). David Enoch uses the term “third-factor explanation” to refer to the same possibility. See Enoch (2011: 167-68).

³⁶ These are the only two possibilities that Street considers in crafting her Darwinian Dilemma for evaluative realists. The success of any specific indirect tracking hypothesis would thus enable realists to avoid either horn of her dilemma.

know whether or not someone is trustworthy. And lo, that feature of the world is also morally significant.

In reality, there are many specific indirect tracking hypotheses, each seeking to establish the moral importance of a property that it would be adaptive to reliably discern. Rather than survey all relevant possibilities, consider just two recent examples. We could adopt David Enoch's strategy for dealing with Darwinian critiques, and assume that our survival is good. Once we make this assumption, and assume, just as plausibly, that believing in the goodness of survival is adaptive, then we have a way of establishing the reliability of a doxastic practice that generates adaptive evaluative beliefs.³⁷ Or we might adopt Knut Skarsaune's suggestion that pleasure is good and pain is bad.³⁸ People who are disposed to make such judgments are obviously in a better position to pass along their genes than those who fail to be disposed in this way. Thus is a link between evaluative truth and epistemic reliability forged. This would enable us to reject the allegation of a massive coincidence that is at the heart of C3.

I am discussing third-factor/indirect tracking hypotheses because I think that they are the most obvious path to resisting C3. But there is a serious worry about strategies that rely primarily on such hypotheses, a worry that at the same time raises concerns about the structure and tenability of genealogical critiques. The success of any indirect tracking hypothesis relies on enlisting at least one substantive moral or evaluative claim. Such claims effectively serve as bridge principles that link a doxastic genealogy with the moral truth. These claims are meant to provide the explanation of the reliability of our moral faculties, given causal origins whose confirmed effect is something other than lighting on the moral truth. But is the realist really entitled at this stage of the dialectic to the substantive assumption (e.g.) that survival is good, or pain bad? If we grant such entitlement, then it seems the genealogical critic has lost the argument before it has even begun. For substantive evaluative or moral claims, if they are able to do the needed work, must enjoy a degree of positive epistemic status that is immune to genealogical critiques. And yet if we are able to identify claims that enjoy such immunity, then realists are in the clear. They would have all they need to vindicate the Natural Reply.

So: Are realists entitled to rely on substantive moral or evaluative claims in addressing the genealogical critiques? Or must they forbear from such reliance? Let us consider the matter briefly before wrapping things up.

³⁷ "Survival (or whatever) is good; so behaving in ways that promote it is (pro tanto) good; but one efficient way of pushing us in the direction of acting in those ways is by pushing us to believe that it is good to act in those ways. ... So the normative beliefs this mechanism pushes us to have will tend to be true." Enoch (2011: 169).

³⁸ Skarsaune (2010).

6. The Use of Substantive Assumptions in Assessing Genealogical Critiques

There are three relevant scenarios for determining the propriety of enlisting substantive moral claims in attempts to reply to the Darwinian Argument.

In the first scenario, we are convinced of the soundness of the Darwinian Argument. We are convinced, in other words, that if realism were true, then it would be massively coincidental were our moral faculties reliable. In that situation, realists are rationally required to suspend judgment about all moral claims. On the assumption of realism's truth, the warrant for our moral beliefs would have been defeated. It would therefore be illicit to rely on those beliefs in seeking to restore that warrant.

In the second scenario, we have undercut the plausibility of the Darwinian Argument, perhaps by means of arguments of the sort that I have given in this paper. Once the putative genealogical defeater has itself been defeated, we are certainly within our epistemic rights to invoke substantive moral claims to bolster the impression of the reliability of our moral faculties.

The third scenario is the most interesting. This is the one in which we are as yet unsure of the power of the genealogical critique. What do we do, for instance, if we find such a critique tempting, and are uncertain about its soundness? If we are in the midst of trying to evaluate its plausibility, must we at that point suspend judgment about all moral claims, or are we allowed to rely on some of them to reinforce our antecedent confidence in our moral faculties?

Genealogical critics will plump for the first horn of this dilemma. Realists who deploy third-factor/indirect tracking hypotheses to subvert genealogical critiques will opt for the second horn. Here I think that the genealogical critics have the upper hand. In this third scenario, we are worried about the reliability of our moral faculties. For all we know, they might be drastically unreliable. It would be illicit at that point to introduce moral beliefs to bolster our confidence in our moral faculties, as any doubts about the faculties themselves ought to be transmitted to the beliefs they generate.

To see this, consider an analogy. You have been spun in circles. Once the spinning stops, you are asked to point to the west. You are still quite dizzy. Though you ordinarily feel confident about your sense of direction, you are now unsure of how reliable it is. Suppose that you nonetheless believe that west is currently to your left. That belief cannot rightly reestablish your confidence in your sense of direction. To the extent that you now doubt the reliability of your sense of direction, the beliefs that are based on that sense alone do not seem able to restore our confidence in its reliability.

Of course, you could manage this restoration were you to offer an independent confirmation of your directional beliefs – say, by pointing to a landmark whose location you are antecedently sure of, and inferring the direction on that basis. With such beliefs justifiably in hand, you can begin the process of assessing the faculty's reliability. What is crucial here is the exist-

ence of a source of warranted true beliefs that is different from the source whose reliability is now subject to doubt.

But this crucial element seems absent in the moral case. On the assumption that we must trace the warrant of our moral beliefs to the exercise of our moral faculties, any doubts about those faculties will be inherited by those beliefs. And further assuming that nonmoral beliefs cannot by themselves justify positive moral beliefs, we are left adrift if we are unable to answer the debunking critiques. So long as we are deeply puzzled about the reliability of our moral faculties, we are unable to restore confidence in them by enlisting beliefs generated by the now-questionable faculties. Neither can we restore confidence by bringing nonmoral beliefs to bear. So, it seems, we cannot restore confidence at all.

There may be a plausible way through the horns of this dilemma. Call it the *formal strategy of faculty vindication*. The success of any version of this strategy would solve the problem of restoring warranted confidence in our moral faculties once we become aware of a criticism that causes us to worry about their reliability.³⁹ And not only that. It would also provide a compelling case against the third premise of the Darwinian Argument:

D3: There is no independent confirmation of the reliability of our moral faculties – their reliability can be confirmed only by showing that they have generated moral beliefs in which we have a high degree of warranted confidence.

We can understand the nature of a formal strategy by contrasting it with two familiar models of faculty confirmation. An *internal coherence strategy* seeks to establish the reliability of a doxastic faculty F by showing that it generates many F-beliefs in which we have great confidence. An *independence strategy* seeks to vindicate the reliability of a doxastic faculty F by showing that F-beliefs are confirmed by various non-F beliefs in which we have great confidence.

However, if we have doubts about F's reliability, and if F ranges over an autonomous domain of inquiry, neither of these familiar strategies will work.

Enter the formal strategy, which seeks to vindicate the reliability of a doxastic faculty by showing that it is identical to, or a species of, a (kind of) doxastic faculty in which we independently have a high degree of warranted confidence. When it comes to our moral faculties, the likeliest candidate to serve as the basis of a successful formal strategy is the faculty of generating a priori beliefs. It may be that there is one such faculty, or an interrelated set of them, that ranges over a wide variety of belief contents. Perhaps the mental

³⁹ To be clear: I do not think that evolutionary critiques should cause us permanent worry. I hope to have shown that their central charge – that of massive coincidence – is as yet unpersuasive on any of the most plausible ways of construing that charge. To the extent that I have succeeded, the formal strategy is here only to bolster the arguments that have come before, and to assuage the concerns of those who still find evolutionary critiques deeply troubling.

operations that generate nonmoral synthetic a priori knowledge are the very ones responsible for generating our a priori moral beliefs. Were that so, we would have good, albeit defeasible, reason to consider our moral faculties reliable, at least with regard to the a priori moral beliefs they generate. And confidence in this reliability would not have been gained by assuming the truth of some substantive moral beliefs, or by seeking to confirm such beliefs by means of nonmoral beliefs.

This may be a good route for realists to take. Realists inclined in this direction would of course have to vindicate belief in the synthetic a priori (or defend a different candidate faculty to use as the basis for the formal strategy). This is not easy, but examples of such propositions (see n. 33 for a few plausible candidates) seem to me more compelling than the premises invoked in arguments attempting to show the impossibility of this kind of knowledge. Realists would also need to do much more than they (and antirealists) have done to explain the operations and identity conditions of our doxastic faculties. So the jury is out on how feasible a formal strategy is.

Suppose that formal strategies of faculty vindication will not work for morality. Then what? If we are convinced by evolutionary critiques, then the realists among us are saddled with moral skepticism. If we are convinced that those critiques are less powerful than their adherents have thought, then we can adopt an internal coherence strategy of faculty vindication and enlist substantive moral beliefs to validate the reliability of our moral faculties. But if we are unsure of the power of such critiques, then, as I have conceded, it seems illicit to rely on substantive moral beliefs to resuscitate our confidence in our moral faculties. In that case, internal coherence strategies are out. Independence strategies are out (given the autonomy of the moral). If formal strategies are also off the table, then it seems that all options for vindicating our moral faculties will have failed.

Thus until such time as we can mount a successful defense of a formal strategy for vindicating the reliability of our moral faculties, moral realists will have to focus directly on trying to undermine at least one premise of the Darwinian Argument. That is what I have tried to do here.

7. Conclusion

Consider two claims at the heart of debates about the merits of evolutionary critiques:

- (1) On the assumption of moral realism's truth, selective pressures are likely to have generated reliable moral faculties.
- (2) On the assumption of moral realism's truth, selective pressures are unlikely to have generated reliable moral faculties.

It is important to note that the moral realist does not have to vindicate (1) in order to fend off evolutionary critiques. All the realist needs is to show that (2) is inadequately supported.

We can see this if we consider the nature of the dialectic. Even Darwinian critics begin with a recognition that (at least) paradigmatic, uncontroversial moral beliefs are strongly presumptively warranted. Many moral beliefs appear to be intrinsically plausible and are very widely endorsed across times and cultures, are embedded in a broad network of mutual doxastic support, survive extended reflection and possess a high degree of credibility as implied by plausible general theories of warrant. The initial positive epistemic status of such moral beliefs is extremely high.⁴⁰

The Darwinian Argument is intended to undermine this status. It can do so only if we have an extremely high degree of warranted confidence in each of its premises. We don't seem to. D5 is an uncertain empirical hypothesis. The status of D3 is also uncertain, and depends on whether realists can succeed in developing a formal strategy for vindicating the reliability of our moral faculties. The best arguments for D1 – or at least, the best that I can think of and locate in the literature – also fall short of their intended aim.

To the extent that I have succeeded in raising doubts about the crucial premises of the Darwinian Argument, we lack compelling reasons to accept (2). The absence of compelling reasons to accept (2) should *not* yield a state of *aporia*. Rather, it should leave us where we began – with a set of highly presumptively warranted moral beliefs whose positive epistemic status is thus far undefeated.

Though realists do not need to vindicate (1), it is interesting to think of what is required to succeed at that task. As a general matter, and setting formal strategies aside for the moment, we can vindicate the likely reliability of doxastic origins in a given domain only by utilizing standards of critical assessment appropriate for that domain. So we are able to know whether the beliefs caused by adaptive moral faculties are likely to be true only by utilizing the standards of critical assessment appropriate to the moral domain.

Those standards are controversial, of course, but at the very least they will include endorsement of paradigmatic moral commitments. Indeed, given the truth of any plausible inference barrier from is to ought, we will have to rely on substantive moral or evaluative claims in order to positively establish the existence of a correlation between selective pressures and the reliability of our moral faculties.

If all has gone well thus far, we are now in a position to do just that. We are only now entitled to rely on indirect tracking hypotheses, for instance,

⁴⁰ Karl Schafer ((2010) “Evolution and Normative Skepticism,” *Australasian Journal of Philosophy* 88: 471-88, at 476) offers another reason to support the presumptive warrant of normative beliefs generally: “[T]he data that supports evolutionary theory is itself based on our perceptual experiences. And so evolutionary theory itself rests on a tacit assumption that we are not in one of the traditional skeptical scenarios. Thus in order for an evolutionary vindication of our perceptual faculties to get off the ground, we must start off by assuming that we possess some sort of initial entitlement *a priori* to assume that our perceptual faculties are operating reliably. ... if we do have an *a priori* entitlement to trust our perceptual faculties, the realist is perfectly entitled to assume that the same is true with respect to our normative faculties as well.”

because we have done the work needed to cast doubt on a serious potential defeater of our moral beliefs. That leaves us in a position to confidently rely on paradigmatic moral beliefs. Some of these can be utilized to forge the needed bridge between our moral faculties' evolutionary origins and their reliability. That would suffice to establish (1).

Until better arguments for the central premises of the Darwinian Argument come along, or an entirely different form of antirealist genealogical critique makes its appearance, we are right to regard moral realism as perfectly compatible with moral knowledge and warranted moral belief. Further, once realists are able to defeat genealogical defeaters, they are then in a position to advance positive arguments that, for all we now know, are sufficient to establish the reliability of our moral faculties.⁴¹

Russ Shafer-Landau
University of Wisconsin
Department of Philosophy
shaferlandau@wisc.edu

⁴¹ For help with previous drafts, I would like to thank my terrific colleagues John Bengson, Elliott Sober and Mike Titelbaum, as well as Alex Byrne, Dave Chalmers, Christian Coons, Terence Cuneo, Tyler Doggett, Bill FitzPatrick, Uriah Kriegel and Benjamin Wald – only one or two of whom will recognize in the current version anything closely approximating the draft that they spent so much time with. (*Mea maxima culpa.*) I am grateful for the excellent comments provided by two anonymous reviewers for *JESP*. I would also like to express my gratitude for kind invitations to present some of these ideas at conferences on philosophical methodology and philosophical realism in Cologne and Urbino, and in seminars and talks at Bowling Green, Carleton and St. Olaf colleges, the University of Chicago, the University of Iowa, Northwestern University, the University of Toronto and the University of Illinois at Chicago.