

Do the evolutionary origins of our moral beliefs undermine moral knowledge?

Kevin Brosnan

Received: 3 October 2009 / Accepted: 1 November 2010 / Published online: 17 November 2010
© Springer Science+Business Media B.V. 2010

Abstract According to some recent arguments, (Joyce in *The evolution of morality*, MIT Press, Cambridge, 2006; Ruse and Wilson in *Conceptual issues in evolutionary biology*, MIT Press, Cambridge, 1995; Street in *Philos Studies* 127: 109–166, 2006) if our moral beliefs are products of natural selection, then we do not have moral knowledge. In defense of this inference, its proponents argue that natural selection is a process that fails to track moral facts. In this paper, I argue that our having moral knowledge is consistent with, (a) the hypothesis that our moral beliefs are products of natural selection, and (b) the claim (or a certain interpretation of the claim) that natural selection fails to track moral facts. I also argue that natural selection is a process that could track moral facts, albeit imperfectly. I do not argue that we do have moral knowledge. I argue instead that Darwinian considerations provide us with no reason to doubt that we do, and with some reasons to suppose that we might.

Keywords Evolutionary ethics · Moral realism · Moral epistemology

In what follows, I will not question the plausibility of the hypothesis that our moral beliefs are products of natural selection. My interest is rather to investigate the implications that are drawn from this claim. Since these implications are said to affect moral realism, let me state briefly state how I will understand this view. There are different varieties of moral realism and my interest here is not to settle which of them is the most plausible (Boyd 1988; Brink 1989; Railton 1986; Sturgeon 1985; Shafer-Landau 2003). It is enough for present purposes to characterize it as the view that what makes any moral proposition true is independent of what anyone believes

K. Brosnan (✉)
University of Cambridge, Cambridge, UK
e-mail: kb407@cam.ac.uk

or says. For example, the claim that torturing others is morally wrong is not made true by the beliefs or attitudes of anyone. If this claim were true, it would be true even if no actual person believed it. I will follow the authors I discuss in this paper by talking about moral facts. Moral facts as I will be using this term are instantiations of moral properties. Our moral beliefs are true or false in virtue of whether or not they correspond to these facts. Moral realism is also committed to the view that some of our moral beliefs are true.

An evolutionary argument against moral knowledge

Having clarified some terminology, we can now turn to the question I want to address: if our moral beliefs are products of natural selection, does it follow that we do not have moral knowledge? In defense of an affirmative answer to this question, consider this argument, versions of which are defended by Ruse and Wilson (1995), Street (2006), and Joyce (2006):

- (1) If our moral beliefs are products of a process that fails to track moral facts, then we do not possess moral knowledge.
- (2) Our moral beliefs are products of evolution by natural selection.
- (3) Evolution by natural selection is a process that fails to track moral facts.
- (4) Our moral beliefs are products of a process that fails to track moral facts.
- (5) Therefore, we do not possess moral knowledge.

In this section of my paper, I will consider the arguments offered in defense of the inference stated in premise (1). In the following section, I will argue against the inference from (2) and (3) to (4). In the final section, I will present an argument against premise (3).

Street (2006) and Joyce (2006, Chap. 6) accept for the sake of argument that moral facts exist. But they argue that our moral beliefs are products of a process (natural selection) that fails to track these facts. From this tracking failure, Street infers that our moral beliefs are probably false, while Joyce infers that they are unjustified. On the plausible assumption that we cannot have moral knowledge unless our moral beliefs are both true and justified, either inference, if correct, would suffice to undermine moral knowledge. To see whether either inference is correct, the arguments offered in defense of each one must be examined. Consider Street's argument first, the heart of which relates to the first horn of the following dilemma:

On the one hand, the [moral] realist may claim that there is *no* relation between evolutionary influences on our evaluative attitudes and independent evaluative truths. But this claim leads to the implausible skeptical result that most of our evaluative judgments are off track due to the distorting pressure of Darwinian forces. The realist's other option is to claim that there *is* a relation between evolutionary influences and independent evaluative truths, namely that natural selection favored ancestors who were able to grasp those truths. But this account, I argue, is unacceptable on scientific grounds. Either way,

then, realist theories of value prove unable to accommodate the fact that Darwinian forces have deeply influenced the content of human values. (Street 2006: 109)

Street refers in this passage to a relation, the *relata* of which are, (a) the evolutionary processes (e.g., natural selection) that govern which moral beliefs we have, and (b) moral facts (this is how I am construing Street’s “independent evaluative truths”). The case in which there is “no relation” between (a) and (b) may be illustrated as follows. Let ‘*p*’ represent the claim that cooperation is morally good. Street’s ‘no relation’ idea means that the probability of our believing that cooperation is morally good is independent of whether cooperation is in fact good. So the ‘no relation’ thesis is satisfied precisely when:

IND: $\Pr(\text{We believe that } p \mid p) = \Pr(\text{We believe that } p \mid \text{not } p)$

The alternative case, where there *is* a relation between (a) and (b), means the probability of our believing that cooperation is morally good is not independent of whether cooperation is in fact good, but is rather correlated with this fact. So the ‘relation’ thesis is satisfied precisely when:

CORR: $\Pr(\text{We believe that } p \mid p) \neq \Pr(\text{We believe that } p \mid \text{not } p)$

One way to see the difference between **IND** and **CORR** is in terms of the *tracking* relation.¹ When the process responsible for our believing that *p* fails to track *p*, our believing that *p* will be independent of *p* in the manner expressed by **IND**—so **IND** results from *tracking failure*. Alternatively, when the process responsible for our believing that *p* succeeds in tracking *p*, our believing that *p* will be dependent on *p* in the manner expressed by **CORR**—so **CORR** results from *tracking success*.

Given this clarification of the relation that Street’s dilemma invokes, it is worth noting the respects in which Street’s way of framing this dilemma is perfectly general.² So long as our moral beliefs are caused by *some process or other*, this process will either track presumptive moral facts or it will not. Evolutionary considerations are irrelevant to this point. The other respect in which Street’s dilemma is quite general relates to the epistemic implications of tracking failure. It doesn’t matter whether the causes of our moral beliefs are evolutionary, or cultural, or whatever: if these causes fail to track, the beliefs they affect are likely thereby to be tainted epistemically.³ Causes that fail to track are “distorting” causes on Street’s view (Street 2006, p. 121). This view may seem plausible. For example, if I

¹ Copp (2008) discusses Street’s dilemma in terms of the tracking relation as well.

² Harman (1977) defends an argument similar to Street’s, except for the fact that Harman talks about “upbringing” while Street talks about evolution. Harman and Street share the view that moral facts aren’t needed to explain our moral beliefs. Harman takes this to imply that there are no such facts, in which case we cannot have moral knowledge, while Street argues that, even if there were such facts, we would still not have moral knowledge. Wielenberg (2010) has also noted the similarities between Harman’s argument and Street’s.

³ In a recent discussion of how one of Darwin’s contemporaries—Henry Sidgwick—deals with the skeptical challenge presented by an evolutionary account of our moral beliefs, Lillehammer notes the following: “It is widely, if not universally, accepted that tracking failure would impugn the epistemic

come to believe that there are 12 people in my neighbor’s house by counting the change in my pocket, then I’ve acquired my belief by a process that fails to track the relevant facts (assuming the aim of my belief is truth). This tracking failure seems to imply that my belief is probably false. If tracking failures have this consequence generally—I will argue in a moment that they never do—then the first horn of Street’s dilemma goes through: our moral beliefs are probably false because they are the result of a process that fails to track moral facts.

Tracking failures do not support Street’s conclusion, however, as the odds version of Bayes’ theorem reveals:

$$\begin{array}{ccc}
 \text{Ratio of Posterior Probabilities} & \text{Ratio of Priors} & \text{Likelihood Ratio} \\
 \frac{\Pr(p|\text{We believe that } p)}{\Pr(\text{not } p|\text{We believe that } p)} & = \frac{\Pr(p)}{\Pr(\text{not } p)} \times & \frac{\Pr(\text{We believe that } p|p)}{\Pr(\text{We believe that } p|\text{not } p)}
 \end{array}$$

The likelihood ratio measures how sensitive an organism’s cognitive mechanisms are to varying states of the world. It is thus a world-to-head measure: if the world is in this state rather than that, the likelihood ratio measures how probable it is for an organism to respond accordingly. The tracking relation illustrates the different values that the likelihood ratio can have, with tracking failures at one end of the scale and tracking successes at the other end. Consider tracking failures first. A cognitive mechanism that fails to track the relevant states of the world is like a weather vane whose arrow is stuck at one position on its axis; it points the same way irrespective of the wind’s blowing north, south, east, or west. Tracking failures and stuck weather vanes are mechanisms whose output is in this way independent of the relevant facts. This means that in the case of a *tracking failure*, the value for the likelihood ratio is 1. Now consider tracking success. A cognitive mechanism that succeeds in tracking the relevant states of the world is like a weather vane that is not stuck but rather moves freely in response to the wind’s direction. Tracking successes and well-functioning weather vanes are mechanisms whose output is in this way dependent on the relevant facts. So in the case of a *tracking success*, the value for the likelihood ratio is greater or less than 1. These two cases are summarized in the following table:

<u>Likelihood Ratio</u>	<u>Tracking Failure</u>	<u>Tracking Success</u>
$\frac{\Pr(\text{We believe that } p p)}{\Pr(\text{We believe that } p \text{not } p)}$	= 1	≠1

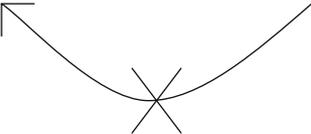
How does this discussion show that Street is mistaken to infer from a tracking failure that our moral beliefs are probably false? Well claims about tracking failures

Footnote 3 continued

credentials of our ethical beliefs and that evolutionary considerations are in some sense relevant to whether or not our ethical beliefs are vulnerable to it” (Lillehammer 2010, p. 5).

and successes are reflected in the likelihood ratio and nothing else. Claims about the probable truth of what we believe are reflected in the posterior ratio and nothing else. As we can see, no value for the likelihood ratio is sufficient to determine the value for ratio of posteriors.

Ratio of Posterior Probability	Ratio of Priors	Likelihood Ratio
$\frac{\text{Pr}(p \text{We believe that } p)}{\text{Pr}(\text{not } p \text{We believe that } p)}$	$\frac{\text{Pr}(p)}{\text{Pr}(\text{not } p)}$	$\frac{\text{Pr}(\text{We believe that } p p)}{\text{Pr}(\text{We believe that } p \text{not } p)}$
$= \quad \times \quad =$		



To obtain posteriors, one must have priors

Of course, one could assign values to the ratio of priors for various moral propositions so as to achieve a value for the posterior ratio. But what would justify one assignment over another? Subjective routes aren't helpful in answering this question, since moral realists would have very different assignments than would anti-realists. Objective routes aren't promising either, because objective estimates for prior probabilities are generally based on frequency data or well-confirmed empirical theories. These sources of justification aren't available in the present case. So in the absence of a justified estimate for the priors, nothing follows about the probable falsity of our moral beliefs from any claim about how the process responsible for our coming to have them fails to track moral facts.⁴

If our moral beliefs are subject to a tracking failure, it does not follow that they are probably false. But it is worth considering whether a tracking failure is sufficient to undermine our being *justified* in holding the moral beliefs we do. Suppose we knew that natural selection fails to track moral facts. Assuming that natural selection causes us to have the moral beliefs we do, does it follow that we are probably unjustified in maintaining these beliefs? Joyce (2006) argues that it does:

We have an empirically confirmed theory [evolution by natural selection] about where our moral judgments come from (we are supposing). This theory doesn't state or imply that they are true, it doesn't have as a background assumption that they are true, and, importantly, their truth is not surreptitiously buried in the theory by virtue of any form of moral naturalism. 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

⁴ The genetic fallacy is generally regarded as a fallacy that applies to deductive inference only. According to this fallacy, facts describing the history of a belief's acquisition do not *entail* any conclusions about that belief's truth-value. As we've just seen, these historical facts do not *probability* any such conclusions either. Thus, if the genetic fallacy is a fallacy, it applies to both deductive *and* inductive inference.

grounds one way or the other for maintaining these beliefs (Joyce 2006, p. 211).

If we have no grounds for maintaining our moral beliefs once we discover this fact about their origins, then we are not justified in continuing to hold these beliefs. The plausibility of this inference depends crucially on the following claims: (1) that tracking success is necessary for justification; (2) that natural selection is the *only* cause of our moral beliefs⁵; and (3) that natural selection is a process that fails to track moral facts. Since each of these claims is necessary for Joyce's conclusion, a successful argument against any one of them would suffice to undermine this conclusion. I do not wish to argue against claim (1), but will simply grant its plausibility. I will, however, now argue against claim (2), and in the final section, against (3).

The role of rational reflection

If the epistemic credentials of our moral beliefs are in some sense undermined by their evolutionary origins, it is worth considering whether processes of rational reflection can restore these credentials.⁶ Street considers and rejects this possibility. On her view, if our moral beliefs result from tracking failures, then they are probably false. And if our reflective process of belief revision starts with false beliefs (garbage in), it must end with false beliefs too (garbage out). Street expresses this worry as follows:

For what rational reflection about evaluative matters involves, inescapably, is assessing some evaluative judgments in terms of others. Rational reflection must always proceed from some evaluative standpoint; it must work from some evaluative premises; it must treat some evaluative judgments as fixed, if only for the time being, as the assessment of other evaluative judgments is undertaken. It follows that all our reflection over the ages has really just been—a process of assessing evaluative judgments that are mostly off the mark in terms of others that are mostly off the mark. And reflection of *this* kind isn't going to get one any closer to evaluative truth, any more than sorting through contaminated materials with contaminated tools is going to get one closer to purity. (Street 2006: 124)

One problem with this argument is that it is based on the inference that if our moral belief forming mechanisms fail to track moral facts, then our moral beliefs are probably false. I have argued that this inference is not valid. Another objection to Street's argument here is based on the possibility of getting closer to moral truth by reflecting *on*, rather than merely *with*, moral beliefs that *are* false. As an illustration of this possibility, consider the following two moral judgments:

⁵ Joyce and Street agree that (2) is false. They accept that causal processes other than natural selection influence our moral judgments. But they do not see these processes as altering the basic fact that our moral beliefs are products of processes that fail to track moral facts.

⁶ Carruthers and James (2008, p. 241) pursue this line of argument against Joyce. .

J₁: One ought to help kin, but not members of one's community who aren't kin

J₂: One ought to help unrelated members of one's community, but not outsiders

Suppose that both of these judgments are false, because it is a moral truth that we ought to help others regardless of their genealogical or physical proximity to us. How might J₁ and J₂ be used to arrive at this truth? Well, even though J₁ and J₂ are false, they do contain kernels of truth. One might discover these kernels by reflecting on the reasons that one ought to help anyone. These reasons may have to do with the fact that others would suffer in the absence of our help and that our help can reduce this suffering. By reflecting in this manner, one might discover that the obligation to help others is grounded less in genealogical and physical proximity than it is in suffering and one's ability to lessen it. In so doing, one might come to judge that one ought to help others, irrespective of whether they are kin or outsiders.⁷

The idea that we can attain moral progress by reflecting on moral beliefs that are false should not be taken to imply that such progress is attainable using beliefs that are *radically* false.⁸ If our moral judgments are way off the mark, it's difficult to see how they can be used to get us closer to the truth. For instance, it's hard to imagine a reflective process by means of which we can arrive at the judgment that others ought not to be tortured by reflecting on a judgment according to which they always should be.⁹ My only point is that the falsity of a moral belief does not automatically exclude it from its being used effectively in a reflective processes aimed at moral progress. So even if evolutionary influences have a corrosive effect on justification, the judicious use of our reflective capacities might be more than enough to counteract this effect.¹⁰

Let's suppose, however, that natural selection is the only influence on our moral beliefs. The epistemic implications of this claim depend on whether natural selection is a process that succeeds or fails to track moral facts.

⁷ The idea here isn't terribly dissimilar to Darwin's views about moral progress:

[A]s man gradually advanced in intellectual power, and was enabled to trace the more remote consequences of his actions; as he acquired sufficient knowledge to reject baneful customs and superstitions; as he regarded more and more not only the welfare, but the happiness of his fellow-men; as from habit, following on beneficial experience, instruction, and example, his sympathies became more tender and widely diffused, so as to extend to men of all races, to the imbecile, maimed, and other useless members of society, and finally to the lower animals—so would the standard of his morality rise higher and higher (Darwin 1871, I: 103)

⁸ I thank an anonymous referee for pressing this point.

⁹ One possibility for doing so might be as follows. If one thought that torturing others was the best way of making them happy, then one could discover this belief to be false by attending to the actual effects of one's action on others; seeing, e.g., that even though they've been tortured, they are miserable rather than happy.

¹⁰ Sterelny (2010) develops a similar point against certain moral nativist views. Sterelny argues that nativists typically overemphasize our reflexive moral judgments and underemphasize the role that conscious reflective reasoning plays in moral cognition.

Does evolution by natural selection track moral facts?

Street and Joyce argue that evolution by natural selection is a process that does not track moral facts. Using Sober's (1984) selection *off*/selection *for* distinction, this claim could be interpreted to mean (a), that there was no selection *for* a tracking capacity, or (b), that there was no selection *of* a tracking capacity. On the assumption that natural selection is the *only* influence on our moral beliefs, the question that needs to be addressed is this: does either (a) or (b) imply that our moral beliefs are products of a process that fails to track moral facts? In this section, I will argue that (a) does *not* support this implication. In the final section, I assume that (b) does have this implication, but will argue that there are good reasons to suppose that (b) itself is false.

Street and Joyce argue in favor of the hypothesis that our moral beliefs are products of natural selection. However, they are both careful to note, correctly in my opinion, that the influence of natural selection on our moral beliefs is *indirect* (Street 2006, p. 119, Joyce 2006, Chap. 4). According to Joyce, it is a *tendency* to make moral judgments that natural selection favored, rather than particular moral judgments with this or that content. This tendency is realized in our brains by certain innate mechanisms (Joyce 2006, p. 137). The content of the moral judgments individuals make is influenced by what these mechanisms do and by various environmental factors, such as cultural learning (Joyce 2006, p. 137). On Street's view, the mechanisms implicated in moral judgment "serve to link an organism's circumstances with its responses in ways that tend to promote survival and reproduction" (Street 2006, p. 127). For example, the mechanism implicated in the moral judgment that we ought to return favors is represented as "effecting a pairing between the circumstance of one's being helped and the response of helping in return" (Street, p. 127). So natural selection favors moral beliefs by favoring the psychological mechanisms that typically produce them.

If these mechanisms evolved by natural selection, we can ask what caused them to evolve. Did they evolve *because* of their capacity to track moral facts? Street and Joyce answer this question negatively. I am willing to grant for the sake of argument that this answer is correct: our moral beliefs are products of psychological mechanisms that were not selected because of their capacity to track moral facts. In the language of the selection *off*/selection *for* distinction, this claim asserts that there was no selection *for* this capacity. But what does this claim imply? Does it imply that that our moral beliefs are products of psychological mechanisms that *lack* a tracking capacity¹¹? It does not. That this capacity was not selected *for* does not imply that it wasn't selected *of*. Bones were not selected *for* being white, but their so being was nevertheless selected. So if our moral beliefs are products of psychological mechanisms that were not selected *for*

¹¹ Street describes two alternative explanations for why our moral belief forming mechanisms evolved, which she calls the "tracking account" and the "adaptive link account" (Street 2006, pp. 125–135). The tracking account is committed to the idea that these mechanisms were selected *for* their capacity to track moral facts. Street argues that this idea is implausible. From this, Street infers that we should accept the "adaptive link account," according to which our moral beliefs forming mechanisms were selected *for* their capacity to promote adaptive behavior. I'm not arguing against anything in this picture. My point is that facts about selection *for* do not settle questions about selection.

their capacity to track moral facts, it does not follow that they are products of a process that fails to track moral facts. That a tracking capacity was not selected *for* is consistent with it having been selected *of*, because the causes of selective processes do not settle how the products of selection operate. Of course, this claim provides no reason to suppose that a tracking capacity *was* selected. So let's now consider the question of whether it was selected, or at least could have been selected.

Was there selection of tracking?

Joyce (2006), Ruse and Wilson (1995), and Street (2006) argue that there was no selection *of* a tracking capacity, and for similar reasons, as these passages illustrate.

Whether we assume the concepts *right* and *wrong* succeed in denoting properties in the world, or whether we think that they suffer from a referential failure that puts them on par with the concepts *witch* and *ghost*, the plausibility of the hypothesis concerning how moral judgment evolved remains unaffected (Joyce 2006, p. 183, emphasis in original).

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 that we do. And surely, redundancy is the last predicate that an objective morality can possess (Ruse and Wilson, 1995, pp. 186–187)

The moral beliefs that evolved by natural selection are “the very same judgments we would expect to see if our judgments had been selected on those [fitness] grounds alone, regardless of their truth[.]” (Street 2006: 132)

We can represent these ideas as follows. Let x be the process that governs the moral beliefs we have (e.g., x could be evolution by natural selection [Joyce 2006; Ruse and Wilson 1995; Street 2006], or it could be our upbringing [Harman 1977]). Let p represent some moral proposition.

Given that x occurs, we'll believe that p whether or not p is true.
Therefore, our believing that p does not track whether p is true.

Another way of representing this argument is to say that that, conditional on x , our believing that p is independent of p . On this construal of the argument, x screens-off p from our believing that p . This is taken to imply that our believing that p is independent of p . In the language of conditional probabilities, we can formulate the argument this way.

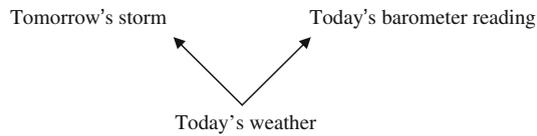
Screening off thesis:

$$\begin{aligned} \Pr(\text{We believe that } p \mid x \text{ occurs and } p) = \\ \Pr(\text{We believe that } p \mid x \text{ occurs and not } p) \end{aligned}$$

Therefore (Tracking Failure),

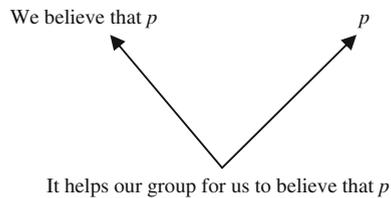
$$\begin{aligned} \Pr(\text{We believe that } p \mid p) = \\ \Pr(\text{We believe that } p \mid \text{not } p) \end{aligned}$$

We can see that this inference is invalid by considering other screening-off relations, for example the relation between weather, barometer readings, and storms.



The weather today screen's-off the storm tomorrow from today's barometer reading. Yet storms and barometer readings are correlated because they trace back to a common cause.

Reichenbach (1956, pp. 158–160) has shown that if A raises the probability of B and of C, and if A screens-off B from C, then B and C will be correlated. This means that we can derive a tracking success from the screening-off thesis as follows. Consider the belief that cooperation with others is morally good. Suppose that this belief was favored by natural selection because it enhanced our capacity for helping behavior; individuals who believe that cooperation is morally good are more likely to help others than are individuals who lack this belief. The helping behavior that this belief generates has two effects: it promotes fitness, and it promotes wellbeing. The former effect is what explains why it might evolve by natural selection. The latter effect is part of what may explain why it is that cooperation in fact is morally good. If what's morally good has to do with behaviors that promote rather than hinder wellbeing, then part of what makes cooperation good is that it typically has this effect. This picture can be represented as follows ('*p*' represents the claim that cooperation is morally good; the arrows represent a relation of probability raising)



If the following two conditions are satisfied—

(1)

$\Pr(\text{We believe that } p \mid \text{It helps our group for us to believe that } p) >$

$\Pr(\text{We believe that } p \mid \text{It does not help our group for us to believe that } p)$

(2)

$\Pr(p \mid \text{It helps our group for us to believe that } p) >$

$\Pr(p \mid \text{It does not help our group for us to believe that } p)$

—it follows that

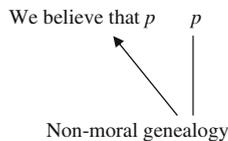
(3)

$\Pr(\text{We believe that } p \mid p) >$

$\Pr(\text{We believe that } p \mid \text{not } p)$

Notice that (1) is not committed to the implausible claim that if our believing that p helps our group, then this belief must or even will evolve by natural selection. Notice also that (2) is not committed to the equally implausible claim that if our believing that p helps our group, then p is true, or is probably true. Contrary to the argument stated above, the screening-off thesis does not support a tracking failure, but actually entails a tracking success, given claims (1) and (2).

Joyce (2006, p. 189) considers another way (which he regards as the *only* way) of deriving a tracking success from the screening-off thesis. He claims that if moral facts are reducible to the natural facts that cause our moral beliefs to evolve, then our moral beliefs will correlate with these moral facts. Joyce presents this scenario as follows; the arrow represents a causal relation, while the straight line represents a reductive relation.¹²



Joyce rejects this picture because he is skeptical that moral facts can be reduced to natural facts of any sort, evolutionary or otherwise (Joyce 2006, pp. 190–199). Any such reduction, on his view, must show how the natural world is capable of providing the “*inescapable authority* we apparently expect and require of moral values” (Joyce 2006, p. 191, italics in original). This inescapable authority is due the idea that moral norms on Joyce’s view prescribe actions that we would want to do, or would have reason to do, *independently of our interests*. But whatever reasons we have to do things, these reasons are to one degree or another necessarily tied to our interests (Joyce 2006, pp. 196–198). So the natural world lacks the type of practical reason required of it to serve as a reductive base for moral facts. For this reason, Joyce rejects the claim represented in the figure above.

I have no quarrel with any part of Joyce’s argument against efforts to reduce moral facts to natural ones. However, I think he is mistaken to regard the success of such efforts as the *only* way to derive a tracking success from the screening-off thesis. As I have shown, deriving a tracking success from the screening-off thesis does *not* require that the natural world possess the type of practical reason required of it to serve as a reductive base for moral facts. The relation between the moral and the natural in my example is not a reductive relation, as it is on Joyce’s; it needn’t be a causal relation either, or for that matter, a supervenience relation. My example only requires that natural facts are capable of raising the probability of moral facts. This is a very modest requirement. I conclude that natural selection is a process that *could* track moral facts.¹³

¹² Joyce construes reduction “in the broad ontological sense of the term” (Joyce 2006, p. 188). Moral facts on this construal reduce to natural facts in the same way the term *cat* reduces to the language of physics and chemistry. .

¹³ My proposal is related to what David Enoch (2010) calls a “third-factor explanation” for the correlation between our moral beliefs and moral truth. Enoch’s idea is roughly that if there is no causal

One objection to my proposal is related to Hume's argument that a moral *ought* cannot be derived from any argument with exclusively descriptive (or *is*) premisses. For example, the following argument is invalid according to Hume.

- (1) Wide distribution of mosquito netting in areas with high rates of malaria is cheap and easy to do and will prevent millions of people from dying (a descriptive *is* claim).
- (2) Mosquito netting should (morally) be widely distributed in these areas.

Call the rule that precludes the validity of any such inference "Hume's rule". Hume's rule is generally regarded as applicable to deductive inference only, leaving open the possibility that descriptive premisses can raise or lower the probability of normative claims. So even though (1) does not entail (2), it *may* raise its probability. Whether (1) does this or not, however, depends crucially on background normative assumptions. For instance, if a consequentialist moral principle is assumed, then (1) clearly does raise the probability of (2). But in the absence of *any* normative assumption, it is hard to see how this relation (or any probabilistic relation) between (1) and (2) can be sustained. Sober (1994) takes this to imply that Hume's rule applies as much to inductive inferences as it does to deductive ones. As he puts this point, "*purely is-premisses cannot, by themselves, provide non-deductive support for an ought-conclusion*" (Sober 1994, p. 109, italics in original). I entirely agree with Sober on this point. This means that the example I used to illustrate how natural selection tracks moral facts won't work unless I assume the truth of some moral principle. I don't see this as a problem, however. It may appear to beg important questions in the context of the evolutionary debunking arguments of Street and Joyce, but it does not.¹⁴ Joyce and Street both assume for the sake of argument that moral facts exist. If moral facts exist, then the inductive version of Hume's rule presents no obstacle to my argument. Moreover, it's as reasonable an assumption as any that among these facts is this one: what's morally good has to do with actions that promote rather than hinder wellbeing.

Another concern with my proposal is that it may exclude non-natural varieties of moral realism. Street is careful to note that the target of her argument is primarily the moral realism of just this variety (Street 2006, p. 112). This is the view, very crudely, that moral facts are not reducible to natural facts and play no role in causal explanation. For example, Street suggests that natural selection cannot track non-natural moral facts because "a creature can't run into such truths or fall over them or be eaten by them" (Street 2006, p. 130). But if moral facts are isolated from the

Footnote 13 continued

relation between our moral beliefs and the moral facts in virtue of which these beliefs are true or false, as realists maintain, then whatever positive correlation there is between the two must be explained by a "third factor" that is responsible for both. Enoch's third-factor is the goodness of our survival. Wielenberg (2010) and Skarsaune (2009) defend similar accounts, but offer different third factors. For Wielenberg, the third factor consists of certain cognitive faculties, while for Skarsaune, it consists in the claim that pleasure is usually good and pain is usually bad. My own account differs from these in two ways: (1) it offers a different third factor, and (2) the relation between this factor and moral facts is neither causal nor logical, as it on the other accounts. I have developed my account independently of these authors.

¹⁴ I thank an anonymous referee for raising this objection. .

world of creatures in these ways, it doesn't follow that natural selection cannot track them. Consider one such fact: whatever promotes rather than hinders wellbeing is morally good. Natural selection can track this fact by tracking facts about behaviors that vary in the degree to which they do promote rather than hinder wellbeing.¹⁵ If cooperation hadn't been good, i.e., if it had hindered rather than promoted wellbeing, it may well *not* have been selected. There's nothing mysterious about this. Moral facts can make a difference to what natural selection favors even if they are non-natural; even, that is, if organisms cannot be eaten or run over by them.

Of course, there are many behaviors that promote fitness but not wellbeing, just as there are many that promote wellbeing but not fitness. So the process of natural selection does not track moral facts perfectly. But perfect tracking isn't needed to avoid the adverse epistemic consequences implied by tracking failures. So long as there is some positive correlation between our moral beliefs and moral facts, the proper use of our reflective capacities may provide enough of an epistemic correction.

Conclusion

If moral knowledge requires that our moral beliefs are true, then evolutionary considerations do not provide us a reason to doubt that they are. If moral knowledge requires satisfaction of a tacking condition, then evolutionary considerations do not undermine this requirement but rather show how it might be satisfied. So far from undermining moral knowledge, Darwinian considerations can help to show how this knowledge might come about.

Acknowledgments I would like to thank Elliott Sober for his invaluable help with this paper, provided promptly at each of its many stages. For reading earlier drafts and for their helpful feedback, I'd like to thank Kim Sterelny, an anonymous referee from this journal, and: Alexander Bird, Alex Broadbent, David Copp, Daniel Guevara, Russ Shafer-Landau, Hallvard Lillehammer, William A. Rottschaefer, Mark van Roojen, and Ralph Wedgwood. I would also like to thank Sorin Bangu, Simon Blackburn, Steve John, Jonathan Ellis, and Richard Otte.

References

- Blackburn S (2010) Sharon street on the independent normative truth as such. Available at <http://www.phil.cam.ac.uk/~swb24/PAPERS/Meanstreet.htm>
- Boyd R (1988) How to be a moral realist. In: Sayre-McCord G (ed) *Essays on moral realism*. Cornell University Press, New York, pp 181–228
- Brink D (1989) *Moral realism and the foundations of ethics*. Cambridge University Press, Cambridge
- Carruthers P, James S (2008) Evolution and the possibility of moral realism. *Philos Phenomenol Res* LXXVII:237–244
- Copp D (2008) Darwinian skepticism about moral realism. *Philos Issues* 18:186–206
- Darwin C (1871) *The descent of man, and selection in relation to sex*. Princeton University Press, Princeton, 1981

¹⁵ Simon Blackburn (2010) defends a similar point against Street. See his “Sharon Street on The Independent Normative Truth as such”, at <http://www.phil.cam.ac.uk/~swb24/PAPERS/Meanstreet.htm>.

- Enoch D (2010) The epistemological challenge to metanormative realism: how best to understand it, and how to cope with it. *Philos Studies* 148:413–438
- Harman G (1977) *The nature of morality*. Oxford University Press, Oxford
- Joyce R (2006) *The evolution of morality*. MIT Press, Cambridge
- Lillehammer H (2010) *Methods of ethics and the descent of man: Darwin and Sidgwick on ethics and evolution*. Biol Philos
- Railton P (1986) Moral realism. *Philos Rev* 95:163–207
- Reichenbach H (1956) *The direction of time*. University of California Press, California
- Ruse M, Wilson EO (1995) Moral philosophy as applied science. In: Sober E (ed) *Conceptual issues in evolutionary biology*, 2nd edn. MIT Press, Cambridge, pp 421–438
- Shafer-Landau R (2003) *Moral realism: a defense*. Oxford University Press, Oxford
- Skarsaune K (2010) Darwin and moral realism: survival of the fittest. *Philos Studies*
- Sober E (1984) *The nature of selection: evolutionary theory in philosophical focus*. MIT press, Cambridge
- Sober E (1994) *From a biological point of view*. Cambridge University Press, Cambridge
- Sterelny K (2010) Moral nativism: a sceptical response. *Mind Lang* 25(3):279–297
- Street S (2006) A Darwinian dilemma for realist theories of value. *Philos Studies* 127:109–166
- Sturgeon N (1985) Moral explanations. In: Zimmerman D, Copp D (eds) *Morality, reason and truth*. Rowman and Allanheld, Totowa
- Wielenberg E (2010) On the evolutionary debunking of morality. *Ethics*