

Dissolving a dilemma: why Darwinian  
considerations don't confront moral realism with hard choices.

Kevin Brosnan  
University of Cambridge

Suppose that moral facts (*MF*) exist. Now consider the process (*P*) by means of which our moral beliefs are formed. *P* operates either, a) independently of *MF*, or b) dependently on *MF*. Street (2006) and Joyce (2006) argue that option b) is implausible. On their view, the moral beliefs we possess are a result of evolution by natural selection. And evolution by natural selection is a process, on their view, that operates independently of moral facts. Against this claim, I will argue that evolutionary processes and moral facts can be positively correlated. But suppose that I am mistaken and that Street and Joyce are correct. This leaves us with option a). From a), Street infers that our moral beliefs are very probably false, while Joyce infers that they are unjustified. I will argue that a) does not support either of these conclusions.

The target of Street's argument is moral realism, which she characterizes as the view that the truth-values of moral beliefs are independent of anyone's beliefs or desires that they be true or false (Street 2006: 110). While her primary concern is with non-natural moral realism—the view that moral facts are not reducible to natural facts and play no role in causal explanation (Shafer-Landau 2003)—she maintains that her arguments apply to naturalistic varieties of moral realism as well—the view that moral facts are identical with or are constituted by natural facts and do play a role in causal explanation (Boyd 1988, Brink 1989, Railton 1986, Sturgeon 1985).

Philosophers who draw metaethical conclusions from evolutionary theory generally emphasize the speculative status of their evolutionary hypotheses (Gibbard

1990, Joyce 2006, Kitcher 2005). Street is no exception: “I try to rest my arguments on the least controversial, most well-founded evolutionary speculations possible. But they are speculations nonetheless [...]” (Street 2006: 112). While she admits that the details of her proposed explanations are speculative—details as to whether our moral beliefs are adaptations, by-products of selection, or a result of drift—she is confident that her evolutionary hypotheses are in general plausible and that they mean trouble for moral realism.

#### An evolutionary explanation of our moral beliefs -- the “adaptive link” account

The strategy that Street adopts in defending her view, which she calls the “adaptive link account”, is to identify an adaptive behavior or action (A) and infer that a particular moral belief (B) is an important cause of A’s production. Consider the belief that “the fact that something would promote one’s survival is a [moral] reason to do it.” (Street 2006: 115) Street’s explanation for why this belief is a product of natural selection is that “creatures who possess this general evaluative tendency tended to do more to promote their survival than those who, say, had a tendency to view the fact that something would promote their survival as counting against it, and so the former tended to survive and reproduce in greater numbers.” (Street 2006: 116)

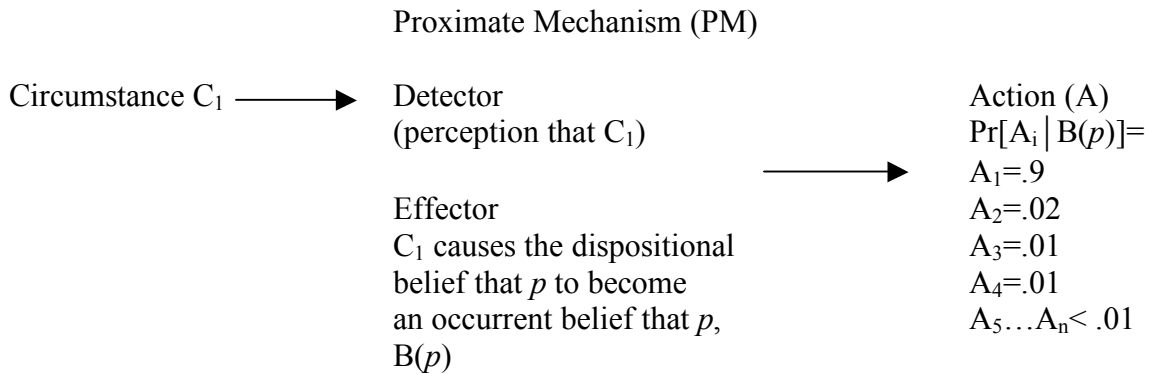
To develop this account, it is useful to distinguish between dispositional and occurrent beliefs. Suppose that Josh believes that dogs are better than cats. Since he rarely encounters either animal, he rarely entertains this proposition. However, on the rare occasion when asked whether he prefers one to the other, he answers that he believes firmly that dogs are superior. On this occasion, Josh has the belief occurrently. When

the belief is not in the forefront of his mind, he possesses it dispositionally. According to representational warehouse models of memory and belief, “a subject dispositionally believes P if a representation with the content P is stored in his memory or ‘belief box’ . When that representation is retrieved from memory for active deployment in reasoning or planning, the subject occurrently believes P.” (Schwitzgebel 2006: 12).

In light of this distinction, we should understand Street’s evolutionary argument as follows: organisms with the dispositional belief that P were favored by natural selection because, a), in the appropriate circumstances C, this dispositional belief tended to become occurrent, and b) this occurrent belief played an important causal role in the production of a fitness enhancing behavior. The following passage illustrates Street’s view:

As a result of natural selection, there are in living organisms all kinds of mechanisms that serve to link an organism’s circumstances with its responses in ways that tend to promote survival and reproduction. A straightforward example of such a mechanism is the automatic reflex response that causes one’s hand to withdraw from a hot surface, or the mechanism that causes a Venus’s-flytrap to snap shut on an insect. Such mechanisms serve to link certain kinds of circumstances—the presence of a hot surface or the visit of an insect—with adaptive responses—the immediate withdrawal of one’s hand or the closing of the flytrap. Judgments about reasons—and the more primitive, “proto” forms of valuing that we observe in many other animals—may be viewed, from the external standpoint of evolutionary biology, as another such mechanism. They are analogous to the reflex mechanism or the flytrap’s apparatus in the sense that they also serve to link a given circumstance with a given response in a way that may tend to promote survival and reproduction. Consider, for example, the evaluative judgment that the fact that someone has helped one is a reason to help that individual in return. Just as we may see a reflex mechanism as effecting a pairing between the circumstance of a hot surface and the response of withdrawing one’s hand, so we may view this evaluative judgment as effecting a pairing between the circumstance of one’s being helped and the response of helping in return. Both of these pairings of circumstance and response, at least if the evolutionary theory of reciprocal altruism is correct about the latter case, are ones that tended to promote the reproductive success of ancestors who possessed them. (Street 2006: 127-128)

This view can be represented as follows:



Although Street does not state explicitly that occurrent beliefs make it more or less probable that certain actions will occur, rather than determining that they will, it is clear from what she does say that this is what she has in mind. This is as it should be, since interfering forces can always alter the normal functioning of any proximate mechanism. For example, the reflex that normally causes my hand to withdraw from a hot stove could fail to do so if the peripheral nerves in my hand were damaged.

There are a couple of issues that Street does not explicitly address. For example, if two or more proximate mechanisms produce the same behavioral output, why did selection favor one over the other? And, why wouldn't selection favor having both?<sup>1</sup> She does consider the question of why selection favored a mechanism that involves beliefs, rather than a brute reflex mechanism that does not, just to set it aside. In regards to the role of desire as it relates to the production of behavior, she assumes, but does not explicitly state or argue, that S's desire or motivation to do action A is more probable if S believes there is a moral obligation to do A than if S believes there is no such moral obligation (or has no views about its morality at all).

One problem affecting Street's account can be brought into view by considering the reason she offers in support of her claim that certain beliefs evolved by natural selection. Street recognizes that a trait can so evolve only if it was fitter than a set of alternative traits. So if the moral beliefs we currently possess evolved by natural selection, we need to identify a set of alternative traits with which they competed. These could be alternative beliefs, or some non-cognitive trait that performs a similar function. The alternatives that Street considers are the following alternative beliefs: "(1') The fact that something would promote one's survival is a reason *against* it; (2') The fact that something would promote the interests of a family member is a reason *not* to do it" (Street 2006: 116, italics added). Let's assume that Street is correct that individuals with these beliefs would be less fit than individuals with those beliefs that Street argues evolved by natural selection. This does *not* provide a good reason to support her hypothesis, even if we assume that natural selection was the only important evolutionary force at work. The reason for this is that selectionist explanations concern the range of *actual* variation in trait T that existed in a particular population. On Street's view, the range of variation one needs to consider in defending an adaptive hypothesis is "the universe of [the] logically possible" (Street 2006: 133). But surely this casts too wide a net. Relative to this hypothetical range of variation, the superior fitness of a trait T that we actually observe is almost impossible to establish. Consider an adaptive explanation for the speed and camouflage of zebras that cited their role in predator defense. Since it is logically possible for zebras to defend against predators with machine guns (Krebs and Davies 1981), it would be very surprising that selection produced camouflage and speed as defense mechanisms rather than the more effective alternative just described. Or

consider an adaptive explanation for the white color of snow owls. According to Street's proposed methodology, one could argue as follows: since white owls would have been fitter than orange ones, whiteness is clearly a product of selection. However, the fact that white snow owls would have been fitter than orange ones is no reason to maintain that whiteness is a product of selection, if orange snow owls never actually existed. The logically possible should be considered only as a necessary condition for membership in the set of variants that was present ancestrally, rather than one that is sufficient. I am not claiming that traits such as the belief that 1' or 2' could not have been among the variants that selection acted upon. But to assess Street's explanatory claims, we need to address the question of whether these traits actually *did* exist. This question is not settled by considering only logical possibility.

Since the dilemma Street poses for the moral realist is independent of this part of her evolutionary explanation, I won't consider any further the extent to which her explanatory claims can be made to work.

### An alleged dilemma for moral realism

The dilemma Street poses for the moral realist is as follows:

On the one hand, the 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)

The claim that evolutionary forces bear no relation to moral facts is analogous, according to Street, to the claim that the direction and magnitude of ocean currents bear no relation to human desires. (Street 2006: 121) One way to represent this idea is in terms of the idea of probabilistic independence; Street’s idea that there is “no relation” between O and Mf then can be formulated as the thesis that

$$\Pr(O \mid Mf) = \Pr(O \mid \neg Mf)$$

Here O = S believes that action A is right, and Mf = action A is right. Street’s “no relation” thesis means that the relation of probabilistic independence (of zero correlation) obtains. Street argues that if our moral beliefs are the result of a process— natural selection—that operates (probabilistically) independently of moral facts, then our moral beliefs are probably false. I will try to show that this does not follow.

The epistemic implications of independence:

Consider the processes by means of which two individuals (Bill and Carl) form their respective beliefs about the number of marbles inside an urn. Bill empties the urn of its marbles and, with all his psychological capacities functioning normally, carefully counts them. Carl writes a number between 1 and 100 on 100 strips of paper, tosses them into a hat, draws one at random, and concludes that the urn contains whatever number of marbles is written on it. Given the process that led Bill to form his belief, it does not follow deductively that his belief is true; after all, careful counting is not an *infallible*

process. Similarly, given the process that led Carl to form his belief, it does not follow deductively that his belief is false; after all, guessing *sometimes* gets it right. While conclusions about the truth or falsity of a belief cannot be *deduced* from considerations as to how that belief was formed, it is worth investigating whether such considerations have epistemic implications that are *probabilistic*.

Sober argues that they do. On his view, the following claim is *false*: “conclusions about the truth of a proposition cannot be *inferred* [non-deductively] from premises that describe why someone came to believe the proposition.” (Sober 1994: 106). To reach a conclusion about the truth status of a proposition believed requires information concerning the relation between how the belief that  $p$  [ $B(p)$ ] was formed and the facts ( $f$ ) in virtue of which  $p$  is true or false. On Sober’s view, if the process leading to  $B(p)$  is *independent* of  $f$ , then  $p$  is probably false. Alternatively, if the process leading to  $B(p)$  is *dependent* on  $f$ , then  $p$  is probably true (Sober 1994: 106). These (genetic) rules of inference suggest that Carl’s belief is probably false, while Bill’s is probably true. They also support Street’s view that, if our moral beliefs were formed by a process that operates independently of moral facts, then our moral beliefs are probably false<sup>2</sup>.

While these genetic rules of inference seem plausible in the context of the urn example, there are other contexts in which they seem mistaken. Consider Ned, who believes that the earth is round because he believes that roundness is the perfect shape and that the earth is perfect<sup>3</sup>. Ned’s belief-forming process is independent of the relevant facts, just as in Carl’s case. However, unlike the case of Carl, this fact does nothing to undermine our confidence in the truth of Ned’s belief. Alternatively, consider Fred, who carefully counts the marbles inside an urn. Suppose that he counts  $n$  marbles. However,

he then comes to believe that the urn contains  $\pm n^2$  marbles. Fred's belief forming process is dependent on the relevant facts, just as in Bill's case. But unlike Bill's case, this fact does *not* suggest that his belief is probably true. It suggests just the opposite.

As these examples suggest, facts that describe how someone came to believe a proposition do not *by themselves* have (probabilistic) implications concerning the truth status of that proposition. To reach a conclusion about the truth status of  $p$  from facts that describe how someone came to believe that  $p$  requires additional information. We can identify just what this additional information is by using Bayes' Theorem to illustrate how the process whereby an agent came to believe that  $p$  [ $B(p)$ ] is related to the probability that  $p$  is true or false, as follows<sup>4</sup>. Specifically, Bayes' theorem entails the following relationship among three ratios.

<u>Posterior Probability</u>	<u>Priors</u>	<u>Likelihood</u>
$\frac{\Pr [p   B(p)]}{\Pr [\neg p   B(p)]}$	$\frac{\Pr (p)}{\Pr(\neg p)}$	$\frac{\Pr [B(p)   p]}{\Pr [B(p)   \neg p]}$
=	x	

Notice that different belief-forming processes provide different values for the likelihood ratio (LR). When the process is such that what you believe is *independent* of what is true, LR=1, and when what you believe is *dependent* on what is true, LR is less than or greater than 1. Notice that no value for the likelihood ratio (other than 0) provides a value for the ratio of posterior probabilities without requiring the input of a value for the ratio of the prior probability of  $p$ . The fact that likelihoods alone provide no value for posteriors is what explains why it is a mistake to infer that a belief is probably

false (or probably true), given that it was formed by a process whose output is independent of the truth of the proposition believed. This also explains why it's wrong to infer that a belief is probably true, given that it was formed by a process whose output is *dependent* on the truth of the proposition believed. It may seem as if a belief forming process whose output  $[B(p)]$  is sensitive to  $p$  is likely to yield true beliefs.<sup>5</sup> For example, the process of carefully counting the number of marbles in an urn would seem likely to produce a true belief about the number of marbles it contained. Street seems to accept this idea, for she says that beliefs formed *dependently* on the relevant facts are probably true (Street 2006: 125). But this is mistaken for the same reason it is a mistake to infer that beliefs formed independently are probably false: both inferences rely on an illegitimate leap from likelihoods to posterior probabilities.

These considerations suggest that the only epistemic implication of independence, absent information about priors, is ignorance as to whether the belief in question is probably true. So if our believing various moral propositions is driven by a process that operates independently of presumptive moral facts, this leaves completely open whether or not any of our moral beliefs are true, or what the probability is that they are true, which is very different from the conclusion that they are very probably false.

How damaging to moral realism is the claim that we have no idea whether our moral beliefs are probably true? Strictly speaking, moral realism is a view about the truth-determinants of moral belief, as Street recognizes. It is perfectly consistent with moral realism to maintain that the moral beliefs we possess are false or are very probably false. Nevertheless, Street maintains that moral realists should accept this conclusion only as a last resort, though I don't see why realists should agree with Street on this.

Moreover, even if Street's assessment in this regard is correct, perhaps the conclusion that we have no reason to deny or affirm the truth of our moral beliefs is more acceptable. Why? Well, consider that any process responsible for the content of our moral beliefs has implications for the value of the LR. But without a value for the ratio of priors, this tells us nothing about how probable it is that a given proposition is true, given that one believes it. Of course, the prior probability of a proposition can be estimated in many cases. For example, suppose that the sensitivity (to the presence of cancer) of cancer-testing device provided a high LR value. If we use this device to test a randomly selected individual (*S*) from a particular population (*P*), knowledge of the cancer rate in *P* would allow us to estimate a value for the ratio of priors. This would allow us to determine the probability that *S* had cancer, given a positive test result. But how could we obtain a value for the prior probability of a moral principle? Suppose that we cannot. Well, if natural selection operates independently of moral facts, and a) it is the *only* influence on our moral beliefs, and b) there is no *other* evidence that bears on our moral beliefs, then this implies that we have no reason to affirm or deny the truth of our moral beliefs. However, we should consider the possibility that there are influences on our moral beliefs *other than natural selection*, for example the method of reflective equilibrium. It may be that natural selection explains the *origin* of our moral beliefs, but this does not imply that it also explains why we continue to hold or reject them, or why we modify them in one way or another. We need to distinguish between the process that initially *produces* a belief and the factors that influence whether that belief is *maintained*.

Joyce agrees with Street that our moral beliefs were formed by natural selection and that natural selection is a process that operates independently of whatever moral facts

there happen to be. The implication he draws from this, however, is that our moral beliefs are unjustified.

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

I interpret Joyce's use of the term "independence" to mean probabilistic independence, just as I did with Street's. Notice that the inference Joyce makes in the last sentence of this passage relies on natural selection being the *only* influence on our moral beliefs. The mere fact that our moral beliefs evolved by natural selection may provide no reason to maintain that they are justified, as Joyce claims, but it doesn't rule out the possibility that there are *other* reasons for thinking they are justified. Suppose that storms are caused only by barometric pressure changes. This doesn't mean that the only information we can have about storms is by knowing what the barometric pressure is. Eye-witnesses and historical records, for example, could provide evidence about storms.

In support of his claim that independence renders beliefs so formed unjustified, Joyce relies on the following example (Joyce 2006: 179). Suppose there exists a belief pill that makes one believe Napoleon lost at Waterloo, and another that makes one believe he won. Now suppose you were given the former pill (at random) without your knowledge at an early age, and as a result possess the belief that Napoleon lost at Waterloo. Since the process that gave you this belief is independent of the relevant fact (in this case, the fact that Napoleon did lose Waterloo), should we conclude that your belief is unjustified? It doesn't seem that we should, as the following Gettier-style

example illustrates. Jones has a very accurate clock that has provided him with the correct time over many days. One day, with all of his perceptual systems functioning normally, he looks at the clock and it indicates that's it's 3 pm, and indeed it is 3 pm. However, since the clock is broken, it would say 3 pm regardless of the actual time. Jones' belief was thus formed independently of the relevant facts, yet his belief clearly is justified. He based his belief on reliable indicators (the clock and his vision), so on externalist grounds he's justified. He had, suppose, clear access to his mental representation of the clock's reading, so on internalist grounds he is also justified. If these judgments are right, as I think they are, we should conclude that independence does not render the beliefs so formed unjustified.

One objection to this point would be to argue that the clock example misses the point. After all, the clock has been reliable in the past, and then it changes. This is not so in the belief pill example. True, but my point is only that 'independence' *alone* does not suffice to show that beliefs so formed are unjustified. The belief pill and clock example are similar in the relevant way, since both represent cases in which this 'independence' condition is satisfied. I suspect that what makes the bill pill example compelling is the existence of evidence that could render the belief it produces justified. To what extent this is also true of our moral beliefs remains to be seen.

We should now consider the question of what to think about cases in which one *discovers* that one's beliefs were independently formed. Does Jones' belief become unjustified once he discovers that the clock is broken? Joyce answers this question affirmatively. If you discover that your belief about Napoleon's defeat at Waterloo is a result of swallowing the belief pill, you should view this belief as unjustified. While his

may be so with respect to this particular example, the inference does not seem to generalize across cases. Suppose Sol has good evidence that he has malaria. He then gets a blood test and it indicates that he does have it. However, Sol is then informed that the testing device is broken such that it would have yielded a positive reading whether malaria was present or not. The likelihood ratio is 1, relative to the observed test outcome, and Sol knows this. Does this render Sol's believing that he has malaria unjustified? No, it clearly does not. The explanation for this verdict is all of the other evidence Sol has for his belief. So if natural selection does operate independently of moral facts, and we know this to be the case, we should *not* conclude that our moral beliefs are therefore unjustified. To reach this conclusion, we would need to consider what *other* influences there are on our moral beliefs and to evaluate the extent to which they might serve to justify them.

There are many questions about why people have the moral beliefs they do that the theory of natural selection does not address; for example, why views about the permissibility of slavery changed in the 18<sup>th</sup> and 19<sup>th</sup> century. It's unlikely that individuals who believed slavery to be morally wrong were biologically fitter than those who believed otherwise, and that, as a result of this superior fitness, the frequency of this belief increased in particular regions over this period of time. Street admits that there are influences on our moral beliefs other than natural selection. As she says, "[a]ny full explanation of why human beings accept the evaluative judgments we do would need to make reference to the large influence of rational reflection" (Street 2006: 15). So even if natural selection operates independently of moral facts and thus provides us no reason to accept (or reject) the moral beliefs we possess because of natural selection, rational

reflection might be a means by which we can assess the acceptability of other beliefs that have a different etiology. And even if a belief is caused by natural selection, this does not mean that selection is the only cause, or that the belief isn't checkable via rational reflection of some kind. Street considers and rejects this possibility:

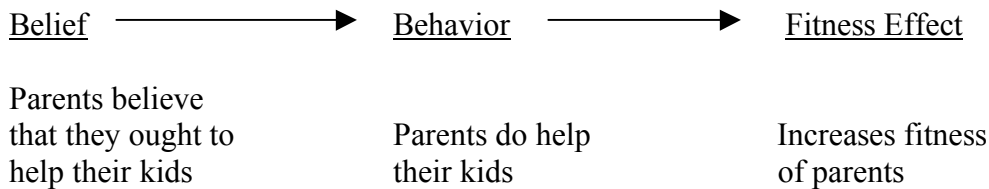
For what rational reflection about evaluative matters involves, inescapably, is assessing some evaluative judgments in terms of others. Rational reflection must always proceed from some evaluative standpoint; it must work from some evaluative premises; it must treat some evaluative judgments as fixed, 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)

Street is surely correct that the method of reflective equilibrium always relies on evaluative judgments, though it is arguable whether this method involves a search for mere consistency (or some type of coherence) among whatever moral beliefs did evolve by natural selection. Consider the following example of how reflective equilibrium might proceed. If a mother sees her child drowning in a shallow pond, what should she do? Save her of course. From this particular judgment, we might construct a more general one: if a mother can save the life of her daughter at little cost to herself, she should do so. If we change this principle by substituting 'a person' for the mother, and 'an unrelated child' for her child, keeping everything else the same, would our judgment about what that person should do change? No. This would allow us to re-formulate the principle as follows: if one can save the life of an innocent child at little cost to oneself, one should do so. Similar changes to the example produce similar results. For instance, we could substitute 'unrelated adult' for child, with the revised principle reading as follows: if one can save the life of a person at little cost to oneself, one should do so.

Iteration of this method enables us to evaluate the moral relevance of different properties. For example, in situations where an individual can be saved at a small cost, it is morally irrelevant whether that individual has the property of being a child or the property of being an adult. This process should not be thought of as a method merely for identifying sets of moral beliefs that are internally consistent logically. For example, the proposition that mothers ought to save their daughters is logically consistent with the proposition that mothers ought *not* to save unrelated children. If the method of reflective equilibrium is to play a role in justifying our moral beliefs, then the relevant notion of consistency has to be something other than logical consistency. Instead, we should think of reflective equilibrium as a process by means of which we accept or reject moral propositions insofar as they are consistent with a particular (proto) moral theory, for example a theory as to what constitutes moral standing. For instance, consider a theory according to which moral standing is directly proportional to biological relatedness. The proposition that mothers ought to save their daughters may be consistent with this theory, while the proposition that mothers ought to save unrelated children may not be. Alternatively, both propositions may be consistent with a different moral theory according to which the cost/benefit ratio of an action (measured in terms human happiness) determines our moral obligations. The process by means of which we establish a view as to what matters from a moral point of view is by no means clear. However, what is clear is that we cannot do so by looking just for logical consistency among our moral beliefs.

The second horn of Street’s dilemma: why presumptive moral facts are ‘unrelated’ to evolution by natural selection

Street offers the “adaptive link” account and the “tracking account” as competing explanations for why certain moral beliefs evolved by natural selection. However, there is a lot that the “adaptive link” account and “tracking account” have in common. For instance, on both accounts, certain moral beliefs evolved by natural selection because they played important roles in the production of behaviors that promoted fitness. This general idea can be represented as follows:



If the tracking account and the adaptive link account agree with this general picture, what is it they disagree about? According to Street,

[t]he tracking account obviously posits something extra that the adaptive link account does not, namely independent evaluative truths (since it is precisely these truths that the tracking account invokes to explain why making certain evaluative judgments rather than others conferred advantages in the struggle to survive and reproduce). The adaptive link account, in contrast, makes no reference whatsoever to evaluative truth; rather, it explains the advantage of making certain evaluative judgments directly, by pointing out how they got creatures who made

them to act in ways that tended to promote reproductive success. Thus, the adaptive link account explains the widespread presence of certain values in the human population more parsimoniously, without any need to posit a role for evaluative truth. (Street 2006: 129)

According to Street's preferred explanation, the probability that selection will produce a particular evolutionary outcome is the same whether moral facts exist or not. The moral beliefs that evolved by natural selection are, on Street's view, "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). This situation can be represented as follows. First some notation:

O=S believes that action A is right  
Mf=action A is right  
Nf=selection favored action A

Street's view is that natural facts screen-off moral facts from evolutionary outcomes, which can be represented as follows:

$$\Pr(O | Nf) = \Pr(O | Nf \& Mf)$$

Street further maintains that this screening-off thesis implies that moral facts and evolutionary outcomes are probabilistically independent.

To see why this does not follow, consider the relation between a drop in barometric pressure at  $T_1$  (Dbp), a low barometer reading (Lbr) at  $T_1$ , and a storm (St) at  $T_2$ . Notice that Dbp screens-off Lbr from St:  $\Pr(\text{St} | \text{Dbp}) = \Pr(\text{St} | \text{Dbp} \& \text{Lbr})$ . However, it does not follow from this that storms and barometer readings are independent of each other. A drop in barometric pressure is a common cause of storms and low barometer

readings, which explains how the latter two events can be correlated with each other. Reichenbach's (1956) principle of the common states the probability relations that must obtain if two events are correlated in virtue of having a common cause. Specifically, he proved 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. We can apply this straightforwardly to the relation between O, Mf, and Nf. From the claim that Nf screens-off Mf from O— $\Pr(O | Nf) = \Pr(O | \neg Nf)$  & Mf—if we add

- A<sub>1</sub>:  $\Pr(O | Nf) > \Pr(O | \neg Nf)$ , and
- A<sub>2</sub>:  $\Pr(Mf | Nf) > \Pr(Mf | \neg Nf)$ ,

it follows that

$$\Pr(O | Mf) > \Pr(O | \neg Mf)$$

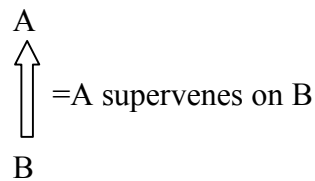

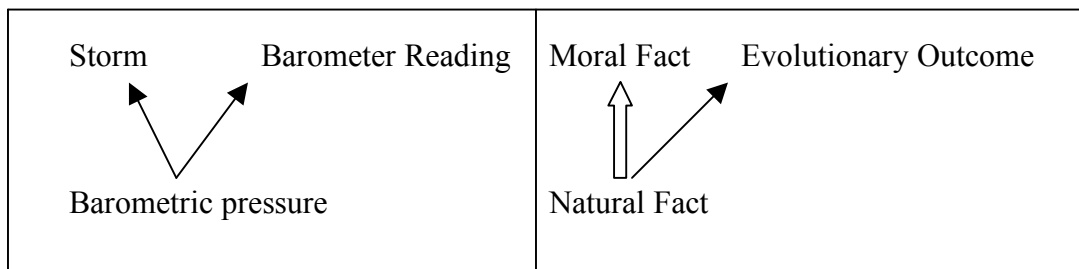
This demonstrates that it is possible for Mf and O to be correlated, in spite of the fact that Nf screens-off Mf from O. As an illustration of this possibility, consider the following example:

- O=parents provide care to their children.
- NF=parental care was favored by selection.
- Mf=it is morally right for parents to care for their children.

A<sub>1</sub> states that  $\Pr(O | Nf) > \Pr(O | \neg Nf)$ . The probability of parents providing care to their children is greater if this behavior was favored by selection than it is if this behavior was not favored in this way. This is a reasonable claim. A<sub>2</sub> states that  $\Pr(Mf | Nf) > \Pr(Mf | \neg Nf)$ . If it is right for parents to care for their children, then there is some natural fact that this moral fact supervenes on, assuming that moral facts supervene on natural

facts. Is the care that parents provide to their children a supervenience base of this moral fact? It seems reasonable to think that it is. I conclude that this example illustrates a reasonable case in which a moral fact and an evolutionary outcome are positively correlated.

The parallel between the moral case and the barometer example can be represented as follows



Barometric pressure screens-off barometer readings from storms, yet barometer readings and storms are correlated in virtue of having a common cause. Natural facts screen-off moral facts from evolutionary outcomes, yet moral facts and evolutionary outcomes may be correlated, if moral facts supervene on natural facts. X's screening-off Y from Z does not imply that Y and Z are independent, contrary to what Street claims.



Are moral facts epiphenomenal with respect to evolutionary outcomes?

Street may be read as offering an argument for the conclusion that moral facts are epiphenomenal with respect to evolutionary outcomes. The claim that moral facts are epiphenomenal in this way means that they do not cause differences in fitness. Suppose that the (dispositional) belief that one ought to help one's children is a result of natural selection. According to Street, what explains this belief's advantageousness has nothing to do with its actually being right to help one's children. On her view, the explanation for this belief is as follows: this dispositional moral belief was connected appropriately with circumstance C, such that, given C, the belief tended to become occurrent and, once occurrent, the act of believing made highly probable certain advantageous behaviors, *whether or not it is right to care for one's children*. This argument is similar to an argument that Shapiro and Sober (forthcoming) discuss in connection with the question of whether psychological properties are causally efficacious in the production of behavior. Shapiro and Sober call this argument the "master argument of epiphenomenalism:"

Given that any instance of a mental property X has a physical micro-supervenience base MSB(X), it would appear that X has no causal powers in *addition* to those that MSB(X) already possesses. The absence of these additional causal powers is then taken to show that the mental property X is causally inert (Shapiro and Sober forthcoming: 7).

To see whether X has causal powers in *addition* to those possessed by MSB(X) requires one to compare what would happen to an effect term Y in the following two cases: 1) where MSB(X) and X are present, and 2) where MSB(X) was present but (X) was not. But if X supervenes on MSB(X), 2) describes an impossible situation. The fallacy of the

“master argument”, according to Shapiro and Sober, is claim that “if X causes Y, then X must have an impact on Y that is additional to the impact on Y that MSB(X) has”.

(Shapiro and Sober forthcoming: 7) So if moral facts have no causal impact on evolutionary outcomes in addition to the impact on such outcomes delivered by the natural facts on which they supervene, it does not follow that moral facts are epiphenomenal.

The objection I’ve developed here against Street parallels an objection that Sturgeon (1988) raises against Harman’s (1986) view that moral facts aren’t needed to explain the moral beliefs we hold. According to Harman, sociological and psychological facts suffice to explain why an individual believes that some action or other is morally right or wrong. These sociological and psychological facts screen-off moral facts from moral beliefs. Harman mistakenly infers from this that moral facts are causally inert with respect to moral beliefs. Consider Harman’s example in which, observing a group of children ignite a gasoline soaked cat, one immediately judges that what they’re doing is wrong. If moral facts are causally inert in the way Harman claims, then one would have this moral reaction even if the children’s behavior was as described but it were not in fact wrong. Sturgeon argues that if the wrongness of this action supervenes on the children’s behavior, then this counter-factual claim represents an impossible situation (Sturgeon 1988: 247). If a natural fact, whether about fitness or burning cats, screens-off a moral fact from a moral belief, it does not follow that moral facts are either independent of moral beliefs or that they are causally inert with respect to them.

Another way to show that screening off does not entail causal inertness is to consider a causal chain that goes from a distal cause ( $D_c$ ) to a proximate cause ( $P_c$ ) to an

effect E. For example, suppose my dialing your telephone number causes your phone to ring, and the ringing of your phone causes you to answer it. The proximate cause screens off the distal cause from the effect; if your phone rings, the probability of your answering it is the same regardless of whether it is I who has dialed your number (you don't have caller ID). But this hardly shows that my dialing your number does not cause you to answer your phone.

Returning to Street's evolutionary arguments, I am not suggesting that moral facts *do* cause differences in fitness. But the theory of natural selection does not require that traits that are fitter be fitter *because* they cause difference in fitness. The claim that trait T is fitter than not-T does not imply that T is a *cause* of this difference in fitness. The fact that T is fitter than not-T is consistent with T's being causally inert with respect to fitness, and even with T causally detracting from fitness. To see why, we need to keep in mind that there is an important difference between selection *of* and selection *for* (Sober 1993). If there is selection *for* a trait T, then possessing T causes an advantage. But if there is selection *of* T, T may be causally inert with respect to fitness, or T may even detract from fitness. Evolution by natural selection increases both types of properties. Selection *for* applies just to properties that cause higher fitness, while selection *of* applies to properties that are associated with higher fitness, whether they cause it, or not. The following example illustrates the difference between selection *of* and *for*. Suppose that in a population of beetles, there is selection for being fast, since fast beetles evade predators better than slow ones. Suppose further that the number of spots on any beetle's wings is selectively irrelevant—there is no adaptive problem in the beetle's environment that spot number affects one way or another. Now suppose that all and only the fast beetles have

12 spots on their wings. It follows that if fast is selected, so is having 12 spots. Because of the perfect positive correlation between these two traits, their fitness values are identical. The trait of having 12 spots is fitter than the trait of having some number other than 12, despite the fact that it is causally inert with respect to fitness. Being fast is selected *for*, having 12 spots is selected *of*. So moral beliefs can evolve by natural selection even if having this or that moral belief doesn't cause a fitness difference.

### Conclusion

One horn of the dilemma that Street poses to the moral realist assumes that our moral beliefs were produced by a process that operates independently of moral facts. The implication that Street draws from this assumption does *not* depend on the claim that our moral beliefs were produced by natural selection. With respect to any moral belief forming process, if this process operates independently of moral facts, we can always ask what the epistemic credentials are of the beliefs it produces. Street's answer is that our moral beliefs are probably false. I have argued that this answer does not follow. If all we know is that our moral beliefs were produced by a process that operates independently of moral facts, then we don't know enough to say whether they're probably false or unjustified. The second horn of Street's dilemma assumes that our moral beliefs were produced by a process that is dependent on moral facts. Whether it is reasonable to suppose this depends on what this process is and on what the moral facts are. Street assumes for the sake of argument that moral facts exist, but doesn't say much about what they are. She does say that it is unreasonable to suppose that evolutionary outcomes and moral facts can ever be correlated. I have tried to show that this too does not follow.

Whether moral facts supervene on natural facts or not, moral facts and evolutionary outcomes may be positively correlated. Positive correlations can be induced because there is a common cause, or because of a common factor that is the supervenient base of one thing and the cause of another. Neither of these cases requires us to assume that moral facts *cause* fitness differences.

## References

- Boyd, R. (1988). How to be a Moral Realist. (In G. Sayre-McCord (Ed.), *Essays on Moral Realism*. Cornell University Press.)
- Brink, D. (1989). *Moral Realism and the Foundations of Ethics*. (Cambridge University Press)
- Field, H. (1990). Narrow' Aspects of Intentionality and the Information Theoretic Approach to Content. (In E. Illanueva, (Ed.) *Information, Semantics, and Epistemology*. Blackwell)
- Gibbard, A. (1990). *Wise Choices, Apt Feelings*. (Harvard University Press)
- Harman, G. (1986). Moral explanations of Natural Facts—Can Moral Claims be Tested against Moral Reality? *Southern Journal of Philosophy*, supp. 24, 57-68.
- Joyce, R. (2006). *The Evolution of Morality*. (MIT Press)
- Kitcher, P. (2005). Biology and ethics. (In D. Copp (Ed.), *The Oxford Handbook of Ethics*. Oxford University Press)
- Krebs, J. and Davies, N. (1981). *An Introduction to Behavioral Ecology*. (Sinauer)
- Railton, P. (1986). Moral Realism, *Philosophical Review*, 95, 163-207.
- Reichenbach, H. (1956). *The Direction of Time*. (University of California Press)
- Sayre-McCord, G. (Ed.) (1988). *Essays on Moral Realism*. (Cornell University Press)
- Shafer-Landau, R. (2003). *Moral Realism: A Defense*. (Oxford University Press)
- Shapiro, L. and Sober, E. (forthcoming). Epiphenomenalism—The Do's and Don'ts. (In G. Wolters and P. Machamer (Eds.), *Studies in Causality: Historical and Contemporary*. University of Pittsburgh Press)
- Singer, P. (1972). Famine, Affluence, and Morality. *Philosophy and Public Affairs*, 1, 229-243.
- Sober, E. (1993). *The Nature of Selection*. (University of Chicago Press)

- Sober, E. (1994). *From a Biological Point of View*. (Cambridge University Press)
- Sober, E. and Wilson, D.S. (1998). *Unto Others: The Evolution and Psychology of Unselfish Behavior*. (Harvard University Press)
- Street, S. (2006). A Darwinian Dilemma for Realist Theories of Value. *Philosophical Studies*, 127, 109-166.
- Sturgeon, N. (1985). Moral Explanations. (In D. Copp and D. Zimmerman (Eds.), *Morality, Reason and Truth*. Rowman and Allanheld.)
- Sturgeon, N. (1988). Harman on Moral Explanations of Natural Facts. (In G. Sayre-McCord (Ed.), *Essays on Moral Realism*. Cornell University Press.)

---

<sup>1</sup> Sober and Wilson (1998) argue that considerations of efficiency, reliability, and economy should guide our inferences about the evolution of proximate mechanisms.

<sup>2</sup> Sober agrees with this inference. On his view, if “[t]he processes that determine what moral beliefs people have are entirely independent of which moral statements (if any) are true”, then “*the moral beliefs we currently have are probably untrue*” (Sober 1994: 107, emphasis his).

<sup>3</sup> I borrow this example from Sober (1994: 94).

<sup>4</sup> Sober uses Bayes’ theorem to illustrate this point as well. In addition, he argues, as I do, that this Bayesian format shows clearly that information about belief forming processes *alone* does not have any implications for whether the proposition believed is true or false.

<sup>5</sup> I’m assuming there that the correct way to represent belief acquisition devices is in terms of likelihoods and that likelihood is a measure of sensitivity. We can think of sensitivity as a ‘world-to-head’ measure. The posterior probability represents reliability, which we can think of a ‘head-to-world’ measure. I borrow the terms ‘world-to-head’ and ‘head-to-world’ from Field (1990).