

Ethics after Darwin
Richard Joyce

Penultimate draft of paper appearing in M. Ruse (ed.), *The Cambridge Encyclopedia of Darwin and Evolutionary Thought* (Cambridge University Press, 2013): 461-467.

Through most of the 20th Century, the influence of Darwin on the philosophical field of ethics was negligible. Things changed noticeably in the last couple of decades or so of that century, and now “evolutionary ethics”—which had lain dormant since Darwin’s contemporary Herbert Spencer—is a lively and hotly debated topic. There are several Darwinian theses that might have bearing on moral philosophy.

- i. Humans are the product of natural selection.
- ii. (i) + humans have been forged by that process to be social organisms.
- iii. (ii) + among the mechanisms that govern that human sociality is an innate moral sense.

The first two are beyond serious question, but the last—moral nativism—can be reasonably doubted. It is a plausible counter-claim that the human tendency to engage in moral assessment (of oneself and others) is not a discrete psychological adaptation but a learned cultural trait that depends on psychological capacities that evolved for other purposes. Darwin himself, however, endorsed all three theses; we find him advocating (iii) in *The Descent of Man*:

I fully subscribe to the judgment of those writers who maintain that of all the differences between man and the lower animals, the moral sense or conscience is by far the most important. ... [A]ny animal whatever, endowed with well-marked social instincts, the parental and filial affections being here included, would inevitably acquire a moral sense or conscience, as soon as its intellectual powers had become as well, or nearly as well developed, as in man. (1879/2004: 120)

What bearing might these theses have on ethics? It is important to start out distinguishing two programs passing under the name “evolutionary ethics.” The first is the *empirical* enterprise of exploring the evolutionary origins of the human moral sense, drawing evidence from primatology, developmental psychology, evolutionary biology, and so on. But though often called “evolutionary ethics,” this is not a field of *ethics* in the traditional sense (any more than the investigation of the origin of the human musical sense is a kind of musical production). By contrast, *philosophical* evolutionary ethics proposes that facts about human evolution can help address certain perennial problems in moral philosophy, such as how we ought to act or whether our moral judgments are justified. When we ask what bearing the above theses might have on ethics, we are asking what impact they may have on ethics as a philosophical subject.

In 1876 (in the first issue of the new academic journal *Mind*), the great utilitarian Henry Sidgwick declared that “the theory of Evolution ... has little or no bearing upon ethics” (1876: 54). Around the turn of the 20th century two influential attacks on philosophical evolutionary ethics seemed to settle the matter in Sidgwick’s favor.

Giving the 1893 Romanes Lecture at Oxford, Darwin’s “bulldog” Thomas Huxley argued that even if the human moral sense is the product of natural selection, this affords it no particular justification. “Goodness and virtue,” he proclaimed, demand self-restraint and the helping of one’s fellows, whereas the process of natural selection demands “ruthless self-assertion” (1893/2009: 82). Moral considerations require that we *combat* the activity of the “gladiatorial

theory of existence” provided by Darwin. We shall see shortly that Huxley seriously underestimated the extent to which natural selection can produce cooperative traits (which is curious, given that he acknowledges that moral sentiments are themselves the product of evolution).

A decade later the Cambridge philosopher G. E. Moore drew attention to what he called the “naturalistic fallacy”—an error supposedly committed by any attempt to derive ethical conclusions from scientific data. Moore’s presentation of the naturalistic fallacy doesn’t place a restriction on deriving normative claims from empirical data *per se*; rather, it asserts that the quality of *goodness* is indefinable and therefore any attempt to define it in some other terms (including evolutionary terms) is doomed to failure. Moore says that same thing about the quality of *yellowness*. But why should we agree that *goodness* (or *yellowness*) is indefinable? Moore seeks to convince us with his “Open Question Argument,” which is as follows. Suppose we try to define goodness by reference to some natural property—let’s say some evolutionary property *E*. Thus when we ask of something, *x*, (1) “Is *x* good?” we are asking (2) “Does *x* have *E*?” Suppose we are inclined to answer the last question affirmatively; we can then sensibly ask a further question: (3) “Is it good that *x* has *E*?”—which, accordingly, would be the same as asking (4) “Does the fact that *x* has *E* itself have *E*?” And at this point Moore throws up his hands and pronounces that in asking (3) we clearly do not mean anything “so complicated” as (4).

Moore’s view was influential for decades, but is both widely misunderstood and dubious. A comprehensive assessment of this argument cannot be pursued here, but the standard objection should be mentioned. Goodness might be identical to some naturalistic property—including some property pertaining to human evolutionary origins (we’ll continue to call it “*E*”)—while this is unobvious to competent speakers. If an ancient Greek, ignorant of molecular chemistry, asks whether *x* is water, he should not be interpreted as asking whether *x* is H₂O—yet this observation doesn’t undermine our confidence that water is identical to H₂O. In the same way, the fact that the question “Is *x* good?” should not be interpreted as “Does *x* have *E*?” (and “Is it good that *x* has *E*?” should not be interpreted as “Does the fact that *x* has *E* itself have *E*?”) doesn’t undermine the possibility that goodness is in fact identical to the property *E*.

The naturalistic fallacy is frequently confused with the claim that one cannot validly derive an “ought”-claim from a set of premises that are purely descriptive (“You can’t get an *ought* from an *is*”)—an injunction that is also widely assumed to sink evolutionary ethics. But much the same objection applies. The evolutionary ethicist may claim that goodness (say) is identical to *E* without supposing that conclusions about what is good can be logically derived from premises that mention only *E*—any more that one can validly derive the conclusion “*x* is H₂O” from premises couched entirely in “water” terms.

The impact of Huxley’s and Moore’s arguments had a lot to do with the virtual abandonment of philosophical evolutionary ethics through the first half of the 20th century (and, in Moore’s case, the abandonment of moral naturalism more generally). There was a glimmer of interest in 1943 when Huxley’s grandson, Julian Huxley, gave a much more positive account of evolutionary ethics in *his* Romanes Lecture, but the later Huxley’s view was not very influential, due no doubt in part to a certain obscurity surrounding his reasoning and, indeed, his intended positive thesis. (Huxley’s view is effectively dissected by C. D. Broad in a critical notice of the following year.) The philosophical advances that really helped break the spell cast by Moore came in the form of mid-century progress in conceptual understanding of identity statements and analyticity. By the 1980s robust forms of moral naturalism were being offered by philosophers, the advocates of which felt entirely unhindered by Moore’s worries. These changes modified the

landscape of moral philosophy in a way that rendered it much friendlier towards the prospects of evolutionary ethics in its philosophical sense.

At the same time, advances in evolutionary biology were rendering theses (ii) and (iii) more plausible—that is, promoting *empirical* evolutionary ethics. Despite the attention that Darwin paid to the natural selection of social traits in general, and to the human moral sense in particular, the profusion of cooperation evident in nature continued to be seen as a challenge for Darwinian thinking. Natural selection, one might be tempted to assume (as did Thomas Huxley), is a process that will always favor self-serving behavior over self-sacrifice. Yet when we look around us we find a natural world teeming with examples of helpful organisms: from the bee's suicidal sting to vampire bats sharing blood. This challenge has been referred to as the "paradox of altruism."

It was not until William Hamilton's work on kin selection in the 1960s that a comprehensive solution began to crystallize. Kin selection essentially presupposes a gene's eye view on evolution, appreciating that a gene carried by organism O1 might further its reproductive chances if O1 sacrifices its interests for the advantage of organism O2, provided that O2 also carries a copy of that gene. Hamilton's theory of kin selection was complemented a few years later by Robert Trivers' work on reciprocal altruism. In this case, O1 acts in a helpful and seemingly self-sacrificing manner towards O2 because there is a high probability of O2 repaying the favor at a later date, with net gain for both parties. ("I'll scratch your back if you scratch mine.") Darwin appreciated both evolutionary forces, though only vaguely. Writing of helpful and inventive individuals in prehistoric tribes, he points out: "Even if they left no children, the tribe would still include their blood-relations" (1879/2004: 154). A few paragraphs later he writes that as reasoning powers increase, "each man would soon learn that if he aided his fellow-men, he would commonly receive aid in return"—a tendency that, he makes clear, may be inherited. Reciprocity may also be indirect, where O1 helps O2 and receives a proportionally greater benefit from O3 (and others). Darwin's frequent acknowledgements of the importance of *reputation* (our love of praise and dread of blame)—which he says clearly was "originally acquired ... through natural selection" (1879/2004: 156)—is in effect an appreciation of the importance of indirect reciprocity in the evolution of human sociality.

Kin selection and reciprocal altruism are by no means the end of the story of the evolution of cooperation—both theories have been refined and complemented by descriptions of further evolutionary processes leading to cooperation (e.g., mutualism)—but it's fair to say that by the mid 1970s it had become accepted that the abundance of cooperative behavior observed in nature poses no major difficulty when it comes to providing a Darwinian explanation of the mechanisms productive of those behaviors.

Against this background, E. O. Wilson's *Sociobiology* had a major impact when it appeared in 1975. Wilson undertakes to explain how natural selection leads to cooperation—a perfectly reputable ambition when applied to ants and zebras, but one that proved incendiary when applied to humans. The leading concern seems to have been that in providing an evolutionary explanation for human traits—including such things as aggression and sexual preferences—one somehow provides a *justification* for these behaviors. The fear was frequently expressed, though never properly explained, that sociobiological theories underwrite certain political systems.

While Wilson focused on the idea of *behaviors* as adaptations, subsequent thinkers came to focus on the psychological mechanisms underlying those behaviors. The shift in emphasis is important, for organisms with the same suite of psychological adaptations may behave very differently if placed in different environments. This change of emphasis heralded a change in

name in the 1990s, from “sociobiology” to “evolutionary psychology.” Evolutionary psychology was pioneered by psychologist Leda Cosmides and anthropologist John Tooby, whose preferred case study was the hypothesis that the human mind contains a “cheater detection module” for governing social exchanges (Barkow, Cosmides & Tooby 1992). Because of this decision to focus on such a “moralistic” human trait, the growth of evolutionary psychology encouraged work in empirical evolutionary ethics, which in turn stimulated discussion in philosophical evolutionary ethics.

Philosophical evolutionary ethics can be divided roughly into two antagonistic programs. First, it has been argued that Darwinian thinking applied to humans can serve to *vindicate* morality—either morality in general or some specific set of moral norms. Second, one might draw the opposed conclusion that moral nativism in fact *undermines* morality, providing grounds for some form of moral skepticism. These will be discussed in turn.

Suppose moral nativism is true. This shows that morally assessing aspects of one’s environment (and oneself) enhanced the reproductive fitness of our ancestors. And from this one might draw the conclusion that morality is useful, and thus justified. (See Campbell 1996.) But such an argument is invalid and the conclusion is, in any case, misleading. It is invalid because of a fallacious tense shift: From the fact that morality *was* useful it doesn’t follow that it *is* useful. If one wants to show that morality *is* practically justified then examining the ways in which moral thinking was useful in the Pleistocene may provide some insight, but it is strictly superfluous; rather, contemporary data are needed. More importantly, the conclusion that morality is practically useful is not the kind of justification in which moral philosophers are typically interested. Metaethics is concerned with whether moral judgments are *epistemically* justified, not whether they are *instrumentally* justified. When we seek epistemic justification for a belief we inquire into the grounds for holding the belief to be true. Ernie’s holding a certain belief might bring him reassurance and happiness, and thus might be instrumentally justified—but if Ernie holds this belief irrespective of any supporting evidence, then it is not epistemically justified.

One might object that moral judgments are not beliefs, and thus the question of their *epistemic* justification does not arise—that instrumental justification is the only kind that matters for morality. But then the line of reasoning from nativism to vindication would require supplementation by a preliminary argument demonstrating that moral judgments are not beliefs (i.e., an argument for noncognitivism). Metaethics has debated the merits and pitfalls of such arguments for decades. One might, however, think that Darwinian considerations can be pressed into service here, to settle the metaethical debate over whether moral judgments are beliefs. If it were shown, for example, that moral judgment emerged in our ancestral lineage because of a pay-off that relied on emotional arousal (e.g., guilt or punitive anger), then one might suppose that noncognitivism is corroborated. But the success of such an argument, while it cannot be excluded, faces serious challenges. After all, cognitivism is not the view that emotions play *no role* in moral judgment. Moral judgments may be prompted by emotion, may produce emotion, may have evolved precisely because of their emotional components, and yet for all that moral judgments may be beliefs. (Compare the hypothesis that we have an innate fear of snakes. Fear is an emotion, but it doesn’t follow that the associated judgment “This snake is dangerous!” is anything other than a belief.) So evidence that morality evolved because of its *emotional* adaptiveness cannot be taken as evidence supporting noncognitivism.

Putting noncognitivism aside, the kind of justification of morality in which metaethicists are interested in *epistemic*. Can Darwinism applied to humans help supply such justification? Several attempts have been made.

If moral nativism is true, then certain mechanisms pertaining to moral judgment have evolutionary functions. This allows one to speak of these mechanisms fulfilling their functions “well” or “poorly,” of what they are “supposed” or “ought” to do. Philip Kitcher, for example, argues that the evolutionary function of morality is to encourage social cohesion when natural altruistic sentiments fail (Kitcher 2011). If this is correct, then moral systems (and moral beliefs) can be assessed according to whether they fulfill or deviate from this function.

The Aristotelian virtue ethicist will make a similar teleological claim, but not necessarily one that pertains to the proper functioning of an innate moral sense, but rather one that pertains to the flourishing of a human understood as a complex organism. Just as biology provides understanding of what it is to be a flourishing frog, as opposed to a diseased or unhealthy frog, so too it can in principle provide the same with respect to humans. Given that humans are social creatures (i.e., given above thesis (ii)), the virtues, it is claimed, are those character traits that are conducive to, or constitutive of, human flourishing. Thus the virtue ethicist takes a Darwinian premise about what kind of evolved creatures we are, and strives to produce a normative output—one that favors such things as friendliness, benevolence, and so forth—and thus hopes ultimately to provide epistemic justification for such claims as “One ought to be friendly,” etc. (Casebeer 2003).

The principal problem with such attempts to vindicate morality using evolutionary data is that whatever normative language legitimately follows appears to be the wrong sort to underwrite the kind of practical guidance we require of *morality*. Consider: The function of a hammer is to bang in nails, but if I find it convenient to use a hammer to prop open the garage window, there is nothing fishy about my action; it is not even an instance of the new function I have assigned the hammer *overcoming* requirements imposed by the hammer’s “real” function. The hammer’s real function may license assertions like “A good hammer bangs in nails well” and “This hammer is supposed to bang in nails,” but it turns out that this normative language is really quite toothless when it comes to making claims upon our practical deliberations independent of our standing interests. In the same way, moral systems may have the evolutionary function of promoting cooperation when altruism falters, and this may thus be what moral systems are “supposed” to do, but if a society chooses to use its moral system for some other end (in support of militaristic imperialism, say), then the “real” function of moral systems carries no weight *per se* to cast doubt on that decision. (That’s not to say that there’s nothing wrong with such a decision, but that we need to look somewhere other than biological functions in order to locate grounds for criticism.)

The same point applies when we consider someone electing to cultivate personality traits other than those virtues conducive to flourishing. Such a person needn’t have given up on the aim of flourishing, but rather has chosen a vision of flourishing other than that laid down by biology. Perhaps this person has embraced the kind of flourishing that goes along with being a Buddhist monk (which presumably diverges spectacularly from what it took to be a fine human specimen in the Pleistocene), and thus cultivates the kind of character traits necessary for this end. That we could legitimately say of this person that her vision of flourishing is not the one she is “supposed” to be pursuing (qua human organism) might sound impressive, but it is not obviously any more of a genuine criticism of her behavior or character than is the observation that in propping open the window with a hammer one is using the object in a manner for which it

is not intended. The problematic consequences of this failure to derive genuine normative criticism from evolutionary function becomes apparent when we note that it applies as much to the person who chooses a life of violent crime as it does to the Buddhist.

Another possible route from nativism to epistemic vindication is via epistemological reliabilism. True beliefs are far more likely to enhance reproductive fitness than false beliefs; therefore on those occasions that natural selection produces some discrete belief-forming mechanism it is likely that the resulting beliefs will be true (Carruthers 1992: 111 ff). Thus beliefs that are fixed or pre-wired by natural selection can be considered the product of a *reliable* process; and hence are, according to the theory of process reliabilism, epistemically justified.

The prospects of any such attempt to vindicate morality are only as good as the prospects of the theory of process reliabilism upon which it depends—and such theories are controversial. One of the problems of reliabilism is that it is difficult to specify precisely *which* process any given belief is the product of, for invariably it is simultaneously the product of numerous. If moral beliefs are the output of some kind of “moral sense,” then it’s natural to assume that when we try to identify “the process” that produced them, we should not look to natural selection in the general sense, but rather the particular evolutionary trajectory of the innate faculty in question. It might be correct that *in general* we are better off with true beliefs than false, but it need not be correct that *when it comes to moral beliefs* we are better off with truth than falsity. A false belief about the value of benevolence may be adaptive in a way that a false belief about the behavior of predators is not. Indeed, this observation segues naturally into discussion of the second program of philosophical evolutionary ethics: that moral nativism *undermines* morality.

Contemporaries of Darwin already felt uneasy about the possible undermining influence his views might have on moral authority. One called his position “dangerous” and expressed concern that moral nativism “aims ... a deadly blow at ethics” (Cobbe 1872: 10). Another wrote that if Darwin’s views on moral nativism were true, “or should they come to be generally accepted, the consequences would be disastrous indeed! We should be logically compelled to acquiesce in the vociferations of [those] who would banish altogether the senseless words ‘duty’ and ‘merit’” (Mivart 1871/2008: 204). The general worry is that if humans assess the world in moral terms only because doing so helped our ancestors produce more babies than their competitors, then these judgments appear not to carry the binding authority over our actions that we usually think they do. Moreover, if it is true not only that evolutionary origins deprive moral judgments of their authority, but also that invoking such authority is the whole point of having a moral system—indeed, that such authority is a necessary feature of our basic moral concepts—then it appears that moral nativism reveals our moral concepts to be bankrupt: they imply an authority that they cannot supply.

Darwin himself gives no hint of having ever been tempted by such skeptical thoughts. He held that moral thinking is both practically necessary for human society and one of the most striking of human adaptations. He considered that any social creature granted sufficient intelligence would evolve a moral sense—though he also conceded that the content of that morality may differ dramatically among species. But one looks in vain for any satisfying metaethical statement from Darwin (understandably enough); rather, one is forced to infer from his seemingly untroubled attitude towards morality that he was unaware of, or had little patience for, the possibility that moral nativism might debunk morality.

In recent years, the debunking argument of evolutionary ethics has been explored by several philosophers. According to Michael Ruse (1986), in order for moral judgments to serve their evolutionary function (roughly, encouraging cooperation) they must be imbued with objectivity.

This is a thesis about the *content* of moral judgments; it doesn't follow that the actions in question (or any other actions) *are* objectively required. In fact, Ruse thinks, moral nativism provides grounds for doubting that any actions are objectively morally required, for supposing that they are so is entirely unnecessary. Nativism may explain why humans make judgments about moral objectivity, but to go further—to suppose that the judgments are *true*; that is, that some actions *are* objectively required—involves populating our conception of the world with properties that play no explanatory role. (It is not merely that they aren't needed to explain our moral judgments, but they aren't needed to explain *anything*—for what explanatory role could they play independently of anyone making a moral judgment?) Humans have been set up by natural selection to believe in objective moral properties (Ruse thinks)—and have been set up to do so irrespective of whether there are any such properties—so there are no grounds for believing in them at all.

Ruse here wields Ockham's Razor to cut objective moral properties from our conception of the world. In fact, we should distinguish two razors (following Sober 2009). Suppose our evidence fails to discriminate between "X exists" and "X does not exist." The Razor of Denial states that we should deny the former and affirm the latter; the Razor of Silence states that we should suspend judgment about both. The debunking argument of evolutionary ethics is more plausible when construed in the latter manner. It is unlikely that moral nativism can show that our moral judgments are all *false*; but that it might show them to be all *unjustified* is an argument with more promise. This argument is pressed by Richard Joyce (2006).

It is important to note that this argument has promise only if moral nativism is understood in a certain way—namely, that the ancestral adaptiveness of moral judgment was secured independently of any truth-tracking relation between these judgments and moral facts. A comparative illustration may help. Suppose that humans are pre-wired by natural selection to divide their social environment into in-groups and out-groups. The supposition implies that such thinking was reproductively useful. But *why* was it useful? The only plausible answers presuppose that our ancestors' environment did actually contain in-groups and out-groups. It is important to see that the hypothesis of moral nativism may be crucially different in this respect. The most plausible accounts of why it was reproductively useful to our ancestors to categorize aspects of their social world as *good*, *bad*, *evil*, *obligatory*, and so on (e.g., that such categorization strengthened social cohesion) *nowhere presuppose* that the environment contained such things as goodness, badness, evil, and obligatoriness.

However, even if the nativist hypothesis nowhere explicitly mentions any actual moral properties, it remains possible that such properties are identical to, or supervene upon, those properties that are explicitly mentioned. Ruse's use of Ockham's Razor, for example, seems to assume that to allow the existence of objective moral properties would admit an extra ontological layer into the world (and thus should be disallowed if unnecessary). But this isn't obviously so. An opponent can counter that objective moral properties were implicitly present all along in the evolutionary worldview accepted by Ruse, just as H₂O was implicitly present in ancient Greek explanations involving water. This debate then moves to the question of whether it is plausible to claim that objective moral properties may be identical to, or supervene upon, those naturalistic properties recognized by science. It is noteworthy that Joyce's evolutionary debunking argument has to be supplemented with an attempt to undermine the prospects of moral naturalism on purely metaethical grounds (2006: chapter 6)—an undertaking to which Darwinian thinking has no obvious contribution to make.

Sharon Street (2006) comes to a similar conclusion to Joyce, though starting out with a different understanding of moral nativism: Whereas Joyce is willing to speculate that natural selection left the *content* of the moral faculty pretty much open, Street supposes that the content of morality has been “deeply influenced” by Darwinian forces. She then poses the moral realist with a dilemma focused on the relationship between these evolved evaluative tendencies and objective moral values. Either (A) there is no relation at all—in which case the chances that natural selection has guided us to approximately correct evaluative judgments are vanishingly small; or (B) there *is* a positive relation: Our evolved moral faculty “tracks” real moral properties. The problem with the latter is that it is an empirical doubtful claim; as noted earlier, the most plausible accounts of the evolution of the moral faculty see its adaptiveness in terms of enhancing social bonds, not in tracking truths.

Street’s argument targets moral *realism*, understood as the thesis that moral truths hold independently of our attitudes. But she is not targeting moral truths *per se*; she leaves open the door to moral facts that are in some sense constructed by us. Similarly, Ruse’s conclusion is that we have no reason to believe in *objective* moral properties—seemingly allowing the possibility of *non-objective* moral properties. “[T]he illusion lies not in morality itself, but in its sense of objectivity” (Ruse 1986: 253). By contrast, Joyce’s skeptical attack is leveled at moral facts *tout court*—subjective as much as objective.

Evolutionary ethics—both the empirical and the philosophical programs—barely existed for the best part of the century following Darwin’s death. In the last few decades it has mushroomed into a rich interdisciplinary field concerned with both the explanation and justification of a fundamental aspect of the human organism.

REFERENCES:

- Barkow, J. H., L. Cosmides, & J. Tooby (eds.), 1992. *The Adapted Mind*. Oxford: Oxford University Press.
- Broad, C. D. 1944. “Critical notice of Julian Huxley’s *Evolutionary Ethics*.” *Mind* 53: 344-367.
- Campbell, R. 1996. “Can biology make ethics objective?” *Biology and Philosophy* 11: 21-31.
- Carruthers, P. 1992. *Human Knowledge and Human Nature*. Oxford: Oxford University Press.
- Casebeer, W. D. 2003. *Natural Ethical Facts: Evolution, Connectionism, and Moral Cognition*. Cambridge, MA: MIT Press.
- Cobbe, F. P. 1872. *Darwinism in Morals and Other Essays*. London: Williams and Norgate.
- Darwin, C. (1879) 2004. *The Descent of Man, and Selection in Relation to Sex*. London: Penguin Books.
- Hamilton, W. 1964. “The genetical evolution of social behavior” I and II. *Journal of Theoretical Biology* 7: 1-52.
- Huxley, J. S. 1943. *Evolutionary Ethics*. Oxford: Oxford University Press.
- Huxley, T. H. (1894) 2009. *Evolution and Ethics*. Princeton: Princeton University Press.
- Joyce, R. 2006. *The Evolution of Morality*. Cambridge, MA.: MIT Press.
- Kitcher, P. 2011. *The Ethical Project*. Cambridge, MA.: Harvard University Press.
- Mivart, G. J. (1871) 2008. *On the Genesis of Species*. BiblioLife.
- Moore, G. E. (1903) 1948. *Principia Ethica*. Cambridge: Cambridge University Press.
- Ruse, M. 1986. *Taking Darwin Seriously*. Oxford: Basil Blackwell.
- Sidgwick, H. 1876. “The theory of evolution in its application to practice.” *Mind* 1: 52-67.

- Sober, E. 2009. "Parsimony arguments in science and philosophy: A test case for naturalism_p." *Proceedings and Addresses of the American Philosophical Association* 83: 117-155.
- Street, S. 2006. "A Darwinian dilemma for realist theories of value." *Philosophical Studies* 127: 109-166.
- Trivers, R.L. 1971. "The evolution of reciprocal altruism." *Quarterly Review of Biology* 46: 35-57.
- Wilson, E.O. 1975. *Sociobiology: The New Synthesis*. Cambridge, MA.: Harvard University Press.