

10 Evolutionary Ethics: A Theory of Moral Realism

Robert J. Richards

Introduction: Darwin's Theory of the Evolution of Morality

In *The Descent of Man and Selection in Relation to Sex* (1871, 1: 72-73), Charles Darwin produced a theory of moral conscience and an explanation for ethical behavior. His theory had several elements. To develop a conscience that recognized morally right and wrong action, an organism had to have (1) a repertoire of social instincts – those urging parental care, group cooperation, altruistic response; (2) sufficient intelligence and memory to make practical judgments of a complex sort; (3) language to codify behavior and communicate desires, requests, and other information to conspecifics; and (4) habits to help shape behavior. Darwin thought the crucial feature of his account to be the social instincts, those innate expressions of cooperative and altruistic response. But he knew there was a problem in providing a natural-selection explanation for behavior that seemed to benefit the recipient rather than the agent. He believed he had the answer:

It must not be forgotten that although a high standard of morality gives but a slight or no advantage to each individual man and his children over the other men of the same tribe, yet an enhancement in the standard of morality and increase in the number of well-endowed men will certainly give an immense advantage to one tribe over another. There can be no doubt that a tribe including many members who from possessing in a high degree the spirit of patriotism, fidelity, obedience, courage, and sympathy, were always ready to give aid to each other and to sacrifice themselves for the common good, would be victorious over most other tribes, and this would be natural selection. (Darwin, 1871, 1: 166)

There are many details of Darwin's proposal that would require extended discussion: for instance, though he based his conception of human community selection on altruistic response in the social insects – in which natural selection operated not on the individual but on the entire hive or nest – one might ask whether the social instincts in humans were produced by group selection or individual selection. Michael Ruse and I disagree about Darwin's answer to this particular question – profoundly. Ruse argues that selection in ancient, clan-groups was roughly equivalent to kin selection as we now

conceive it, and thus a kind of individual selection, while I regard Darwin as having generalized his notion so that it applied to human tribes in which there was no expectation that all individuals would be related, thus a matter of group selection (Richards and Ruse, 2016). Ruse believes that as proto-human communities expanded, individuals began to aid one another under the expectation of reciprocity. He assumes that this original attitude would become cemented into our genes by a process of individual selection, since such an attitude would be individually beneficial. Ruse forgets that Darwin called reciprocal altruism a “low motive” (Darwin, 1871, 1: 163), being at root selfish. Darwin claimed his theory of community selection removed “the reproach of laying the foundation of the most noble part of our nature in the base principle of selfishness” (Darwin 1871, 1: 98). Strong evidence, I believe, that he advanced a theory of group selection. Moreover, if it were only a matter of individual selection, the most advantageous trait would be one in which the bearer feigned cooperation while defecting. But in the group selection scenario, feigning cooperation while defecting would be deleterious for the whole group and the group would be selected against. Thus in the long run, the advantage goes to groups in which most of its members are sincere altruists. There is further evidence that Darwin was a group selectionist; but for the purposes of this chapter, the question of the exact character of Darwin’s conception – whether supporting group or individual selection – is not immediately crucial. Both Ruse and I agree that Darwin was fundamentally correct when he held that human beings have evolved to act altruistically toward one another, even if sometimes failing to do so.

Despite our differences about the target of community selection, Ruse and I both endorse Darwin’s assumption of an innate attitude toward cooperation and toward helping those in need. The empirical evidence for such an attitude is fairly strong, but we are both interested more in the philosophical consequences of this assumption. So in this chapter, I will simply assume that something like the Darwinian theory of morality is our best account of impulses and judgments we have come to call moral. That is, I will assume we have evolved to be altruists, that altruism evokes a special attitude of approval, and that analysis would show that fundamental moral principles – e.g., the golden rule, the Categorical Imperative, etc. – amount to altruism.

A further feature of Darwin’s account has Ruse’s and my allegiance: that the altruistic instinct needs the guidance of rational considerations. We may feel obliged to act for the community good, but only a rational analysis based on experience and acquired knowledge can determine for us who are members of our community and what really does serve the community good. The Hippocratic physicians who administered to victims of the Athenian plague treated their patients by bleeding them and giving emetics, which could only hasten their deaths. However, the physicians did so under a mistaken medical

belief about the efficacy of those treatments; they served their patients at considerable risk of becoming infected themselves and many did become infected. The development of medical science over the centuries has supplied effective treatments for plague victims. Contemporary doctors who harbor comparable altruistic instincts now have a rational strategy for successfully requiting those instincts. We would today, nonetheless, look back and judge those ancient physicians to be moral heroes, since their motives and intentions were altruistic, even if their beliefs were invincibly defective.

Most critics who object to an evolutionary ethics would, I think, be willing to grant these assumptions, since the potent philosophical objections do not lie along the lines of empirical adequacy. The deeper question I would like to explore in this chapter is the validity of the most powerful objection raised by philosophers to the general theory of human moral evolution, namely, that it fails to be an account of authentic moral behavior. The objection has been developed in two distinct but related ways: that an evolutionary theory of morality commits the naturalistic fallacy and that it fails to be objective, that is universally binding, and by default must embrace only a subjective relativism. These specifications of the general objection are connected.

Objections to an Evolutionary Theory of Morality: The Naturalistic Fallacy

G. E. Moore claimed, in his classic formulation of the naturalistic fallacy (1902), that it was a mistake to identify the moral good with any natural property. He particularly had in mind Herbert Spencer's identification of "more evolved" with "ethically better" (Moore, 1903, 48-54). Such identification would be liable, Moore argued, to the open-question objection: we could always ask whether this altruistic act, or this more evolved behavior, was, after all, really good, and the answer, yes or no, should enlighten us. If moral goodness were simply identified with the natural trait of altruistic impulse, the question would amount to, is this altruistic act really altruistic? The answer would seem hardly to enlighten us. Moore believed we could intuitively perceive the moral good as a nonnatural property accruing to different actions in particular contexts; it was a real property, the perception of which provided reason to approve of actions to which it accrued or supplied reason to pursue such actions. Moore's objection to Spencer's moral theory wasn't that it failed to be a realistic theory, but rather that it was not the right kind of realism.

Thomas Henry Huxley introduced the objection of the naturalistic fallacy even before Moore formally named it. Huxley's lecture "Evolution and Ethics" (1893) was also directed at the ethical theory of Spencer, his erstwhile friend, but it told equally against Darwin's theory of moral evolution. Huxley recognized a disjunction between the biologically fit and the morally fit: "Social

progress means a checking of the cosmic process at every step and the substitution for it of another which may be called the ethical process; the end of which is not the survival of those who may happen to be the fittest, in respect of the whole of the conditions which obtain, but of those who are ethically the best" (Huxley, 1893, 9: 81). Huxley conceded that our so-called moral sentiments had evolved no less than our aggressive sentiments; but this outcome, simply considered, furnished no moral rules, no moral goals. "Cosmic evolution," he cautioned, "may teach us how the good and the evil tendencies of man may have come about; but, in itself, it is incompetent to furnish any better reason why what we call good is preferable to what we call evil than we had before" (Huxley, 1893, 9: 80). In nuce, Huxley deployed the open question: Why ought I act on this altruistic attitude rather than on that vengeful attitude? Huxley's analysis reduces to the more common way of stating the naturalistic fallacy, namely, that one cannot justify an "ought proposition, a norm, by appealing to a fact – no "ought" can be derived from an "is."

The naturalistic fallacy has lost its sting. Several considerations mitigate, really abolish, its force (Richards, 1986, 1989, 1993, 2016). Take the simple assertion that facts cannot justify norms. In science, favorable experimental outcomes, thus empirical facts, are used in the justification of natural laws, and laws carry normative force – e.g., this copper, if dunked into sulfuric acid, ought to produce cupric sulfate. And counterfactually, if all copper and sulfuric acid suddenly disappeared from the world, it would still be true that if sulfuric acid were added to copper it should produce cupric sulfate. So if facts cannot justify norms, science as we have come to understand it, must collapse. Justification is not always or even usually based on a deductive relationship. Of course, we feel no moral outrage if a well-justified natural law fails on further testing. The normative implications for natural laws are indeed different from moral principles that mandate altruism.

What work does the concept "ought" perform? In itself it has minimal semantic meaning; it connects a particular causal matrix or property with an expectation. It achieves more definite semantic content from the particular causal matrix or property that serves as the basis for the expectation. So there is no moral ought or instrumental ought devoid of a particular basis. Consider the kind of "ought" the moral imperative is in practice. It cannot be a categorical "ought," as Kant believed, that is, an "ought" without a basis. After all, we don't admonish a two-year-old with a categorical "ought-not" about swiping cookies, but recognize that the individual must be in a certain condition (i.e., human, of sufficient age and mental capacity) to be morally admonished. So the structure of the real, moral imperative is, if you are a human being, with the full capacities of such – the factual basis for the moral judgment – you ought not lie, cheat, or steal. In light of the evolutionary assumption I've made, we would say, since you are a human being who has evolved to have

altruism as an essential component of your natural fabric, then you ought to act altruistically. The ought carries moral weight, as opposed to, say, only instrumental weight, because of the evolved capacity for altruism, which we have identified (as part of the granting of empirical fact) as the moral inclination. This evolved capacity, I have premised, comes with a very strong attitude of approval for altruistic behavior and a comparable negative attitude when a violation of such behavior has occurred. Because the evolutionary process has deeply embedded the altruistic capacity in the human group, we are warranted in the expectation – a kind of prediction – that in the requisite circumstances an individual ought to act altruistically. Such an expectation, of course, is modified by other circumstances. When, for example, a known thief is about to strike again, we recognize that this criminal is likely to ignore his better impulses. If those impulses are altogether absent – say, in the case of a profound psychopath – would we think such a person, if we were aware of his character, *ought* not steal? We would certainly want to be protected against the psychopath in our midst; but if such a person were simply, say, born that way – born with a kind of insanity or profound mental deficiency – could we hold such an individual morally responsible? We may want to hold the psychopath legally responsible and be protected against the dangers he or she might pose, but on careful consideration I do not believe we could hold him or her morally responsible; for we cannot morally demand something of one who is quite incapable of the action, or of one who sincerely does not recognize a moral requirement.

At this juncture, let me examine two kinds of objection to the sort of analysis I've offered of the naturalistic fallacy, that of Moore and that of Richard Joyce. Moore urged an open-question rebuttal to the likes of Spencer's evolutionary ethics. To consider the force of the objection, let us presume not only that we have evolved to approve altruistic acts, that is, recognize them and regard them as special, but also that we carry the attitude as a reflective criterion. So if we reflectively ask ourselves, but is altruism really a special good, then the criterion with which evolution has endowed us would certify our immediate judgment. One could keep asking the reflective question, but if we have evolved in the way supposed, the answer would be boringly the same. This case would be similar to a Kantian immediately perceiving various examples of causal relationships, but then, in a reflective mode asking himself or herself: does every event really have a cause? The answer must be "of course," since every such reflective question evokes the same category of judgment – there is no other way to think about the world of our experience. Would the reflective person learn anything? Initially, very likely; he or she would learn that the first impulse conforms to the criterion that reveals itself as inescapable. This consideration holds for the moral reflection: if we have evolved in the way specified, there is no other criterion to apply, certainly

not one that would make sense of our general moral experience. The moral criterion, if generally imposed by evolution, is inescapable. So in this scenario, Moore's open question swings closed.

I have distinguished two factors in a moral judgment – the recognition of community needs and a natural sentiment, a sentiment that has evolved over generations and that approves acting in respect of those needs. Richard Joyce argues that pro-social feelings are irrelevant to moral judgment: the difference between doing something “because you want to do it is very different from doing something because you ought to do it” (Joyce, 2006a, 50). To use Joyce's example, I might feel love for someone – that is, I don't want to harm the person – but lack any sense of obligation not to harm her (Joyce, 2006a, 51). What then does the “ought” of obligation mean? Joyce suggests it means “inescapability and authority (Joyce, 2006a, 62).” I have just given an account of the inescapability of the evolutionary criterion. But whence the authority? The authority of a moral judgment gives it “practical clout,” a reason for following a moral precept, such as acting altruistically. Joyce doesn't think we can reduce the authority to a sentiment or feeling – a want – since we presume moral values “bind people irrespective of their desires or interests” (Joyce, 2006a, 192). Herein lies his debunking or skeptical attitude: there is no adequate naturalistic account of the practical clout of a moral judgment. Maybe, however, Joyce hasn't tried hard enough.

Moral judgments can be usefully distinguished into two types: our immediate response to someone in distress, which elicits a desire or need to help – the archetypal example of impulsively jumping into a river to save a drowning child – and a more reflective and calmer judgment, when, for instance, I take out my credit card to respond to an Internet plea from Doctors Without Borders. In the first instance, no reasons are given, but the urge to action is immediate, impulsive, and undeniably pushed by feeling; nonetheless, we usually think of someone who has so acted, without calculation of the dangers, as deeply ethical and morally courageous. Such a person displays what an Aristotelian would call the *habitus* for ethical behavior. The altruistic desire, honed over millions of years of evolutionary selection, is of a different caliber than the desire for pistachio nut ice cream. In the second instance – that of taking out my credit card – I have time to reflect on my reasons for so acting; but I've evolved, so our assumption allows, to approve of altruism and weigh it against the set of other desires. Its authority is simply that of the special attitude that makes it stand out against other momentary needs and fluctuating impulses. Another desire may be momentarily stronger – e.g., to use my extra money to buy expensive tickets to a play rather than send it off to doctors who are short of bandages. The inescapability of the altruistic attitude supplies the urgent desire that has become biologically entrenched. Under this scenario, the altruistic attitude forms part of what it means to be

human; the subjective desire for pistachio nut ice cream lacks the authority of millennia of selection, and it has no special claim to being part of our humanity. I'll elaborate on this notion of human identity in a moment.

Objectivity of Evolutionary Ethics

This rejoinder to the problem of the naturalistic fallacy leads to the other specification of the general problem of the authenticity of moral response under the Darwinian theory: are moral values given to us in experience the same way objects and events are so given – that is, objectively – or are we coerced by evolution into accepting a faux morality, one that is essentially a subjective preference? Some who advance the so-called debunking argument believe we are subjectively constrained, not objectively compelled. Michael Ruse and E. O. Wilson, as well as Sharon Straight, offer nice examples of this kind of argument. Joyce and Horn furnish other examples in this volume.

Ruse and Wilson provide a succinct set of considerations which imply that the Darwinian moral judgment fails to be objective, and thus not universally binding. Human beings, the authors maintain, “are deceived by their genes into thinking that there is a disinterested, objective morality binding upon them, which all should obey” (Ruse and Wilson, 1986). The instinct to reciprocate cooperation, which Ruse and Wilson regard as the result of individual selection, is represented by epistemic rules that have become innate. These innate principles dispose “us to think that certain courses of action are right and certain courses of action are wrong.” They give us “the illusion of objective morality.” They are, however, entirely subjective, relative to the contingencies of environmental circumstances during the course of our peculiar, human evolutionary history. Ruse and Wilson contend that these rules are simply the “idiosyncratic products of the genetic history of the species and as such were shaped by particular regimes of natural selection.” To make this assertion vivid, they contrast our human genetic history with that of “an alien intelligent species evolving rules its members consider highly moral but which are repugnant to human beings, such as cannibalism, incest, the love of darkness and decay, parricide, and the mutual eating of feces” (Ruse and Wilson, 1986, 186). This alien species could have been us. So what we take as the universally binding and objective rules of morality could have been otherwise. We prefer feeding the hungry and aiding the poor; they prefer dining on feces and throwing their fathers under a bus. There but for the grace of our peculiar selection history go we.

Ruse and Wilson believe our moral attitudes are the result of the evolutionary process. What they deny is that, precisely because of this process, our ethical rules could meet an ideal of what such rules ought to be, namely objective and universally binding. (One should note that Ruse and Wilson are assuming

an objective standard of morality to deny that evolution could provide one. But whence this objective standard?) Thus for the evolutionary theorist, there is no reason to regard the moral rules we adhere to as any different than personal preferences, preferences like that for pistachio nut ice cream. Before responding to this assault on our ordinary moral attitudes, let me make some distinctions relative to the terms “subjective” and “objective.” I believe Ruse and Wilson glide over important discriminations.

The pair of terms “subjective” and “objective” could be given an ontological meaning or an epistemic meaning. A proposition, an idea, or an attitude is held, of course, in the mind of an individual, a subject; these mental events are thus ontologically subjective. An event that occurs extra-mentally, in the natural world, would be objective in the ontological sense – a solar eclipse, for instance. Epistemically, however, a subjective belief would be one that is the result of personal idiosyncrasy or prejudice, while an epistemically objective belief – the kind characteristic of scientific propositions – would be one intersubjectively verifiable, capable of confirmation by others acting without bias. Ruse and Wilson slide around between the ontological meaning of these terms and the epistemic.

Ethical propositions are certainly ontologically subjective. After all, they are held by individual subjects – they are held in the minds of individuals. They needn’t be epistemically subjective, however. Consider logical propositions (e.g., the principle of noncontradiction, the principle of modus ponens, etc.) or mathematical propositions (e.g., the Pythagorean theorem). Except for Platonists, such propositions are usually taken to be ontologically subjective, but are also considered epistemically objective – that is, they follow from accepted standards that are ultimately intersubjectively verifiable. Ruse and Wilson reject the idea of “morality as a set of objective, eternal verities,” “extrasomatic guides,” which for their money would constitute an authentic morality (Ruse and Wilson, 1986, 186). But they seem to be looking for a set of propositions chiseled in stone and lying on a mountain top – ontologically objective entities. Certainly, moral codes exist in the minds of individuals; they are ontologically subjective. But this does not mean they must also be epistemically subjective. The validity of a given syllogism or even a natural law is objective if justified by accepted standards, which standards are ultimately intersubjectively verifiable. A given argument or proposition can be epistemically objective without being ontologically objective.

Can ethical arguments be epistemologically objective? Consider the history of our species. Most evolutionist will maintain that our rational ability is the result of the evolutionary process. Our ability to assess arguments – say, for the validity of evolutionary theory itself – is not impugned by its origin as an evolved capacity. Those of our hominid predecessors who regarded the saber-toothed cat in their paths as both gentle and dangerous have left a very short

line of descendants, while those who evolved cognitively to avoid contradictions have prospered. We have come to hold the principle of noncontradiction because of natural selection. Yet, the application of the principle to establish mathematical or logical systems can hardly be rejected as merely a subjective preference. Mathematical and logical propositions can be perfectly objective without being derived from “extrasomatic guides.” The claim that extrasomatic guides are necessary for objectivity is to confuse epistemic objectivity with ontological objectivity. Moreover – and this I’ll explore further – the fact that our cognitive faculties have survival value does not at all imply that they are unfit for evaluating truths about the world. Rather they have survival value because, quite generally, they are reliable guides in determining truths about the world.

Ruse and Wilson suggest that an ideal of objective morality cannot be realized, since we could have evolved otherwise than we have. They seem to believe that there could be an alien species, in all essentials like us, but one with a moral code not based in cooperation and altruism. But if that species were intelligent and social, thus like us in that respect, then it would have to be moral like us as well. For what would a social group be like if it were not bound together by ties of cooperation and altruism? Ruse and Wilson’s scenario of social aliens that were non-altruistic and noncooperative would be comparable to a story about an intelligent species that did not observe the principle of noncontradiction. What kind of hippogriff-like species could that be? Such an alien species could not be us. Nor could it even evolve as a social species. To put it crudely: early in our evolutionary line a sharp right turn could have been taken, producing a species with very small brains and non-social attitudes. But then, we would not be talking about human beings and their moral judgments. Ruse and Wilson’s debunking argument simply fails.

Debunkers sometimes cite an imaginary case proposed by Darwin himself that seems to endorse the utter subjectivity of moral judgments. In the *Descent of Man*, Darwin supposed that if humans had evolved much like hive bees, then unmarried females might regard it “a sacred duty to kill their brothers, and mothers would strive to kill their fertile daughters” (Darwin 1871, 1: 73). This imaginative scenario seems to imply these bee-men would have evolved a moral sense quite different from ours. Thus our human moral sense would appear to be merely a contingent fact of our particular history, quite relative to the ancient environment of our evolutionary descent. These bee-men would exactly mirror Ruse and Wilson’s alien species. Two responses seem apposite. First, I believe Darwin’s scenario would be just as impossible as that of Ruse and Wilson. It’s hard to imagine that if you knew – had rationally determined – your sister would be gunning for you, and your mother was going to stab your sisters, that you could be a member of a stable social community. Bees have a workable cooperative community, since they also have a very dim

rational capacity. But second, we need to recognize two features of moral judgment: the moral (altruistic) motive and the rational consideration of how to enact that motive – both are essential for a practical moral judgment, and Darwin thought both necessary for the development of conscience. We might properly judge the Aztec priest, who sought to sacrifice young maidens to make the corn grow and thus save the community – and who did so with the community's approval – perfectly moral but with defective agricultural knowledge. In Darwin's case of the bee-men, if we abstract from the impossibilities of his scenario, we might say that they have the same moral motive as we do – desire to act for the community welfare – but that they have unwarranted beliefs about what would advance the community welfare. They would not be an alien species, but one with us, yet one that needed enlightenment and needed rational knowledge, much like our own ancestors.

Sharon Street (2006) contrives another strategy for the debunking argument, though ultimately bases her approach on ideas similar to those of Ruse and Wilson. She begins with the assumption that realist theories of value must posit truths existing independently of our evaluative judgments. By contrast, the adaptive non-realist theory – which she believes to be the most reasonable theory in light of evolution – simply regards our value judgments as adaptations whose principal function is that of reproductive success. So the issue of realism versus adaptational relativism reduces to the following: do we hold certain values because they are antecedently and independently true, or do we hold them as true because of our contingent evolutionary history and then merely assume them to have been antecedently and independently true? Street argues the latter is the case: they are adaptations that could have been otherwise, and so lack objectivity. The basic structure of Street's argument, then, is comparable to that of Ruse and Wilson.

Street argues that those who adopt the realistic stance concerning moral values face a dilemma. Either they must deny there is any relation of moral values to natural-selective forces or they must posit a relationship. If there is no relation, then any impact of contingent evolutionary forces on our value system, she believes, could only be accidentally enforcing, but more likely distorting. On this horn of the dilemma, any adaptive cognitive response would likely mislead us in regard to any supposed real moral values. Taking the other horn of the dilemma, if there is a relationship, then adherents of this view must regard our evolved capacities as designed to recognize moral truths and to act on them. So on this horn of the dilemma, the adherents maintain that we have evolved rational faculties of evaluation, which allow us to recognize certain evaluative truths, like the value of altruism, and on the basis of recognizing these facts or truths, we judge certain acts as morally appropriate and others as moral transgressions. But Street thinks this latter hypothesis of an evolved evaluative faculty must fail in light of a stronger one, which she calls

“the adaptive link account.” The adaptive link account simply explains our adoption of certain moral truths because “they forged adaptive links between our ancestor’s circumstances and their responses to those circumstances, getting them to act, feel, and believe in ways that turned out to be reproductively advantageous” (Street, 2006, 127). She believes the adaptive link account to be superior because it’s more parsimonious: it cuts directly to the ultimate cause, namely, reproductive success, without positing the superfluous notion of an evolved general capacity to recognize evaluative truths. At this juncture, Street’s argument joins that of Ruse and Wilson, as well as that of Joyce and Horn, and all slide into the slough of implausibility.

Street maintains that a general capacity to recognize factual truths would actually be of “no advantage or even a disadvantage” (Street, 2006, 130). She asks: is there any survival advantage to recognizing “truths about the presence or absence of electromagnetic wave-lengths of the lowest frequencies” or, presumably, most of the truths of physics, cosmology, and higher mathematics? Of course not. There might be a selective advantage in directly recognizing the dangers of fire, for example, but she perceives no advantage in a general rational capacity for evaluation. So if a general evaluational capacity carries no reproductive advantage, natural selection could not account for it and thus it could not evolve as part of our repertoire of traits. It’s crucial for Street’s argument that she dismisses an evolved capacity to make evaluations in many domains. Because, if we have such an evolved capacity, it would show two features of human evolution: that an evolutionary adaptation is not likely to be distorting our truth judgments and, concomitantly, that evaluations might be more than contingent preferences, might be revelatory of facts in the world.

Two simple rejoinders seem decisive in nullifying Street’s objection to an evolved capacity for making general evaluations. First, having a large collection of separate and independent instincts, each designed to respond to a particular scenario in the environment – e.g., innate responses to the dangers of fire, flood, lightening, dark Chicago alleys at 2:00 AM, as well as the infinite number contingencies that might assail us – would be hugely more costly in evolutionary terms than a general capacity to learn and make evaluative assessments in several domains. But secondly, we do have a rational and evaluative capacity. Where did that come from? Did it fall from a tree? Straight’s argument is similar to that of Alfred Russel Wallace, who maintained that natural selection could not account for man’s excessively big brain and his capacities for music and mathematics, since these had no survival value. Wallace contended that this kind of objection to a natural-selection explanation made room for and furnished evidence for positing a spirit world, where higher powers would confer these traits on mankind (Richards, 1987, 176–184). Is that the implication of Street’s own position?

Street's further objection, and an apparently more telling one, concerns the kind of natural entities evaluative facts might be and their use in making evaluative judgments. She asks, "exactly what natural fact or facts does the evaluative fact that one should care for one's offspring reduce to?" (Street, 2006, 131) Or comparably, when we judge pains to be bad, "we need make no reference whatsoever to the fact that they are bad; we need only point out how [they] tended to promote reproductive success to take them to be bad" (Street, 2006, 151). In these cases, Street's assumption is that we immediately endorse care of offspring and shun pain because these evaluative attitudes have resulted from their aid in the reproductive success of our ancestors. We don't judge these situations to be good or bad because we have first perceived the fact that they are good or bad; rather the perception of an infant in need or the feeling of pain immediately evoke from us an evaluative feeling and that feeling has had reproductive advantage for our ancestors. We make the evaluation, according to Street, and then call it a fact, rather than initially perceiving an independent truth that caring for infants is morally good and pain is physically bad, and then so decide to endorse the one and avoid the other. Thus there is no independently existing moral fact in the world, only our own valuing some attitude or behavior, which we subsequently call a fact. So Street's analysis is similar to that of Ruse and Wilson: we are deceived if we think there are independent moral facts; rather our moral judgments are really about survival.

Street's version of the debunking argument sounds like it furnishes a plausible objection to moral realism, or it does until we examine the two fundamental assumptions on which it is based: that facts in the world stand ontologically apart from us as independent entities and that a trait naturally selected for, and thus part of our evolutionary acquirements, must be epistemically directed only to survival in the first instance. As is the case of the Ruse-Wilson analysis, both of these assumptions confuse ontological objectivity with epistemic objectivity.

Consider my judgment that the ball on the floor is red – a prototypical judgment about a fact in the world. Do I judge the ball to be red simply because I first recognize an independent fact in the world, a red ball, and on that basis make the true judgment "the ball is red"? This is the paradigm Street's discussion suggests. But I think it's the wrong paradigm. Isn't it rather that the experience of the object evokes from me an evaluative attitude – that is, a trait that has been selected for – that leads me to call the ball red? What kind of fact is a red ball, after all? The best physical and physiological theories would say something roughly like this: in the world there are clumps of atoms whose electrons have been excited by electromagnetic energy to give off photons that strike the cones in the retina – the color receptors in the eye – in a certain way; and, further, our neural reaction is an adaptive response

that leads us (who possess normal trichromatic vision) to have the sensational quality of red in our visual field. Absent a normal human eye (or trichromatic vertebrate eye), there are no red balls in the world. Most all of an object's qualities that we take to be existing independently in the world of our experience depend on the physiology of our nervous system, a nervous system that has evolved over millions of years. Moreover, being able to make basic color discriminations by way of qualitative sensation has survival value, and so such a response can only be the result of an adaptive link causally forged by evolution, one which unites our qualitative sensations to a certain kind of physiological stimulation that usually stems from a physically independent source. Thus our judging the ball to be red, a quintessential factual judgment, is a function of our evolutionary history, though it's not epistemically about that history, but about the fact of a red ball. Is my judging that torturing innocent children to be morally bad really any different than judging the ball to be red? Both judgments are functions of an evolutionary history and both are immediate, virtually a part of the perception itself and both are typically regarded as factual. If "badness depends in an important sense on our evaluative attitudes (Street, 2006, 151)," so does redness. Both are the result of the interaction of independent structures (e.g., atoms excited by a light source or the social structures of cooperation and altruism) with our neuro-mental evaluations. Those evaluations are part of the experience of facts in the world. Moreover, when we reflect and ask ourselves "but is the ball really red," we presumably do so by applying an acquired internal standard of the quality red, without which the reflective judgment would not be possible. The same is true when we reflectively ask, but is that act really morally sanctioned. What assures us that our judgments are objective is that they can be tested and intersubjectively verified. After all, we sometime discover that the ball only looks red because of the nonstandard character of the environment, just as we might be mistaken in judging the moral character of an act (e.g., not torture but a doctor injecting a vaccine).

We do think it a fact of the world that red balls might be found rolling around on the floor. We also believe the obvious fact that torturing innocent children is morally bad and, consequently, that it should be condemned. But in neither case do we initially assert to ourselves that it is a fact that there is a red ball before me (or that torturing children is bad) and, having established the fact, evaluatively judge there to be a red ball before me (or that child torture is wrong). Both are immediate, perceptive judgments about the world – the world we live in – and both are functions of our evolutionary history. So Straight's argument that the evolutionary process precludes a realistic view of the world either defeats realism not only in morality but in all areas of epistemic concern (e.g., red balls), or her analysis fails for want of understanding what "fact in the world" could mean. I would suppose neither option would be

acceptable to her or to the others who advance a debunking argument. That means the debunking argument itself ought to be rejected. That classes of cognitive judgments have an evolutionary genesis is thus no reason per se for rejecting the objectivity of those judgments, and thus their factual character.

Our access to a reality beyond mind can only be indirect. This is the message of realism – critical realism, that is. To make sense of our experience, we have to make a number of assumptions. We test and use those assumptions especially in our commerce with the world. Those assumptions were shaped during the course of the history of science, some assumptions having been rejected by reason of evidence and others added. Gradually these assumptions have become more firmly established and refined through the historical development of reason, of science. Developing science provides ever better evidence for constructing a map of reality. And that map constitutes the factual terrain of the world, at least one part of the factual world. But that terrain is not something we directly perceive. We well understand how the visual system produces color discriminations, and also how this likely arose through the long history of natural selection. In the same way, we understand why certain social behaviors have been perceived as morally good or morally bad, and also why these perceptions likely arose through natural selection operating on our ancestors. Together the extra-mental structures and their representations in our experience constitute the relational facts of our world. The extra-mental structures are postulated by the best scientific evidence we have, such as electromagnetic energy of a certain frequency and the effectiveness of cooperation and altruism for social creatures. As creatures with an evolved nervous system and cognitive apparatus, we perceptively evaluate both, one giving us something as factually red and the other as giving us something as factually morally good (or bad). Our experiential judgments are thus about complex, real structures of the world, one aspect more remote and dependent on our best theories, the other more perceptively immediate. Both aspects are sustained by the criterion of intersubjectivity. They are thus perfectly objective.

Conclusion: The Darwinian Theory of Morality

The Darwinian theory of morality supposes (1) that humans have evolved to be altruists and that fundamental principles of morality – the Categorical Imperative, the Golden Rule, etc. – can be reduced to altruism and (2) that evolution has instilled in mankind the social instincts of caring for the young and acting for the benefit of others, just as it has instilled in the human vertebrate the qualitative sensations produced by various neural impulses. These social instincts, of course, have survival value, as do our qualitative, sensory discriminations. As human beings, we have evolved to be intrinsically social

creatures. That means we have wrapped into the very fabric of our being inclinations that are powerful and regarded as special. In the same way, we have evolved a rational, evaluative capacity. Both have had survival value. You might even conclude our ancestors were social long before they became rational. Of course, in the human species, both traits have a distribution of more or less. In some few instances, these capacities may be completely absent. While there may be some imaginative scenarios in which a population, nominally called human, could have evolved into mindless and completely asocial creatures, that could not be our population. As a result of our moral and rational capacities, we make claims and formulate principles based on these capacities, which have the sanction of intersubjective affirmation – that is, they are objective and universal. That simply is who we are as human beings. Aristotle assumed humans were moral animals. Darwin demonstrated it.