

Zed Adams

Genealogies of Ethics*

Abstract: There have been many genealogies of ethics. Philip Kitcher's *The Ethical Project* stands apart in its ability to incorporate the insights of earlier genealogies while avoiding their oversights and mistakes. In this essay, I compare and contrast Kitcher's genealogy of ethics with two contemporary alternatives, those offered by Frans de Waal and Richard Joyce. Comparing Kitcher's genealogy with these alternatives makes it easy to highlight his most useful contribution to our understanding of the origin of ethics: the idea of ethics as a social technology. I conclude by identifying an oversight of Kitcher's own genealogy, a significant way in which the function of ethics-as-a-technology has been transformed from its origin to today.

1. Introduction

Quarantining a philosophical question is rarely good for its health. Many of the most rewarding works in ethics, for instance, simultaneously deal with issues in metaphysics and epistemology, philosophy of mind and language. That said, it is pointless to try to deal with *all* of the possible issues raised by a specific philosophical question. Some issues are tangential, after all. The difficulty is finding a middle ground between these two pitfalls: underreaching and overreaching. Philip Kitcher's *The Ethical Project* is a masterful example of a work that achieves this middle ground. His combination of curiosity and discretion results in a work that explores some of the most difficult questions surrounding the genealogy of ethics without either neglecting the larger philosophical problems raised or getting bogged down in tangential considerations.

The genealogy of ethics that Kitcher offers is both novel and plausible. In this short comment, my aim is not to dispute what he says about the origin of our ethical lives. My aim is rather to draw attention to a dimension of ethics that Kitcher does *not* discuss. Such a criticism might seem unfair, given what I have just said about the discretion Kitcher exhibits in refraining from addressing issues that are irrelevant to the underlying questions at hand. My ultimate goal is to show that in spite of Kitcher's enviable discretion, he nonetheless neglects an *essential* feature of our ethical lives.

* For helpful comments and suggestions, I am indebted to Jake Browning, Ben Chan, Jay Elliott, Jason Fissette, Adam Gies, Nat Hansen, and Eliot Michaelson.

There are four parts to this commentary on Kitcher's *The Ethical Project*. First, I situate Kitcher's genealogy of ethics in the current literature by comparing and contrasting it with two competing genealogies, those of Frans de Waal and Richard Joyce.¹ Second, I spell out what I take to be the core innovation of Kitcher's genealogy, the idea that ethics is a social technology. Third, I draw upon John Stuart Mill to argue that Kitcher's genealogy neglects an essential aspect of the ethical project. Fourth, I conclude with some diagnostic comments about what might have motivated Kitcher to neglect this aspect of the ethical project.

2. Genealogies of Ethics

There have been many genealogies of ethics. In order to see what is distinctive about Kitcher's genealogy, it will help to see how it compares and contrasts with two of the most prominent sorts of genealogies offered today (which are usually seen as exhausting the genealogical options): namely, those that aim to either *vindicate* or *debunk* ethics on evolutionary grounds.²

2.1 The Vindication of Ethics

Frans de Waal is a clear representative of the sort of genealogy that aims to vindicate ethics (see de Waal 2006). Drawing upon his own and others' studies of animal behavior, de Waal has argued that many of the supposedly distinctive features of human ethical life find surprising parallels in the animal kingdom. In particular, de Waal has claimed to find parallels for justice and altruism. Consider the following two cases:

"Many highly cooperative nonhuman species seem guided by a set of expectations about the outcome of cooperation and the division of resources. [For instance] the brown capuchin monkey responds negatively to unequal reward distribution in exchanges with a human experimenter. Monkeys refused to participate if they witnessed a conspecific obtain a more attractive reward for equal effort, an effect amplified if the partner received such a reward without any effort at all." (Brosnan/de Waal 2003)

"Rhesus monkeys trained to pull chains for differential rewards discover that pulling the more rewarding chain causes another monkey to be shocked. After witnessing this shock, two-thirds of the subjects prefer the nonshock chain, receiving half as many rewards. Of the remaining third, one stops pulling the chains altogether for 5

¹ I have not picked these authors at random. Beyond being two of the most prominent, empirically-informed contemporary authors working on the evolution of ethics, Kitcher himself cites each of them multiple times. It seems clear that he takes them to be two of his primary interlocutors.

² Joyce (forthcoming) is a clear representative of the assumption that these are the only two options.

days and another for 12 days, starving themselves to avoid shocking another monkey.” (Masserman/Wechkin/Terris 1964; cited in Preston/de Waal 2011)

In the first case, de Waal takes the capuchin monkeys’ aversion to unequal treatment to indicate an expectation of equal treatment. And he takes this expectation of equal treatment to be surprisingly similar to our own sense of fairness, in that fairness involves expecting an equal reward for equal effort. Given that a demand for fairness underlies our own sense of justice, de Waal proposes that the capuchin monkey’s behavior might shed light on the evolutionary origins of justice.

In the second case, de Waal takes the rhesus monkeys’ preference for giving up half or all of their possible rewards in order to avoid shocking other monkeys to indicate a willingness to sacrifice one’s own well-being so as to avoid harming another. And he takes this sort of self-sacrifice to be similar to our own sense of altruism. As such, de Waal takes this behavior to be another possible evolutionary building block for ethics.

For de Waal, discovering these evolutionary precursors to central parts of our ethical lives vindicates ethics because they undermine what he takes to be the central challenge to ethics: the worry that ethics is unnatural, an artificial imposition that does not have any real basis in human nature. By revealing surprising similarities between central parts of our ethical lives and non-human animal behavior, de Waal hopes to show that there is nothing unnatural about ethics.³ As we will see, Kitcher agrees that there is nothing unnatural about ethics but he does not think that merely noting similarities between the behavior exhibited by capuchin and rhesus monkeys and human ethical life suffices to give a complete genealogy of ethics.

2.2 The Debunking of Ethics

Richard Joyce is a clear example of an opposing sort of genealogy of ethics, one that aims to debunk ethics (see Joyce 2006). Joyce begins by agreeing with de Waal that it is possible to provide an evolutionary explanation for the origin of our ethical sense.⁴ But Joyce takes the possibility of such an explanation to raise problems for ethics, not solve them. Joyce points out that giving an evolutionary explanation for ethics only shows that having an ethical sense was useful for our ancestors’ genes, and he argues that this sort of usefulness falls far short of the sort needed to justify believing in ethics today. It falls short because it only provides an instrumental justification for having an ethical sense, whereas what is needed is an epistemic justification for our ethical beliefs—one that warrants thinking that they are true, not merely instrumentally useful for our ancestors’ genes.

³ For criticism of this argument, see Adams 2007b.

⁴ Although Joyce and de Waal are in broad agreement on this point, Joyce clearly thinks that de Waal has overstated the extent to which he (de Waal) has in fact provided such an explanation, as well as underestimated the difficulties involved in successfully doing so. See Joyce 2010.

As a way of showing how our ethical beliefs lack epistemic justification, Joyce draws a distinction between two different sorts of causal explanations for why we hold certain beliefs. The distinction is between (i) explanations that explain the benefits of holding certain beliefs only on the condition that the beliefs are true and (ii) explanations that explain the benefits of holding certain beliefs without presupposing that the beliefs are true. Call the first sort a *truth-dependent* explanation, the second sort a *truth-independent* explanation. Joyce offers mathematical beliefs as a case for which a truth-dependent explanation is correct. With regard to the belief that $1 + 1 = 2$, for instance, Joyce observes that “we have no grasp of how this belief might have been selected for, how it might have enhanced reproductive fitness, independent of its truth” (Joyce 2006, 182). This is why mathematical beliefs are not debunked by the existence of an evolutionary explanation for our tendency to think about math. The situation is different with regard to truth-independent explanations. Consider the following sort of case: many people are disgusted by the thought of drinking apple juice from a brand new and immaculately clean bedpan because they believe the bedpan is contaminated by some kind of invisible and irremovable contaminant. The most plausible explanation for the widespread acceptance of this belief about invisible and irremovable contaminants is truth-independent—it does not depend on there being such contaminants. Joyce thinks that if it is possible to provide a truth-independent explanation of the origins of a belief, then that belief’s epistemic standing is undermined. And since Joyce he thinks he can provide just such an explanation of the origin of our ethical beliefs, he thinks ethics is debunked.⁵ Kitcher agrees with Joyce that it is possible to give a truth-independent explanation for the origin of the ethical project. As will become clear, however, Kitcher does not think that being able to give a truth-independent explanation for our ethical sense suffices to debunk ethics.

2.3 The Debunking of Ethics

As noted, Kitcher’s genealogy is similar to de Waal’s in that he aims to naturalize our ethical sense, as a way of making us feel “at home with [our] ethical propensities” (*EP*, 279). And it is similar to Joyce’s in that he gives a truth-independent explanation for the origin of ethics.

The biggest difference between Kitcher’s genealogy and those of de Waal and Joyce is that Kitcher does *not* “suppose ethical practice is already present, at least in embryo, in our evolutionary cousins or our hominid progenitors” (*EP*, 9). Unlike de Waal and Joyce, Kitcher does *not* think that the origin of ethics can be explained in terms of biological evolution, so it will not do merely to identify the evolutionary building blocks of ethics or the inherited dispositions that our ethical sense depends upon. These will not do because ethics involves cultural developments that have happened much too recently to be explained in terms of biological evolution. For Kitcher, ethics is a human invention, not a biological inheritance. It is, as he puts it, a social technology, akin to other tools that we have invented to cope with the difficulties of life (such as chairs and computers).

⁵ For criticism of this argument, see Adams 2007a.

3. Ethics as a Social Technology

For Kitcher, ethics is a technological innovation designed to solve the problem of living together peacefully in large groups. It draws upon our disposition for psychological altruism, which Kitcher takes to be part of our biological inheritance, but is designed to cope with some of the limitations of that disposition: simply put, psychological altruism is not very reliable (i.e., it is prone to breakdown) and does not extend very far (i.e., it works best in small groups). Kitcher refers to all such limitations as ‘altruism failures’. In response to altruism failures, Kitcher thinks that we have come up with a more effective solution for the problem of living together: to command ourselves to obey norms that facilitate social stability. Here is how Kitcher describes the origin of this technology:

“The cumbersome peacemaking of our original hominid [progenitors was] replaced by a new device, one preempting rupture rather than reacting to it, and in principle capable of operating in a wide variety of contexts. That device is necessary for what we think of as ethical practice. I shall call it a ‘capacity for normative guidance’.” (*EP*, 69)

Kitcher’s genealogy begins with an observation about chimpanzees: they spend *a lot* of time engaging in “mutual grooming and other forms of physical reassurance” (*EP*, 67) as a way of maintaining social stability.⁶ Kitcher then assumes, plausibly, that something similar was true for our evolutionary progenitors. What’s distinctive about us is that we went on to invent a more efficient—i.e., less time-consuming and more reliable—way of achieving this same goal. We began to command ourselves to get along with each other. The benefits of this invention are manifold: it anticipates and prevents possible breakdowns in social stability, rather than merely restoring it after it has broken down, and allows us to engage in much more complex cooperative projects (such as living in cities). In short, this invention is what makes a distinctively human form of life possible.

Kitcher’s account of the origin of ethics exhibits a pair of features that are rarely combined: it is both novel and plausible. Here are two aspects of its novelty: first, it shows how giving a truth-independent explanation for the origin of our ethical sense need not debunk ethics, and, second, it does so without committing the naturalistic fallacy.

Kitcher’s genealogy is truth-independent because it explains the origin and subsequent development of ethics in terms of fulfilling a social function, rather than in terms of discovering truths. Here is how Kitcher puts this point:

“One way to break the stranglehold of the idea that [ethical] progress consists in accumulation of truth is to consider an area in which advances are understood differently. Technology serves as a paradigm. Our world is full of instruments, machines, and devices that improve

⁶ “During some periods in the recorded histories primate troops, particularly when social tensions are running high, the animals devote three to six hours per day to plucking and smoothing one another’s fur.” (*EP*, 57)

on previous efforts. [...] Progress with respect to these artifacts, and in the domain of technology generally, is readily understood as functional refinement. [...] Functional refinement consists in satisfying the wish more reliably or more completely, and doing so in ways that generate fewer problems for potential users.” (*EP*, 218–9)

Unlike Joyce, Kitcher does not take such a truth-independent genealogy to debunk ethics, because he thinks this genealogy merely involves getting clear about what we have been doing all along when we have engaged in ethical discussion and debate. As evidence for this historical claim, Kitcher looks at a series of paradigmatic instances of ethical progress and argues that it is false to think that these moments involved anything like the discovery of new ethical truths. For instance, he considers the case of John Woolman, the pioneering American abolitionist. Woolman kept a detailed journal that details the emergence of his anti-slavery views. But, Kitcher notes,

“Woolman’s *Journal*, the most revealing document we have, contains no mention of an occasion when he saw the plight of slaves in a new way (a sudden revelation at the battered bodies before him, that—somehow—transmitted ethical insight). [...] [B]lack slaves do not appear as individuals in Woolman’s narrative—nor do individualized slave owners whose ‘corruption’ Woolman might suddenly ‘perceive’.” (*EP*, 184)

Kitcher thinks that this is generally true about ethical progress, that such progress does not take the form of discovering new ethical truths. Rather, it involves refining our existing ways of preventing altruism failure. Why does this historical discovery not debunk ethics? Simply because the idea that ethical progress *must* involve discovering new ethical truths is a philosophical illusion. In other words, it is only if one is already committed an illusory philosophical picture of what ethical progress *must* be like—that it *must* involve the discovery of new ethical truths—that one would be inclined to think that giving a truth-independent genealogy of our ethical sense suffices to debunk it.⁷ If we free ourselves from this philosophical picture, Kitcher thinks we will be able to welcome his genealogy as an illuminating account of what we have always been doing in ethics and not as a skeptical debunking.

A second novelty of Kitcher’s genealogy is that it shows how the historical origins of ethics might inform our current practices, while avoiding the naturalistic fallacy altogether. That is, it provides a way of moving from a premise about what we *have* done in the past to a conclusion about what we *should* do now, without assuming that such a move is always justified (because it is not—there are obvious counter-examples). It does this by identifying a problem for our ancestors that we are still faced with today: the problem of altruism failure. If the ethical project worked as a solution to the problem of altruism failure in the past, and we still have this problem today, then the ethical project should still

⁷ A significant portion of Joyce’s evolutionary debunking of ethics involves saddling ordinary ethical thought with this sort of philosophical picture of ethical progress.

work as a solution. Granted, if there were a better way to solve this problem, then we could abandon the ethical project. But, plausibly, there is not a better way.⁸ Hence, we are justified in continuing to engage in the ethical project.

4. Mill's Question

In spite of the novelty and plausibility of Kitcher's genealogy, I worry that it is significantly incomplete. One way to bring out this incompleteness is to draw upon a famous episode from John Stuart Mill's *Autobiography*. At the age of twenty, when Mill had come to believe that Jeremy Bentham's utilitarianism said everything there was to say about ethics, Mill asked himself the following question:

“Suppose that all your objects in life were realized; that all the changes in institutions and opinions which you are looking forward to, could be completely effected at this very instant: would this be a great joy and happiness to you?” And an irrepressible self-consciousness distinctly answered, ‘No!’” (Mill 1989, 112)

The reason Mill found himself answering ‘No!’ to this question, and thereby discovering the incompleteness of Bentham's utilitarianism, was that he realized that all of the actions recommended by Bentham were *merely* of instrumental value—that none of these actions were worth engaging in for their own sake. Given this, the fulfillment of all of Bentham's utilitarian goals would thereby deplete the world of any value; there would be nothing left worth doing. This led Mill to the disturbing realization that Bentham thinks nothing is really worth doing *for its own sake*.

Now let us ask a similar question of Kitcher's ethical project. Imagine that you came to believe that Kitcher's genealogy has said everything there is to say about ethics. Now ask yourself what you would think if all altruism failures were prevented: would this be a great joy and happiness to you? I submit that if one really believed that Kitcher's genealogy was complete, one's reaction to this question would be similar to Mill's reaction. If you took preventing altruism failures to be *the* goal of the ethical project,⁹ and you imagine this goal fulfilled, then there would be nothing left worth doing because—on such a view—nothing is really worth doing for its own sake.¹⁰

⁸ Kitcher himself seems to assume a rather limited set of possible solutions to this problem. He appears to assume that there are only two options: (i) mutual grooming and other forms of physical reassurance or (ii) ethics (cf. *EP*, 278). One could think that this is a false dichotomy while still agreeing with Kitcher's larger point that the ethical project is, in fact, the best available solution to the problem of altruism failure.

⁹ Kitcher says that this is *the* function of the ethical project too many times to note every instance. For example, see *EP*, 222, 225, 240, 269, and 339.

¹⁰ There are two further parallels between Mill's dissatisfaction with Bentham and a dissatisfaction one might have with Kitcher. First, Bentham's flattening of any qualitative distinctions between pleasures is echoed in Kitcher's anti-elitist comments, which similarly seem to preclude the possibility of any qualitative distinctions between ends (cf. *EP*, 314–5). Second, another

To be fair to Kitcher, he does leave room in his genealogy for the invention of a technology to generate new functions of its own (cf. *EP*, 237–241). He even acknowledges that these new functions might lead to the emergence of a “vastly enriched notion of the good life” (*EP*, 137). But his discussion of this possibility is far too undeveloped. Most importantly, he does not discuss how certain human actions could come to be ethically worthwhile for their own sake. Nor does he discuss how one’s very identity as an agent could come to be informed by engaging in these actions. In short, Kitcher’s genealogy is significantly incomplete because it does not adequately explore the manifold ways in which the ethical project has come to be concerned with addressing needs that were simply unimaginable for our hominid progenitors.

5. Diagnosis

Kitcher is too intelligent a thinker for the incompleteness of his genealogy to be merely an oversight. I think there is a deeper reason for it. More than once, Kitcher emphasizes that on his account of the ethical project, *there are no ethical experts*.¹¹ Given the rest of his view, this is a rather peculiar thing to insist upon. After all, in every other domain of technology, there are experts. For example, consider the first technology that Kitcher mentions (when he introduces the idea that ethics is a social technology): chairs. I myself once tried to build a chair. It came out horribly, largely because of my lack of expertise in building chairs. If Sam Maloof had tried to build the same chair, it would have come out wonderfully. This point about technological expertise is not specific to chairs: I cannot think of an example of a technology for which there are not experts. Why, if ethics is also a technology, should it be any different?

Given these mundane facts linking technology and expertise, there has to be something else motivating Kitcher’s insistence that there are no ethical experts. I think the best explanation for Kitcher’s insistence in this regard is that he implicitly recognizes the incompleteness of his own genealogy of ethics. That is, I think Kitcher’s rejection of ethical experts reflects an implicit recognition that the main intellectual problems for ethics today have been fundamentally transformed. If the main intellectual problem for ethics today were figuring *how* to remedy altruism failures, then it is hard to see why this would not be a question for experts. This is, I submit, because it is generally true that when we know what the function of a technology is, it is just a fact of life that some people are better than others at inventing things that fulfill that function. Perhaps the most difficult question for ethics today, however, is *not* how to fulfill an antecedently-identified function. Perhaps the reason that we tend to think of ethics less as an engineering problem, and more as an inquiry into who we are,

source of Mill’s dissatisfaction with Bentham is the sense that Bentham effaces any distinction between projects that are valuable *because* they are mine, as opposed to projects that are valuable as a result of having been brought up to care about them. Although Kitcher clearly recognizes the importance of this distinction, his genealogy does not give us the resources to understand *why* it is important. I return to this point in the concluding section.

¹¹ For example, see *EP*, 8, 251, and 286.

is because the most difficult ethical question today is figuring what, if anything, has *become* of the function of ethics. It is with regard to this question—the question of figuring out what our ethical needs have become—that there are no experts. This, I think, is the reason Kitcher’s genealogy is incomplete: we do not, as yet, know what the ethical project has made of us.

Bibliography

- Adams, Z. (2007a), Review of *The Evolution of Morality*, in: *Ethics* 117(2), 363–369
 — (2007b), Review of *Primates and Philosophers*, in: *Ethics* 117(3), 552–555
 Brosnan, S. F./F. B. M. de Waal (2003), Monkeys Reject Unequal Pay, in: *Nature* 425, 297–299
 de Waal, F. B. M. (2006), *Primates and Philosophers*, Princeton
 Joyce, R. (2006), *The Evolution of Morality*, Cambridge/MA
 — (2010), Review of *Primates and Philosophers*, in: *Quarterly Review of Biology* 85, 100
 — (forthcoming), Ethics after Darwin, in: Ruse, M. (ed.), *The Cambridge Encyclopedia of Darwin and Evolutionary Thought*, Cambridge
 Kitcher, P. (2011), *The Ethical Project*, Cambridge/MA
 Masserman, J. H./S. Wechkin/W. Terris (1964), ‘Altruistic’ Behavior in Rhesus Monkeys, in: *The American Journal of Psychiatry* 121, 584–585
 Mill, J. S. (1989), *Autobiography*, ed. John M. Robson, London
 Preston, S. D./F. B. M. de Waal (2011), Altruism, in: Decety, J./J. Cacioppo (eds.), *The Handbook of Social Neuroscience*, New York