

*Cailin O'Connor/Nathan Fulton/Elliott Wagner/
P. Kyle Stanford*

Deus Ex Machina: A Cautionary Tale for Naturalists*

Abstract: XXX

“In every system of morality, which I have hitherto met with, I have always remark’d, that the author proceeds for some time in the ordinary way of reasoning [...] when all of a sudden I am surpriz’d to find, that instead of the usual copulations of propositions, *is*, and *is not*, I meet with no proposition that is not connected with an *ought*, or an *ought not* [...] as this *ought*, or *ought not*, expresses some new relation or affirmation, ‘tis necessary that it shou’d be observ’d and explain’d; and at the same time that a reason should be given; for what seems altogether inconceivable, how this new relation can be a deduction from others, which are entirely different from it.” (David Hume, *Treatise on Human Nature*, 3.1.1)

1. Introduction: Chekhov’s Gun

Naturalistic approaches to ethics face a persistent challenge, formulated perhaps most influentially by Hume in the passage above. The worry at the heart of this ‘is-ought’ problem can be formulated as a question: how could one even possibly reason from empirical facts about what is the case to normative conclusions about how things ought to be? The challenge has been advanced with particular vigor against those philosophers who appeal to the evolutionary history of our own species to inform our understanding of the character, status, or content of our own ethical practices, with critics insisting that only a ‘naturalistic fallacy’ could lead us to think that such considerations might have any relevance for our own efforts to reach substantive ethical conclusions. Despite the increasing influence of such naturalistic and evolutionary approaches to the subject, no response to Hume’s challenge has come to be broadly accepted. Most philosophers still firmly believe that *is* cannot be used to justify *ought*.

* Our thanks to Simon Huttegger and Tucker Lentz for helpful suggestions and discussion of this material.

With *The Ethical Project*, Philip Kitcher becomes the latest naturalist to join the fray. His book offers a novel approach to these foundational issues that consists of three central elements. The first is an admittedly speculative genealogy, or ‘how possibly’ story, of how the ethical project emerged in our evolutionary ancestors and how we came to have it in the form that we do. Second is a set of metaethical proposals centered on a novel and distinctive account of ethical progress. And the final element is a particular normative stance consisting of a specific set of principles, values, and commitments that he argues should guide our revision and extension of the ethical project into the future. These three elements may seem strange bedfellows, but each plays an important role in Kitcher’s account. Roughly speaking, Kitcher seeks to finesse the is-ought problem by using the middle element to connect the other two, appealing to his innovative conception of ethical progress to try to establish the relevance of his somewhat speculative genealogy to the substance of our own ethical deliberations.

We think there is much to admire in Kitcher’s approach, both for those who take naturalistic approaches to ethics seriously and for those who wonder why they should. We are especially enthusiastic about his proposed reconception of ethical progress. But at the end of the day we do not think that Kitcher’s intriguing metaethical proposals can serve as the needed glue between the speculative genealogy he sketches on the one hand and the substantive normative recommendations he seeks to advance on the other. As will become clear in what follows, this judgment represents a substantial repudiation of Kitcher’s central ambition: like so many before him, we suggest, Kitcher has tried and failed to somehow wring an *ought* from an *is*. But to see how and why, we must begin by describing the ingenious inversion of the usual understanding of the relationship between truth and progress in the ethical domain that lies at the very heart of Kitcher’s project.

Typically, ethical progress is understood in terms of truth—progress consists in coming to believe a higher proportion of freestanding ethical truths. Kitcher suggests that we should instead treat progress itself as the more fundamental notion. On this picture, truth is understood in terms of progress: truth is simply ethical belief that is preserved indefinitely through progressive ethical changes. The remaining question, of course, is how one understands progress. According to Kitcher, ethics represents a historically situated set of practices that first emerged in response to challenges plaguing early hominid social life. The ethical enterprise then *itself* evolved as new challenges emerged from our efforts to navigate the existing challenges successfully. Ethics thus represents a social technology with not only an original function but also further, generated functions that have arisen in the course of our efforts to fulfill the original one. As a consequence, ethical progress is best conceived as a matter of sophisticating that social technology to better fulfill these original and derived functions rather than apprehending a larger proportion of a static body of independently certified ethical truths.

Kitcher thus disavows the need for, along with the possibility of, any transcendent or transcendental backing for the authority of ethical claims. Whatever

authority such claims have will be rooted solely in the extent to which they can solve the challenges of allowing us to live together while collectively pursuing the various projects that actually matter to us. This also means that the sort of progress we find in ethics is of a distinctive variety: it is progress *from* an existing set of conditions that is achieved by better fulfilling the original and derived functions of the ethical project, rather than progress *to* the apprehension of a fixed and timeless body of ethical truths.

This conception has the potential to genuinely reorient much of our usual thinking about ethics in ways that promise to improve both our understanding of the ethical project and our ability to carry it out. Perhaps most importantly, this conception of progress is alive to the genuine possibility of fundamental conflict between the different functions that have emerged as the ethical project itself has evolved. But the very real prospect of such conflict also threatens to undermine the possibility of distinguishing genuine ethical *progress* from ‘mere change’ in our ethical views. This leads Kitcher to conclude that his metaethical account remains incomplete without the addition of some specific normative stance that can serve to pick out some changes as genuinely progressive even in cases of fundamental functional conflict. The particular normative stance he goes on to embrace advocates a return to what he sees as the original function of the ethical project. Kitcher then proceeds to make use of this standard both to evaluate the progressiveness of past transformations in our own history of ethical practice and to propose how the existing ethical project may be extended and modified in genuinely progressive ways.

The remainder of our paper will proceed as follows. First, we will discuss Kitcher’s speculative genealogy, outlining what we take to be several important problems with it and arguing that Kitcher’s intention to reform the ethical project by returning it to what he regards as its original function renders these problems especially acute. We will then go on to suggest that Kitcher’s attempt to anchor a renewed normative vision of our collective ethical project in a naturalistic account of its genesis and genealogy founders on the very rock that Hume threw in its path long ago, despite Kitcher’s protests to the contrary. We argue that although Kitcher manages to pick out the particular normative stance he advocates using criteria of coherence and continuity, these criteria are in no way sufficient to justify the privileged position he both wants and needs to assign to that normative stance. In the last section of the paper, we will present a diagnosis of where we think Kitcher’s central ambitions for the very project of understanding our ethical practice in naturalistic terms have gone wrong, and we will offer an alternative vision of what we think naturalized ethics should aim to accomplish instead.

2. Problems with Kitcher’s Genealogy

As noted, Kitcher’s approach is rooted in what he admits is a speculative genealogy of ethical practice, a ‘how possibly’ story intended to illustrate how the ethical project as we know it even *could* have emerged from the conditions of

early hominid life. Kitcher is particularly concerned to articulate such a genealogy for those skeptics of ethical naturalism who do not see how any such genesis for the ethical project is possible and who therefore feel compelled to embrace alternatives like divine authority as a counsel of despair. Quite a number of features of Kitcher's genealogy strike us as suspect, but we will restrict our attention to just a few that are of special importance because they make fundamental contributions to the content of the normative stance that he ultimately seeks to defend.

First, Kitcher argues that the original function of ethics as a social technology was the prevention of 'altruism failures', i.e., instances in which hominids were insufficiently sensitive to the desires and needs of fellow group members to allow for effective group living (see *EP*, 222ff.). Although we think Kitcher is quite right to focus our attention on the selective advantages conferred by the ethical project, we also think that the prevention of altruism failures constitutes an implausibly narrow description of the ways in which such advantages were probably achieved. Although enhancing the altruistic dispositions of its participants was undoubtedly *among* the early functions of the ethical project, cooperative efforts that seem to have little to do with altruism also seem likely to have played an important role in improving fitness in ethically minded hominid groups. Such cooperative efforts may arise when all group members simply pursue their own interests rather than some group members sacrificing for the sake of others or even modifying their own behaviors in response to their perceptions of the needs and desires of others. At times, Kitcher seems to feel the pull of such an expanded menu of functional possibilities for the ethical project, but instead of broadening his conception of the original function of ethics he unfortunately stretches the meaning of 'altruism failures' to implausibly encompass nearly any type of cooperative or coordinated social interaction, going so far as to claim at one point (*EP*, 126) that even a failure to fully develop one's own innate talents may be seen as a form of altruism failure because of its potential negative impact on the welfare of the group. If the ethical project originally enhanced fitness through, for example, non-altruistic cooperation, then its original function ought to be understood more broadly as something like the promotion of social coordination.

A further problem is that Kitcher's genealogy embodies the presumption that in coming to be guided by normative considerations hominids actually sat down in the 'cool hour' of the evening to seek agreement concerning specific rules of conduct (e.g. *EP*, 5–6; 104f. and ch. 3 generally). In other words, Kitcher suggests that ethical conventions arose explicitly, as the result of substantive ethical discussion, rather than implicitly as the result of daily group interactions not specifically aimed at determining ethical rules. One of the great triumphs of recent work in evolutionary game theory, we think, is that it has shown how easily and naturally rule-governed behavior can emerge without any such explicit agreements. Work in this tradition¹ has shown that in situations where coordination is desirable—such as those that characterized the beginnings of the ethical project—behavioral conventions can arise through repeated interaction

¹ See, for example, Young 1993; Skyrms 1996; 2004; Binmore 2005; Alexander 2008.

with variable outcomes on the part of the participants. No extensive ‘cool hour’ discussion is necessary to determine ethical practice in a particular society, as features of day-to-day life and interaction will often allow such practice to arise naturally. This would not mean that ‘cool hour’ discussions concerning rules and norms of conduct do not occur or are not important, but simply that we have little reason to think that this is how such rules would have become established (or even revised) in the first place.

A final significant problem with Kitcher’s genealogy concerns his treatment of the emergence of normativity itself (*EP*, ch. 2). Simple rule-governed behavior may evolve with relative ease in a wide range of cognitively unsophisticated creatures, but even when such behavior is mediated by an explicit awareness of the consequences of noncompliance with the rule it remains a far cry from a full-fledged *normative constraint*, in which an individual or a group commits to the idea that one *should* or *should not* behave in certain ways. We think Kitcher has evaded the most important work in explaining any such transition, and that the evasion becomes particularly evident in his discussion of the emergence of punishment.

On Kitcher’s view, the emergence of normative guidance proceeded stepwise, and the implementation of punishment by early hominid groups constituted a key step in its development (*EP*, 87f.). Punishment for rule-breaking was eventually internalized, leading to full, self-directed normative guidance, and in this way widespread normative guidance gained a foothold in hominid society. The problem with this account is that it equivocates between a thin conception of punishment that simply represents the unpleasant practical consequences of non-compliance and a much more normatively robust conception involving something like a *justified* evil delivered upon someone who *deserves* it because she has violated a *legitimate* rule of conduct. This difference is perhaps especially easy to see when we ask whether the requisite non-interference of other group members (see *EP*, 89; 91) in a particular case of ‘punishment’ is motivated only by the practical constraint of comparative power (or advantage) or also involves such considerations of justification, desert, and legitimacy. We need not accept the existence of the sort of special ‘ethical point of view’ devoid of emotion that Kitcher rightly rejects as a myth (*EP*, 80–81) in order to see these cases as importantly different. Moreover, we agree with Kitcher’s suggestion that it is pointless to ask when ‘real’ punishment emerges in the sort of stepwise evolutionary progression he envisions (*EP*, 89–90) and that the development of robust normativity will involve the recruitment of a wide array of different entangled motivational and emotional states (including fear, reverence, solidarity, and prudential calculation; *EP*, 81–86ff.). Nonetheless, his repeated slide (e.g. *EP*, 89; 91) from the brute modification of others’ behavior through positive and negative reinforcement to a more normatively robust conception of punishment for wrongdoing simply obscures the need to show how distinctively normative judgments of the sort that we do at least *sometimes* experience could emerge from even the most heterogeneous collection of rule-governed practical motivations and social expectations. He thus either assumes that some sort of normativity is present from the very beginning, or else simply refuses to say when and how the

most important step in the emergence of ‘normative guidance’ as we ourselves know it has been taken.

There are a number of further concerns and reservations we have about the particulars of the genealogy of ethical practice that Kitcher offers. We single out these three in particular because they are among those that matter most for Kitcher’s larger project. If Kitcher is right in thinking that he can generate substantive moral guidance from genealogy in the way that he imagines, it is all the more important to get the relevant genealogy right, or at least make it as plausible as possible. It will make a substantial difference to see the weight of history as driving us to enhance cooperation and coordination generally rather than simply prevent failures of altruism. If explicit early deliberations in the cool hour of the evening about what the rules for living together should be are largely a genealogical fantasy, then their role in the original project cannot be any part of the authority on which we use them to shape or model contemporary ethical deliberation. And if Kitcher’s story does not actually manage to get any characteristically ethical sort of normativity off the ground it cannot reassure those who wonder whether any naturalistic story can do so.

3. From Genealogy to Morality

Although we have suggested that these concerns regarding the details of Kitcher’s moral genealogy are especially important ones, surely the most pressing question is whether and how Kitcher can actually manage to derive substantive conclusions about what our own approach to ethics should be from any such genealogy. Indeed, the later sections of *The Ethical Project* are chock-full of substantive claims concerning an extraordinarily wide variety of ethical issues, ranging from the ethically permissible size of the human population, to the ideal distribution of educational opportunities, to the circumstances in which the abortion of fetuses with abnormal genetic conditions should be permitted (*EP*, 316–317). How could Kitcher’s distinctive metaethical apparatus permit such normative recommendations to emerge from any naturalistic account of how we came to have the ethical project in the first place? Does he really believe he has gotten an *ought* from an *is*?

Return for a moment to Kitcher’s suggestion that progress rather than truth is the fundamental metaethical notion and that such progress must be understood as the improvement of an existing state rather than advance towards some independently specified ideal. Note that this conception works well as long as we restrict ourselves to just a single function (perhaps the original) or a set of consistent functions: progress then simply consists in innovations of method that allow us to fulfill this function or set of functions more effectively. But the conception seems powerless to help when faced with conflicts between functions that emerge as the ethical project itself evolves over history, or with questions about whether the adoption of some new function (or the abandonment of an existing function) would itself be progressive, or with questions more generally about what the function or functions of the ethical project *should be*.

Kitcher seems to clearly recognize at least some dimensions of the problem. He allows, for instance, that new functions will be generated as we seek to fulfill the original or existing functions of the ethical project in new and better ways, and he recognizes the possibility of genuine and fundamental conflict between modifications that better pursue some of these functions at the expense of others.² A pure devotion to the notion of ‘progress from’ might seem to demand that in such circumstances we simply privilege the original function of ethics (or in the case of conflict between derived functions, the earlier of the two), but Kitcher wisely refrains from embracing any such arbitrary solution. Instead, as we noted above, he allows that his metaethical apparatus remains incomplete without the addition of a particular, substantive *normative stance* specifying how the existing ethical project *ought to be revised* and thus allowing it to resolve such functional conflicts. Such a stance consists in commitment to a particular conception of the good, particular methods of ethical deliberation, and (as consequences) the many above-noted specific recommendations concerning abortion, education, population size, and much else besides.

On its face, this further appeal to a normative stance is curious, for Kitcher seems to simply announce his preference for the original function of the ethical project over the derived (cf. *précis*, 14). The only virtue Kitcher is ultimately able to claim for this normative stance is coherence: it is not nothing, he insists, that the various parts of his normative stance fit together naturally with one another. But *bare* coherence is presumably easily achieved by normative stances, and is in fact achieved by any number of simplistic and strident ethical codes under which few if any of us would care to live. It is a challenge to achieve such coherence only while also seeking (as Kitcher does) to preserve or even maximize *continuity* with the ethical project as it has come down to us through history. The central issue, then, is whether it *counts in favor* of including a particular element as part of our own normative stance or part of the ethical project going forward that it has been an element of that project in the past and/or present. Citing the influence of Dewey and other pragmatist thinkers, Kitcher is clear that he is concerned with ethical problems only as they arise in our lives as they are actually lived, which perhaps precludes massive discontinuities in the evolution of the ethical project, but that is not the issue. The question is whether such continuity is merely a practical necessity or instead a virtue that we should value *as such* and strive to maintain *for its own sake*.³ Kitcher thus picks out

² His clearest example seems to us not very clear (see, e.g., *EP*, 248–249; cf. *précis*, 10–11). It is the conflict between the egalitarian distribution of goods favored by the original function of preventing failures of altruism and permitting the inequalities in such distribution that emerge as a consequence of adopting methods that expand the resources available to the community as a whole.

³ This also highlights why Kitcher’s response to various forms of skeptical challenge is so unsatisfying—he repeatedly contents himself with pointing out that such skepticism is not a practical possibility for us. It is true that the genuine skeptic cannot be silenced (*EP*, 271), that skeptical concerns are typically posed by real people who wish simply to be reassured in the deference they pay to their own moral convictions (279), and that characters like Plato’s Thrasymachus, Hume’s sensible knave, and Nietzsche’s free spirit are cardboard cutouts rather than actual interlocutors (274–280). But the central issue here concerns the status or character of our moral convictions, and these responses no more help with concerns about justification

the normative stance he advocates by simply grafting together criteria whose authority or appeal do not seem to fit together: a requirement of coherence whose normative credentials we may (perhaps) treat as unproblematic but which is otherwise fairly undemanding, and a demand for continuity which radically restricts the range of possible normative stances but whose claim on us looks to be more a matter of practical necessity than normative force. This curious combination of constraints may allow Kitcher to pick out just a single normative stance, but it does not seem to have done nearly enough to justify the privilege.

We suspect that Kitcher will welcome our criticism up to a point. He will remind us of his insistence that the proper role of philosophy is simply that of midwifery, inviting us to ‘join the conversation’ and propose a competing normative stance if we have one to offer. But this is not responsive to our real concern, which is the conflation of continuity as a practical necessity with continuity as a normative virtue or justification. Nor does it help to appeal to the idea that ethical progress is always ‘progress from’ and thus embedded in our ethical history, for the very question is whether we should regard continuity as a normative virtue in deciding how to extend the ethical project progressively from here forward. If we ask why we should regard continuity as a virtue rather than a practical necessity, Kitcher seems to have no answer to give, nor even one to propose that coheres with the normative stance he himself has advocated. That is, Kitcher’s *own* normative stance would seem to give us no reason whatsoever to value continuity with the existing ethical project for its own sake, nor to think that pursuing internal coherence under the constraints of continuity is the best way to modify the existing ethical project. The suspicion arises that Kitcher has simply found a number of his *own* values reflected in some (but only some) historical and genealogical aspects of the ethical project and tried to claim a special form of authority for them on that basis.

This consideration suggests a further way to see why the curious combination of demands that Kitcher uses to pick out his preferred normative stance remains dubious even if we concede that ethical progress must be ‘progress from’ an existing background and that coherence is indeed a virtue. Kitcher draws our attention to several dramatic shifts in ethical practice to motivate the intuition that ethical progress (rather than ‘mere change’) is possible: the abandonment of *lex talionis*, the expansion of the civil rights of women, ending chattel slavery, and the acceptance of the legitimacy of homosexual love. But surely we do not think that these dramatic shifts were progressive because they produced a more coherent ethical practice (presumably they did so in some ways and not in others), still less because they maintained continuity with other aspects of the existing practice. We would view these particular changes as progressive even if we decided that in fact they reduced the coherence of the ethical project overall

than the fact that inductive skepticism is not a practical possibility for human beings serves to answer Hume’s problem of induction. Thrasymachus cannot be recruited as an ally (274) who has a mistaken view of the genesis and genealogy of the ethical project, the sensible knave will indeed shrug at being diagnosed as inhuman (276), and the free spirit will reply that taking “a leap into a completely unknown, and doubtfully human, form of life” (282) is precisely what he had in mind. Real people who do in fact “feel the tug of ethical commands” (279) should be no more reassured about the tug of those commands than they were before.

and departed from existing practices in fundamental ways (as of course they all did in important respects). We judge these changes to have been progressive, simply because they generate a practice that better accords with the most important values that we ourselves embrace. It thus seems illegitimate for Kitcher to suggest that his favored normative stance is somehow ratified or recommended by the very structure or history or character *of the ethical project itself*. Kitcher may be able to pick out his preferred normative stance by grinding the existing ethical project as he finds it against the constraints of coherence in a particular way, but if we *endorse* the resulting normative stance this is not ultimately because of its continuity or its coherence, but instead the fact that it agrees with the ethical values that we ourselves hold most dearly.⁴ And without that agreement, Kitcher has no way to defend the reality of ethical progress (rather than ‘mere change’) in any of the difficult cases or when questions of what the functions of the ethical project *should be* genuinely arise: continuity and coherence drift away as irrelevant, and any defense of such a normative stance would ultimately have to rest on its own bottom.

Perhaps all this is just what Kitcher means to concede in allowing that philosophy can do no more than propose such normative stances and serve as midwife to their emergence. But if so the form of that concession is exceedingly obscure. It seems for all the world as if Kitcher takes himself to *support* the normative stance he advocates by appealing to its combination of coherence and (selective) continuity with the historical ethical project.⁵ In fact, we should stress that the challenges we have raised constitute a problem for Kitcher because he sees himself as answering rather than dismissing or ignoring such concerns about justification.⁶ From the very earliest descriptions he gives of his own undertaking he notes that “[c]onvincing proposals come with some form of support” and insists that by demanding coherence while preserving the original function of the ethical project “the deepest challenge—the most important accusation against naturalistic ethics—is turned back” (*EP*, 9).⁷ He identifies worries about the ‘naturalistic fallacy’ with the challenge issued by Hume with which we began our own paper, and insists that “the Humean challenge can be tackled head-on”

⁴ On this view it is perhaps unsurprising that, as Kitcher documents in fascinating detail, the grounds on which earlier thinkers were moved to advocate policies like the abolition of slavery or educating women often stand in stark contrast to the grounds on which we ourselves would advocate them.

⁵ Perhaps another way to frame the concern is to suggest that Kitcher constantly slides back and forth between an ‘internal’ perspective in which judgments of progress are implicitly relativized to a presupposed (but inadequately justified) specific scheme of values and an ‘external’ perspective in which no such scheme is presupposed but in which there is no distinction to be made between genuine progress and ‘mere change’. But Kitcher does not seem to be aware of, much less advertise, any such vacillation.

⁶ He occasionally seems tempted by the latter option instead: he explicitly sets aside questions about whether ancient changes in our ethical or social arrangements that are now firmly entrenched were progressive in some global sense as “idle”, for example, and suggests that a pragmatist is not “moved to settle the question” in part because it “does not leave us uncertain about how to go on from where we are” (244).

⁷ The “deepest challenge” here is a form of Hume’s own: Kitcher’s text references his response (in Section 56) to a challenger who asks why she herself should be bound by the ethical project in the form it has taken historically.

(258). Deploying his distinctive metaethical apparatus to show ethical transitions can be genuinely progressive, he argues, will “remove the sting from the challenge”, leaving the ethical inferences he champions “exempt from the mysteries Hume rightly queried” (258). And he goes on to insist that, “[o]nce ethics is viewed as a social technology, directed at particular functions, recognizable facts about how those functions can better be served can be adduced in inferences justifying ethical novelties. The mystery that worried Hume disappears.” (262) We hope that it is by now clear why we disagree.

4. Diagnosis and Prescription

We think we can offer a more general diagnosis of what has gone wrong. Kitcher has rightly seen that the authority of his (and our) substantive normative convictions cannot be grounded in the benediction of a powerful being, or the structure of rationality itself, or any such Archimedean point. He has rightly noted the profound differences between the achievement of ethical insight and processes like scientific discovery by which we learn about persisting and independent features of the external world. And he has rightly understood that our substantive moral convictions must gain any authority they have from the role(s) that they play in human lives as they are actually lived. But Kitcher nonetheless seems to keenly feel a need for some *proxy* for the transcendent or transcendental backing he has dismissed if his own convictions are to amount to anything more than the latest ethical fashions. He then quite literally relocates God in the machine: reimagining ethics as a social technology, he recasts the normative authority of his *own* ethical convictions as a matter of what it is to be using that technology *correctly*—in (some of) the ways it was originally used, to fulfill (some of) the functions for which it developed more efficiently and effectively—and does so in ways that serve to write his own substantive convictions into the very rules of the game.

Let us be clear that we support many of the details of Kitcher’s broadly left-leaning 21st Century neo-liberal politics, and we certainly view developments such as the abandonment of *lex talionis* and chattel slavery as well as the expansion of civil rights for women and homosexuals as paradigmatic instances of ethical progress. But we do not think these commitments are or can be supported by anything more than the sorts of first-order argumentative considerations that we would ordinarily give in favor of them. In particular, they are not and cannot be supported by metaethical claims about the history, genealogy, or character of the ethical enterprise itself. There is no ‘getting behind’ the ordinary sorts of reasons we would give concerning the impact of particular policies on the lives of those we think entitled to our moral concern, no appeal to values with a special standing built into the very fabric of the ethical enterprise, and (consequently) no special way to make those commitments compelling for agents whose fundamental values are sufficiently and radically at odds with our own. Such opponents can be convicted of cruelty, or unfairness, or insensitivity, but they cannot simply be waved away on the grounds that they are making a

fundamental error about the very nature of the ethical project that somehow favors our own views over theirs: no substantive moral views are privileged in light of a special relationship they bear to how we came to have the ethical project in the first place. Those views will have to be defended by showing how they can allow as many of us as possible to live together in tolerable peace, security, grace, dignity, or whatever states matter most to us.

Ultimately, we think the role of naturalism in the ethical project should be quite different than what Kitcher seems to envision. Instead of finding some replacement for a transcendent or transcendental backing for the authority of our own ethical convictions, we suspect naturalists would be better served by embracing the difficult process of learning to live without any such backing in the first place. We want naturalism to give us a sufficiently clear-eyed view of the status, role, and functioning of the substantive values we presently hold to be willing to defend them while recognizing that no such transcendental backing nor even a proxy for one grounded in the history or origins of the ethical project itself is ultimately possible. As Kitcher's own discussion helps to illustrate, a naturalistic genealogy can help us to better understand how particular functions are served by particular substantive principles and commitments. And our confidence in the importance of continuing to fulfill those functions can be reaffirmed by seeing that connection more clearly, even if that importance cannot be defended against a dissenter whose existing values are sufficiently at odds with our own. We can come to recognize that these are merely *our* values without consequently losing sight of the fact that they *are* our values. In short, we seek a naturalistic explanation of the ethical project that allows us to stand up for the values we actually hold, while also making room for any sensible limits we ourselves recognize on the legitimacy of imposing those values on others. We thus seek a naturalistic explanation of the ethical project that allows us to look its origins straight in the eye without losing any of our enthusiasm for carrying it out.

Bibliography

- Binmore, K. (2005), *Natural Justice*, New York
Hume, D. (1978), *Treatise of Human Nature* (2nd ed., edited by L. A. Selby-Bigge and P. H. Nidditch), Oxford
McKenzie Alexander, J. (2008), *The Structural Evolution of Morality*, Cambridge
Skyrms, B. (1996), *Evolution of the Social Contract*, Cambridge
— (2004), *The Stag Hunt and the Evolution of Social Structure*, Cambridge
Young, H. P. (1993), The Evolution of Conventions, in: *Econometrica* 61, 57–84