

FUNCTION AND TRUTH IN ETHICS

Michael Smith

In the first part of *The Ethical Project* Philip Kitcher provides a naturalistic account of the function of ethical practice. Though he conceives of this as an unashamedly empirical task—Kitcher denies that the function of ethical practice is to get at the ethical truth, or to get people to behave in accordance with the ethical truth—it is not one that he thinks has a debunking upshot. Instead he thinks that his account of the function ethical practice, which is based on his reading of the practice's origins in the behaviour of our hominid ancestors, provides us with all of the materials required for a pragmatic naturalist account of ethical truth. More familiar a priori attempts to ground ethical truth in the spirit of Kant are given short shrift, as are the supposedly more empirically adequate non-cognitivist accounts of ethical judgement which hold that such judgements express emotions, commitments, intentions, desires, or non-cognitive attitudes of some other kind.

Though there is great deal to admire in *The Ethical Project*, especially for someone like me who was so antecedently ignorant of the relevant science and history, I am afraid that I remain unconvinced. My discussion is divided into three main sections. In the first I consider Kitcher's account of the function of ethics and his pragmatic naturalist account of ethical truth. If ethical truth really were fixed by the function of ethical practice in the way he suggests, then this would make the truth of ethical claims turn on issues that seem to me, and I would have thought to Kitcher too, to be quite irrelevant. It would also mean that many of our ethical beliefs, including some of those required to maintain ethical practice so understood, are false. This doesn't entail that his account of ethical truth is mistaken, but it does suggest that something might have gone wrong, and I suggest what that something might be. In the second section I argue that Kitcher's account of the function of ethical practice does not commit him to a pragmatic naturalist account of ethical truth. Non-cognitivism remains an alternative, notwithstanding his reservations. Since it does a much better job of shoring up our ethical commitments than his own pragmatic naturalist account, I suggest that Kitcher would be better off embracing some form of non-cognitivism. In the third and final section I briefly consider Kitcher's objections to Kantian approaches to ethical truth, and I explain why at least one version of this approach, the one that I myself prefer, is not vulnerable to his objections. This too remains an option for someone who accepts his account of the function of ethical practice.

1. Kitcher's accounts of the function of ethical practice and ethical truth

According to Kitcher, ethical practice has its origins in the biologically-driven need our hominid ancestors had to live in cooperative communities. This need led to their having a direct concern not just for their own kin, but also for those who were members of the cooperative coalitions to which they belonged, as their limited altruistic concern gave rise to the cooperative foraging behavior that enabled the members of these coalitions to do better than those who lacked such concern.

Since altruistic concern wasn't always strong enough to override other dispositions the members of these coalitions had to please themselves, a tendency to engage in elaborate and time-consuming peace-making behaviours developed alongside such concern.

Peace and mutual tolerance are typically hard-won. Precisely because of this, observations of chimpanzee societies disclose periods of intense social interaction,

lengthy bouts of grooming undertaken to reassure friends who have been disappointed by recent behavior. At times of great tension within a group, chimpanzees can spend up to six hours a day huddled together, vastly longer than any hygienic purpose demands. (*The Ethical Project* p.73)

Though limited altruistic concern and elaborate peace-making strategies were sufficient for cooperation among individuals so long as the groups remained small, by the time our ancestors were using tools that required a trade in raw materials that were mined and transported large distances, and even more so by the time they lived in cities of a thousand or more, "strategies of peacemaking through face-to-face reassurance...[were]...no longer applicable", so there had to be "a system of agreed-upon rules for forestalling potential conflicts and for dealing with people who are relative strangers" (*The Ethical Project* p.116). Individual behavior thus came under the control of what Kitcher calls "normative guidance" (*The Ethical Project* pp.67ff).

What is the distinctive feature of normative guidance? According to Kitcher, there is normative guidance whenever the "cumbersome peacemaking of our original hominids is 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" (*The Ethical Project* p.69). As I understand it, this consists in the possibility of our having and conveying thoughts with normative contents, normative thoughts that themselves become a source of motivation.

The simplest—and original— form of normative guidance consists in an ability to transform a situation that would otherwise have been an altruism failure, by means of a commitment to following a rule: you obey the command to give weight to the wishes of the other. (*The Ethical Project* p.74)

Normative guidance is thus the hallmark of ethical practice as we know it, as in large groups cooperation is underwritten by thoughts about the behaviours required by rules and commands. Ethical practice, so understood, has been on the scene for much longer than we might have thought: "[b]y thirty thousand years before the present, the enterprise of framing rules for life together, the ethical project, must have been quite well developed" (*The Ethical Project* p.118).

If the function of ethical practice is to remedy altruism failures, then we can evaluate particular moments in the history of ethical thinking by that very standard. For this reason, Kitcher sees certain transformations in the ethical thinking of our ancestors and ourselves as examples of ethical progress: the move from an eye-for-an-eye conception of retribution, which allowed for those other than a wrongdoer to bear the burden of punishment for a wrongdoer's actions, to a conception that makes the wrongdoer himself bear that burden; the move from a conception of wrongdoing that emphasizes a loss of honor to one that emphasizes the common good; the end of slavery; the extension of legal and political rights to women; and the partial but still continuing shift towards more accepting attitudes among heterosexuals towards homosexuals (*The Ethical Project* pp.138ff). In each case, those who were candidates for altruistic concern all along, but were excluded from it, have been or are being brought within its scope.

Should we suppose that the ethical judgements that favour extending altruistic concern in these ways are *true*? Kitcher's answer is that we should suppose that they are true because the concept of ethical progress can itself be used to define a conception of truth for ethical claims:

Descriptive counterparts of ethical rules count as true just in case those rules would be adopted in ethical codes as the result of progressive transitions and would be retained through

an indefinite sequence of further progressive transitions. There is no prior conception of ethical truth, so that people make ethical progress when they discover (or stumble on) independently constituted ethical truths. Progress is the prior notion, and descriptive counterparts of rules *come to count as true* in virtue of the fact that they enter and remain in ethical codes that unfold in a progressive sequence... (*The Ethical Project* p.246).

We didn't just change the way we treated women, women deserved to be treated better than they had been treated in the past, and this proposition—the proposition that women deserved better treatment—is true. It is true, or so we have reason to believe, because the rule requiring us to treat women in this way would be retained in every further progressive transition in our ethical thinking.

Attractive though this picture is, there are, I think, some hidden problems. The main source of these problems lies in Kitcher's characterization of the function of ethical practice. To repeat, he tells us that that function is to remedy altruism failures, or alternatively to satisfy the "endorsed desires of all" (*The Ethical Project* p.241). But, given his own description of ethical practice, neither of these is quite accurate. A better characterization would be *to remedy those altruism failures that would interfere with cooperation*, or *to satisfy the endorsed desires of all to the extent required for cooperation*. This is because progress itself isn't characterized ethically, but is rather a matter of *ongoing cooperation*. Kitcher himself is of course aware of this, but he doesn't always bear it in mind when he explains how ethical progress has and will occur.

Consider, for example, his explanation of why it would count as ethical progress to extend altruistic concern to non-human animals.

Is anything similar available in the case of nonhuman animals? Apparently so. Just as the late Paleolithic witnessed first occasions of transient association among neighboring groups and later increases in band size, so, at the very end of this period, people began to set up more regular patterns of association with some kinds of nonhuman animals. The practice of domestication creates something like a society, one including some nonhuman members. That practice refines the ethical function of satisfying the endorsed desires of all. Yet, because the animals newly included differ in some important properties from the people across the river, obvious questions arise. Does this expansion really create anything like a society? If so, to what extent can the rules adopted within the local group be carried over to the new members?... Just as participants in the ethical project have learned, over tens of thousands of years, to deal first with nongroup members in a limited range of contexts, to form larger societies governed by a common framework of rules, and ultimately to apply an ethical code to people with whom interactions are negligible or even nonexistent, so too we can understand a progressive sequence of changes identifying altruism failures in inflicting, or even tolerating, pain in animals whose anatomy, physiology, and behavioral reactions resemble those already embraced within the ethical framework. ... The idea that it is good to relieve pain, wherever it occurs, is a natural extension of ideas about the good, progressively elaborated during the evolution of the ethical project. Accompanying that idea is a related thought about human beings: it is not good for people to be insensitive to pain, whether it occurs in other people or in nonhuman animals. Inflicting pain—or even permitting it—produces human beings who are debased, whose characters and lives are less good than they might be. As the ethical project evolves, views of the good human life become richer, and engaging in conduct that causes unnecessary pain to others, including nonhuman animals, comes to appear

detrimental to living well. (*The Ethical Project*, pp.307-309)

What Kitcher says here sounds like a credible bit of ethical reasoning, but we aren't officially supposed to be in the business of evaluating the thought that it is good to relieve pain, wherever it occurs, in terms of its intrinsic plausibility. That would be for progress to be ethically characterized. Instead we are supposed to evaluate it in terms of its power to serve the function of ethics, where that function is to fix the content of rules whose adoption would remedy altruism failures that would interfere with cooperation.

Seen in this light, it seems quite implausible to suppose that the rule we have to adopt, for ethical practice to serve its function, is the injunction to relieve pain, wherever it occurs. It is quite implausible because we are evidently able to ignore vast quantities of pain, both human and non-human, without that having any affect at all on the levels of cooperation we enjoy. A failure to extend altruistic concern to non-human animals quite generally would thus seem to leave cooperation largely intact. Moreover, to the extent that a failure to extend our altruistic concern towards non-human animals in some limited way would have an affect on cooperation, the explanation of this would have to be that a failure to extend such concern would interfere with cooperation, not the truth of the claim that it is good to relieve pain, wherever it occurs. In the first instance what is important is the cooperation of other humans, not the alleviation of non-human animal pain. Indeed, I take it that the claim that it is good to relieve pain, wherever it occurs, thus turns out to be false, on Kitcher's way of understanding ethical truth. What is true is rather that it is good to relieve pain, wherever it occurs, *to the extent that a failure to do so would interfere with cooperation*, a claim that does not have the same ring of intrinsic plausibility as the original. There is thus a huge gap between the rules we need to adopt, in order for ethical practice to serve its function, which is ongoing cooperation, and the ethical beliefs which many of us have and give voice to insofar as we talk with each other about these rules, Kitcher included.

A similar problem arises when Kitcher explains why the extension of legal and political rights to women counts as ethical progress, though in this case he is more up front about what's at stake.

What was discovered? Factual knowledge advanced: people learned that, under different conditions of socialization, women wanted things traditionally denied to them; that they found satisfaction in attaining some of these things; that fulfillment of the wishes did not thwart desires previously seen as central to female nature—public life combined more or less satisfactorily with family life. Increased factual knowledge proliferated desires for access to public life, fostering acceptance of the desires as prevalent and no longer pathological. Recognition of the suppression or frustration of those desires aroused sympathy, recruiting male as well as female allies for the reform movement. Like the early elaboration of normative guidance, in which particular altruism failures cause too much trouble, the increase of sexual egalitarianism occurred partly because, in the end, traditionalists wanted a quieter life. (*The Ethical Project*, p.153)

As Amia Srinivasan says in her commentary on this passage:

This may or may not be a sound account of how the feminist movement originated, but the point is that Kitcher interprets the question 'Why do we believe that women and men have equal rights?' as a request for causal explanation, not as a demand for justification. Feminism is justified not because women want to be treated like men and because it so happens that granting them equality doesn't cause too much damage to family life: it is

justified because of the equality of men and women. (Srinivasan 2012)

Indeed, in the terms in which I put it earlier, on Kitcher's way of thinking about ethical truth, it turns out to be false that women are the equals of men. What turns out to be true is rather that women are the equals of men *to the extent that men's failure to treat them as equals would interfere with cooperation*. Once again, this claim doesn't have the same ring of intrinsic plausibility.

To repeat, what these examples bring out is the gap that exists between the rules we have to adopt, if ethical practice is to serve the function Kitcher says it has, and the ethical beliefs that many of us, Kitcher included, naturally have and express insofar as we take these examples to be examples of ethical progress. Nor should that be surprising given that the idea of cooperation that is doing the work in Kitcher's account of ethical practice has to take the empirical facts as given. To stick to the example of relations between men and women, the explanation of their cooperation must take the preferences of men and women as given, and it must also take as given their relative positions in a power hierarchy, where that power hierarchy might itself be kept in place in part by the fact that certain ethical beliefs that are false, or anyway false according to Kitcher's pragmatic naturalist account of ethical truth, are widely shared. With that idea firmly in mind, consider a somewhat more cynical explanation of how men came to extend their altruistic concern to women.

Let's assume that though men had more power than women, and that though many of them would have preferred to continue to dominate women in the ways they always had, they did not have sufficient power to continue to do so while retaining their cooperation. They therefore had to compromise. How much men needed to compromise is, however, an open empirical question, a question whose answer is settled by where the equilibrium points lie in the satisfaction of men's and women's preferences, given the differences in their power relations. Ethical beliefs take on a special significance in this context, given the effect that they can have on cooperative behaviour. If, for example, both men and women could come to believe that women are the equals of men *simpliciter*, and the effect of their acquiring this belief was to limit the extent to which men had to compromise with women—perhaps women would no longer agitate for all of the changes that the truth of the belief would demand if they believed (falsely) that, because men were on board, they would make those changes willingly—then widespread possession of the ethical belief that women are the equals of men *simpliciter*, a belief which to repeat looks is false according to Kitcher's pragmatic naturalist account of ethical truth, precisely because it would limit the need for men to compromise with women, and allow them to maintain their superior position in the power hierarchy, would be a crucial part of what made the "ethical progress" Kitcher identifies possible. There would, however, in this case be a very substantial gap between the content of the rules that men and women adopt, rules which limit their cooperation, and the content of the ethical beliefs that make their adoption of those rules stably support their cooperation. It would be no surprise if we so easily fall into a more intrinsically plausible commentary when we explain ethical progress.

Of course, this more cynical story wouldn't be interesting if it didn't seem to be a more or less accurate description of what has in fact happened in relations between men and women, insofar as men's altruistic concern has been extended to women. But precisely because it does seem to be so accurate, the interest of the story is plain. It suggests that if we were to accept both Kitcher's account of the function of ethical practice and his pragmatic naturalist account of the truth of ethical claims, then his pragmatic naturalist account of truth may well end up debunking the truth

of many of our ethical beliefs. To function well in remedying failures of altruism to the extent required for cooperation, ethical practice may itself cause us to have ethical beliefs whose possession is cooperation-conducive, but ethical beliefs that turn out to be false, given Kitcher's account of the truth of ethical claims. This, at any rate, is what seems to be the case as regards ethical relations between men and women, and a similar story looks like it might be right regarding ethical relations between humans and non-human animals as well.

It is perhaps worth noting that what I have said here is in much the same ballpark as Kim Sterelny's criticisms of *The Ethical Project* (Sterelny 2012). There are, however, important differences. Like me, Sterelny thinks that an account of ethics like Kitcher's might well make many of our explicit ethical beliefs turn out to be both crucial for cooperation and false, given that they are so ripe for manipulation by power elites. However Sterelny's argument for this conclusion begins by taking issue with Kitcher's claim about the role played by normative guidance in ethical practice. In his reply to Sterelny, this is the issue on which Kitcher focuses (Kitcher 2012 pp.172-174). By contrast, I have accepted that normative guidance plays the role that Kitcher says it does in ethical practice; I have argued that this means that Kitcher himself needs to distinguish between the rules that we adopt and act on, on the one hand, and the explicit ethical beliefs that we express when give our commentary on these rules, on the other, ethical beliefs that are themselves cooperation-conducive; and I have argued that, to the extent that normative guidance plays the role that Kitcher says it does, there will often be a gap between these two. If ethical truth is defined in terms of the rules we adopt, in the way Kitcher supposes, it follows that he himself must admit that the explicit ethical beliefs many of us have, ethical beliefs that may well be crucial to the stability of ethical practice, could well turn out to be false.

2. Kitcher's account of the function of ethical practice and his rejection of non-cognitivism

Is it possible to accept Kitcher's account of the function of ethical practice, but reject his account of ethical truth? Or, to put the question in the slightly different terms suggested by Srinivasan, is it possible to disconnect questions of causal explanation from questions of justification, and to concede that while Kitcher answers the causal questions correctly, his answers to the questions of justification are incorrect? If so then that would leave us free to admit that though we have in fact internalized certain rules as the result of transitions in ethical thinking that are cooperation-conducive, and though these rules are—let's stipulate—rules that would remain internalized through an indefinite sequence of further progressive-in-Kitcher's-sense transitions, our acting in accordance with those rules might lead astray ethically. One obvious way in which to do this would be by accepting some kind of non-cognitivism.

Consider once again ethical relations between men and women. Perhaps I notice something peculiar about myself. I find that when I reflect on relations between men and women, I can see no ethically relevant distinctions between them: women are, I say, the equals of men *simpliciter*. But I also notice that the rules we have all internalized and live by, men and women alike, myself included, only lead us to treat women as equals to the extent required to elicit their cooperation. I therefore notice that there is a cooperatively stable future in which I am able to retain my relatively higher position in the power hierarchy, and so control the extent to which I give up benefits so that additional benefits can flow to women, notwithstanding my ethical views. With all of this understood, what could I do when I judge that women are the equals of men *simpliciter*?

The only answer Kitcher can give is that I am expressing a false ethical belief. Non-cognitivists, by contrast, can say that I am expressing my endorsement of a norm that requires me to treat women as equals *simpliciter*. In other words, I am condemning the rules that I know myself to have internalized, and recommending the internalization of alternative rules, rules that would lead me to treat women as the equals of men *simpliciter*. I could well understand that the transformation I endorse would lead to a rockier future, one in which willing cooperation between the sexes may never be fully forthcoming. Given our stipulation, I could also have good reason to believe that that transformation will never come about. But I could still coherently endorse the transformation, as I could think that that future, rocky though it would be, would be a better future than one in which men and women stably cooperate but only because they don't act on their ethical belief that women are the equals of men *simpliciter*.

On this non-cognitivist way of thinking, the mere fact that men and women would still stably cooperate without women being treated as the equals of men *simpliciter* puts no pressure on us to suppose that women are not really the equals of men *simpliciter*. Similarly, to switch examples, the mere fact that people would stably cooperate with each other without their internalizing a rule requiring them to relieve the pain of both human and non-human animals whenever they are able to do so puts no pressure on us to suppose that that rule doesn't capture the ethical truth of the matter. Nor should it be thought that, as non-cognitivists, we would have to shy away from talking of ethical truth. Allan Gibbard, among others, has shown how and why we would still treat ethical judgements as truth-apt even if we were to deny, as many non-cognitivists do, that such judgements are truth-apt in any robust sense (Gibbard 1990). Given that Kitcher himself seems to have ethical views about the treatment of women, and non-human animals, and homosexuals, that are at some distance from the ethical truth about these matters, given his pragmatic naturalist account of ethical truth, and given that these views that are more intrinsically plausible than the views that that account of ethical truth commits him to, the question is why doesn't he embrace some form of non-cognitivism.

Chrysostomos Mantzavinos asks Kitcher this question in his commentary on *The Ethical Project* (Mantzavinos 2012, p.38). Here is Kitcher's reply.

For non-cognitivism to succeed, however, would require an account of progress for the emotions taken to be expressed—an account of what is meant by saying that the emotions elicited at later stages are more 'apt' than those appearing at earlier times. Despite my admiration for Allan Gibbard's work, I cannot find any such account in his writings, nor in those of any other non-cognitivist I know. So, if the challenge is to be met, some serious work will have to be done. (Kitcher 2012, p.179)

But once we have distinguished questions of causal explanation from questions of justification, Kitcher's demand that the non-cognitivist tells us what it is for an emotion to be 'apt' can be seen to be ambiguous.

If progress is understood in terms of the removal of impediments to on-going cooperation, then an 'apt' emotion is simply one which grounds our internalization of rules which would be internalized in any transformation of the rules whose internalization would similarly result in on-going cooperation. The question thus concerns causal explanation, and the non-cognitivist should borrow Kitcher's answer. Importantly, however, the non-cognitivist is not committed to endorsing such on-going cooperation—remember again our discussion of ethical relations between men and women. But if progress is understood in terms of aligning our ethical

commitments with the ethical truth, then an 'apt' emotion is simply one whose expression would result in our making true ethical claims. The question concerns justification, and the non-cognitivist should answer it in his own terms, something he would presumably do by expressing his endorsement of emotions whose expression would result in ethical claims that are much the same as the ethical claims he makes when he expresses the emotions he in fact has.

The upshot is that there is no challenge here that the non-cognitivist lacks the resources to meet. It therefore seems to me that Kitcher should abandon his account of ethical truth and open himself up the possibility of embracing some form of non-cognitivism. Doing so would allow him to argue in favour of our internalizing rules without regard to whether or not the internalization of those rules would result in stable cooperation. Given that stable cooperation might itself depend on the maintenance of unjustifiable discrepancies in power relations, this is surely just as it should be.

3. Kitcher's rejection of Kantian approaches to ethical truth

Embracing non-cognitivism isn't the only way in which someone could accept Kitcher's account of the function of ethical practice while rejecting his pragmatic naturalist account of ethical truth. Another alternative would be to embrace an account of the truth of ethical claims in the spirit of Kant. Kitcher himself is, however, no fan of Kantian approaches.

Consider some approaches found in or inspired by Kant. Ethics expresses the requirements of pure practical reason: to deny the moral law reason generates is to fall into a mode of irrationality, in which one contradicts oneself. Or it might be said: ethics consists in a set of principles ideally rational agents would agree to under ideal circumstances, so failing to abide by its precepts is to violate conditions of rationality. Tough-minded skeptics would hardly be brought to silence by either of these dicta or by any plausible emendations of them. The skeptic speaks: "You can call the procedures you use to generate the rules you favor 'pure practical reason,' if you like, and suppose those who don't go along with them are involved in some sort of contradiction, but the mere label doesn't frighten me, and the effects you envisage don't appear particularly dreadful. If I reject these rules, I am hardly doing something similar to asserting a statement and its negation—and even if I were, it's not obvious anything I care about would be compromised by doing so. The history of science contains episodes in which people have worked quite well with internally inconsistent ideas (think of the Bohr model of the atom), and I have no reason to think my 'practical irrationality' will pose difficulties for me. Nor am I much moved by the thought of rules hypothetically ideal people (supposedly better than me) in some fictional situation would agree to. Why should I be bound by what they would decide? There is no argument for thinking the purposes I care about would be ill served by flouting any such precepts."...Maybe, though, the skeptic's answers have exposed something wrong with him: he has failed to meet ideal conditions of rationality. If that is so, the criterion for a successful reply to the skeptic is modified. He does not have to be silenced; one must merely have an account of why his responses are problematic. Given this understanding, however, pragmatic naturalism can do just as well as the allegedly superior nonnaturalistic approaches. Where Kantians and contractarians see failures of ideal rationality, pragmatic naturalism diagnoses an inability to appreciate how central the ethical project is to human life. (*The Ethical Project*, pp.272-3)

How should Kantians respond? The first thing Kantians should say is, I hope, apparent from the earlier discussion.

According to Kitcher, the Kantian's response to the skeptic reduces to the claim that there is "something wrong with him", and this means that the Kantian and the pragmatic naturalist are on an equal footing as regards the skeptic. They agree that there is something wrong with him, but disagree about what that thing is: the Kantian says that the skeptic is practically irrational, whereas the pragmatic naturalist says that he has "an inability to appreciate how central the ethical project is to human life." The issue between should therefore be decided by figuring out whose account of what's wrong with the skeptic is better supported by the evidence. But as the earlier discussion made clear, there isn't obviously anything wrong with someone's having the inability that the pragmatic naturalist identifies. To be sure, the ethical project, as Kitcher describes it, is central to human life to the extent that it sets the terms for stable cooperation. But given that stable cooperation might itself be ethically problematic—remember again our discussion of what might underwrite the stable cooperation we find between men and women—that doesn't count in favour of anyone's appreciating that kind of cooperation, in the sense of their being disposed to bring it about. We should therefore all be skeptical of the ethical project, insofar as we agree that it has the function Kitcher identifies.

This leaves us with Kitcher's disparaging remarks about Kantian approaches at the very beginning of the passage. Kitcher's skeptic says: "You can call the procedures you use to generate the rules you favor 'pure practical reason,' if you like, and suppose those who don't go along with them are involved in some sort of contradiction, but the mere label doesn't frighten me, and the effects you envisage don't appear particularly dreadful." That isn't just the skeptic talking, that's Kitcher talking. As I understand it, he is saying is that the Kantian doesn't have a plausible story to tell about what the practical irrationality of the skeptic consists in. Kantians claim that the skeptic is practically irrational, but this claim is pure bluff. The question is whether anything can be said to convince Kitcher that he is wrong about this.

The Kantian's idea, at least as I understand it, is that we can derive a substantive account of the reasons that agents have, where the substance of those reasons is recognizably ethical, from the concept of an agent. Such a derivation is possible, it seems to me, if the following premises are all true:

- (1) that agents *qua agents* have the function of being desire-realizers (analytic truth about agents);
- (2) that it follows from (1) that ideal agents—that is, agents who perform their function optimally—must have and exercise the capacity to realize their desires no matter what their content and know what the world is like, no matter what it is like, at least insofar as the world's being the way that it is bears on their realizing their desires (the modal conception of the ideality of a functional kind);
- (3) that it follows from the potential for conflict in the optimal possession and exercise of these two capacities—think of the conflict present in an otherwise ideal agent who desires to believe that *p* whether or not it is true—that ideal agents must have certain dominant conflict-resolving desires: specifically, ideal agents must have a dominant desire not to interfere with their exercise of their capacities to know their world or realize their desires in it (on condition that those desires wouldn't lead them to interfere), whether in the present or the future, and they must also desire to do what they can to help ensure that they have the capacities to know their world and realize

their desires in it, whether in the present or the future (novel premise whose truth depends on (2), the fact that agents are temporally extended, and a pain-staking examination and rejection of alternative ways in which the conflict identified might be resolved);

(4) that there is no consistent way for an ideal agent to restrict his desires to help and not interfere to the current and future stages of himself, and hence that ideal agents have expanded versions of these two dominant desires ranging over all agents: that is, they desire to help *everyone* and not to interfere with *anyone* (a premise inspired by Nagel (1970) and Parfit (1984) that teases out the implications for the theory of practical rationality of the fact that each agent is (potentially) just one agent among many);

and

(5) that there is an analytic tie between facts about an agent's reasons for action, on the one hand, and facts about which of that agent's options their ideal counterpart has desires for, on the other (analytic truth about reasons for action).

Though this is not the place to argue that these premises are all true, I should say that I have argued for their truth elsewhere (Smith 1994, 2011, 2012, 2013). What is crucial for present purposes, however, is not so much the truth of these premises as a conditional claim. If (1)-(5) are all true, then all agents have dominant reasons to help and not interfere, which are substantive reasons for action with recognizably ethical content, and those agents who fail to help, or who interfere, act in ways that they have dominant reason not to act, which is by all accounts one way in which an agent could be practically irrational. Though Kantian accounts of the content of ethical claims and their truth might still be mistaken—one or more of premises (1)-(5) might be false—it thus seems to me that a plausible story can be told about why the skeptic is practically irrational. The Kantian's claim that the skeptic is practically irrational thus isn't pure bluff.

What should Kantians say about Kitcher's account of the function of ethical practice? Kantians can agree with Kitcher that the function supported by the practice's origins in the behaviour of our hominid ancestors is to preempt and remedy altruism failures that interfere with cooperation; they can agree that ethical practice achieves this through getting individuals to internalize certain rules; and they can agree that there are therefore certain rules whose internalization would, as a matter of empirical fact, lead people to stably cooperate. But Kantians should insist that it is an open empirical question whether people have reasons to perform the acts that would result from their internalization of these rules, even if their internalization would lead them to stably cooperate. They should say that that depends on whether the scheme of cooperation is one in which people are helping and not interfering to the extent that they can. If it is such a scheme, then they should insist that there is a happy coincidence. But if it isn't, then the Kantians, like the non-cognitivists, should insist that we have good reasons to reject such a stable scheme of cooperation in favour of a scheme in which people do help and don't interfere to the extent that they can.

Here once again we see the relevance of Srinivasan's distinction between questions of causal explanation and questions of justification. If our focus is on purely causal questions, then there is nothing wrong with supposing that the function of ethical practice is the function Kitcher identifies. But if we switch to questions of justification, then, at least according to the version of the Kantian approach that I myself prefer, we should suppose that the function of ethical practice *should be* to get people to believe and behave in accordance with the ethical truth as spelled out in (1)-(5). Alternatively put, insofar as we think that our function is desire-realization, and think

that it follows from this that our function is to be sensitive to reasons, we should suppose that the function of ethical practice is to get at the truth about what the reasons are, and then to get people to behave in accordance with that truth.

REFERENCES

- Gibbard, Allan 1990: *Wise Choices, Apt Feelings* (Cambridge, MA: Harvard University Press).
- Kitcher, Philip 2011: *The Ethical Project* (Cambridge, MA: Harvard University Press).
- _____ 2012: "Afterthoughts. Reply to Comments" in *Analyse & Kritik* (1) pp.167-189.
- Mantzavinos, Chrysostomos 2012: "The Ethical Project: A Dialogue" in *Analyse & Kritik* (1) pp.21-38.
- Smith, Michael 1994: *The Moral Problem* (Oxford: Wiley-Blackwell).
- _____ 2011: "Deontological Moral Obligations and Non-Welfarist Agent-Relative Values" in *Ratio*, (24), pp.351-363
- _____ 2012: "Agents and Patients, Or: What We Learn about Reasons for Action by Reflecting on Our Choices in Process-of-Thought Cases" in *Proceedings of the Aristotelian Society*, (122), pp.309-331
- _____ 2013: "A Constitutivist Theory of Reasons: Its Promise and Parts" in *LEAP: Law, Ethics, and Philosophy* (1), pp.9-30
- Nagel, Thomas 1970: *The Possibility of Altruism* (Princeton: Princeton University Press).
- Parfit, Derek 1984: *Reasons and Persons* (Oxford: Oxford University Press)
- Srinivasan, Amia 2012: "In the Long Cool Hour" in *London Review of Books* Vol.34 No.23, December, pp.17-18.
- Sterelny, Kim 2012: "Morality's Dark Past" in *Analyse & Kritik* (1) pp.95-115.