

On a Better Moral Theory and a Theory of a Better Morality¹

by

Geoff Sayre-McCord

(draft for discussion, not citation)

Introduction

Normally, there is a sharp distinction between a better theory of X and a theory of a better X. That the theory of a better X is a theory according to which things are different from the way one's (so far) best theory says they are is (normally) no reason whatsoever to think one's (so far) best theory is wrong, just reason to wish X were different (and, if it is possible, reason to work to change X). That it would be better if everyone were treated as equals is no reason whatsoever to think that they are; that it would be better that death came quickly, painlessly, and late in life is no reason whatsoever to think it does; that it would be better if we could fly with the snap of a finger is no reason whatsoever to think that we can...

In contrast (I claim) when the subject matter is normative, this normally sharp distinction is elided. There being a difference between one's theory of the best X (the best morality, the best standards of inference, the best rules of justification...) and one's (so far) best theory of X, necessarily provides a reason (though perhaps not a decisive reason) to think one's (so far) best theory is wrong.

The elision plays an essential role in a range of arguments concerning morality, practical rationality, and theoretical rationality, a few of which I discuss. Yet it smacks of depending crucially and unacceptably on wishful thinking – on supposing that the fact that things would be better if only they were a certain way provides some reason to think they are that way. As a result, it invites invocation of “Wouldn't it be nice that p, therefore p” reasoning.² of a sort Thomas Nagel suggests there might be and David Enoch sketches. I think that the invitation should be resisted. The elision is to be defended, I argue, not as an instance of defensible wishful thinking but as a reflection a constraint on acceptable normative theories that is itself explained by a distinctive characteristic of normative concepts that sets them all apart from descriptive concepts.³

A Normally Sharp Distinction

Normally, as I have said, there is a sharp distinction between a theory being a better theory of X and it being a theory of a better X. Suppose, for instance, that we

¹ Thanks to Joshua Blanchard, Richard Holton, Barry Maguire, and David Plunket for thoughts and feedback on the earlier versions of this paper..

² Thomas Nagel suggests that such reasoning is sometimes legitimate and David Enoch offers an intriguing sketch of a defense of such reasoning.

³ The distinction between normative and descriptive concepts is of course a term of art that, I am thinking, is vindicated (if it is) by capturing a theoretically valuable contrast. In what follows I try to identify an important feature of that contrast.

develop a theory of what the laws of (say) North Carolina are. This is of course no easy task. We would need to take account not just of the laws on the books, but whether and how they were enforced, as well as the nature and grounds of past judicial decisions (and the weight of *stare decisis*), the relevance of common law and the distinction between that and statutory and regulatory law, of which “violations” were actually punished, and which just “taxed,” of whether various candidate laws conformed to the standards of legal legitimacy in force in the state, of how state law is related to federal law, etc..⁴ Imagine, though, that we (tentatively, but with some confidence) arrive at a theory of what the laws actually are.

Next, suppose we turn to the project of figuring out what the best set of laws for North Carolina might be. Again, this is of course no easy task. Empirical and normative complexities abound. We will need to take account, presumably, of whether and how easily they might be understood, of what enforcing them would involve, of their compatibility with Federal laws, of their capacity to win willing compliance, of the effects of putting them in place, of the degree and nature of their infringement on liberty, of whether they are fair and fairly enforceable, etc..⁵ Still, imagine that we (tentatively, but with some confidence) arrive at a theory of what the best laws would be.

Imagine too that our two theories do not coincide perfectly – that on our understanding of what the laws actually are, the laws of North Carolina would be better if only some of them were different. (Not an implausible conclusion.) Would this give us any reason to think that our original theory of what the laws are is mistaken. Normally not⁶; rather, it would give us reason to wish the laws were different and reason to work to change the laws (assuming they can be changed).

The same is true if we were to switch from a theory of the laws of North Carolina to one concerning the workings of the human mind, or geological forces, or social conventions. In each case, our coming to think things would be better were they different than we take them to be is, it seems clear, *no reason at all* to think that they are different – only reason to wish they were different and to change them if we are able. There is, in other words, a sharp distinction between a better theory of the human mind, of geological forces, and of social conventions, on the one hand, and a theory of better human minds,

⁴ Here I mean to leave room for whatever sorts of considerations people think bear on a proper understanding of what the laws actually are.

⁵ Here I mean to leave room for whatever sorts of considerations people think bear on a proper understanding of what the laws should be.

⁶ The “normally” is important here, since if our theory of the nature of North Carolina’s laws includes the claim that certain normative requirements are a part of the law the shortfall in the law as we take it to be may actually serve as grounds for thinking we have misunderstood what the law actually is. More about this in what follows. For now, assume that no normative elements appear in our positive theory of the laws of North Carolina, or that the ones that do are not elements that tell in favor of the laws we have identified as better.

of better geological forces, and of better social conventions, on the other. The first tells us how things are, the second how it would be good for them to be.

The Distinction Elided

So much is true, *if* the theories are of non-normative matters, but not, I believe, otherwise. Suppose, for instance, that we develop a theory of what the standards of morality are. This is of course no easy task. We would need to take account not just of what is intuitively plausible, but (presumably, though this is somewhat contentious) of what might be brought under general plausible principles that are mutually consistent, of what is metaphysically tractable, of what is epistemically accessible, etc.⁷ Imagine, though, that we (tentatively, but with some confidence) arrive at a theory of what the standards of morality actually are.

Next, suppose we turn to the project of figuring out what the best standards of morality might be.⁸ Again, this is of course no easy task. Complexities abound: at least all those already mentioned concerning our attempt to develop a theory of morality.⁹ Still, imagine that we (tentatively, but with some confidence) arrive at a theory of what the best standards of morality would be (by the lights of our original theory). Imagine too that our two theories do not coincide perfectly – that on our understanding of what morality’s standards actually are, a different set of standards would be morally better. Would this give us any reason to think that our original theory of what the standards of morality are is mistaken. I think it would; we would be finding ourselves with reason to think we had originally gotten the standards wrong, not with reason to wish the standards of morality were different and reason to work to change those standards (if only that were possible). In this case, we would have found there was reason to guide our judgments, actions, and reactions by the morally better standards, rather than what we had originally taken to be the standards of morality.

Significantly, the legal example switches sides, so to speak, if (i) our best theory of law incorporates moral principles as central to what the laws are and (ii) it is in light of

⁷ Here I mean to leave room for whatever sorts of considerations people think bear on a proper understanding of what the principles of morality actually are.

⁸ The quality of a moral or legal system is a direct reflection, I am assuming, of the quality of the standards or laws that comprise it. But, I am also assuming, the quality of those standards or laws may well depend in part on how they play out in practice – say, on whether the moral standards can secure allegiance and whether the laws of a legal system can be consistently enforced. To what extent how things would play out in practice matters to the evaluation of the standards or laws in question is of course a contentious issue. For my purposes, it is to be settled by our best theory of morality or of the law, respectively, since the relevance of the theory of a better X to the best theory of X depends on the best theory including the standards in light of which the other theory counts as better.

⁹ If I am right that the project of developing a theory of the best morality is part and parcel of developing the best theory of morality, the complexities involved in doing the latter will involve all those involved in the former. Specifically, the burden of evaluating how the standards we arrive at fare, morally, in comparison with alternatives, which I highlight in discussing which morality is best, has a place in determining which theory of morality is the best.

those principles that our theory of a better set of laws counts as better. This may happen, for instance, on Ronald Dworkin's theory of law according to which equal concern and respect are central to legal legitimacy in the United States. In that case, if our best theory of North Carolina's laws include a prohibition on same-sex marriage yet a commitment to equal concern and respect tells in favor of allowing such marriages, we would have reason (albeit not necessarily conclusive reason) to think our original theory might be mistaken about the prohibition on same sex marriage. Evidence against the prohibition is provided by the standards that, according to our best theory of the law, are a part of the law and serve as constraints on other laws.¹⁰

Some Examples Concerning Morality

To put a little flesh on the bones of what I have in mind, let me mention several examples of arguments that seem to work (if they work) by taking advantage of the elision between a better theory of morality and a theory of a better morality.

One example can be found in Adam Smith's discussion of fortune's effect on judgments of praise and blame. Smith starts out with a crisp and elegant argument for thinking that "To the intention or affection of the heart... all praise or blame, all approbation or disapprobation of any kind, which can justly be bestowed upon any action, must ultimately belong."¹¹ What turns on luck (good or bad) cannot justly, Smith maintains, have an impact on whether one is praised or blamed, approved or disapproved. For what turns on luck is not under our control and not properly attributable to us. Strikingly, though, Smith then immediately turns to canvassing all the ways in which our praise and blame, approval and disapproval, are actually influenced by fortune. He finishes with a robust defense of this practice. This defense plays out against the following observation

If the hurtfulness of the design, if the malevolence of the affection, were alone the causes which excited our resentment, we should feel all the furies of that passion against any person in whose breast we suspected or believed such designs or affections were harboured, though they had never broken out into any actions.¹²

And it turns crucially on the claim that "if the baseness of the thought which had given birth to no action, seemed in the eyes of the world as much to call aloud for vengeance as the baseness of the action, every court of judicature would become a real inquisition." Allocating praise and blame on the basis of actions and their effects, rather than merely

¹⁰ Similarly, if the best theory of law is as Natural Law theorists claim, then in discovering that what we initially took to be the law was, in fact, not acceptable in light of Natural Law, we would have reason to think we had gotten the law wrong.

¹¹ Third paragraph in section on the effects of fortune in *The Theory of Moral Sentiments*. p. 133. Get proper cite.

¹² *ibid.* pp. 151-152.

on an agent's (supposed) intentions and designs, is, Smith argues, essential to the protection of privacy.

Smith never returns to his original argument against moral fortune nor to his claim that

The consequences which actually, and in fact, happen to proceed from any action, are, if possible, still more indifferent either to praise or blame, than even the external movement of the body. As they depend, not upon the agent, but upon fortune, they cannot be the proper foundation for any sentiment, of which his character and conduct are the objects.¹³

So we never get a clear idea of what, in his original argument, Smith would reject. Nonetheless, Smith seems to think that the importance of privacy, of avoiding an inducement to inquisition, and of providing an incentive, however malevolent one's attitudes, not to act on them, all recommend standards for praise and blame that take actions and their effects into account (even as intentions too are considered).¹⁴ In effect, then, having first set out a ban on moral fortune as a part of the best theory of morality, Smith offers apparently moral grounds for thinking a better morality would allow the influence of fortune in the allocation of praise and blame.

Appealing to different grounds for the same conclusion, Margaret Urban Walker argues that leaving room for moral luck is morally important.¹⁵ Specifically, she argues that acknowledging moral luck is the only way to accommodate the possibility of a robust kind of agent, an "impure agent," whose agency is ineluctably bound up with causes and effects of her actions that are often not under her control. Such an agent takes responsibility, and can be held responsible, for more than what is directly under her control. Yet allowing that there are such agents involves, on Walker's understanding, rejecting the "control condition" on responsibility that seems so strongly recommended by our best theory of morality. Reject it, though, she thinks we should, given that absent impure agency there is no room for the notable moral virtues of dependability and integrity. In effect, then, having acknowledged that the "control condition" is understandably a central aspect of our best moral theory (at least until we see that there is a morally better alternative), Walker offers moral grounds for thinking a better morality would allow the influence of fortune in determining who is responsible for what. And she takes that as grounds for thinking that in fact there is moral luck.

Using the same form of argument, albeit in support of the opposite conclusion, people regularly argue against moral luck on the grounds that holding people responsible

¹³ *ibid.*, pp. 133-134.

¹⁴ It is important to note that Smith is defending a focus on actions and their effects not just within our judicial system but across the board in our judgments of praise and blame. His reference to "every court of judicature" pretty clearly is meant to cover every context in which we judge of one another (and ourselves), and not restricted to the legal system.

¹⁵ See her "Moral Luck and the Virtues of Impure Agency," *Metaphilosophy* 22 (1991), pp. 14-27.

for what is beyond their control is unfair. The crucial claim is that a better morality – a fairer morality – is one that restricts an agent’s responsibility to what she has control over.¹⁶

Switching from moral fortune to moral dilemmas, I have offered an argument along these lines as well, holding that for any dilemma-allowing moral theory there is a non-dilemma-allowing theory (albeit not the same other theory for each) that is morally better, by the first theory’s own lights. What makes the latter theories better, I maintain, is that they avoid demanding the impossible of agents and do so without losing any of the virtues of the dilemma-allowing theories. And since every (plausible) dilemma-allowing theory is bettered by one that does not allow dilemmas, I argue, we have reason to think there are no dilemmas (even as there are wrenching choices that may well be between pressing but incommensurable demands).¹⁷

Turning to a different issue, many have maintained, against certain moral theories, that they are too demanding. The claim is not just that the theories... in requiring that people sacrifice everything for the greater good, or that they act impartially... demand a great deal of people, or that the demands are difficult, psychologically, to meet. Rather, the objection is that the demands are unjustified when and because they involve sacrificing things that weigh morally against the demands in question. Certain personal commitments, or the relations we bear to friends and family, or the extent that a meaningful life develops within, and only within, communities, people have argued, provide reason to reject any set of principles that treats these commitments, relations, or communities, as if their significance either depended on their contribution to the greater good or to acting, successfully, according to principles that are impartial.¹⁸ Williams, Nagel, Scheffler, plus Erik Zhang’s phrase: there is room for “reasonable partiality.”

Warren Quinn argues that people have rights because “it is fitting that they should.”¹⁹ In other words (I think): it is good (that is fitting) that people have rights, therefore we have reason to think that they do. Or, to put it in still another way: a moral theory according to which people have rights is a theory of a better morality, a more fitting morality, so (he infers) a better moral theory is one according to which people do have rights. Inspired by Quinn, Thomas Nagel embraces this line of argument, to the same end, defending basic moral rights that reflect our status as inviolable on the grounds that a morality that acknowledges such rights is a better morality than one that doesn’t. “The argument is that we would all be worse off if there were no rights... – ergo, there

¹⁶ Walker of course can acknowledge this cost while arguing that on balance making room for the possibility of impure agents more than compensates.

¹⁷ In a “Moral Argument Against Moral Dilemmas,” unpublished manuscript.

¹⁸ This objection has long been offered against utilitarianism and other forms of consequentialism, usually with consequentism’s standard focus on maximizing the overall good in view. But, starting with Bernard Williams, the worries about demandingness took as their target the emphasis on impartiality, shared by consequentialism and many non-consequentialist theories alike. [Talk about supererogation, satisficing, agent centered perogotives...?]

¹⁹ Morality and Action, (Cambridge: Cambridge University Press, 1991), p. 173.

are rights.”²⁰ According to him, this is because “one morality is more likely to be true than another, because the former makes for a better world than the latter — not instrumentally, but intrinsically.”²¹

Finally, to take one last example, Michael Slote defends our right to pursue our personal aims and projects rather than, and at the expense of, the public good, on the grounds that the “moral autonomy” such a right allows is intrinsically valuable. Any theory that rejects such a right – act consequentialism, for example – offers standards that are worse than those of alternatives that recognize the right. So a theory of a better morality is a theory of a morality that affords people the right to pursue their private aims rather than the public good.²² And from that he infers that a better theory of morality is a theory according to which morality affords such a right.

Our interest here, of course, is not in the soundness of any of these particular arguments but in the argument form they share. Should we think that a difference between our (currently) best theory of morality and our theory of a better morality gives us reason to revisit our understanding of morality? If so, why?

Before turning to this, a few observations are in order.

The first is that such arguments are plausible at all, it seems, only if the dimension along which one morality is supposed to be better than another is recognizably moral. That one rather than another is financially, or self-interestedly, or legally, better is no reason to think that morality is the actual one (unless, perhaps, these improvements are morally important).

Moreover, in order for the theoretical availability of a better morality to provide *us* reason to revisit our (current) best theory of morality the value at stake must be one that we can be brought to recognize. The most elegant case of this is when the value that is in play is already recognized as morally important by our (current) best theory.²³ But

²⁰ This is true, he maintains, whether or not the rights are violated.

²¹ “Personal Rights and Public Space,” *Philosophy and Public Affairs* 24 (1995), p. 92. Nagel emphasizes that the value in play here is, specifically, not instrumental. He is not arguing that people having rights has, or will have, good consequences, but rather that their having rights is intrinsically valuable.

²² See Common-Sense Morality and Consequentialism. (London: Routledge and Kegan Paul, 1985), Chapter 2.

²³ Smith’s, Walker’s, and my arguments all work with considerations – the importance of privacy, of integrity, and of fairness – that pretty much all plausible moral theories recognize as being of moral importance. Indeed, when it comes to my argument against moral dilemmas, many who defend dilemmas willingly acknowledge that they are unfair but argue that their unfairness is the essence of the tragedy they engender. In contrast, Quinn’s, Nagel’s, and Slote’s arguments, which appeal to the fittingness or intrinsic value of rights, or the inviolability or autonomy they make possible, are likely to meet opposition from the start from those who embrace various influential versions of consequentialism, though not all. Even in these cases, of course, if these things are morally valuable then there will be reason, even if consequentialists do not recognize that reason, to judge the alternative theories better, at least in certain respects.

the arguments presumably get their force (assuming they have any) not from our recognizing that one (conception of) morality is morally better, but from it being so. And any justification of the argument should, as a result, trace back to a link between the actual standards of morality and their moral credentials.

Finally, that we have a theory of a morality that is morally better than that represented by our (currently) best theory *in some respect* is compatible with that morality being morally worse in other respects. And those other respects might well outweigh (or undermine) the reason we have for thinking our (currently) best theory is wrong. When it comes to each of the arguments I have mentioned, their force depends not just on the moral value of what is at stake, but on what, if anything, of moral significance is sacrificed by the alternative morality on offer.

Still, keeping these considerations in mind, it seems as if sometimes we might well find ourselves thinking that our (currently) best theory of morality does not, after all, align with our theory of the best morality. In these cases, I am inclined to think, we will have found a reason (albeit not always a conclusive reason) to think our best theory has gotten morality wrong. What is going on?

Not Just Reflective Equilibrium

An important question is whether this elision is simply reflective equilibrium in action. Certainly, I am insisting that in cases in which the implications of a theory are morally worse than those of some other theory that is evidence against that theory. That sounds simply like reflective equilibrium. But it is worth keeping in mind that the objection I am pointing is not that there is some original considered moral judgment that the theory is getting wrong. The problem is not that it implies something we otherwise think is false, but that embracing the standards on offer lead us to unjustified judgments – because the standards we are relying on, given the alternatives, are not themselves justified.

In appreciating the difference, it is worth keeping in mind the contrast among (i) having general principles that entail more particular conclusions, (ii) having general principles that help to **explain** more particular conclusions, and (iii) having general principles that help to **justify** more particular conclusions.²⁴

Part of what I am trying to bring out is that a condition for the acceptability of a “conclusion in light of a principle,” is that the principle is itself justified (or at least, not unjustified). A test of that of course is whether its implications are acceptable. But another test is whether there is an alternative principle that is relevantly better, but the

²⁴ In thinking about (iii) I have often been tempted to think of it as an instance of (ii) where what is being explained is the goodness, rightness, or whatever of the more particular conclusions. I now think though that (iii) is importantly distinct precisely because the explanation working as a justification depends on the normative credentials of the considerations that are doing the explaining.

first principle's own lights.

If it would be good were something good, does that show that it is good?

One possible explanation, suggested by David Enoch, is that at least some moral terms express operators that respect a rule analogous to the following concerning possibility: that if it is possible that it is possible that p , then it is possible that p ($\diamond\diamond p \rightarrow \diamond p$).²⁵ The idea is that, for instance, “it is good that...” might be such that if it is good that it is good that p , then it is good that q ($GGp \rightarrow Gp$).²⁶ If so, then we could show the actual moral value of something by showing the moral value of it being morally valuable. Assuming the rule is in place for “it is good that...” we would have, if not quite an argument from wishful thinking, an argument of the form “Wouldn't it be morally good that p , for a moral p (specifically a moral p to the effect that something is morally good),” to “therefore p (that is, “therefore something is morally good.”).

It is a little tricky to represent the various arguments mentioned above as fitting this schema smoothly, if only because the target claims are about whether certain moral claims – that there is moral luck, that we face true moral dilemmas, that people have rights – are true, not whether they are or would be good. Presumably, though, these claims may be treated as roughly equivalent to claims about something else, say holding people responsible even when they have no control, or treating people as inviolable, being morally good.

If only “it is good that p ” entailed that p , in the way “A knows that p ” entails p – that is, if only value claims were factive – then $GGp \rightarrow Gp$ would follow straight-away; just as Gp would entail p , so too would GGp entail Gp . But normative claims, including those concerning what is good, or justified, or obligatory, are not in this sense factive. That p is good, or justified, or obligatory (or right or virtuous or ...) is, sadly, perfectly compatible with $\sim p$.

In this respect, at least, talk of something being good is analogous to talk of something being possible – no grounds at all for thinking what is good or possible is actual. So the question is whether there is some explanation of why $GGp \rightarrow Gp$ even though $\sim(Gp \rightarrow p)$. As with possibility, concerning which it is similarly true that $\sim(\diamond p \rightarrow p)$, some account is needed for why the relevant operator, iterated, entails its application.

Yet when it comes to “it is good that...” and other moral operators, there is reason to doubt any such explanation is available, because there is reason to doubt the entailment. To take just one case, consider the apparent fact that people are not only less than perfect, some people are not morally good at all. Considering just those people, it seems clear that it would be morally good if they were morally good. But that neither entails nor even gives us any reason to believe that they are morally good after all. So

²⁵ “Wouldn't It Be Nice If p , Therefore, p (for a moral p),” *Utilitas*, vol. 21, no. 2, 2009, p. 222-224, p. 223.

²⁶ At least if we hold fixed that it is moral goodness that is in play.

GGp (where Gp=the people in question are morally good) is true, though patently Gp is not. Other cases, it seems, are easily generated.²⁷ That *it would be good for something or someone to be good* does not entail (or even reason to think, so it seems) *that it or she is good*.

This version of a counter example doesn't really work since 'good' is not being used consistently as an operator. Unfortunately, in English, it is a bit hard to get an example going that sounds grammatical. Though I am hopeful that I can. If we switch to talk of 'ought' (and the suggestion that it iterates) I think the counter-examples to O(Op) → Op (where "O" stands for "It ought to be that...") are easier to generate. For instance, suppose that it ought to be that it ought to be that Alex visits Barbara, say because (for some reason) it ought to be that Alex promises Barbara that he will visit her (in a context where making that promise will make it true that it ought to be that he keeps it). In such a case we cannot reasonably conclude that he ought to visit Barbara since he may well have not made the promise and it might be that whatever reason there was to make the promise was not, independent of the promise, enough to establish that it ought to be that Alex visits Barbara.

In whichever way these cases work out, there are two things to keep in mind. One is that if (what we might call) the self-establishment thesis is true of some moral operators – so that the goodness of something being good establishes that it is good, or that something's being such that it is obligatory that it is obligatory establishes that it is obligatory – we need an explanation of *why* these operators work in this way. The other is that any such explanation needs to take account of the fact that the form of argument seems to go through not just in cases where it would be good (or it ought to be) that certain things are good (or ought to be), but also in cases where it is good (or it ought to be) that certain things are bad (or forbidden).²⁸

It is worth highlighting here that the better world argument is different from the one I offer. What is at stake is not the world being better (or worse), instrumentally or

²⁷ See Rob van Someren Greve's "Wishful Thinking In Moral Theorizing: Comment on Enoch," *Utilitas*, vol. 23, no. 4 (2011), pp. 447-450, for arguments along these lines. There are epi-cycles we might go through here, say by wondering whether any alteration for the better in one person's character will be more than off-set by a moral loss somewhere else. But this would just be a distraction, since the problem with the inference doesn't turn on the particular facts but with the intelligibility of the case. Van Someren Greve considers the possibility that if we limit *p* to something that is itself a moral claim, say "that human happiness is good," a version of the inference from GGp to Gp might go through. As van Someren Greve points out, it is not at all clear that the restriction would facilitate the inference, and even if it did, it would leave us with a claim about a moral claim – that it is or would be good if human happiness were good, but not with the moral claim that we swapped in for *p*: that human happiness is good. See p. 450.

²⁸ There seems to be no easy way to define what is bad in terms of something else being good. This means it would be hard for the self-establishment thesis to play a role in explaining an argument that compared two theories of what is bad and argued in favor of one, against the other, on the grounds that it was a theory of a (morally) better morality. Things are easier if the theories are of what is forbidden and one is advanced on the grounds that it morally ought to be that certain things are forbidden, since the claim that things ought to be forbidden can plausibly be translated into claims that their negations ought to be.

otherwise, but of morality being better (or worse) – specifically of it being more (rather than less) justified by its own lights.

An alternative...

In any case, I think the truth lies closer to, though only closer -- not in -- a defense of the inference running the other direction, from something being good (or right or obligatory or justified) to it being good that it is (i.e., $Gp \rightarrow GGp$). The (slightly misleading) idea is that there must be something good about – some justification of – something counting as good if it is actually good. Thus, again a bit misleadingly: GGp is a necessary condition, not a sufficient condition, for Gp .

If I am right, in each of the cases mentioned above, the force of the argument that one morality is morally better than another is that it gives us reason to suspect that the original theory (which we initially took to be the best theory of morality) is of a morality that is not, given the alternative(s), one with which we have reason to comply.²⁹ The arguments work by calling into question a credential that the actual standards of morality must have: justifiably making a claim on our allegiance.

One way to see the force of the demand that the standards of morality stand as themselves morally justified, is to consider a familiar dialectic concerning dispositional accounts of value.³⁰ According to such accounts, what it is for something to be valuable is for it to be such that it would garner approval (or some other positive response) from someone appropriately situated. Such accounts face a familiar challenge: that to be plausible, whatever garners the requisite approval must merit the approval it garners.³¹ According to many people, this challenge reveals that such accounts cannot in the end be adequate since they will need (these people suppose) a distinct standard of value to determine which things merit (and which do not) the approval they might garner.

This is too quick, though, since nothing precludes dispositional accounts redeploying the standard on offer (as set by the reactions of the impartial spectator) to determine whether what garners approval merits that approval; the test is simply whether their garnering approval would itself garner approval. If so, these views hold that the objects that would garner approval merit it – that it is good that they would receive it – and if not, they hold that the objects do not, after all, merit the approval they would garner. This is not a trivial test, since people, properly situated, might well disapprove of their own pattern of approval. Nonetheless, it is also a test that might well be passed and

²⁹ A bit more carefully. What the arguments purport to show that in certain respects the standards on offer by the putatively best theory are one's that on balance, given the alternatives, we do not have reason to comply with. We may well have reason to comply with many of the standards under most circumstances. To the extent this is so, the theory of a better morality will preserve those aspects of the original theory.

³⁰ Variations on this theme were advanced by Francis Hutcheson, David Hume, and Adam Smith, and more recently by Roderick Firth, Michael Smith, and David Lewis, among others. Include cites....

³¹ John McDowell has made this point especially forcefully in "Values and Secondary Qualities, in *Morality and Objectivity: A Tribute to J. L. Mackie*" ed. by Ted Honderich (Routledge, 1985), pp. 110-129.

if it is, those who embrace the account will be in a position to affirm that what garners the approval of those properly situation merits the approval garnered.

Thus, recognition of the requirement that, on such accounts of value, what garners approval must merit the approval does not mark a decisive objection to them. Yet it is a requirement that must be passed if the theory is to be plausible, and it is worth noting why. If a theory of value did not pass, it would mean that – by the theory’s own lights – there was nothing good about the standard it offered for distinguishing what was good from what not.³²

Significantly, in accepting the requirement there is no assumption that passing it is grounds for thinking a standard is the right one (so no commitment to thinking that $GGp \rightarrow Gp$). Rather, there is the recognition that if a standard does not pass, it should be rejected (since $Gp \rightarrow GGp$, and failing the test is a matter of discovering that, on the account in questions, $\sim Gp$).

The same line of argument, it is worth noticing, finds a home in discussions concerning theories of practical and theoretical reason. Thus many argue that the standard maximizing theory of practical reason (according to which the rational choice is the one that maximizes expected utility) needs to be revised in light of its own focus on optimality as a measure of standards of reasoning, since maximizing expected utility predictably results in suboptimal outcomes in various game theoretic situations.³³ Similarly, many people hold that an acceptable theory of induction needs to endorse itself as a measure of what to expect – on pain, otherwise, of judging its own measure of rational expectation irrational.³⁴ In much the same way, others have argued that an acceptable rule for determining (partial) rational credence must by your own lights vindicate your actual credences.³⁵ In each of these cases, the argument is that an acceptable normative theory in the relevant domain must be one that is not relevantly bettered, by its own standards, by an alternative. Again, meeting this requirement does

³² Adam Smith criticizes Hutcheson for failing to recognize the importance of this requirement. See The Theory of Moral Sentiments, p. 323. Hume meanwhile makes a point of emphasizing that “not only virtue must be approv’d of, but also the sense of virtue: And not only that sense, but also the principles from whence it is derived” (in The Treatise of Human Nature, p. 619).

³³ See, for example, David Gauthier’s defense of constrained maximization. Cite.

³⁴ See David Lewis on the requirement of immodesty in “Immodest Inductive Methods,” *Philosophy of Science*, vol. 38, no. 1 (1971), pp. 54-63. His conclusion: “immodesty — in the class of otherwise adequate methods, under an appropriate accuracy-measure, on the total evidence — is a necessary condition of inductive adequacy” p. 63.

³⁵ See James Joyce’s “A Nonpragmatic Vindication of Probabilism,” *Philosophy of Science*, vol. 65, no. 4 (1998), pp. 575-603. He argues that only standards that satisfy the laws of probability theory will meet this requirement and takes this as a vindication of probabilism. J. Predd, R. Seiringer, E. Leib, D. Osherson, H. V. Poor, and S. Kulkarni, in “Probabilistic Coherence and Proper Scoring Rules” offer a proof that “No matter what proper scoring rule is in force, any incoherent forecast [that is, any forecast that is inconsistent with the probability calculus] can be replaced by a coherent one whose [epistemic] penalty is lower in every possible circumstance; there is no such replacement for a coherent forecast” p. 3. The quasi-dominance strategy in play here is analogous to the one I use in my argument against moral dilemmas.

not establish a normative theory as right, but failing to meet it rules a theory out as mistaken.

Tentative Conclusion

The fact that the general argument reaches beyond morality recommends finding an explanation for it that does so as well.³⁶ My suggestion is that the key idea is that our normative concepts (whether moral or not) are concepts that make a claim on our allegiance, that purport to be such that – and here I purposely use a generic characterization – their satisfaction *ipso facto* means someone or other has reason to act or react in certain ways to what satisfies them. Just who has the reasons, and what they have reason to do, varies dramatically according to the concepts in play. Thus our concept of duty, when satisfied by an action available to someone, might imply that the person has decisive reason to perform it, while our concept of value, when applied to some object, might imply that someone (just who, might vary) has some reason to cherish, or promote, or protect it, and our concept of evidence, when applied to some consideration, might imply that people have some reason to believe what it supports.³⁷

This general feature of our normative concepts crucially distinguishes them from non-normative concepts. Non-normative concepts may often be such that their satisfaction in fact provides reason for someone or other to act or react in some way. But such concepts do not purport to be such that their satisfaction *ipso facto* provides reasons. The reasons their satisfaction provides, if indeed there are such reasons, depend on other (conceptually, though not necessarily metaphysically) distinct features of the situation.³⁸ Thus whether we have reason to stop at a particular red light depends on more than whether it is red, for instance on whether we have a desire to conform to the law, or a concern not to get a ticket, or... whereas our having (some) reason to perform a particular action is settled by whether it is valuable to do (or so our concept of value implies).³⁹

³⁶ Ryan Preston-Roedder, in “A Better World,” *Philosophical Studies*, 168 (2014): 629-644, offers a general characterization of several of the moral examples I mentioned above, with special focus on Slote’s defense of moral autonomy and my argument against moral dilemmas. He argues that they are all instances of, what he calls, a “a better world argument.” This is an argument that apparently applies forcefully within morality in light of morality having as a constitutive aim making the world morally better. The idea is that in denying something of intrinsic moral worth a theory is effectively self-defeating. Whether this strategy might be extended successfully to cases outside of morality is a bit unclear. To the extent it does generalize, though, I suspect it is explained by the distinctive feature of normative concepts I discuss below.

³⁷ There is of course plenty of room for argument about which sorts of reasons, for which sorts of actions and reactions, of which people, distinguish one normative concept from another.

³⁸ It may well be that the very thing picked out by a non-normative concept – say the concept of PAIN – also always, and perhaps necessarily, satisfies a normative concept (e.g. BAD). Many think this is true of this example, though I do not.

³⁹ Needless to say, while (if I am right) normative concepts are such that their satisfaction establishes reason for someone to act or react, particular normative concepts may prove bankrupt, in ways that meant they are not actually ever satisfied. This is the view one might well take of, for instance, chastity in the

Moreover, this feature of our normative concepts means that competence with them carries a distinctive burden: that we be sensitive to whether the criteria we rely on in applying them are such that satisfying the criteria *ipso facto* establishes that someone or other has reason to act or react in a certain way. Evidence to the contrary – evidence that satisfying the criteria we are relying on does not *ipso facto* establish reasons – constitutes a direct challenge to the criteria. Such evidence gives us grounds for revising our criteria. That is what is going on, I think, when we discover a theory of a better morality, or a theory of better practical reasons, or a theory of a better principle of induction, or a theory of a better rule for determining what we should believe in situations of uncertainty. In each case, that the alternative is better provides grounds for thinking our original criteria (as captured by our current best theory of the domain) do not actually establish that we have the relevant reasons, at least when those standards conflict with those identified by the theory of a better set of criteria.

Another way to look at, what I take to be, the same phenomena, is to see that competence with our normative concepts, which purport to establish reasons, commits us to seeing arguments that they do not as a challenge to our understanding of the concepts in question.⁴⁰ We in effect are committed to seeing the claims of our normative concepts as justified. Exactly what this involves is a bit hard to say. There are two ways we might think of the situation. On one, and this is how I have just put things, we are committed to seeing the standards we embrace (as criteria for our normative concepts) as themselves justified. On the other, and this is how I put things earlier, we are committed to holding that our normative concepts are such that their satisfaction *ipso facto* establishes reasons (for someone or other do act or react in some way to what satisfies them). In any case, I suspect these are actually deeply connected: we have relevant reasons to act or react in certain ways to the satisfaction of a concept only if the standard it sets is itself a justified standard for determining what reasons we have.⁴¹

eighteenth century. On this view, either chastity is a non-normative concept and the mere fact that some act is chaste gives no one reason to do act or react in any particular way, or else it is a normative concept the satisfaction of which requires that someone has such a reason, and no one does (so nothing actually satisfies the concept).

⁴⁰ It is worth noting that arguments purporting to show that the fact that something is red give us no reason to do anything constitute no challenge at all to our understanding of the relevant criteria for redness – even if they might show that we shouldn't care about what is red.

⁴¹ In thinking about whether certain standards are justified we face two familiar views. what does it take for a standard to be justified. On one view: the standard is justified only if there is positive support for it. On the other view: the standard is (we might say permissively) justified as long as there is not be, on balance, reason against it. The first invites a familiar infinite regress; the second avoids that, but traffics in permissive justification. I am a devotee of the latter, for a broad range of reasons. Right now, though, the point is either conception of what it takes for a standard to be justified leaves us thinking that when there is a relevantly better alternative the standard is not justified. Alternatively, though, my point may be that our normative concepts are concepts that purport to be such that their satisfaction means that there is reason for someone or other to do (to act or react) in particular ways to what satisfies them. The challenge posed by there being a theory of a better morality is that the standards of that morality seem to be the ones the satisfaction of which provide people with reason to act and react in various ways.

February 21, 2024