



ARTICLES

Is Norcross Right about Right?

Shelly Kagan 

Yale University, New Haven, USA
Email: Shelly.Kagan@Yale.edu

Abstract

In *Morality by Degrees*, Alastair Norcross defends a view he calls “scalar consequentialism.” I argue, first, that Norcross does not use the term consistently, since in most passages this seems to refer to a version of consequentialism that rejects all claims about rightness altogether, yet in other passages Norcross claims that scalar consequentialists should nonetheless embrace his favored “contextualist” account of rightness. I also argue, second, that the particular arguments offered by Norcross as to why consequentialists should forgo more traditional consequentialist accounts of rightness (such as maximizing or satisficing) are unpersuasive.

Keywords: Norcross; scalar consequentialism; rightness

1. Scalar consequentialism

In the preface to *Morality by Degrees*,¹ Alastair Norcross tells us that while still a graduate student he suggested in a seminar that “consequentialist theories in general (and utilitarianism in particular) would be better served by jettisoning the maximizing account of rightness, and any other account, and concentrating instead on purely comparative judgments of alternatives” (v). Later in the term he learned a name (coined by Michael Slote) “for this approach.” The name was “scalar morality.” Later still, Norcross and Frances Howard-Snyder wrote a paper “arguing for the idea.” It is an idea, he tells us, that he continued to explore and develop in still other work, an idea which finally finds a book length exposition and defense here.

Given all of this, it would seem reasonable for a reader to conclude that in *Morality by Degrees* Norcross is indeed defending something he calls “scalar consequentialism” and that scalar consequentialism does without any account of rightness at all.

Well, perhaps that is a bit of an overstatement. It would surely be compatible with all of this to think that the scalar consequentialist might go on to make *negative* remarks about rightness – perhaps saying things like “no acts are right, no acts are wrong” or “the best version of consequentialism avoids ever saying of an act that it is right or that it is wrong,” and so on. To the extent that saying something like this is offering an “eliminativist” account of rightness, it is presumably fully compatible with scalar consequentialism to offer such an account. What the scalar consequentialist must not

¹Oxford, 2020. All parenthetical page numbers are references to this book.

do, I suppose, is offer what we might call a “positive” account of right and wrong, something along the lines of “an act is right if and only if such and such is the case.”

And I imagine that a genuinely scalar consequentialist account must not only avoid such (positive) talk about right and rightness; it must also avoid (positive) talk about other deontic terms like “required” and “wrong” and “permissible” and “obligatory” and “duty” and “forbidden” as well. It would, after all, hardly be in the spirit of scalar consequentialism, as described in the preface, to offer a consequentialist theory that studiously avoids offering a positive account of when an act is *right*, but which then goes on to offer such an account of when an act is “permitted” or “required” or “wrong.” (For simplicity, going forward I will drop the qualifier “positive” in discussing the types of accounts scalar consequentialism must avoid.)

Is it indeed the case that Norcross is defending scalar consequentialism, so understood, in his book? As I say, the preface certainly gives us reason to think so. And Norcross certainly regularly suggests that he is indeed defending scalar consequentialism (see, for example, 11 or 141). What’s more, he regularly says the sort of thing that we would expect a scalar consequentialist to say, for example, “consequentialism should not employ the notions of right and wrong” (48), or “a consequentialist should not judge acts to be right or wrong” (103), and “we should reject the notion that morality issues demands at all” (21).

But at other times Norcross says things that seem to me to be incompatible with scalar consequentialism as I have been describing it. For Norcross goes on to offer his own favored account of rightness, a contextualist account (122–127). I’ll give the basic idea later, but for the moment the crucial point is that, according to Norcross, if we accept this favored account (which he wants consequentialists to do), then various claims of the form “such and such act is right” or “such and such an act would be wrong” can be appropriate ones for us to make. Indeed such claims can be *true* (110, 136, 142, 147). All of which is just to say, Norcross defends an account of rightness after all, one which he is encouraging consequentialists to embrace.

So what is going on? Has Norcross actually *abandoned* scalar consequentialism? Is it just that he now believes that more familiar forms of consequentialism, such as maximizing consequentialism or satisficing consequentialism, offer inadequate accounts of rightness, and that he has what he takes to be a *superior* theory of rightness? Has Norcross given up on the scalar consequentialism of his youth?

No, that doesn’t seem to be what’s going on either. For Norcross seems to think that his favored theory of rightness is perfectly compatible with scalar consequentialism (109, 111, 136, 147). What he wants us to accept is a package that consists of both scalar consequentialism *and* his contextualist account of rightness. (I might note, in passing, that there are further contextualist accounts Norcross wants us to accept as well – of good and bad and harm – but they won’t be our concern here.)

As I hope I am making clear, I find all of this baffling. I would have thought that Norcross’s contextualist account is precisely an account of rightness (it certainly appears to be one, and Norcross certainly treats it as such), and yet that scalar consequentialism is precisely a view that “jettisons” all accounts of rightness, that is, offers *no* such account, and so, in particular, cannot be combined with the contextualist account. It looks to me for all the world as though Norcross is telling us to accept a version of scalar consequentialism which *includes* an account of right of just the sort that scalar consequentialism is defined as disallowing (or, alternatively, that he is telling us to accept such an account in *addition* to accepting scalar consequentialism, even though scalar consequentialism disavows all such accounts).

To be sure, I find it virtually impossible to believe that Norcross could be making such an elementary mistake. I have long admired Norcross's work. He normally writes with exemplary clarity and precision, and indeed much of *Morality by Degrees* displays exactly that kind of care. But it surely looks like Norcross is contradicting himself, or proposing an incoherent package of views, or some such.

Well, the more charitable thing to do, no doubt, would be to refuse to put so much weight on a few remarks from the preface, and to look to see the official definition of scalar consequentialism provided in the main body of the book. Perhaps once we have the official definition in front of us, the apparent contradiction will disappear. But unless I have overlooked it, there is no official definition of scalar consequentialism to be found later in the book. That's a bit surprising, since Norcross frequently offers crisp official statements of positions he means to criticize or to endorse. But I can't find such an official definition for scalar consequentialism itself. (And, oddly enough, "scalar consequentialism" doesn't even appear in the index!) What we can find, instead, are occasional passing remarks like this: "consequentialist ethical theories should not be interpreted as theories of either the rightness or goodness of actions, but instead as scalar theories that evaluate actions as better or worse than possible alternatives" (11). Or this: "...a scalar theory, that doesn't issue any demands at all..." (35). Or this: "scalar utilitarianism doesn't tell us what we ought to do" (42). But these of course do still seem to be more or less in line with how I have been interpreting scalar consequentialism, namely, as a view that rejects all accounts of rightness (and of related deontic properties).

Still, there are certainly other passages that do point to, or at least suggest, a different reading of scalar consequentialism, one which does not insist that we must do without any account of rightness at all. Here are a few typical passages: "consequentialists should abandon (at least at the fundamental level) the notions of right and wrong actions" (74). And: "consequentialism is not fundamentally concerned with such staples of moral theory as rightness, duty, permissibility, obligation, moral requirements..." (108). And: "a consequentialist (at least) has no room for the relevant notions *at the fundamental level* of her theory" (110, emphasis in original). (See also 12, 27, 36, 46, 147, 151.)

The crucial thing to notice about the passages just quoted is the presence of the term "fundamental." The thought here seems to be that theories of rightness are not *essential* or *fundamental* to consequentialism. Note that if this is indeed the key idea behind scalar consequentialism, then scalar consequentialism is more modest with regard to rightness than I've made it out to be. It isn't committed to the thought that we should do without any account of rightness at all. On the contrary, scalar consequentialism now seems to be perfectly compatible with putting forward such an account and holding it to be the uniquely *correct* account of rightness (or perhaps the *best* account of rightness, from a consequentialism perspective). The only nonnegotiable point would be that we would have to agree that this account of rightness is not "fundamental" to consequentialism. Apparently, then, a favored account of rightness can be fully embraced as part of our complete moral theory, and doing this is fully compatible with embracing scalar consequentialism, as long as we agree that this account of rightness is not part of our consequentialist theory "at the fundamental level."

Very well, that seems to be progress. But now, of course, we will want to know, what *is* it to be part of the "fundamental level" of a moral theory? And here again, as far as I can see, Norcross never gives us an official account of the notion.

Luckily, however, Norcross does offer us an account of something he calls "core consequentialism." Norcross offers a series of consequentialist claims (3–5) and suggests, plausibly, that together they yield the view that

an action is morally better or worse than available alternatives, and thus there is greater or lesser (moral) reason to opt for it, entirely to the extent that the world containing it is overall better or worse (contains more or less net intrinsic value) than the worlds containing the alternatives.

Suppose we assume that in saying that no account of rightness is part of consequentialism at the *fundamental* level, Norcross is simply pointing out that his series of basic consequentialist claims do not include or entail anything at all about rightness. Then it does seem to be true that accounts of rightness are not “fundamental” to consequentialism in this sense – because of course there is no account of rightness that shows up yet in the set of views Norcross *calls* the core of consequentialism.

But if that’s *all* that Norcross is trying to say, then that would hardly be the sort of controversial claim that I think it clear Norcross is trying to make. So interpreted, after all, all that Norcross would be doing is making a trivial observation, that the particular claims he lists do not yet say or entail anything about rightness. But what of it? Perhaps there are *further* claims, including claims about rightness, that should be *added* to this initial group, and perhaps these further claims are themselves part of the essence or core of consequentialism as well!

What Norcross means, presumably, is that the claims he lists actually *exhaust* the fundamental claims made by consequentialism. Or, at the very least, even if we do decide that some additional claims also belong in the “core” of consequentialism, Norcross clearly means to be suggesting that no particular claims about *rightness* are part of the fundamental core of consequentialist theorizing.

Is that true? Actually, I think it might be a reasonably plausible thing for Norcross to suggest. After all, maximizing consequentialism (where one must do as much good as possible) and satisficing consequentialism (where one need only do “enough” good, a “satisfactory” amount) and even “eliminativist” consequentialism (which makes no claims about rightness at all) all do strike me as having a decent claim to being counted as genuinely consequentialist theories. Yet maximizing and satisficing consequentialism disagree about the nature of rightness, and eliminativist consequentialism wants to reject any account of rightness whatsoever. So if we are going to count all three of these theories as consequentialist theories – and speaking personally, I am perfectly prepared to do that – then it really does follow that no claim about rightness (not even the negative, eliminativist claim that there should be no *positive* claims about rightness) can be essential to being a version of consequentialism.²

So Norcross would be right to say that rightness and claims about rightness are not essential to consequentialism; they are not fundamental to the theory.

So suppose we interpret “scalar consequentialism” in this more minimal way, not as committed to the *rejection* of all talk of rightness, but simply committed to the thought that talk of rightness, if it is ever added, will not be part of the essence of consequentialism, it won’t be *fundamental* to the consequentialist theory.

If we do say this, then Norcross can rightly insist that we can indeed combine scalar consequentialism with a contextualist account of rightness after all, in a perfectly

²Might Norcross have put *too much* into what he thinks of as the core of consequentialism? As he recognizes, in effect, his core consequentialism rules out views like rule consequentialism or motive consequentialism as not being truly consequentialist (5 note, 34 note). I’m inclined to agree with Norcross about this – provided that we are thinking about *normative* or *factoral* consequentialist theories, rather than *foundational* theories. But the question is too complicated to pursue here, and it won’t be relevant for what follows.

coherent fashion. So long as we agree that the contextualist account isn't fundamental to consequentialism, it can still be the correct or best account of rightness, and we can *add* it to the more essential elements of consequentialism to have a fuller and more adequate moral theory.

But if we can do this, if we can combine scalar consequentialism with the contextualist account of rightness, adding that contextualist account later, not at the fundamental level, why can't we do this with a maximizing account instead, or a satisficing one? Why can't a traditional maximizing consequentialist, for example, agree that maximizing isn't a fundamental part of consequentialism – that there are no truths about rightness at the *fundamental* level of consequentialism – but insist for all that, that a maximizing account is the uniquely correct or best account of rightness, and that it should be added to the more fundamental parts of consequentialism to arrive at a fuller and more adequate moral theory? Similarly, of course, for satisficing consequentialists, who can readily agree that the satisficing account isn't part of the common *core* of consequentialism, all the while insisting that it is the *right* account of rightness and so should be added to that common core. (And similarly as well, though less expansively, for eliminativist consequentialists, who can now insist that we should add – again, at the nonfundamental level – a *rejection* of all claims about rightness whatsoever.)

Thus, if we interpret scalar consequentialism in this more minimal way – so that it merely asserts the ideas that lie in the common core of consequentialism, without taking any sort of stand at all, one way or the other, with regard to rightness – then of course it is now coherent for Norcross to propose combining scalar consequentialism and his contextualist account of rightness. But doing this robs scalar consequentialism of most (perhaps not all, but most) of its controversy. Scalar consequentialism no longer tells us to do *without* claims about rightness, it just reminds us that since consequentialists disagree among themselves about the *correct* account of rightness (and whether we should even have such an account), no such account is fundamental to consequentialism in the sense of being part of what all consequentialists have in common. And having said this, there is nothing in scalar consequentialism per se – construed in this minimal, tolerant fashion – that gives us reason to favor Norcross's contextualist account as opposed to some *other* account of rightness.

So which is it? What does Norcross mean by his talk of scalar consequentialism? Is it the minimal view, perfectly compatible with a contextualist account of right, but equally compatible with a maximizing or satisficing (or eliminativist) account? Or is it the more restrictive view, one which rejects all positive accounts of rightness whatsoever (in effect, incorporating eliminativism), and which is thus incompatible with Norcross's defense of his contextualist account? I honestly don't now.

Here are two more pieces of evidence. The subtitle of Norcross's book is "Reasons without Demands." Obviously enough, this strongly suggests that Norcross means to be defending a view which rules out all moral demands whatsoever, full stop. We must avoid talking about what is right, or what is required, or what is forbidden, or what is a duty, and so on. This clearly points to the more restrictive (eliminativist) interpretation of scalar consequentialism.

On the other hand, after having argued for scalar consequentialism, Norcross suggests that there are two positions the scalar consequentialist can in principle take with regard to rightness (109–110). They can choose to endorse the eliminativist position, rejecting all talk of rightness whatsoever, *or*, he says, they can embrace a contextualist account of rightness. I simply do not see how to make sense of this passage unless Norcross is here construing scalar consequentialism in a *minimal* fashion – so that it is

compatible with scalar consequentialism to go on to *offer* an account of rightness, albeit at the nonfundamental level. So perhaps Norcross really does have the minimal account of scalar consequentialism in mind after all. (Although if he does, it isn't clear why Norcross doesn't see that construed minimally, the scalar consequentialist has many other accounts of rightness available as well, not just the contextualist one.)

Is there perhaps some *third* plausible interpretation of scalar consequentialism, one which would allow us to reconcile all of these sundry passages? Is there an account which leaves room for combining scalar consequentialism with contextualism about right (rather than presupposing eliminativism), but which nonetheless rules out combining it with maximizing or satisficing? I imagine that Norcross *must* have some such view in mind. But I cannot figure out what that view might be.

2. Arguments

Let's put aside this exegetical question, as best we can (and important as it is), and instead look at what I take to be the main arguments that Norcross offers against including a more familiar account of rightness within our consequentialist theory. Perhaps they will persuade us to abandon all talk of rightness altogether, or at least convince us that if one *is* going to talk about rightness one should be a contextualist about it, as Norcross suggests.

Demandingness

As Norcross notes (7–10, 15–16), one common argument against maximizing consequentialism is that it is a very demanding theory indeed. If an act is right if and only if no available alternative act would result in a world with more intrinsic value, then few if any of us live morally permissible lives, and morality leaves us with little or no room to pursue our own individual well-being. We are called on to make potentially unlimited sacrifices of our own welfare, provided only that – as might well be the case – making such sacrifices results in a better world overall than the alternatives. Intuitively, however, this has seemed to many to be an implausibly demanding moral view, one requiring sacrifices that are not, in fact, morally required.

Obviously, we can neatly avoid this objection if we eliminate all talk of rightness (as well as of related deontic concepts) altogether. If we accept a version of consequentialism where there simply are no demands – no duties, no requirements, and so on – then it obviously cannot be the case that such a theory is “too” demanding. So perhaps this provides some reason to embrace an eliminativist version of consequentialism. After all, if we *limit* ourselves to saying that it is morally better to make huge sacrifices for the greater good – that there is greater moral reason to do this than to not do this – then what we are saying seems less objectionable. It was only in making the further claim that one is *required* to make such sacrifices that maximizing consequentialism seems to go too far.

On the other hand, even if this objection is thought compelling, it only gives us reason to reject a *maximizing* version of consequentialism. As Norcross recognizes, a satisficing version – where one is only required to do *enough* good – can also avoid the objection (21). Provided that the satisfactory level is set at a low enough point, satisficing consequentialism can still get away with making demands (insisting that it is wrong to do too *little* good) while avoiding the charge of being *overly* demanding. So the demandingness objection – even if we agree with it – cannot by itself provide us with reason to reject all talk of rightness whatsoever; it cannot provide reason to eschew all demands whatsoever. That will take a different argument.

It may also be worth noting that if one *does* find the demandingness objection compelling, as far as I can see that would also give us reason to reject Norcross's favored contextualist account of rightness as well. According to that account the rightness or wrongness of an act is always a matter of comparing it to a relevant alternative, where that alternative is fixed by the conversational context (122–127). (Roughly, an act is wrong on this account if the results are worse than those of the relevant alternative; right, if better than or equal to them.) And no doubt, there will be many conversational contexts for which such an account will be less demanding than a maximizing account would be. But for all that, there will be still *other* contexts where this view yields moral judgments about right and wrong that are every bit as demanding as traditional maximizing (123). That is to say, even the contextualist will hold that in certain conversational contexts it can be literally true to say of a given act that it is morally *required* – even though common intuition would nonetheless insist that the sacrifice in question is *not* actually obligatory at all. Thus, if the overdemandingness objection is a compelling one it will speak against contextualism too, and not merely against maximizing. (Yet it won't speak against satisficing, so it won't speak against all accounts of rightness.)

But *should* we find this objection compelling? I don't myself think so, and I have argued for this claim elsewhere.³ But what is especially interesting for present purposes is the fact that even Norcross doesn't seem to think that the "demandingness" objection is an especially strong one. Norcross spends several pages (16–20) arguing that we have significant grounds for being skeptical about the intuitions that lie behind the complaint. But if that's right, of course, then even maximizing consequentialists may escape this first argument without too much trouble.

In short, if there is a good argument against theories of rightness (whether all of them, or all others except for contextualism), then even Norcross seems to recognize that we will have to look elsewhere to find it.

All or nothing

A second argument offered by Norcross turns on the claim that right and wrong (and related notions) as understood by traditional consequentialist theories (like maximizing and satisficing consequentialism) are "all or nothing" in the sense that they don't come in degrees (11, 14, 22, 46). An act can be right, or it can be wrong, but one act cannot be "more right" than another. Yet, as Norcross rightly insists, it is perfectly clear that we can indeed rank some acts as morally better (or morally worse) than others. For example, if you can do more good by giving your money to charity than by spending it on luxuries for yourself, then it is better to give \$1000 to charity than to give only \$500, and better to give \$500 than to give only \$100, and so on. Norcross argues that instead of thinking of morality as sorting our actions into two groups – right and wrong – we should recognize that there is actually a *scale* here, with acts being ranked as better or worse. (Hence the name for Norcross's favored view, *scalar* consequentialism, and the name for Norcross's book: *Morality by Degrees*.)

I certainly agree with Norcross, of course, that we can rank acts as better or worse and that it is important not to lose sight of this. But what I don't yet understand is why Norcross thinks this fact provides any sort of argument at all against *also* having an account of right and wrong in our moral theory. It isn't as though having one of these (say, an account of rightness) rules out having the other. Take maximizing

³In *The Limits of Morality* (Oxford, 1989).

consequentialism again, classifying an act as permissible if and only if no alternative act open to the agent would have better results. Obviously enough, such a principle does indeed sort all acts into one of two groups: those that are permissible (or right) and those that are impermissible (or wrong, or forbidden). But what of it? Nothing stops the maximizing consequentialist from adding that among the acts that it would be wrong to do, significant differences remain, with some acts being significantly worse than others. Reverting to the example I just gave, if we suppose that giving \$1000 to charity is an optimal act (so that giving less or more would actually do less good), then it might well be true that giving only \$500 or \$100 would both be wrong. But for all that, the maximizing consequentialist can certainly agree that giving only \$100 is even worse than giving only \$500. And of course they can also recognize that under ordinary circumstances it will also be true that killing someone is worse still (that is, worse than giving \$100 to charity), while killing two people is even worse than that.

Admittedly, the maximizing principle does not in and of itself generate the comparative, scalar judgments I've just offered. But again I want to ask, what of it? We've already agreed that any principle offering an account of which acts are right (or permissible, and so on) will be an addition to the common consequentialist core – the set of ideas that lay out what is fundamental and shared by all consequentialist moral theories. And it does seem as though those shared ideas suffice to generate the relevant comparative judgments. Indeed, as we saw above, the summary that Norcross labels “core consequentialism” specifically ranks acts in keeping with the amount of intrinsic value had by the worlds that contain those acts. It is clear that Norcross himself appeals to these core elements of consequentialism so as to produce the scalar judgments that he wants us to retain. But that being the case, I simply do not see why a maximizing consequentialist (or, for that matter, a satisficing consequentialist) cannot similarly avail herself of those same ideas to generate the same comparisons. If the eliminativist consequentialist can make use of core consequentialist ideas to rank actions, and if those who accept the contextualist account of rightness can make use of core consequentialist ideas to rank actions (despite the fact that contextualism also apparently generates “all or nothing” judgments of right and wrong), then it seems clear that other forms of consequentialism – including maximizing and satisficing consequentialism – can do so as well. As far as I can see, then, there is nothing in the observation that consequentialists will want to rank actions, that provides any reason at all to avoid more familiar consequentialist accounts of right and wrong.

(I might also add, though it isn't essential for the point just made, that it is far from evident to me that right and wrong and the other deontic concepts as ordinarily understood actually all *are* “all or nothing” concepts, in the way that Norcross claims. We certainly do talk of one act being more *wrong* than another, for example, and we similarly talk of one act being more strongly *required* than another. So even if turns out to be true that *some* deontic concepts cannot take comparative forms, it also seems true, for all that, that others can.⁴)

Analyses

In what may be intended to be a third argument, Norcross criticizes two particular accounts of right and wrong that have sometimes been defended by moral philosophers. One, inspired by Mill, holds (roughly) that an act is wrong if and only if it would be

⁴Does *right* allow for comparative forms? See Thomas Hurka, “More Seriously Wrong, More Importantly Right,” *Journal of the American Philosophical Association* 5 (2019): 41–58, section 5.

appropriate to impose various sanctions on the agent who performs it. Norcross argues at length that such an account is problematic, especially from a consequentialist perspective (28–34). The other account, the one I want to focus on, holds that “right” just *means* “supported by the strongest (moral) reasons” or “what we have most (moral) reason to do” (34). Norcross notes that if this were right, then the maximizing account of rightness would hold as a trivial conceptual matter, so that anyone who believes that some acts are morally superior but not obligatory must be “confused about the meaning of words” (35) – which is a highly implausible thing to claim.

But why should someone who accepts a maximizing version of consequentialism have to agree that the maximizing claim (that one is required to act in an optimal fashion) holds as a matter of the very *meaning* of the word “right” (or “required” or “permissible” and so on)? Why can’t the maximizing claim simply be put forward instead as a substantive moral *truth*? As far as I can see, maximizing consequentialists can readily agree that the word “right” (for example) doesn’t *mean* “supported by the strongest reasons” – all the while insisting, for all that, that an act *is* right if and only if it is supported by the strongest reasons. (A loose analogy: “water” doesn’t *mean* “H₂O,” but for all that, something *is* water if and only if it is H₂O.) The maximizing account of rightness can still be *correct*, even if it isn’t itself an analytic truth. So we still haven’t been given a compelling reason to abandon it. Similarly, of course, a satisficing account of rightness or a contextualist account of rightness could be correct without being analytic. Even if none of these are true as a matter of the very meaning of the term “right,” it simply would not follow that none of these accounts is *correct*.

That’s not to say that the maximizing account of rightness is *not* a conceptual truth. It might well be. Or it might be that a satisficing account or a contextualist account is a conceptual truth. All that someone who favors one of these accounts as compared to the others needs to insist upon is that if their favored account is indeed a conceptual truth, it is not a *trivial* truth, obvious to anyone at all who understands the meaning of the term “right” (or “required” and so on). Since people clearly disagree about the best account of rightness, it certainly isn’t *obvious* what the right account is; so it will take a fair bit of philosophical argumentation to settle the question. But it could still be, for all that, that there *is* a correct account and it holds as a conceptual matter. (An analogy: it is, arguably, a conceptual truth that there is no largest prime, even though the truth of this claim certainly needn’t be obvious to everyone who understands the very meaning of the term “prime.”)

Smith and Jones

In what may be a fourth argument (22–25), Norcross invites us to imagine an example where Jones is obligated to give 10% of his income to charity, and he is torn between giving 11% and 12%. Meanwhile Smith is torn between giving 9% and 10%. (I think we are also supposed to assume that Smith is also obligated to give 10%, and that Smith and Jones have the same income as well.) If I persuade Smith to give 10%, then he does the right thing. But Norcross suggests that from the consequentialist perspective there is no reason to try to persuade Smith to give their 1% more, rather than trying to persuade Jones to give 1% more instead (in which case Smith will *not* be doing the right thing). Doesn’t that show that “the supposed difference between right and wrong is not significant”? (24)

(In introducing the Smith and Jones example, Norcross claims that it shows that rightness and wrongness isn’t really an all or nothing affair (22). But I’m afraid I

don't actually see how it particularly aims to do that. What it actually seems to me to aim to show, rather, is that the difference between acting rightly and acting wrongly is not, in itself, an especially morally *significant* one. I've interpreted the passage accordingly.)

I have to confess, I find this example somewhat baffling. From the perspective of maximizing consequentialism, at any rate, if Jones is really obligated to give 10% of his income – as Norcross stipulates – then it seems to follow that if Jones gives *more* than this then he will be bringing about a *less* good state of affairs; giving 11% rather than 10% will actually make things *worse*. (Otherwise, it would have been correct to say, rather, that Jones is obligated to give either 10% *or* 11%.) And then it seems likely that giving even more, that is, giving 12% rather than 11%, will make things even worse still! (In which case, shouldn't we be trying to *dissuade* Jones from giving more, rather than trying to *persuade* him to do this?) Yet Norcross is clearly asking us to imagine that the more Jones gives (within these ranges, at least) the better: that's why it is supposed to be a matter of indifference whether we get Smith to give an extra 1% or get Jones to do it. But if Jones really can do more good by giving 12% rather than 11% (and, presumably, more good by giving 11% rather than 10%), then how can it be that Jones is obligated to give only 10%, since this is clearly not his *optimal* contribution? So whatever it is that is going on in this example, it is hard to make sense of from the perspective of maximizing consequentialism.

Might Norcross instead be thinking of the case from the standpoint of *satisficing*, rather than maximizing, consequentialism? Perhaps the more Smith or Jones give the more good they do, and any contribution of at least 10% is "good enough" to meet the satisficing requirement? But then Norcross should have said that Jones and Smith are each obligated to give *at least* 10%. And he doesn't say that either.

As I say, then, I find the example as described confusing. Perhaps, however, we can tweak Norcross's example, and still bring out the point that I think Norcross is after. Imagine that Jones's optimal contribution is actually 15%, while Smith's is only 10%. (Imagine that if either gives more than their respective percentages, the amount of good they do will start to go down.) Then maximizing consequentialism demands that Jones give 15% and that Smith give 10%. Imagine as well that in the relevant circumstances if either gives less than this, that will fail to be doing "enough good" – so that satisficing consequentialism also demands these same two contributions (15% and 10%, respectively). So both maximizing consequentialism and satisficing consequentialism require that Jones give no less than 15% and that Smith give no less than 10%.

Next, suppose that Jones is only choosing between giving 11% and giving 12% (where giving 12% does indeed do more good than giving 11%), and Smith is only choosing between giving 9% and 10%. Finally, suppose that if either gives the 1% more in question this does the *same* amount of extra good, regardless of whether it is Smith who gives it (giving 10% rather than 9%) or Jones (giving 12% rather than 11%).

In this case, if I persuade Smith then I have indeed made the difference between someone doing the right thing and their failing to do so (since Smith will not otherwise meet his obligation). In contrast, if I persuade Jones, then I have not at all made that kind of a difference (since Jones will fail to meet his obligation either way). By hypothesis, the amount of good that I will do is the same, regardless of who I persuade. So now we can ask, does the fact that if I persuade Smith rather than Jones I will have brought it about that someone does the *right* thing give me reason to try to persuade Smith rather than Jones?

Presumably not. If we assume that there is no intrinsic value per se in someone's having performed a right act (and I think it clear that Norcross wants us to assume

a value theory where this is so), then I have no stronger a consequentialist reason to try to persuade Smith than I have to try to persuade Jones. The fact that Smith would be doing the right thing, while Jones would not, gives me no consequentialist reason at all to persuade one rather than the other. So, again quoting Norcross, doesn't this show that "the supposed difference between right and wrong is not significant"?

Well, what it *does* show – given all our assumptions – is that rightness doesn't have the kind of moral significance that would direct me to make one choice rather than another in a case like this. But that should hardly surprise us. Given the common core of consequentialism, for me to have more reason to persuade Smith it would have to be the case that persuading Smith brings about an intrinsically more valuable world than the one that would obtain if I persuaded Jones instead. And we have been at pains to construct a case where that simply isn't so. (We may need to throw in still further details, for example, that it doesn't especially matter to anyone else whether more people are doing the right thing. But let's throw in those further details as well.) So Norcross is quite right (given all our assumptions) that the potential rightness of Smith's giving 10% doesn't provide *me* with a reason to persuade him rather than to persuade Jones (24).

But of course that hardly shows that the notion of rightness has no moral significance whatsoever! It should matter to *Smith* whether he is doing the right thing (just as it should matter to Jones whether *he* is doing the right thing, and to me whether *I* am doing the right thing). For only if Smith meets his obligations (and thus acts permissibly) is he avoiding making a *mistake* with regard to his actions. After all, someone who fails to act as they are required to do is failing to act on a (morally) decisive reason, and as such they are acting *incorrectly*. That, at any rate, is how it will look to those of us who want to make the relevant deontic judgments.

It is not, after all, a mere coincidence that we use the same general language of right and wrong – a set of terms that come into play whenever there is the possibility of error – in moral contexts too. To act wrongly isn't merely to do something *worse* than some alternative; it is to act in a way that morality rejects as unacceptable and mistaken. So if there is the possibility of *erring* in our action – and nothing in this example gives us reason to think otherwise – then it seems to me significant for each of us to note the acts that avoid such errors, as well as flagging the acts that would involve them. And this is precisely what the language of right and wrong (as well as permissible, required, forbidden, and so on) does.

All of this is connected, obviously, to the action guiding function of morality: we want our moral theory to guide us with regard to how to act correctly. Identifying certain acts as right, and other acts as wrong, provides the desired form of guidance. To be sure, the comparative scalar judgments that Norcross emphasizes also provide guidance of a sort, marking some acts as better or worse than others. But only theories that make use of deontic concepts – like right, wrong, permitted, forbidden, and so on – have the conceptual resources to steer us clear of *error*.

Acknowledging the importance of having this form of guidance still leaves room, of course, for disagreement concerning which acts are required, which forbidden. The maximizing consequentialist thinks that only optimal acts avoid all error whatsoever; the satisficing consequentialist thinks instead that it takes more than mere nonoptimality to render an act morally mistaken or unacceptable. The contextualist thinks that what counts as an unacceptable act depends on the conversational context. These are interesting and important disagreements. But nothing in the Smith and Jones example gives us reason to think we are better served by doing without such judgments altogether.

Nor, as far as I can see, does the example give us reason to prefer the contextualist account of rightness over alternative accounts. To be sure, Norcross claims that if we settle on any given account of rightness (I assume he means any fixed account, rather than a contextually sensitive account of the kind he favors) the choice of which acts to count as required will be “arbitrary” (23). But it seems to me that this fails to take the idea of morally wrong acts as being mistaken sufficiently seriously. There is certainly room for disagreement concerning which acts, if any, are indeed mistaken and unacceptable from the moral point of view. But that’s a far cry from showing that any particular answer (other than leaving the answer dependent on context, or refusing to mark any acts at all as unacceptable) must be an *arbitrary* one.

No extra reason

There is one final line of thought that I want to mention. At one point Norcross seems to suggest in passing that the fact that a given act is right doesn’t provide any *extra* reason to do it (24). As we might express the thought: there can be any number of reasons to do a given act – it might promote someone’s welfare, for example, or show respect for their autonomy, or it might involve keeping a promise, and so on – but the very fact that the act is morally *required* is not itself one of the reasons to do it. (Norcross might not agree with all the reasons on this list, but I’m just illustrating the idea.) Extraordinary circumstances aside, if one were naming all of the reasons one had for performing a given act, one would never conclude “and in addition, here’s one more reason to do it: it’s obligatory.”

Obviously, this observation is similar to the one that lay behind the argument based on the Smith and Jones example (and it occurs in the midst of the discussion of that case). But the thought there was that the mere fact that an act is the right act for someone else to perform won’t normally provide *me* with a reason to help bring it about that the person in question actually performs that act. That is, the potential rightness of someone else’s act is not itself something that affects the balance of reasons that *I* have. Here, however, Norcross seems to be going further, suggesting that, strictly speaking, the fact that a given act is the right one for someone to perform doesn’t even provide *that very person* with a reason to do it. Putting the same point in the first person: the mere fact that I have a duty to do something doesn’t actually provide me with a reason to do it! There may well be significant and weighty reasons to do a given act, but its being required doesn’t seem to be among them; indeed, it doesn’t seem to be any reason at all. And shouldn’t *that* fact give us pause about any claim to the effect that being required (or right, or forbidden, and so on) is a significant moral property?

I am not at all sure whether Norcross really does intend this to be an independent argument for jettisoning talk of rightness (or for making do with a “mere” contextualist account). But I presume that some might take it to be one, so let’s consider it a bit further.

The first question, I suppose, is whether it really is true that the fact that a given act is obligatory provides no further reason for doing it. As it happens, on some metaethical accounts of the nature of obligation that claim might be denied. (On some contractarian accounts, for example, the very fact that some act would be required by a principle we could all agree to is itself one further reason for doing that act.) But if we are limiting ourselves to views that accept the common core of consequentialism as Norcross describes it – according to which the only reasons for performing an act are the intrinsic value of the results – and if we assume, once again, that the performance of a right

action need not have any intrinsic value in and of *itself*, then it really does seem to follow that the obligatoriness of an action is not, in itself, a reason for performing it.

But even if it is true that from the consequentialist perspective an act's having the property of being required is not in itself a reason to perform that act, it doesn't follow that the property is not a morally significant one. Knowing that an act is required may not give you an *extra* reason to do it, but it does tell you something crucially important about the *nature* of the reasons you have for doing it. In particular, being morally required is a property that an act has when the various reasons to perform that act are jointly (morally) decisive, settling the question of how someone who perfectly conforms to morality would act. After all, an act can have the property of being *supported* by reasons without those reasons *decisively* supporting it. But when the act is *required*, in contrast, the reasons that support the act *settle* the matter of what one should do. (Similarly, of course, if an act is forbidden, then the reasons that *oppose* doing that act are morally decisive; and if it is permissible, then there is no morally decisive reason to *refrain* from doing it.)

As usual, acknowledging this point does not resolve the question of *which* acts, in fact, are required. Is maximizing consequentialism correct when it claims that an act is required whenever all the alternatives would have worse results? Or is satisficing consequentialism correct when it claims, instead, that even a uniquely optimal act need not be obligatory (if some alternatives, though nonoptimific, are nonetheless good enough)? Indeed, even an advocate of a contextualist account of rightness could accept the thought that when an act is required, there is a decisive reason to do it; they need only insist that *whether* a reason is decisive or not is relative to a given conversational context.

Choosing *among* these rival views about rightness requires further argument. But unless I am overlooking it, nothing in the observation that requiredness is not in itself an *extra* reason gives us any reason to conclude that the concept of requiredness is an *unimportant* one (or one that can only plausibly be adjoined to a contextualist account).

Is Norcross perhaps skeptical about the very *idea* of a reason being decisive – settling the matter of how one should act, or think, or feel? I'm not sure. But in any event, none of the arguments we've surveyed give us reason to *embrace* such skepticism.

So where then are we left, after all of these arguments? Is there compelling reason for the consequentialist to abandon the traditional quest to identify right acts and distinguish them from wrong ones? Should we “jettison” the search for a general criterion of rightness (whether the maximizing one, a satisficing one, or some alternative)? Not as far as I can see. Obviously, it remains a controversial question just what account of rightness is the correct one. But despite everything that Norcross has said, it seems to me that we were right all along to be trying to figure out the answer.⁵

⁵I am grateful to Walter Sinnott-Armstrong for inviting me to contribute to this symposium and for suggesting the title for my essay.