Norcross, Alastair. *Morality by Degrees: Reasons without Demands*.

Oxford: Oxford University Press, 2020. Pp. 176. $50.00 (cloth).

Ethical theories have traditionally been understood as theories of right action. In *Morality by Degrees*, Alastair Norcross urges Consequentialists to reconceive of their ethical theorizing. Consequentialists care about promoting *value*, something that comes in degrees. Accordingly, it makes little sense to try to twist the view into centrally concerning an all-or-nothing property of *rightness*; we may even do better to banish the latter from our fundamental theory altogether. The resulting view – *Scalar Consequentialism* – could plausibly dethrone Maximizing Act Utilitarianism as the paradigmatic form for consequentialism to take. In this short and very accessible book, Norcross draws together his previous papers on scalar morality into a unified whole. His provocative arguments are essential reading for anyone interested in the Consequentialist tradition.

The introductory chapter nicely explains how Norcross conceives of the heart of consequentialism, as captured by the following principle:

> **"Core Consequentialism (CC):** 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 intrinsic value) than the worlds containing the alternatives." (5)

CC is both *comparative* and *scalar*. It speaks of what is better or worse, rather than what is outright good or bad. These evaluations are scalar, admitting of degrees, rather than involving any "all-or-nothing" binary verdicts.

A wide range of consequentialists can agree with CC. Maximizers and satisficers, for example, should accept CC even as they supplement it with further fundamental principles to establish a boundary between right and wrong. What's distinctive about Scalar Consequentialism is the claim that CC exhausts fundamental ethics. Norcross can still speak with the vulgar, given the contextualism about 'right'/'wrong' (and other binaries) defended in chapter 5, but this is mere verbiage: there is no normative property of rightness, on his view. At its base, ethics furnishes us (only) with reasons, not with demands.

Why think that there is no property of rightness? Chapter 2 presents Norcross' central argument: the distinction between right and wrong is supposed to be *significant*, but from a utilitarian perspective, there would seem no basis for judging some (equal-sized) increments in value to be more significant than others. For example, if we're obliged to give 10% of our income to charity, then increasing from 9% to 10% must somehow be more significant than increasing from 11% to 12%, but no utilitarian should think that! The purest form of utilitarianism would posit no such special thresholds, instead simply specifying (as CC does) that the more good we do, the better.

While Norcross contrasts his view with Maximizing Act Utilitarianism, I find it tricky to pin down substantive points of disagreement between the two views. Norcross agrees, of course, that to maximize the good would be best of all (and what we have *most* moral reason to do). He prefers to frame this as "an ideal to which we aspire" rather than "as a standard that *must* be met (perhaps to avoid censure)" (46). But then, no sensible maximizer advocates censuring people merely for falling short of perfect optimality, so that can't really be what's at issue between them. My sense is that the two views are best understood as verbal variants. Our choice between them may then come down to which way of talking better reflects the underlying normative structure, and here Norcross seems on strong ground. It's an important fact about utilitarianism that there's no *extra* significance or reason (out of proportion to the additional value thereby realized) to do what's right or best, and the scalar view

better highlights this structural feature. Norcross can thus be understood as offering a better way to describe what Maximizing Utilitarians were trying to get at all along.

Chapters 3 – 4 argue for the striking conclusion that we should also be skeptical of evaluating actions as 'good', 'bad', or 'harmful', or even of comparing the values of actions that occur at different times—as it turns out there is no principled way to determine *the* value of an action, except in comparison to its various alternatives. The final chapters (5 – 6) then demonstrate how a contextualist analysis of the banished moral terms allows Scalar Consequentialists to accommodate such discourse on the cheap.

It's natural to assume that consequentialists should judge an action good (bad) just to the extent that it makes things go better (worse) than if the act hadn't been performed. But consider *Button Pusher* (62): faced with ten buttons labelled '0' – '9', Agent is told that pressing a numbered button will result in that number of people being killed. If no button is pressed within thirty seconds, ten people will be killed. Our natural account implies that Agent's pushing '9' would be a good act (at least if they are otherwise disposed to do nothing), which seems an inapt evaluation of gratuitously killing nine people. It may be better than letting all ten be killed, but it surely isn't a *good* act given that Agent could just as easily have saved everyone.

On the other hand, sometimes saving just one out of ten people does seem positively good, e.g. if each life saved requires a separate sprint into a burning building (77). So we can't just call anything suboptimal 'bad'. Similar variation is found in our use of 'harm': intuitively, Agent harms nine people by pushing the button that kills them when he could have saved them, even though they still would have died had he done nothing. But the imperfect rescuer does not, of course, harm those that he fails to save. So it seems that whether an act is good, bad, harmful, or the like, is not something that follows simply from a neutral accounting of how the outcome compares to what would have happened

otherwise. This is a striking and important result. Which alternative we take to constitute the relevant

baseline can differ from case to case.

Norcross takes this to motivate a contextualist view on which conversational context selects a

salient alternative as the one deemed "relevant" in that context. But it's an interesting question

whether a more principled determination might yet be possible. For example, we might take the

relevant alternative to be determined by what could be *reasonably expected* of any (minimally decent)

agent. In Button Pusher, we expect any minimally decent agent to push the '0' button to save all lives

costlessly, so anything worse is outright 'bad'. In cases where greater self-sacrifice is involved, any aid at

all might strike us as 'good' in virtue of being more than is minimally expected.

This alternative account depends upon there being an objective threshold of *adequate moral*

concern: a least amount of altruistic motivation an agent must exhibit in order to qualify as minimally

decent. Norcross does not explicitly discuss such an idea, but it seems clear that he would be skeptical. It

certainly goes beyond the conceptual resources that he allows himself. But it's not clear why the rest of

us must feel so constrained.

Our contrasting expectations in the button pusher vs fire rescue cases seem to reflect genuine

normative differences between the cases, not just the arbitrary expectations embedded in

conversational contexts. Against a background where Agent is known to be villainous (such that

everyone expected him to watch all ten die), we might resignedly sigh, "Well, it's a good thing he only

killed nine people this time," as an implicit comparative claim. But I'm still inclined to insist that

costlessly saving all ten seems the *normatively privileged* alternative for determining whether the act

was *absolutely* good (warranting a distinctive kind of pro-attitude on our part, perhaps).

It's worth asking what hangs on Norcross' contextualist analyses. In contrasting his

contextualism to an error theory about the associated terms, Norcross notes that on his reductivist

account, "it is possible, even quite common, to express substantively true or false propositions

involving" these terms (110). But why care about that? Defining 'God' to mean *love*, one could express

substantively true or false propositions involving the term 'God', but they wouldn't have theological

significance. Matching ordinary usage in the assignment of truth values to linguistic strings adds further

constraints, but still doesn't seem all that philosophically significant. We should care less about the

words, I think, and more about their *inferential roles*: what *follows* from calling something good, bad, or

harmful? The answers may push us away from contextualism. If harms warrant resentment, for

example, contextualism about 'harm' would seem to saddle us with the awkward implication that

whether resentment is truly warranted could depend upon arbitrary conversational context.

The superficiality of contextualist analyses seems especially troubling when applied to free will

and determinism (chapter 6). Norcross suggests that "[e]ven if strictly speaking, an agent couldn't have

done otherwise, conversational context may select certain counterpossible alternatives as the relevant

ones with which to compare the action." (134) On the other hand, we're told that in the context of a

philosophy seminar discussing determinism, "there may be no relevant alternatives to an agent's actual

behavior." (135) But what *reasons for action* an agent has depends upon their option set. (Letting five

die could be an excellent choice if the only alternative is killing ten. Other options could render it a

terrible choice, by contrast.) So, to fix the moral reasons (avoiding relativism), we need a principled way

to determine an agent's available options. Conversational context seems ill-suited to this task.

Related objections may be raised against Norcross' rejection of moral rightness. While granting

Norcross' central insight that rightness should not be *central* to our understanding of consequentialism,

I think he is too quick to assume that we lack the conceptual resources to mark *any* significant

distinctions here.

Responses to Norcross' earlier work suggest at least two ways that the boundary between right

and wrong could mark a significant, non-arbitrary distinction, without implying that we have

disproportionate moral reason to bring about right-making increments of value. First, if we accept a

distinction between moral and nonmoral (e.g., prudential) reasons, we might take an act's deontic

status to reflect how these reasons balance out (Rob Lawlor, "The Rejection of Scalar

Consequentialism," *Utilitas* 21(1) [2009]: 100–116, 106). On this view, while some altruistic sacrifice may

be required of us, there is some point at which the personal sacrifice becomes sufficiently great that our

prudential reasons to favor ourselves trump our moral reasons to promote the impartial good.

Alternatively, we might take wrongness to be constitutively tied to *blameworthiness* (Brian McElwee,

"The Rights and Wrongs of Consequentialism," *Philosophical Studies* 151(3) [2010]: 393–412, 400). On

this view, our reasons for action could all be impartial, but not all failures are so great as to warrant

censure. I survey both options in greater depth elsewhere (R. Y. Chappell, "Deontic Pluralism and the

Right Amount of Good," in Douglas W. Portmore (ed.), *The Oxford Handbook of Consequentialism* [New

York: Oxford University Press, 2020], 498-512).

 Readers hoping to learn Norcross' thoughts on these alternative conceptions of rightness will

be disappointed that he does not consider either. (The preface notes that the book does not seek to

engage with published criticisms of his work, preferring to concisely introduce and motivate his

approach. As a result, the book will be of value primarily to readers who aren't already well-versed in

Norcross' past work on the topic.)

 There is a section (2.4) titled "Wrongness as Blameworthiness" (28), but the title is misleading:

Norcross instead addresses the Sidgwickian notion of its being *expedient to express* blame. We know

from the case of belief that whether an attitude is rationally warranted (or fitting to its object) is a

separate matter from whether it is beneficial to possess. The same is true of desire: Norcross speaks of

intrinsic value as what is "*intrinsically desirable*, or worth desiring or pursuing for its own sake" (2),

which is certainly not the same thing as what it would be most expedient to desire. Expediency invokes

the *wrong kind of reasons* for these attitudes, in contrast to the kind of (object-given) considerations

that truly merit or warrant the attitude in question. And so it goes with blame: we can distinguish the expediency of blaming someone from the question of whether they truly *merit* negative reactive attitudes (perhaps for demonstrating ill will, or counting some for less than one, in violation of utilitarian principle). Norcross decisively disposes of the suggestion that the former concept could have anything to do with the wrongness of an act. (It could conceivably be expedient to blame even someone who acts optimally, after all!) But the latter would have been more worth discussing.

(Granted, wrongness should not straightforwardly be identified with blameworthiness, as the latter depends upon the agent's motives in a way that the former does not. But more complex analyses are possible. For example, one might hold that permissibility is a matter of acting in a way that would be *compatible* with adequate moral concern. This renders the agent's actual motives suitably irrelevant to the permissibility of their action.)

So, it remains far from clear that consequentialists can accept no normatively significant conception of rightness. To pursue this debate, it seems to me, we must look beyond the ideological confines of pure utilitarianism. We need to learn whether there is any possible basis for quality-of-will theorists to (non-arbitrarily) specify the boundary between adequate and inadequate moral concern (compatibly with a utilitarian account of our moral reasons for action). Or whether the best theory of practical reason is one that distinguishes moral and non-moral reasons, and if so, what their relative weights might be. Depending on the answers to these questions, there may be room for consequentialists to develop their views in different directions from what Norcross recommends.

Morality by Degrees is an enjoyable read: short, breezy, and often humorous. It offers an accessible presentation of an important philosophical perspective, one that arguably deserves to be regarded as the paradigmatic consequentialist perspective. Even for those of us inclined to bristle against its implicit conceptual constraints, the book contains many important lessons. The central

insight, that consequentialism should not be understood as primarily concerned with rightness, is both compelling and strikingly revisionary. Norcross' arguments against conventional assumptions about 'good', 'bad', and 'harm' are similarly illuminating. Readers disinclined to accept all the details of Norcross' view will be provoked to seek alternatives—hopefully, alternatives more promising than the initial views we might otherwise default to, which this book so decisively refutes. The book is well-suited to assign as part of a graduate seminar on ethical theory. It raises many important philosophical questions, and I look forward to it spurring further discussion on the key points of dispute.

Richard Yetter Chappell

University of Miami