



# Right and Wrong: Assessing Scalar Consequentialism

Brian McElwee<sup>1</sup>

Accepted: 2 November 2022  
© The Author(s) 2022

## Abstract

Demoralising ethical theory involves eschewing the deontic categories of moral obligation, moral permissibility, and moral impermissibility from our ethical thought. In this paper, I evaluate the case made in Alastair Norcross's recent book, *Morality By Degrees* (2020), for a consequentialist version of such demoralisation. Norcross defends scalar consequentialism, a radical variant of consequentialism which restricts fundamental normative verdicts to a scalar ranking of available actions, ordered according to the goodness of the consequences they produce. Following an introductory Sect. 1, I assess the positive case for scalar consequentialism in Sect. 2, concluding that no strong case has been made for the view. In Sect. 3, I assess the case against the view, concluding that while scalar consequentialism may be able to avoid *the action-guidingness objection*, it falls foul of *the force objection*. In Sect. 4, I expand on this critique, showing that Norcross gives an unstable account of how to assess attitudes, such as desires, beliefs and emotions. In Sect. 5, I argue that appeal to a contextualist reductionism does little to make the scalar view appealing.

## 1 Introduction

What is morality? Here are two sorts of answers.

First, an abstract answer. Morality is a particular set of normative truths. It is a task for the moral philosopher to pick out these normative truths, or at least the fundamental ones. We might characterise morality as, in essence, a single fundamental normative truth such as the Categorical Imperative or the claim that we ought to maximise happiness. Or we might characterise morality as a *set* of principles, irreducible to each other: a duty of beneficence, a duty of reparation, a duty of non-maleficence, etc.

Alternatively, we might give a second, more concrete answer. Morality is something we find in the world—a social-psychological phenomenon. Morality consists in beliefs, actions, and feelings which shape our behaviour, reining in anti-social tendencies and promoting cooperation, harmony and flourishing. (Or at least it does this when it functions well.) One part of morality is psychological: an agent acts in ways influenced by morality, when he shrinks, out of moral inhibition, from attacking, stealing from, or cheating others.

<sup>1</sup> University of Southampton, Southampton, UK

Another part is social: people respond to some anti-social behaviour (attacking, stealing, cheating, etc.) with angry protest, calling out the perpetrator's actions as unacceptable. Such protests may serve to shape future behaviour, of the initial miscreant and the wider community, by instilling or strengthening within them moral emotional inhibitions.

Both these answers are quite acceptable ones. Indeed, we might think that unless we give both answers, we leave out something important in characterising morality.

If we ignore the first, we clearly say too little, limiting ourselves to empirical description of human behaviour. This would be inadequate, since there are not just empirical truths concerning morality (sociological or anthropological truths about what people think, feel, and do under the guise of morality). There are also normative truths, about what people *should* think, feel, and do.

But if we ignore the second, we are similarly in danger of saying too little. If the moral philosopher's task is to articulate those fundamental normative truths that are *the moral truths*, but we fail to attend to the real-world practice of morality, and all that it involves—notably, emotional responses—we may err by leaving out some of the fundamental moral truths.

Some philosophers, 'demoralisers', see the primary task of the moral philosopher (as opposed to the moral anthropologist) as giving an account of the fundamental ethical truths. But on their view, the fundamental moral truths are fewer than commonly thought. They are simply the truths about how good, or how reason-supported, actions are, relative to alternative actions. And that's it. Once we have ranked available actions, our work is done, and classification of actions into the traditional moral categories of obligation, permissibility and impermissibility at best adds nothing crucial, and at worse misleads. It is the rejection of these moral categories that characterises the 'demoraliser'.

I argue in this paper that this demoralised picture of the moral landscape is too spare. Morality, as we find it in the world, is fuller, richer, more complex. Consequently, the demoraliser ends up saying too little. If our philosophical task is to articulate those fundamental normative truths that are the moral truths, then simply saying how good, or reason-supported actions are, relative to the alternatives, is inadequate. And it is inadequate because morality is not all about actions.

Return to the second conception of morality: morality as a practice. Central to this picture of morality are emotions: the agent's inhibitory feelings of prospective guilt and self-blame; and the community's feelings of anger, blame, resentment, in responding to anti-social behaviour—feelings expressed in acts of protest against perceived anti-social behaviour.

If we are to fully articulate the fundamental moral truths, we must evaluate not just people's actions, but also these moral feelings, saying when *they* are reason-supported.

I argue that demoralisation is in principle mistaken, because it overlooks the centrality of feelings to morality, and consequently gives an impoverished portrayal of the moral landscape. The core moral deontic concepts of moral obligation, moral permissibility and moral impermissibility essentially concern emotional responses such as blame and guilt.

My focus will be on Alastair Norcross's scalar consequentialist version of demoralisation, defended in his recent *Morality By Degrees*.<sup>1</sup> Scalar consequentialism claims that we should eschew deontic categories (the obligatory, the permissible, the impermissible) from our fundamental moral thought, and limit ourselves to a scalar ranking of available actions,

<sup>1</sup> See also Crisp 2006, Ch. 1; Sinhababu 2018.

ordered according to the goodness of their consequences. In this paper, I evaluate this radical view. In Sect. 2, I argue that no convincing positive case for scalar consequentialism has been made. In Sect. 3, I assess the case against the view, concluding that while scalar consequentialism may avoid *the action-guidingness objection*, it falls foul of *the force objection*. In Sect. 4, I expand on this critique, showing that Norcross gives an unstable account of how to assess attitudes, such as desires, beliefs and emotions. In Sect. 5, I argue that appeal to a contextualist reductionism does little to make the scalar view more appealing.

## 2 The Case for Scalar Consequentialism

First, it's important to make clear what is distinctive about scalar consequentialism. Scalar consequentialists defend a form of consequentialism characterised by four main claims, SC1-SC4:

**SC1.** Consequentialists should provide a comprehensive ranking of available actions from best to worst, as follows: the best action (or the one there's most reason to perform) is the one that leads to the best consequences; the second best action (the one there's next most reason to perform) is the one that leads to the next best consequences, etc....

A corollary of this is:

**SC2.** Consequentialists shouldn't rest content with simply triaging available actions into the deontic categories of the obligatory, the permissible, the impermissible.

**SC3.** Consequentialists can, and should, not just rank actions, but also give an account of how much better any given action is than any given alternative, as follows: if action A has slightly better consequences than action B, and has much better consequences than action C, then A is a *slightly* better action than B, and a *much* better action than C (there is *slightly* more reason to do A than B, and *much* more reason to do A than C).

**SC4.** Consequentialists should eschew (at least from the fundamental level of moral theory) the deontic categories of obligation, permissibility and impermissibility.<sup>2</sup>

Scalar consequentialists like Norcross have done moral theory a service by emphasising SC1, SC2 and SC3. However, we should note that these are not especially distinctive or radical claims for a consequentialist. They can be endorsed by many forms of consequentialism.<sup>3</sup> What *is* distinctive and radical is SC4.

In evaluating scalar consequentialism, it is crucial not to be distracted by the appeal of SC1-SC3 into endorsing the radical SC4.

### 2.1 Demandingness Objections

Norcross begins his case in favour of scalar consequentialism by noting that it avoids the demandingness objections that notoriously face the most well-known version of consequentialism:

<sup>2</sup> Compare what Norcross calls Core Consequentialism: '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.' Norcross 2020, 5. Here Norcross endorses SC1 and SC3. He argues for SC2 and SC4 at pp. 22–28, and across Chs. 1 and 2 more generally.

<sup>3</sup> Note however that standard forms of rule consequentialism reject SC1.

**Maximising Act Consequentialism (Max AC):** We are morally obliged to perform the action that brings about the best (expected) consequences. All alternative available actions are morally wrong.

With preventable suffering widespread (e.g. because of severe poverty), Max AC requires ordinary individuals in affluent countries to devote most of their money, time, and effort to decreasing the suffering of strangers. The theory says not just that typical affluent citizens currently fall short of what they're morally required to do (which is plausible), but that even someone who devotes very large amounts of her money, time and effort—but less than the most she could—acts wrongly. This seems implausibly demanding. It apparently implies that such an agent, who seems so admirable, in fact merits moral blame (given the close link between moral obligation and blame, discussed below). It also seems to leave no room for supererogatory actions, those that *go beyond* what is required by moral duty.

Some consequentialists have been content to bite the bullet, and persist with Max AC. Many, however, have sought to develop a form of consequentialism which avoids such extreme demands. There are perhaps three main alternatives:

**Scalar Consequentialism**, which asserts SC1-SC4.

**Satisficing Consequentialism** We are morally obliged to bring about *good enough* consequences (where the consequences need not generally be optimal in order to count as good enough, but instead must meet a specified threshold).

**Rule Consequentialism** We are obliged to comply with the code of moral rules whose general acceptance would produce the best consequences.<sup>4</sup>

Scalar consequentialism is the most radical of these three alternatives. It avoids the demandingness objections faced by Max AC, not by developing a principled account of more limited demands, but instead by eschewing moral demands altogether! According to the scalar view, nothing is morally obligatory. This seems at least as counter-intuitive as making Max AC's extreme demands.

A caveat is that scalar consequentialists may try to soften their radicalism by insisting that they eschew moral obligations only at the level of fundamental moral theory, and that we may permissibly speak of obligations, even though there really are none; that obligation-talk can be legitimate if relativized to the context of utterance. I assess such moves in Sect. 5.

Satisficing consequentialism is much less radical. If Max AC's demands are too extreme, the obvious alternative is to make more limited demands. Satisficing consequentialists aim to do this in a way that is fully compatible with consequentialism. Why then does Norcross endorse the more radical, less intuitive view? What arguments are there for endorsing full-blown scalar consequentialism (i.e. a consequentialism that asserts not just SC1-SC3, but also SC4)?

Unfortunately, Norcross makes a straw man of his satisficing consequentialist opponent, by burdening her with the denial of SC1 and SC2:

'Satisficing versions of utilitarianism [or consequentialism generally], no less than the traditional ones, assume that the rightness of an action is an all-or-nothing property... Both forms [maximising and satisficing] share the view that a moral miss is as good as a mile.' (22).

Norcross himself doesn't cite any satisficing consequentialists. It's unfortunate that he doesn't engage with recent developments of the view, most notably Richard Yetter Chappell's 'Willpower Satisficing'.

<sup>4</sup> I lay aside assessment of rule consequentialism here. See Hooker (2001) for extended defence.

It seems obvious that satisficers would *not* claim that a moral miss is as good as a mile, and would endorse SC1 and SC2. Dividing available actions into the permissible and the impermissible does not commit one to thinking that all the permissible actions are equally good, or that all the impermissible ones are equally good. Indeed, the very characterisation of satisficing consequentialism seems to presuppose an initial comprehensive ranking of the kind scalar consequentialists propound. The satisficer claims that an agent is not morally required to perform the action at the top of the ranking (drawn up by reference to the value of the consequences of each action), but is instead required to perform one of those actions above a certain threshold in the ranking. But this natural portrayal of the satisficer's view presupposes that there is a comprehensive ranking of actions in place, prior to the drawing of the threshold!

Moreover, one of the key motivations for satisficing consequentialism is to allow room for supererogation, unlike the maximiser. But supererogation would be incoherent on the view with which Norcross burdens the satisficer. Supererogatory actions are ones that are better than some permissible alternative.<sup>5</sup> The very notion of supererogation thus *presupposes* that permissible actions can be ranked as better or worse than each other.

Once we have recognised that the most plausible versions of satisficing consequentialism endorse SC1-SC3, is there any reason to prefer the scalar view? Are there any good reasons to accept SC4?

We can separate out three strands of argument in Norcross's work.

## 2.2 The Gradability Argument

One simple argument Norcross offers is that since the properties which ground moral properties are clearly gradable, we should expect the moral properties themselves to be gradable:

'the property of an act that makes it right or wrong—how much good it produces relative to available alternatives—is naturally thought of as a matter of degree. Why, then, is rightness and wrongness not a matter of degree?... I will argue that, from the point of view of a consequentialist, actions should be evaluated purely in terms that admit of degrees.' (2020, 14).

However, as Rob Lawlor points out, 'It is not at all unusual to derive an absolute property from a gradable one... Consider the example of a being the winner of a long jump competition... Clearly, the length of jump is something that comes in degrees... But we do not conclude that the property of being the winner comes in degrees.' (102)<sup>6</sup>

The bare fact that some grounding property P is gradable doesn't then imply that any other property grounded in P must itself be gradable. Norcross's concern, however, may be that non-gradable properties are superfluous. He could insist that, even if we can acclaim some action as best, and that being best is a non-gradable property, we should not afford any special status to being the best action, aside from the status of there being proportionally

<sup>5</sup> Norcross's own definition of supererogation is misleading: 'A supererogatory act is generally characterized as an act that is not required, but which is in some way better than the alternatives.' (21) A supererogatory act need not be better than all alternatives; only than some *permissible* alternative. A supererogatory act may itself be suboptimal.

<sup>6</sup> More generally, the property of being best is clearly derived from the gradable property of goodness, but is not itself a gradable property. (Lawlor, 102) When Norcross says that actions should be evaluated purely in terms that admit of degrees, he presumably doesn't mean to eschew evaluations that recognise some action as best—yet that is, strictly, a property that does not admit of degrees.

more reason to perform it than the second best action, than the third best action, and so on. Norcross's more appealing claim is that it's odd for a consequentialist, for whom what matters is just goodness and badness, to focus on other properties of an action, in a way that isn't proportional to the marginal value of the action relative to available alternatives.<sup>7</sup>

If the scalar consequentialist's point here is simply a reiteration of SC2, reminding us that graded recommendations are important, and that some moral philosophy has had an overly heavy focus on assigning actions to the deontic categories, this is all to the good. As I'll argue below, however, rejecting the deontic categories altogether (the distinctive aspect of the scalar view) is unjustified. Graded recommendations of available actions need supplementing with the traditional moral categories of moral obligation and moral permissibility, which embody norms for how we should feel about agents' behaviour.<sup>8</sup>

Consequentialists should indeed rank *all* available actions, rather than being content to simply divide them amongst the deontic categories of obligatory, permissible, and impermissible (SC1 and SC2 again). But as we have seen, that is no reason to prefer the scalar view to the satisficing view, since satisficers will happily accept SC1 and SC2.<sup>9</sup>

### 2.3 The Persuasion Argument

A second argument is the Persuasion Argument. (See Lang 2013.)

Norcross considers a putative duty of beneficence, requiring us to provide aid to the needy. Suppose the satisficing consequentialist concludes that, for agents with income level L, giving 10% or more of their income is morally permissible, while giving less than 10% is morally impermissible. (Norcross 2020, 22) Suppose now that I'm in a position *either* to persuade Smith to increase his giving from 9 to 10%, *or* to persuade Jones to increase his giving from 11 to 12%. But I can only persuade one of them. Smith and Jones both have income level L, so their proposed duty of beneficence is the same (to give at least 10%), and the extra increment in each case will confer the same expected benefit.

A good consequentialist, Norcross says, should be indifferent between persuading Smith and persuading Jones. The amount of good produced by persuading Smith is the same as the amount of good produced by instead persuading Jones. But surely if there's the boundary between permissible and impermissible giving that the satisficer claims, it should be more

<sup>7</sup> Thanks to one of the referees for pressing these issues, and for some of the phrasing I've deployed in this paragraph.

<sup>8</sup> Cf. Chappell (2020) which argues that 'Consequentialists should be *deontic pluralists* and accept a maximizing account of the *ought of most reason*, a satisficing account of *obligation*, and a scalar account of the *weight of reasons*.' (Chappell 2020, 498).

<sup>9</sup> Another self-described scalar consequentialist, Neil Sinhababu, seems to endorse the division of actions into deontic categories. He says that right and wrong themselves come in degrees. It would be an improvement to English language, he suggests, if we used 'wronger' and 'righter'. Sinhababu notes that English already draws evaluative distinctions amongst wrong actions: we do talk of some actions as being 'more wrong than' others. Note, however, that we don't talk of some actions being more *right* than others. This asymmetry may not be a coincidence. On the sentimental view defended below, it's not. To say that an action is wrong is to say it's blameworthy (in standard circumstances). Some actions can be more wrong than others when they merit more severe blame than others. But to say that an action is right or permissible is to say that it doesn't merit blame. But not meriting blame is *not* something that comes in degrees. Either something merits blame, and the question arises of how severe that blame is, or else it does not, and no question of degree arises.

important to successfully persuade Smith rather than Jones, since then they would *both* act permissibly, rather than Jones alone?

This argument has faced forceful objections since Norcross first deployed it. Gerald Lang argues that the argument over-generalises. Norcross argues that a good consequentialist persuader will *simply not care* about permissibility facts in deciding where to focus his persuasive efforts, since *all* he cares about is the amount of good produced. Since such a good consequentialist persuader wouldn't care about permissibility facts, such facts can't have any importance.

However, Lang runs the same argument for facts about betterness. Suppose the persuader is instead in a position either to persuade Smith to increase his giving from 9 to 10% (call the resultant state of affairs P1), or to persuade Jones to increase his giving from 10 to 11% (call this P2). Jones's act in P1 is *just as good as* Smith's act in P1. Jones's act in P2 is *better than* Smith's act in P2 (or in P1 for that matter). But Norcross's line of reasoning seems to commit him to saying that a good consequentialist persuader will again *simply not care* about just-as-good-as facts or betterness facts in deciding whom to focus his persuasive efforts on, since *all* he cares about is the amount of good produced. And since such a good consequentialist persuader wouldn't care about just-as-good-as facts or about betterness facts, *they* can't have any importance.

Something has gone wrong for Norcross here, since these are the very sorts of normative facts he claims *do* have importance.<sup>10</sup> So this argument looks unsuccessful as it stands.

It is unfortunate that Norcross does not respond in *Morality by Degrees* to such criticisms of his earlier work. (As he explains in the Introduction, he opted for a punchier book laying out his view, rather than a longer one engaging with critics.)

## 2.4 The Arbitrariness Argument

A third, closely related argument is the Arbitrariness Argument:

'A related reason to reject an all-or-nothing line between right and wrong is that the choice of any point on the scale of possible options as a threshold for rightness will be *arbitrary*.' (23).

This argument does target satisficing consequentialism, and purports to give a reason to endorse SC4, the distinctive claim of scalar consequentialism.

A first response to Norcross is that he fails to examine any rationales propounded by satisficers for drawing the boundary where they do. His conclusion that any such boundary would be arbitrary is overly hasty. Chappell for example argues that we can use our best judgements about when typical agents would warrant feelings of blame as the basis for a principled, non-arbitrary distinction between permissible acts and impermissible ones. In the case of duties of beneficence, if we judge that a typical agent with an income level  $L$  giving  $x$  per cent of his income *would* warrant blame for doing so little, then we are justified in concluding that giving  $x$  per cent is impermissible for such an agent. If a typical agent with income level  $L$  giving  $x+n$  per cent of his income *would not* warrant blame (again consider

<sup>10</sup> Lang 2013. Lawlor likewise critiques the Persuasion Argument, suggesting that Norcross is confusing two issues: the issue of permissibility, and the issue of what consequentialist agents ought to prefer about other agents' actions: 'This tells us something about the consequentialist's attitude to permissibility, but it does not tell us anything about permissibility itself. Lawlor 2009, 105.

the putatively admirable agent mentioned above who acts suboptimally), then we are justified in concluding that giving  $x+n$  per cent is permissible for an agent with that income.<sup>11</sup>

This sort of account illuminates a false attraction of arbitrariness arguments against the satisficing view. It is possible that, for duties of beneficence, there is no very *precise* boundary between permissible and impermissible levels of giving. Confidently declaring that moderately affluent agents are required to give 10 per cent of their income rather than, say, 8 per cent, may have an air of random stipulation, pulling a figure out of the air.

However, all that's required to vindicate the satisficer's framework is that there be clear-cut cases of permissible giving and impermissible giving either side of what may be a fuzzy threshold. A moderately affluent person who gives 50 per cent of his income to the needy clearly acts permissibly, while one who gives a pittance acts impermissibly. The boundary may be somewhat vague, but this is no objection to the satisficing view, given that vague boundaries are everywhere.

The blame account fits well with such fuzzy boundaries. Our best judgements about blameworthiness regarding levels of beneficence may not be so fine-tuned as to pronounce unequivocally on every possible level of giving: following conscientious reflection and discussion, we may find ourselves still unsure as to whether an agent who gives 8% of their income, or 10%, when fully informed about all the relevant facts, would merit feelings of blame.

It might be suggested here that the satisficer is at a disadvantage compared to rival versions of consequentialism. Don't we especially want guidance from a moral theory when our intuitions are unclear, when we are unsure how much can be properly demanded of a reasonable agent?

Two things can be said in reply. First, I think it's a mistake to demand so much from our moral theories. The present discussion focuses on competing versions of consequentialism, but it's worth bearing in mind that moral theories of all stripes must give some account of the morality of beneficence. For any theory, there are limited options: (i) maximising: we are required to be as beneficent as we can (perhaps subject to constraints against harming others in order to do good); (ii) going scalar: denying that we have any duties of beneficence and limiting ourselves to ranking different degrees of beneficence; (iii) satisficing: we're required to help others to some degree somewhat below the most we can.

Any theorist who opts for (iii) then faces the choice between a very precise account and an account that is vague to some degree. On any approach, consequentialist or otherwise, it seems to me that a less precise account looks *more* plausible than a more precise one. Whether consequentialist, Rossian deontologist, contractualist, or virtue ethicist, it's difficult to see how a very precise figure—say 5% or 10% or 20% of one's income - is to be generated. There is no plausible methodology for setting a threshold for adequate beneficence beyond careful reflection, discussion with others, filtering out self-serving motives, making vivid to oneself the plight of the needy and so on. In going through such a process, considering progressively increasing amounts, asking oneself 'Would this amount really be enough?', it would be decidedly odd to expect that there comes a precise 'click' where one moves suddenly from a confident judgement that 'No, this wouldn't be enough' to the contrary judgement 'Yes, this would indeed be enough'. Any theory that somehow manages

<sup>11</sup> This sentimentalism about moral obligation is developed and defended in McElwee 2018 and Chappell 2019.



to conjure a precise figure out of the air looks suspect, rather than at an advantage over its rivals.

The moral theorist should not be rattled by vague boundaries, given that these are completely ubiquitous in the concepts we deploy every day, from heaps to baldness to redness. If the satisficer had to resort to a wholly *arbitrary* boundary, that would indeed be a black mark against her theory. An imprecise but non-arbitrary boundary, by contrast, looks natural.

Second, even if we did judge that there's something unsatisfying about a theory failing to specify a precise boundary (which I don't concede), the satisficer seems to me still to be in a much stronger position than his maximising and scalar rivals.<sup>12</sup> It is much better to be roughly right than precisely wrong. We emphatically are not required to give away as much of our income as would maximise overall utility—this would be enormously, and implausibly, demanding. And equally emphatically, we do, contra the scalar view, have *some* moral duty of beneficence: a typical affluent citizen who does little or nothing to help anyone in need acts wrongly.

It is notable that some of Norcross's own remarks seem to acknowledge the relevance of certain features of actions that might be used to mark a non-arbitrary boundary: the difficulty of performing the act, the degree of risk involved, the cost or sacrifice so acting would impose on the agent. (e.g. Norcross 2020, 123) (Note that appeal to such features fits well with the blame view, since difficulty, risk, degree of sacrifice are obviously relevant to whether blame is a fitting response to an agent.) If such features have some moral relevance, they can provide a basis for a principled, non-arbitrary cut-off point. For instance, the degree of self-sacrifice *adequate* to avoid meriting blame, in light of the goods at stake, guides our judgement about whether some particular level of altruistic giving is permissible or impermissible.

It seems to me then that the arbitrariness argument doesn't establish a convincing case to prefer the scalar view over a satisficing view.

In conclusion, while SC1-SC3 look plausible, no convincing case has been made for the distinctive aspect of the scalar view, namely SC4, the claim that we should *only* endorse scalar verdicts.

### 3 The Case Against Scalar Consequentialism

I will now argue that while scalar consequentialism can avoid *the action-guidingness objection*, it falls foul of *the force objection*.

#### 3.1 The Action-Guidingness Objection

One of the main objections faced by the scalar view is that it fails to guide action. (Lawlor 2009; Lang 2013) One feature of the presentation of the view that may invite such an objection is its sometime framing in terms of the betterness of actions. An agent turning to an ethical theory for guidance may expect to be told *what to do*, or *what he should do*. If instead he is merely told that action A is better than action B, which is better than action C, etc., we might wonder whether he has really got what he's looking for.

<sup>12</sup> Again, I am setting aside rule consequentialist alternatives here.

Speaking in terms of (normative) reasons for action should however be enough to assuage such worries. The scalar consequentialist who, like Norcross, is happy to speak in terms of reasons for action can issue verdicts about what there is most reason to do. Note that this is one of the main senses of the terms ‘should’ and ‘ought’.

In one important sense, scalar consequentialism may in fact do *more* by way of action guidance than some other moral theories. Scalar consequentialism provides, and insists on the importance of, *detailed* action-guidance: it gives a comprehensive ranking of all available actions, telling you how much reason there is for doing each alternative. Why might this be important? In real life, the most choiceworthy action may be unappealing: it may be costly to the agent, or to her loved ones, or it may be difficult to perform. Consequently, the agent may decide not to perform that action. Scalar consequentialism gives such an agent guidance about what’s the next best thing to do. Importantly, it does not withdraw its strongest recommendation from the action with the best consequences. It gives a ‘second-best recommendation’ to the second best action. Compare a customer choice magazine advising consumers. If the magazine only told the consumer the very best car, or hi-fi, she might feel short-changed. She may decide she’s not going to buy the very best car, or the very best hi-fi: it may be beyond her budget. Guidance on what is next best looks useful.

So in at least one important sense of action-guidance, it seems that the scalar consequentialist offers all we could want.

Some readers may have a niggling concern here that there is more to the action-guidance worry than this. They may have in mind an objection to the scalar view that I take to be much more troubling: *the force objection*.

### 3.2 The Force Objection

Scalar consequentialism gives all we could want by way of action-guidance construed in terms of *recommendations* to the agent. It gives a comprehensive scale of recommendations: the strongest recommendation to the action with the best consequences, the next strongest recommendation to the action with the next best consequences, etc.

But what is distinctive about obligation-talk is its distinctive *force*—a force that goes beyond mere recommendation. Claims of moral obligation, and claims of moral wrongness (claims that we have an obligation to refrain from some action) have a must-y force.

The force objection is that the scalar view leaves us unable to say to all that we want to say in ethics because it eschews judgements with this must-y force.

To my mind, this distinctive force of moral obligations is best understood by reference to the moral emotions: one *must* do what is morally required, or else one merits distinctive moral emotional reactions. There are different aspects to this: first-personal and third-personal, prospective and retrospective. When a third party judges that an agent has violated a moral obligation, she is disposed to feel blame towards that agent, and to judge such feelings warranted. This is no coincidence: there is a close relationship between blame and obligation. At least one thing that a judgement that some action A is morally obligatory adds to the judgement that there’s good reason to do A is that one must do A or else one merits such distinctive emotional reactions from others (absent excusing conditions).<sup>13</sup> Similar things can be said about first-personal moral judgements. When I judge that I’ve violated a

<sup>13</sup> On some views, there is a metaphysical connection: an action’s being obligatory *just is* for that action to be one whose non-performance would ordinarily merit feelings of blame. Although it is very widely held that

moral obligation, I will tend to feel guilt, and to judge that such feelings are warranted. And prospective first-personal judgements of moral obligation have a similar emotional element. When I judge that I have an obligation not to do something, I will tend to feel an emotional barrier to performing that action, and moreover to judge that such feelings are appropriate. Picture how you might feel in Singer's pond case, if you contemplated briefly walking away and leaving the child to drown. Or consider your feelings when tempted to cheat a friend or a partner.

Acknowledging a tight connection between moral obligation and blame does not commit us to denying the existence of blameless wrongdoing. An agent's particular circumstances may give them an excuse, getting them off the hook of blame, for performing a morally wrong action (one that a *typically* circumstanced agent would merit blame for performing).

The scalar consequentialist (and a demoraliser of any stripe) goes wrong in leaving out of the picture this emotional core of our moral lives, the powerful moral sentiments that give the concepts of moral obligation and moral wrongness their punchy force.

Such reactions are absolutely central to our moral practice. It is essential to our moral lives that we react in such ways, and central to ethical theory that we form judgements about the norms governing such reactions. A demoralised ethical theory that limits itself to graded recommendations of actions leaves out central aspects of our moral lives: aspects which need normative evaluation just as much as our actions themselves. A full account of moral normativity must ask not just how much reason there is for any action (compared to the available alternatives), but also what reasons there are to *feel* towards agents who act in various ways in various circumstances.

Consequentialists have, of course, often been open to radical revisions of our moral thought and practice. So might the scalar consequentialist cite our moral sentiments as just another part of 'common-sense' morality that should be jettisoned?

Several intertwined thoughts speak against this. First, given the crucial role such sentiments play in reining in anti-social behaviour, it seems highly unlikely that striving to eliminate these sentiments would itself produce good consequences. Second, it seems questionable that we human beings are *able* to stop reacting to each other with such emotions. Given how deep-seated such responses are in our psychology, aiming for wholesale suppression of moral emotions again seems unlikely to promote human flourishing. Third, aren't there just truths about when such emotional reactions are warranted? Isn't outrage simply a fitting response to the actions of Nazi Germany, to callous people-traffickers, to politicians who advance their own careers at the expense of the most vulnerable?<sup>14</sup> Consequentialists have sometimes been sceptical about fitting emotions, but recent work has shown how they can, and plausibly must, on pain of implausibility, make room for these. (Chappell 2012; McElwee 2018)

We might say, of course, that the force objection is really just another aspect to the action-guidingness objection. In some sense, I provide you with 'guidance'—guidance that goes beyond mere graded recommendation—by saying that you must perform some action. This is fine, but we shouldn't overlook that what is left out at the deepest level of analysis is in fact emotion-guidance: the demoraliser fails to tell us *how we should feel* towards someone

---

there's *some* close relationship between moral obligation and the aptness of feelings of blame and guilt, I'm neutral in this paper about the precise nature of this relationship.

<sup>14</sup> See Srinivasan (2018) for a compelling articulation of this line of thought.

who performs particular acts that had the consequences they did, in circumstances where alternative acts with different consequences were available.

Norcross himself does consider whether the centrality of blame to our moral lives might make trouble for his view. As we'll see presently, however, his discussion is led astray by focusing exclusively on blame-*actions*, rather than blame-*feelings*. *Both* blame-*actions* and blame-*feelings* are central to our ethical practices.

Consider again third-party judgements that an obligation has been violated. Typically, the third party will feel blame towards the transgressor, and will often express such feelings in some act of reproach: shouting angrily 'You can't do that!', giving him the cold shoulder, or suspending cordial relations. The common-sensical view of the normativity of such blame-*actions* is that they are default-warranted just when blame-*feelings* are warranted. There are certainly defeaters of such warrant: if performing an *act* of reproach towards the mafia don is likely to leave me dead, then I may have most reason not to perform it, even though there is overwhelming reason to *feel* blame towards him for his brutal actions.

For a fully-fledged consequentialist, the reasons for acts of reproach are derived in just the same way as reasons for any other action: it is just a matter of the good and bad consequences of the reproachful action, relative to the goodness of the consequences of alternative actions. This may be a bit counter-intuitive, but only in the way that we are used to in encountering consequentialists who bite the bullet in sticking to their exclusively consequence-based account of reasons for action.

But being consequentialist about reasons for action leaves open what to say about reasons to feel.

An exclusively consequentialist account of reasons to feel (one that says there is reason for an agent to feel *F* just to the degree that the agent's feeling *F* will have good consequences) looks unappealing. This is a fact that few consequentialists have appreciated the import of. It is notable that consequentialists don't tend to endorse direct consequentialist accounts of reasons for *belief*. And just as well, as they'd be implausibly revisionary.<sup>15</sup> Norcross himself eschews a consequentialist account of reasons for belief. But direct consequentialist accounts of reasons to feel also appear implausibly revisionary. There is sometimes some attractiveness in the thought 'I shouldn't blame him; it won't do any good'. Note, however, that this appeal attaches much more clearly to blame-*actions*. When it comes to blame-*feelings*, consider again the examples given above, of the actions of Nazis, people-traffickers or ruthless politicians. That their actions warrant feelings of outrage seems undeniable. And the appropriateness of such outrage seems not be contingent on the outrage itself producing some beneficial outcome.<sup>16</sup>

*Blameworthiness* is about reasons to feel blame-*feelings*, not about reasons for blame-*actions*. Norcross considers blameworthiness in the context of a possible 'linguistic analysis' of wrongness: 'an action is wrong if and only if it is blameworthy (or ought to be punished)' (28).

<sup>15</sup> Indirect consequentialist accounts of reasons for belief may have more plausibility. See e.g. Macleod 2018 for an interpretation of John Stuart Mill's views on the normativity of belief which fit this description. Mill may also be plausibly interpreted as endorsing an indirect consequentialist account of reasons to feel.

<sup>16</sup> One possible stance is to adopt a hybrid or dual-source model of reasons to feel: some reasons to feel are determined by the fittingness of the feeling to the object of the feeling, while other reasons to feel are determined by the goodness of the consequences of feeling that way. For discussion, see Chappell (2012); Howard (2016); McElwee (2018).

This phrasing involves a misleading conflation. The practical question of whether we ought to punish someone (or to act reproachfully towards them at all) is clearly distinct from the question of whether that person is *blameworthy*. This latter question is plausibly understood in terms of whether it would be fitting to feel blame towards them—whether there are the ‘right’ kind of reasons to feel blame towards them.

Norcross unfortunately goes on to discuss practical reasons to punish someone (construed entirely in terms of the value of the consequences of punishing them), and proceeds as if he has discussed blameworthiness itself. He concludes that there can be no conceptual connection between wrongness and blameworthiness or punishability for the utilitarian, since it is ‘absurd to say that [someone] has done something wrong just in virtue of the fact that it is appropriate or optimific to punish her’. (34)

But the principles regarding wrongness Norcross appeals to in order to reject the claim that wrongness is conceptually tied to practical reasons to punish are every bit as plausible when we replace wrongness with blameworthiness. (McElwee 2010b, 399) This shows that, on the ordinary understanding, blameworthiness amounts to something distinct from there being practical reasons to punish.

Overall, Norcross’s demoralising of ethical theory fails to do justice to the central role of feelings, and blame-feelings in particular, in our ethical lives.

#### 4 Consequentialism and Reasons to Feel

At various points in *Morality by Degrees*, Norcross makes clear that he endorses global consequentialism. This is the view that for all X, the best X (or the right X) is the one that brings about the best consequences. (Pettit and Smith 2000) A natural consequentialist thought is that it’s not just actions that we should evaluate by reference to their consequences. We should evaluate all sorts of other things in the same way: the best political institutions are those that produce the best consequences, the best motives are those that produce the best consequences, the best codes of moral rules are those that produce the best consequences, and so on.

Norcross endorses this:

‘My suggestion is that utilitarianism should be treated simply as a theory of goodness of states of affairs and of the comparative value of actions (*and, indeed, of anything appropriately related to states of affairs, such as character traits, political systems, pension schemes, etc.*).’ (Norcross 2020, 27, emphasis added).

‘It is at least part of the essence of any consequentialist view that the moral status of anything (*whether an action, a character trait, an institution, or anything else*) is entirely determined by the consequences of that thing itself.’ (Norcross 2020, 34, emphasis added. See also p. 108.)

However, there are two problems with Norcross’s endorsement of global consequentialism. First, it conflicts with some of his other core claims, and second, it is independently implausible.

Consider the following, typical, statement from Norcross:

‘Here then, roughly, is the motivating idea behind consequentialist approaches to morality. We care, or at least *we should care*, about how good or bad the world is. To the extent that we can make a difference, it is better, morally better that is, for us to make the world a

better place rather than a worse place, and the bigger the positive difference we can make, the better.’ (Norcross 2020, 7, emphasis added).

Note the phrase: ‘we should care... about how good or bad the world is’.

Why should we so care? Suppose, in some particular circumstances, caring about how good or bad the world is will itself have bad consequences. Should I still care then about the good and bad in the world? The global consequentialist, who purports to give a consequentialist treatment of cares, just as he does with acts, sets of rules, or institutions, is committed to saying I should *not*.

This looks distinctly unattractive. Cares can be fitting or unfitting to their object. There is something correct about being pleased when the world goes well, and being pained when the world goes badly. There is something correct about desiring that the world goes well, and being averse to the world going badly.

So Norcross says on the one hand that we ought to care about, hope for, desire that the good be promoted. But his global consequentialism commits him to the view that such caring is wrong when it has bad consequences.

The best way for Norcross to resolve the tension in his views is surely to reject global consequentialism and allow that we should endorse patterns of caring that involve the caring being fitting to its object: being glad when good prevails, and being pained when bad prevails.<sup>17</sup> (The alternative would be to retract the claim that we should care about how well or badly the world goes, while somehow insisting that *what really matters* is how well or badly the world goes. This position looks difficult to sustain, driving a wedge between what matters and what we have reason to care about. A defender of such a view would at least owe us an account of the concept of *mattering* divested of any link to what merits caring about.)

That suggests that alongside consequentialist evaluation of *actions*, consequentialists should propound *in general* a fittingness account of attitudes (feelings/emotions/sentiments) like desire, and being glad (as well as a non-consequentialist, fittingness account of reasons for belief). The implications of this are extensive. In evaluating attitudes, we should not attend just to the good and bad consequences of having them, but also to whether they are fitting, apt, appropriate, merited, correct. And many of our ethical concepts are bound up in the normativity of such attitudes. Virtues, for example, are character traits that are fitting to admire. A person’s well-being is what we have reason to desire for that person’s sake. Most importantly, for present purposes, morally wrong actions are ones that are fitting to blame agents for performing, and morally obligatory actions are those that are fitting to blame agents for omitting.

Once we see how infused our ethical concepts are with emotions, which have a normativity of their own, the clearer it is that the demoraliser’s stark portrayal of the ethical landscape is inadequate. An ethical theory that limits itself to giving graded recommendations of actions leaves out so much that is central to our ethical lives.

We can see now that scalar consequentialism’s eschewal of moral obligations, which are centred around the normativity of blame-feelings, is just one way in which a demoralising approach, overlooking the crucial normativity of our sentiments, is impoverished.

If fitting blame-feelings are, as some satisficers propose, the marker of the non-arbitrary, independently significant boundary between permissibility and impermissibility, then a satisficing view looks like it can avoid all of Norcross’s objections, while also avoiding the very strong objections to the scalar view. Consequentialists convinced by demandingness

<sup>17</sup> For critique of global consequentialism, see Chappell 2012 and McElwee 2018.

objections to Max AC should get down to the tricky business of working out the best version of the satisficing view, and leave the scalar view behind. (Alternatively, they might retreat to rule consequentialism, or abandon consequentialism altogether, of course.)

## 5 Contextualism to the Rescue?

Scalar consequentialism at first sight appears an extremely radical view: there are no moral obligations; nothing is morally wrong. However, Norcross in the end propounds a qualified version of the theory, which may seem to withdraw some of the radical edge, and make its clash with deeply held ethical views (e.g. that torturing innocent children is morally wrong) less stark. The qualified view says that we should eschew the traditional deontic categories of obligation, permissibility and impermissibility from our fundamental ethical theory, but allows that it may nonetheless be legitimate to make assertions couched in terms of these categories.

Norcross recommends a contextualist treatment of the deontic terms, though his positive case for contextualism about the moral categories is brief.

In determining our obligations, there may of course be relativity to the context *of the agent*: what someone is required to do is impacted by the details of her situation. Take duties of beneficence. In a society where there's an established practice, and social expectation, of giving 10% (e.g. some religious societies), it may be morally wrong to skimp by giving only 5% or 7.5%. But where there is no such practice, or social expectation, and people typically give far less, it's arguable that giving 7.5% may fulfil one's financial duties of general beneficence. Such obligations may also be affected by the extent to which one's government addresses severe need (domestically and abroad) through general taxation.

But the case for endorsing relativity to the *context of discussion or evaluation* is hazier.

Norcross offers the following analysis and elucidation:

**'R-con** An action is right iff it is at least as good as the appropriate alternative.

The idea here is that the concept of right action (and duty, permissibility, obligation and the like) invokes a standard, against which the action in question is judged... the contextualist approach I am suggesting allows the conversational context to affect the standard.' (Norcross 2020, 122–3).

A first worry about this approach is whether Norcross's analysis can distinguish obligation from permissibility. The same schema is offered for each here, but with no indication of how the 'appropriate standards' for permissibility and those for obligation are to be differentiated.

A second worry is a reiteration of the force objection. The proposed schema looks manifestly inadequate as a full analysis. Nothing in the analysis has the punchy, must-y force requisite for claims of obligation. Such force makes obligation-claims clearly quite distinct from mere comparative evaluative claims, like at-least-as-good-as claims. To say that some action is obligatory manifestly goes beyond saying that it's at least as good as some other

option; specifically, it says it *must* be done.<sup>18</sup> Consequently, Norcross's schema looks at best a *partial* analysis, awaiting supplementation with an account of the force of obligation.

A third worry about Norcross's discussion is that in trying to make his contextualism plausible, he appeals to materials he has denied his opponents. For instance, he says:

'It seems likely that most (ordinary) contexts will be sensitive to such factors as difficulty (both physical and psychological), risk, and self-sacrifice in establishing the appropriate ideal.' (123).

First, we should note here that obligations are *not* ideals: they have a must-y force that ideals lack. More importantly, we should ask why it is only now that difficulty, risk and self-sacrifice are relevant? These are precisely the sorts of features that Norcross's satisficing opponent appeals to in order to determine a non-arbitrary threshold between the permissible and the impermissible. They are the features of actions that typically render agents less blameworthy for failing to perform them.

Shifting from one example where the morally optimal action involves significant risk, sacrifice, or difficulty to another example where the morally optimal action is risk-free, costless and easy may indeed prompt us to change our judgement about whether it's morally obligatory to perform the morally optimal action. (As Norcross indicates in his verdicts about 'Button Pusher' and 'Burning Building' in Chap. 5.) But this is simply because we've changed the example by modifying morally relevant features.

By contrast, shifting the 'context of discussion' in which a *single* example is considered should not prompt us to change our verdict about what is morally obligatory in that case.

But this is what Norcross seems to propose:

'In order to get a context that would set the rescue of all ten people [in his Burning Building case, where an agent can save up to ten people, but only by incurring very great risks] as the appropriate ideal, we could imagine a conversation among committed maximizing consequentialists, or perhaps among proponents of a Christ-as-ideal moral theory'. (Norcross 2020, 123)

If some particular moral zealots adopt a hugely demanding conception of the morality of assistance, that surely doesn't set a context in which we should default to their ultra-demanding standards for what's morally permissible. Rather, we should judge for ourselves about whether morality really does demand so much as they claim.

Likewise, in the case of duties of beneficence, it seems there is just a substantive disagreement between those with lax libertarian standards (which Norcross himself seems to 'reject') and more stringent standards. If Norcross is correct to reject such standards, then, even in a context in which only lax libertarians are present, we should regard as straightforwardly false their under-demanding verdicts on our duties of beneficence.

Norcross characterises some supposedly relevantly different contexts (which allegedly require different standards) in terms of which available actions are salient in those contexts. When actions with relatively bad consequences are salient, the standards for acting permissibly are lower than when actions with relatively good consequences are salient. I find it difficult to see why we should think the salience of certain options in a particular context of discussion affects the relevant standard determining the best judgements of permissibility made in that context. If some particular agent (as in Button Pusher) is so dastardly that she doesn't contemplate performing the only remotely humane option, we should emphatically

<sup>18</sup> Similar things can be said about permissibility, which can be inter-defined with obligatoriness: an action is permissible iff it's not obligatory to refrain from it.



not *adopt their perspective* for our context of discussion and resort to describing the lesser of the two options actively under her consideration as morally permissible (or as morally good).<sup>19</sup>

Overall, it appears that contextualist versions of scalar consequentialist are yet to be motivated convincingly. The best that can be said for them is that they avoid the unpalatable eliminativism of more radical forms of the view. But, once we acknowledge that SC1-SC3 can be endorsed by consequentialists of many stripes, rejecting the scalar view altogether, and adopting a satisficing view, looks the consequentialist's best bet.

## 6 Conclusion

I have argued that, even for consequentialists, who may feel less at home than their deontological opponents amongst the traditional deontic categories of obligation, permissibility and impermissibility, de-moralising ethical theory is a mistake. Demoralisers mistakenly speak as if all we wanted from our ethical theory is a list of graded *recommendations* of actions. We impoverish our ethical theory by leaving out the traditional deontic categories. On the interpretation offered here, this is because such categories are the way we go beyond such graded recommendations of actions, and include evaluation of how we should feel about (and more generally react to) agents who act in different ways. Such a picture does much more justice to morality as we find it in the world. The demoralisers' stark perspective leaves out the normativity of emotions. But given how essential sentimental responses are to morality as it appears in the real world, overlooking the normativity of the emotions leaves an ethical theory that is deeply impoverished and inadequate to tasks it needs to address.

**Acknowledgements** I would like to thank Tyler Paytas and Richard Rowland who organised a conference on demoralising ethical theory at the Rome Campus of the Australian Catholic University in 2019, as well as all the conference participants. The rich discussions there informed my thinking about the issues discussed in this paper. I'm also very grateful to two anonymous referees for this journal who provided really helpful comments on an earlier draft.

<sup>19</sup> Neil Sinhababu's version of scalar consequentialism, which likewise appeals to 'linguistic resources', seems vulnerable to similar worries. Sinhababu says: 'To see how the phenomena under discussion call for the application of these linguistic resources, let's again consider Toby, who is giving 40% of his income to charity. Commonsense morality tells us that this is the right thing to do.' (Sinhababu 2018, 3138). Sinhababu claims that 'things change when we attend to a larger range of possibilities'. (Sinhababu 2018, 3138) If Toby 'were choosing between giving 40% and giving 50%, with the additional 10% coming from giving up personal comforts of no significant moral value, there's at least some temptation to say that giving up 50% is the unique morally right choice.' (Sinhababu 2018, 3138). This case does little to support the application of the linguistic resources that Sinhababu and Norcross have in mind. Both the initial verdict on the case, and the purported urge to shift that verdict in light of other options becoming salient seem no part of our ordinary moral thought. Common sense morality would surely *not* say that giving precisely 40% is *the* right thing to do. Instead, it would say it's *a* permissible (indeed supererogatory and admirable) thing to do. As would giving 45%, or 50%. And there is no pressure whatsoever from common moral thinking to retract the claim that giving 40% is permissible when we 'notice' that giving 50% is an option. Rather, common moral thought says unequivocally that both courses of action are permissible, and that giving 50% is morally better than giving 40%. Moreover, in evaluating Toby's decision to give 40%, further options of giving 35% or 45% or 50% are hardly recondit possibilities whose newfound salience has the potential to transform our initial judgement that giving 40% is permissible, supererogatory and admirable. (There is also a worry here about moral hazard. Such a view of moral obligation gives agents an incentive to make morally good options less salient to themselves, so that less is demanded of them.)

## Ethics Declarations

**Conflict of Interest** The author has no relevant financial or non-financial interests to disclose.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Chappell RY (2012) Fittingness: The sole normative primitive. *Philosophical Q* 62(249):684–704
- Chappell RY (2019) Willpower Satisficing. *Nous* 53(2):251–265
- Chappell RY (2020) In: Portmore D (ed) *Deontic Pluralism and the Right Amount of Good*. The Oxford Handbook of Consequentialism, Oxford
- Crisp R (2006) *Reasons and the Good*. Oxford
- Hooker B (2001) *Ideal code. Real world*, Oxford
- Howard C (2016) In defense of the wrong kind of reason. *Thought* 5:53–62
- Lang G (2013) Should utilitarianism be scalar? *Utilitas* 25(1):80–95
- Lawlor R (2009) The rejection of scalar consequentialism. *Utilitas* 21(1):100–116
- Macleod C (2018) Mill on the primacy of practical reason. *Analysis* 78(4):630–638
- McElwee B (2010a) Consequentialism and permissibility. *Utilitas* 22(2):171–183
- McElwee B (2010b) The rights and wrongs of consequentialism. *Philos Stud* 151:393–412
- McElwee B (2018) The Ambitions of Consequentialism. *J Ethics Soc Philos* 17(2)
- Norcross A (2020) *Morality by degrees. Reasons without demands*, Oxford
- Pettit P, Smith M (2000) In: *Morality R, Hooker C (eds) Global Consequentialism*. E. Mason and D. Miller, Edinburgh
- Sinhababu N (2018) Scalar Consequentialism the Right Way. *Philos Stud* 175:3131–3144
- Srinivasan A (2018) The Aptness of anger. *J of Political Philos* 26(2):123–144

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.