

# Consequentialism and Permissibility

BRIAN MCELWEE

*Lincoln College, University of Oxford*

Scalar consequentialism, recently championed by Alastair Norcross, holds that the value of an action varies according to the goodness of its consequences, but eschews all judgements of moral permissibility and impermissibility. I show that the strongest version of scalar consequentialism is not vulnerable to the objection that it is insufficiently action-guiding. Instead, the principle objection to the scalar view is simply that it leaves out important and interesting ethical judgements. In demonstrating this, I counter Rob Lawlor's contention that consequentialists cannot consistently care about permissibility and impermissibility.

Powerful objections to maximizing act-consequentialism, in particular claims that it is overly demanding, have seen defenders of consequentialism attempt to devise alternative forms of the theory which might avoid such objections. One candidate, championed by Alastair Norcross, which has recently had much attention, is scalar consequentialism, which holds that the value of an action varies according to the goodness of its consequences (perhaps relative to the goodness of the consequences of the alternatives), but which eschews altogether judgements of permissibility and impermissibility.<sup>1</sup>

In 'The Rejection of Scalar Consequentialism', Rob Lawlor comes to three main conclusions:<sup>2</sup>

1. Norcross's positive arguments in favour of scalar consequentialism are unconvincing.
2. Scalar consequentialism cannot be the strongest form of consequentialism because it is not sufficiently action-guiding.
3. Consequentialists in general cannot (consistently) care about moral permissibility.

In this article, I will show that Lawlor rejects scalar consequentialism for the wrong reasons. I will dispute 2 and 3, showing that scalar consequentialism (henceforth SC) can indeed be sufficiently action-guiding, and explaining how consequentialists in general can and

<sup>1</sup> Alastair Norcross, 'The Scalar Approach to Utilitarianism', *The Blackwell Guide to Mill's Utilitarianism*, ed. H. West (Oxford, 2006). See also Michael Slote, *Common-Sense Morality and Consequentialism* (London, 1985).

<sup>2</sup> Rob Lawlor, 'The Rejection of Scalar Consequentialism', *Utilitas* 21.1 (March 2009).

should care about moral permissibility. In arguing the latter point, I will indicate how Norcross's SC falls short, and suggest how it might be modified or supplemented to provide the strongest form of consequentialism.

### CAN SCALAR CONSEQUENTIALISM BE ACTION-GUIDING?

While much of Lawlor's paper focuses on disabling the specific arguments which Norcross adduces in favour of SC, Lawlor's main positive argument against SC is that it fails to be action-guiding in the way we expect of a moral theory. I will argue that the judgements of SC can indeed guide action, and can in fact do so more comprehensively than moral theories framed solely in terms of permissibility.

To illustrate where he takes SC to go wrong, Lawlor refers to an example given by Tim Mulgan, where an agent, Achilles, can, at no cost to himself, produce optimal consequences by pressing button *n*.<sup>3</sup> Lawlor analyses the respective responses of maximizing and scalar consequentialists:

Both accounts tell Achilles that he has reason to press *n*, but the maximizer says more than this. The maximizer states that Achilles has a *conclusive* reason to press *n*. As a result, he tells Achilles that he ought to press *n*, and that it is *impermissible* to do anything else. Clearly, this offers significantly *more guidance* than a theory that merely tells you that you have a reason to press *n*, but says no more.<sup>4</sup>

The first thing to note here is that it is clear that SC need not, and should not, limit itself to claiming that there is *a* reason to press *n*. Scalar consequentialists can, and should, say that Achilles has *most* reason to press *n*. As Lawlor notes, making a claim couched in terms of most reason may appear equivalent to making an *ought* claim. And Norcross, for one, claims to eschew ought claims. However, it is clear that what Norcross really wants to jettison is not simple claims about what there is most reason to do, or what is the most choiceworthy thing to do, but rather oughts understood as moral *requirements*. An ought claim of this latter kind has the implication that in failing to comply, one acts in a way that is morally wrong or impermissible. It is clearly this ought of moral requirement that Norcross believes our ethical thought and practice would be better off without. My aim here in any case is not to defend Norcross himself, but rather to show that the strongest form of SC is not vulnerable to Lawlor's action-guidingness objection. And it is clear that a scalar consequentialist can make ought claims construed as judgements about what there is most reason to

<sup>3</sup> Tim Mulgan, *The Demands of Consequentialism* (Oxford, 2001), ch. 5.

<sup>4</sup> Lawlor, 'The Rejection', p. 113.

do, as they do not commit him to any sharp cut-off points, or to the supposedly all-or-nothing moral notions of wrongness, permissibility and obligation. If 'ought' is understood as synonymous with 'has most reason to', and has none of the paradigmatically moral implications (of wrongness, impermissibility, blameworthiness and so on), then scalar consequentialists can accept ought claims. To say simply that one has most reason to bring about the best available consequences is not equivalent to saying that one is *morally required* to do so, on pain of meriting blame, guilt or serious criticism. For the purposes of clarity, I will express SC in terms of claims about what we have *most reason to* do, or what is the most *choiceworthy* thing to do.

Once we concede that SC can assert that Achilles has strongest reason to press n, is there any respect in which maximizing (or satisficing) consequentialism provides more action guidance? In particular, does a theory which tells us what we have *conclusive* reason to do, or what it is *impermissible* not to do, give any more guidance to the agent in which act to perform? Obviously, if we take 'conclusive reason' to be synonymous with 'strongest reason', it can say nothing more. But perhaps instead, to say that one has conclusive reason is to say that one has overwhelmingly strong reason to perform this action rather than one of the alternatives. But this, again, is just the sort of claim that a scalar consequentialist is able and willing to make: if pressing n has much better consequences than all the other alternative actions, then there is *much* stronger reason to press n than do anything else.

But perhaps to say that one has conclusive reason to press n is just to say that doing anything else would be *impermissible*. This brings us to another frequent criticism of SC, namely that it often seems clearly correct to say not just that there is very strong reason to do X, but that it would be *wrong* not to. What do we add by claiming that an action is not just unchoiceworthy, but wrong? Clearly, what we leave out is assessment of the agent in performing that act, a judgement about what response the agent merits. In saying that an action X is wrong, we do not simply claim that there was very strong reason to do otherwise, but that if one does X, one merits serious criticism, condemnation, or feelings of blame and guilt. This may indeed be a strong criticism of SC (I return to this below), but it is emphatically not the criticism that SC is insufficiently action-guiding in virtue of its being limited to claims about the strength of reasons to act. To add that one would act wrongly, that one would merit blame or serious criticism if one did X, is not to say any more by way of guidance about how worth doing an action is. Whether a theory is sufficiently action-guiding or not is determined solely by its ability to provide an account of how much reason there is to perform various alternative actions.

Lawlor claims that in order to be sufficiently action-guiding, a theory must 'rule out various bad options...[but that]...Scalar consequentialism does not rule out *any* options – not even the worst actions, such as killing and eating old ladies'.<sup>5</sup> But it is clear that SC does 'rule out' such actions in the only sense relevant to guidance of action: it pronounces such acts extremely unchoiceworthy relative to the alternatives. It does not, of course, 'rule out' such actions in the sense of claiming that such actions are morally wrong, blameworthy, and deserving of the most serious criticism and disapproval. The refusal of SC to make such judgements may ultimately be unconvincing, but this is not in any way in virtue of a failure to be sufficiently action-guiding.<sup>6</sup>

Indeed, one might precisely see action-guidingness as a feature in virtue of which SC is in fact at an advantage over many traditional moral theories couched in terms of right, wrong, permissibility and obligation. SC actually provides us with more detailed guidance than such theories. In the Achilles example, for instance, SC tells us precisely how strong the reason is for pressing n, and how strong the reason is to do each of the alternatives (in terms of the goodness of the consequences of each alternative). Thus, SC tells Achilles not just what he has strongest reason to do, or what there is conclusive reason to do, but also *how* conclusive the reason is! It details how much more reason there is to press n than to do anything else. If anything, it appears that scalar consequentialism in fact gives more, rather than less, by way of guidance to agents in that it does not just offer one piece of advice: 'Do this'. Rather it gives an assessment of the relative choiceworthiness of all options.

Why might this be a good thing? One reason is that in real life the most choiceworthy option, the action that we have the most reason to perform, is often one that we will not perform, because it is difficult, costly or demanding. In this case, SC tells us what the next most choiceworthy thing to do is. And the next most choiceworthy after that, and so on. (Crucially, SC does not withdraw its initial judgement as to what the best thing to do is.) Given that I am not going to live the life of a 'moral saint', it is useful to have a ranking of the relative choiceworthiness of other, less demanding lifestyles. Another way in which scalar views may be superior in guiding action to

<sup>5</sup> Lawlor, 'The Rejection', p. 109.

<sup>6</sup> One might argue that an ethical theory is insufficiently action-guiding if it fails to distinguish moral reasons from egoistic reasons, and offer an account of how they interact. (Lawlor has suggested this in personal communication.) For an account of how we should view the relationship between egoistic reasons and impartial reasons, see Brian McElwee, 'Consequentialism, Demandingness and the Monism of Practical Reason', *Proceedings of the Aristotelian Society*, vol. CVIII, pt. 3.

views simply framed in terms of permissibility and impermissibility is that the former give a comprehensive ranking. The latter structure divides actions into three categories – ‘obligatory’, ‘permissible, but not obligatory’ and ‘impermissible’. But within each category, some alternative actions will have more to be said for them than others. In the case of beneficence towards the distant needy, for example, it may be permissible to give anything from 10 per cent of our income upwards, and forbidden to give anything less, but clearly a theory is more informative if it says that we have more reason to give 20 per cent than to give merely 10 per cent. We might consider a moral theory to be deficient if it only tells us what is optimal, and what the bare minimum is. Agent A may be prepared to act in a way that is better than the bare minimum required to fulfil his moral obligations, by giving, say 20 per cent. A theory which tells us that she has more reason to do this than to stick to the minimum is more informative than one that does not. Likewise, A might not be willing to perform the optimal action, say sacrificing her life by diving on the grenade, so to know that a theory gives a ‘second-best’ recommendation to some other action (e.g. not fleeing) is informative too. A morality which limits itself to judgements about the cut-off point of permissibility and impermissibility would appear to do less by way of guiding the agent.

The common objection that SC cannot adequately guide action seems misplaced then. We may of course still be left with the feeling that SC fails by eschewing all talk of wrongness and obligation. We do want to say more than simply assess the relative strength of reasons for alternative actions; we also want an assessment of the agent himself. As Lawlor argues, Norcross at least has not done enough to motivate this thinned-down ethical view, and SC appears to leave out, without justification, a large part of our ethical thought.

## CONSEQUENTIALISM AND IMPERMISSIBILITY

I now turn to the role that the moral notions of permissibility and impermissibility should have in consequentialist ethical thought. Lawlor argues that consequentialists cannot (consistently) care about impermissibility. His conclusion is in fact that the only thing SC has going for it is its honesty about the irrelevance of categories of permissibility and impermissibility in consequentialist ethical thought. This is mistaken: the reason why SC is not the strongest form of consequentialism is because it eschews judgements of moral permissibility and impermissibility, while other forms of consequentialism can and do accommodate such judgements, as I will illustrate.

Lawlor's claim that consequentialists cannot consistently care about permissibility arises in response to Norcross's 'no significant difference' argument for SC.<sup>7</sup> Norcross argues that, for a consequentialist, it can be no more important that one increase from 9 per cent to 10 per cent the portion of one's income one gives to the needy, than it is to increase one's contribution from 11 per cent to 12 per cent, even if we were to judge that a contribution of 10 per cent is what is morally required. Norcross's central idea here is simply that the importance of any given improvement should be measured according to the amount of good it does. If we stipulate that each increase in giving will produce the same amount of good, the consequentialist will judge that each increase is equally important. Given this stipulation, if one were to be in a position to persuade a 9 per cent giver or an 11 per cent giver to up their contribution by the same amount (assume they have the same income), the consequentialist is committed to saying that one has equally strong reason to persuade either, assuming that they are equally open to persuasion.

From this, Lawlor concludes that consequentialists cannot consistently care about permissibility (since all they care about is the consequences of different courses of action). This is wrong, for two reasons. First, it generally has good consequences to treat people differentially depending on whether they have acted permissibly or impermissibly. Second, consequentialists, like anyone else, can care about all kinds of interesting ethical facts, not just those concerning what we have reason to do. It is precisely the denial of such interesting facts (rather than its account of what reasons for action we have) that makes SC so counter-intuitive.

Let us look at the first reason, then, why Lawlor is wrong to say that 'the distinction between the permissible and the impermissible seems to be irrelevant from the consequentialist point of view'.<sup>8</sup> Take the case of punishment, which is often thought to be especially problematic for consequentialism. One might argue that consequentialism already has a full account of when we should punish (in terms of the goodness of the consequences of punishing or not punishing), one which leaves no room for whether the agent acts permissibly or not, and so permissibility can have no role to play. But this is seriously misleading. There is a very important sense in which permissibility plays a role in consequentialist accounts of when we ought to punish people. It is a good thing from a consequentialist point of view if we adopt a general practice of punishing those, and only those, who have acted impermissibly, and indeed taking very stringent measures to ensure that only those whom

<sup>7</sup> See Norcross, 'The Scalar Approach', p. 220.

<sup>8</sup> Lawlor, 'The Rejection', p. 105.

we have very strong reason to believe are guilty are punished. The main reason why this is so important from a consequentialist point of view is, as Mill reminds us, the huge importance security has in contributing to human welfare.<sup>9</sup> If we are forever in fear of being punished for crimes we have not committed, our well-being deteriorates very significantly; much of the human misery in totalitarian states is attributable to this very insecurity. We can thus see that a consequentialist criterion of the value of actions, and of the justification of punishment in particular, does not preclude caring about whether agents act permissibly or not. This is because a practice of punishing only those who are guilty generally has overwhelmingly better consequences than one of punishing the innocent.

It is of course a standard thought-experiment in considering the plausibility of consequentialism to ask whether we should punish an innocent person if we *know* that doing so will have the best consequences. And consequentialists are committed to biting the (perhaps uncomfortable) bullet that this is the right thing to do in such circumstances. But it does not follow from this that consequentialists do not care about whether agents act in a morally permissible way or not. Adopting a practice of only punishing the guilty is one of the things consequentialists think we should do, precisely because doing so brings about such good consequences.<sup>10</sup> Consequentialism is in fact not nearly so far from common-sense morality in this respect as is often implied. According to common-sense morality, it is indeed sometimes right to punish an innocent person for the sake of good consequences. Perhaps the most obvious example is the punishment of young children, presumably morally innocent, as a means to a good end, namely, their learning to act well. Similarly, there are many circumstances in which common-sense morality judges that we ought not to punish the guilty, because of the bad consequences of doing so. If trying a guilty politician will predictably lead to civil war and the deaths of many innocents, a strong strand in common-sense morality tells us that we should forego the punishment of the guilty man.

However, we are seldom in a position to know for sure that punishing someone innocent, or sparing someone guilty, will produce the best consequences. Indeed, we very rarely have strong reason to believe that doing so will have good consequences. Adopting a practice of punishing only the innocent is a strategy that plausibly has overwhelmingly better consequences than trying to work out, on a case-by-case basis,

<sup>9</sup> John Stuart Mill, *Utilitarianism*, ed. R. Crisp (Oxford, 1997), ch. 5.

<sup>10</sup> Obviously, there is a significant distinction between acting in a way that is morally impermissible and acting in a way that is legally impermissible. I will not discuss this complication here as it provides no special problem for the consequentialist.

whether punishing an innocent person might just turn out to produce better consequences. This consequentialist rationale for punishment is hostage to empirical contingencies, but for the consequentialist, and to a lesser degree for common-sense morality too, this is just as it should be. The difference between the two moral outlooks is a matter of degree as regards the extent to which consideration of consequences should mitigate considerations of guilt or innocence in punishing people. My aim here is not to offer a full-scale defence of consequentialist accounts of punishment, but we might also note that the large role within common-sense morality for retributivist rationales in the justification for punishment should make us pause in holding common sense to be the last court of appeal in such matters in any case.

Let us return now to the case of persuading people to meet their obligations of beneficence (e.g. to give 10 per cent of their income). The consequentialist claims that in determining who we should try to persuade, only the consequences of the available acts of persuasion are relevant. As Lawlor points out, a consequentialist might happily concede that there is, contra Norcross, a genuine fact about how much one is morally required to do to help the needy. But, says Lawlor, if he does concede this, he is committed to regarding facts about whether moral requirements have been met as irrelevant to the question of whom we should try to persuade. Lawlor's conclusion is that Norcross's example has provided no argument against the claim that there are genuine cut-off points, but instead shows that if there are such cut-off points, consequentialists do not or should not care about them.

This is perfectly correct in one sense. The sense in which we might wish to qualify it is, again, that it may have good consequences to adopt a practice or strategy of persuading or forcing those who fail to act permissibly to do so. If it is *stipulated* in a particular example that urging A to give more will be a better bet in producing good than persuading B, and *that there are no further relevant consequences*, then the mere facts about who has acted permissibly are indeed irrelevant, from the consequentialist's point of view, *in determining what I have reason to do*. But in general, of course, such facts *are* relevant. Holding people to account for failing to reach what we deem a minimally acceptable standard of behaviour itself has very significant (and overwhelmingly positive) consequences. Thus, it is very clear that Lawlor's claims are too strong when he says,

this just highlights the fact that the consequentialist does not care about fairness or the distribution of burdens, and does not care about permissibility. This tells us something about the consequentialist's attitude to permissibility, but it does not tell us anything about permissibility itself.<sup>11</sup>

<sup>11</sup> Lawlor, 'The Rejection', p. 105.



Lawlor may be right that Norcross has done nothing to impugn the common-sense idea that some courses of action, and levels of beneficence, are permissible while others are impermissible. However, we can be much more precise than Lawlor with regard to what the consequentialist is committed to. He is committed to the claim that whether acts are permissible or not can only have an effect on what one has reason to do if facts about what is permissible have a bearing on the consequences of various alternative actions. They generally do, in that it has better consequences to adopt a practice of treating people who meet their moral requirements differently from those who do not.

There is a second, more fundamental reason why it is deeply misleading to say that consequentialists cannot consistently care about permissibility and impermissibility. This is simply that facts about permissibility are interesting ethical facts. Not all ethical facts are about reasons for action. Nor is it true that an ethical fact can only be interesting to the extent that it bears on what reasons for action there are. In recognizing this, we see where SC falls short. In limiting itself to claims about reasons for action (or, as in Norcross's case, about the comparative value of different actions), SC ignores interesting ethical facts that we wish to discover. We want to know not just that we have very strong reason not to perform an action, but also whether it is morally wrong. In the case of certain types of action, for example, a failure by Achilles to save many lives when it will cost him nothing, as Tim Mulgan says, 'It is surely not enough to say that these characters could have done better. They... behave wrongly.'<sup>12</sup> So it seems wrong to deny, as Norcross does, that any positive assertion that acts are permissible seems wrong or impermissible can be correct.

The fact that our moral theories have trouble determining the extent of our moral obligations as regards, for example, beneficence to the needy does not seem to undermine the fact that there are indeed obvious cases of moral wrongness. Someone who deliberately inflicts severe pain on an innocent person simply for his own amusement acts morally wrongly. A scalar consequentialist can, of course, say that the person has very strong reason to act otherwise. But to say that the person acts wrongly is to say more than this. The scalar consequentialist does indeed seem to be leaving out an important moral judgement. What is it, then, that the scalar consequentialist leaves out in limiting himself to a judgement couched in terms of value or reasons? It seems to be evaluation of the agent's performing that act, or an account of the appropriate response to the agent. We do not simply want to say that the agent could have acted in ways which were much better than the

<sup>12</sup> Mulgan, *The Demands*, p. 143. See also Brad Hooker, 'Right, Wrong and Rule-Consequentialism', *The Blackwell Guide to Mill's Utilitarianism*, ed. H. West (Oxford, 2006).

way he did act, but also that he merits serious criticism, feelings of blame and resentment for having done so. We wish to say not only that Achilles has very strong reason to press *n*, but that if he fails to do so, he merits serious criticism and feelings of blame.

This need not have any bearing, apart from the indirect sort of bearing I outlined above, on what we have reason to do. Indeed, allowing facts about moral permissibility and wrongness any greater role in determining what we have reason to do might appear to bespeak a fetishism for the moral categories. It is quite plausible to hold that the mere fact that an action is impermissible itself provides no reason not to perform it. To say that an action is impermissible or wrong may well *imply* that there is reason not to do it, but does not itself seem to provide additional reasons. The fact that an act will cause unnecessary pain, for example, gives us reason to not to perform that act, and will also (at least partly) explain the wrongness of the act. But the wrongness of the act does not then provide some extra reason not to perform it. And such independent judgements of wrongness are quite compatible with a thoroughgoing consequentialist account of what provides reasons for action.

Aside from the ‘indirect’ bearing that permissibility facts have on what we have reason to do, answering these questions about permissibility, fairness, and so on are just independently interesting. Not all ethical questions are about what we have reason to do. Consequentialism is precisely a theory of what reasons for action there are, or about the betterness and worseness of different available actions. This may be, as Norcross, suggests, the ‘fundamental moral fact about an action’.<sup>13</sup> But much more than this must be said to establish that ‘Once a range of options has been evaluated in terms of goodness, all the morally relevant facts about those options have been discovered.’<sup>14</sup> Questions about whether someone has acted morally wrongly, questions about blameworthiness and praiseworthiness, and questions about fairness are also important ethical questions, ones that are distinct from the question that consequentialists should directly answer, that of what there is reason to do.

The consequentialist does indeed have a ready answer to what we should do in situations where one person is doing his fair share, while the other is not. We should do whatever will produce the best consequences. This consequentialist judgement about the reasons for action an agent has does not preclude making other ethical judgements, in particular judgements about moral obligations and about fairness. Suppose person A gives 20 per cent of her income to the needy,

<sup>13</sup> Norcross, ‘The Scalar Approach’, p. 228.

<sup>14</sup> Norcross, ‘The Scalar Approach’, p. 228.

while person B gives nothing. We might say that B has failed in his moral obligations, while person A has met her moral obligations. It is unfair that B does not do his bit. These judgements are each perfectly compatible with a fully-fledged consequentialist account of reasons for action. So, in response to Lawlor's suggestion that (consistent) consequentialists do not care about permissibility, there are two responses: they care about it in so far as questions about permissibility have an (indirect) bearing on the ways in which it is good to treat people. And also they may care about the permissibility facts for their own sake, as independently interesting facts which are compatible with a consequentialist account of reasons for action.

### CONCLUSION

Perhaps the major motivation to adopt scalar consequentialism is the need to avoid the demandingness objections faced by traditional maximizing act-consequentialism, which implausibly morally requires us to do the best thing available. To object to maximizing act-consequentialism on these grounds is not, however, to impugn consequentialism as a theory of what makes actions better and worse, or as a theory of what there is (most) reason to do. In fact, consequentialism understood this way has much appeal, hence the apparent attraction of SC. However, in saying that producing the best possible consequences cannot possibly be morally required, the scalar consequentialist appears to lurch to the opposite extreme by claiming that, in fact, *nothing* is morally required of us. This appears at least as counter-intuitive as maximizing act-consequentialism. But it should be very clear that these are not the only options. In endorsing the scalar consequentialist's ranking of actions from best to worse, we do not commit ourselves to the claim that nothing is morally wrong.

Scalar consequentialism can be seen as combining two types of claim, one positive and one negative:

Positive claim: The best available action is the one which produces the best consequences. All actions can be ranked from best to worst according to the goodness of the consequences they produce. The value of an action, and how much reason there is to perform the action, is determined by how good its consequences are relative to the alternatives.

Negative claim: There are no further moral truths beyond these claims about best, betterness and worseness. In particular, no actions are morally permissible or impermissible.

Consequentialists ought to endorse the positive claim, but reject the negative claim. This leaves the question for consequentialists of how we are to determine what our moral obligations are. One of Norcross's concerns is that any dividing line between the permissible and the

impermissible would be wholly arbitrary. But, as Lawlor rightly points out, this is incorrect, since the level of self-sacrifice that a course of action requires of an agent is relevant in determining whether it is obligatory or not. It may instead be that, in the case of duties of beneficence, our moral obligations are vague. Such vagueness is not sufficient to motivate SC, since there are clear-cut cases of moral permissibility and moral impermissibility. Giving 75 per cent of our income to good causes (even if we could give more) is clearly morally permissible, while doing nothing whatsoever to help those less well-off than ourselves in any way is morally impermissible. Actively harming people seems to be an even clearer case of morally impermissible action.

In determining what our moral obligations are, we are clearly able to pick out the relevant features, namely the effort, sacrifice, cost, difficulty for the agent. These features, which are the focal point of our sentiment of blame, are precisely what we turn to as a rationale for rejecting maximizing act-consequentialism. Thus, some form of cost-sensitive consequentialism seems clearly preferable.

Within this framework, a number of possible approaches are available to the consequentialist, each distinct from the implausibly demanding theory that we are morally obliged (on pain of meriting blame, guilt and serious criticism) to bring about the best consequences we can, but which nevertheless recognize (unlike SC) that we do have moral obligations. Some sort of 'hybrid' account, which affords us moral permissions to favour ourselves to some degree, along the lines suggested by Samuel Scheffler,<sup>15</sup> is one appealing consequentialist approach to moral requirements.<sup>16</sup> My point here is that there is no reason why a defender of consequentialism about what we have reason to do should be forced to adopt either the moral extremism of maximizing act-consequentialism or SC's rejection of all moral requirements.

We have seen that moral theories are expected to do (at least) two main jobs: to guide our actions (tell us what we have reason to do), and to tell us what we *must* do, what is required, on pain of meriting blame and serious criticism. Scalar consequentialism in fact has the tools to perform the first task more fully than does a theory which simply divides actions into the two categories 'permissible' and 'impermissible'. The real objection to scalar consequentialism is that in eschewing talk

<sup>15</sup> See Samuel Scheffler, 'Prerogatives without Restrictions', *Philosophical Perspectives* 6 (1992), pp. 377–97. For other possible approaches, see Slote, *Common-Sense Morality*; Liam Murphy, *Moral Demands in Non-Ideal Theory* (New York, 2000); Brad Hooker, *Ideal Code, Real World* (Oxford, 2000), Mulgan, *The Demands*, pt. IV.

<sup>16</sup> My own favoured approach is one which endorses an account which is less systematic than Scheffler's in determining what moral requirements we face.

of moral requirements, it leaves out too much. The consequentialist should drop Norcross's negative claim, and supplement his ranking of actions with some sort of account of what we must do, on pain of being blameworthy. The result would be a moral theory which, to my mind, has the potential to retain all that is best in Norcross's scalar consequentialism, while avoiding its most striking faults.<sup>17</sup>

brian.mcelwee@lincoln.ox.ac.uk

<sup>17</sup> I am grateful to Rob Lawlor and Gerald Lang for useful discussion of the issues raised in this article.