Should Utilitarianism Be Scalar?

GERALD LANG
University of Leeds

Scalar utilitarianism, a form of utilitarianism advocated by Alastair Norcross, retains utilitarianism's evaluative commitments while dispensing with utilitarianism's deontic commitments, or its commitment to the existence or significance of moral duties, obligations and requirements. This article disputes the effectiveness of the arguments that have been used to defend scalar utilitarianism. It is contended that Norcross's central 'Persuasion Argument' does not succeed, and it is suggested, more positively, that utilitarians cannot easily distance themselves from deontic assessment, just as long as scalar utilitarians admit – as they should do – that utilitarian evaluation generates normative reasons for action.

I. INTRODUCING SCALAR UTILITARIANISM

Utilitarianism has traditionally been regarded as a theory of right action.[1] The different forms of utilitarianism offer varying accounts of what makes an act right, but they tend to share the formal characteristic of *having* an account of what makes acts right. Traditional versions of utilitarianism will also typically insist that rightness is productive of duty, or obligation: agents are obliged to do what is right, and to refrain from doing what is wrong.

For not unrelated reasons, traditional maximizing utilitarianism has also been exposed to two complaints, among many others. First, it is extremely demanding, or at least has the strong potential to be, given the amount of unrelieved suffering our world contains.[2] Second, it makes no room for supererogatory actions, which are actions that are morally praiseworthy, but not morally required.

Alastair Norcross has argued for an intriguing variant of utilitarianism, 'scalar utilitarianism', which preserves traditional utilitarianism's ranking of states of affairs from better to worse, based on how much utility is instantiated in each of those states of affairs, but dispenses with traditional utilitarianism's categories of rightness and wrongness.[3] Other prominent deontic or action-guiding categories,

[1] I will most often refer to utilitarianism, rather than consequentialism, though little turns on the distinction between them in this debate.

[2] This is most obviously true of act-utilitarianism, of course.

[3] See Alastair Norcross, 'The Scalar Approach to Utilitarianism', *The Blackwell Guide to Mill's Utilitarianism*, ed. H. West (Oxford, 2006), pp. 217–32; and Norcross, 'Reasons Without Demands: Rethinking Rightness', *Contemporary Debates in Moral Theory*, ed. J. Dreier (Oxford, 2006), pp. 38–53. This is slightly rough; I provide a more precise characterization of scalar utilitarianism at the beginning of sect. II.

© Cambridge University Press 2013
Utilitas Vol. 25, No. 1, March 2013
doi:10.1017/S0953820812000295

https://doi.org/10.1017/S0953820812000295 Published online by Cambridge University Press

such as the notions of duties, requirements, obligations and demands, are similarly dismissed by scalar utilitarianism.

According to Norcross, scalar utilitarianism's dismissal of rightness and wrongness allows it to escape from the complaints sketched above – from the 'Demandingness Objection' and the 'Supererogation Objection', as I shall refer to them. Why should this be? Let us go through the objections in turn.

According to the scalar view, utilitarianism will not make *any* demands on the agent, so it surely cannot be accused of being *too* demanding. That is enough to tackle the Demandingness Objection.

The Supererogation Objection is treated in a similar way: since there is no particular act which it says is required, scalar utilitarianism can hardly fail to accommodate acts which produce more utility than those which would be required by a non-scalar form of utilitarianism. A small complication arises here. This line of argument may not pass muster if we want a reply to the Supererogation Objection which continues to accommodate the special *value* of supererogatory actions; for, if we think that the *value* of supererogatory acts consists, precisely, in their going *beyond duty*, we may be presupposing that there are real duty-setting standards which supererogatory acts in some sense improve upon. Norcross displays an awareness of this line of objection, but counters it with the suggestion that '[t]he intuition that drives the belief in supererogation can ... be explained in terms of actions that are considerably better than what would be expected of a reasonably decent person in the circumstances'.[4]

Though Norcross's primary target is that of maximizing utilitarianism, he is also opposed to satisficing utilitarianism, which identifies right action with actions that produce *enough* utility.[5] In eschewing maximization, satisficing utilitarianism fields fairly obvious solutions to the Demandingness Objection and Supererogation Objection. It still comes within the ambit of Norcross's attack, however, because it shares with maximizing utilitarianism a commitment to the significance of rightness.

These possibilities alone give utilitarians – and, of course, interested critics of utilitarianism – ample reason to carefully investigate the scalar version of utilitarianism. Scalar utilitarianism raises a large number of complex theoretical issues, such as the general relationship between the evaluative and the deontic, and I cannot hope to address all of them in this article. I will instead address what are in my

[4] Norcross, 'Reasons Without Demands', p. 44.

[5] For a fuller discussion, see Michael Slote, 'Satisficing Consequentialism', *Proceedings of the Aristotelian Society*, Supplementary Volume 58 (1984), pp. 139–63.

https://doi.org/10.1017/S0953820812000295 Published online by Cambridge University Press

view the most interesting arguments in play for and against scalar utilitarianism, and attempt to make further progress with them.

II. THE PERSUASION ARGUMENT

To secure scalar utilitarianism, Norcross pursues a two-step strategy: first, to demonstrate that, by utilitarianism's own lights, rightness and wrongness cannot be an all-or-nothing matter, but rather matters of degree; and second, to demonstrate that, since the categories of goodness and badness are already matters of degree rather than all-or-nothing, the categories of rightness and wrongness are, in effect, redundant. I will not have anything to say about Norcross's second step here, but will focus squarely on his first step instead. Since I believe that his first step fails, the question of whether his second step succeeds does not arise in any acute way.[6]

Before going any further, we should pause to see what kind of dismissal of all-or-nothing rightness and wrongness Norcross is recommending. There would seem to be two main candidates. First, Norcross might be urging an *eliminativist* approach to rightness and wrongness, according to which no acts can be truthfully described as right or wrong. Second, and more weakly, he might be urging a *silencing* approach to rightness and wrongness, the upshot of which is that utilitarians should simply pay no attention to rightness and wrongness, even if the existence of these properties need not be strictly denied.

As I see it, it is the silencing approach, rather than the eliminativist approach, which matters more to scalar utilitarians. Presumably Norcross and other scalar utilitarians are not in a position to stop someone from *defining* all-or-nothing rightness in traditional utilitarian terms if she wants to, and that is going to make the eliminativist approach hard to sustain. The arguments he offers suggest instead that utilitarians ought to *pay no deep attention* to facts about rightness and wrongness, and that would appear to indicate his allegiance to the silencing approach.

Norcross's argument against the relevance of the all-or-nothing nature of rightness and wrongness turns on what he regards as the *arbitrariness* of the thresholds required to record the distinction between rightness and wrongness. For reasons that will soon become plain, I will refer to this argument as the *Persuasion Argument*.

[6] But see Rob Lawlor, 'The Rejection of Scalar Consequentialism', *Utilitas* 21 (2009), pp. 100–16, at 101–2, for suggestive criticism of the second step.

https://doi.org/10.1017/S0953820812000295 Published online by Cambridge University Press

Norcross outlines the Persuasion Argument in the following passage:

> Suppose Jones is obligated to give 10 percent of his income to charity. The difference between giving 8 percent and 9 percent is approximately the same, in some obvious physical sense, as the difference between giving 9 percent and 10 percent, or between giving 11 percent and 12 percent. Such similarities should be reflected in moral similarities. A moral theory which says that there is a *really significant* moral difference between giving 9 percent and 10 percent, but *not* between giving 11 percent and 12 percent, looks misguided . . . To see this, suppose that Jones were torn between giving 11 percent and 12 percent and that Smith were torn between giving 9 percent and 10 percent. The utilitarian will tell you to spend the same amount of time persuading each to give the larger sum, assuming that other things are equal. This is because she is concerned with certain sorts of consequences, in this case, with getting money to people who need it. An extra $5,000 from Jones . . . would satisfy this goal as much as an extra $5,000 from Smith . . .[7]

I shall now immediately introduce some more assumptions and terms which will play a role in the assessment of the Persuasion Argument. First, assume that Smith and Jones command identical income levels – call that amount N – so that percentile differences in what they donate to charity have the same monetary value, and therefore the same utility. Also imagine there to be a *utility spectrum*. Since Norcross stipulates that the right action is that which coincides with giving 10 per cent of N, we will take that point in the utility spectrum to define the *R-point*. Acts that consist of giving less than 10 per cent are wrong – thus acts that lie to the left of the R-point on the utility spectrum lie in the 'wrongness zone', or the *W-zone* for short. Acts that lie to the right of the R-point on the utility spectrum fall within the 'rightness zone', or the *R-zone* for short. I will also describe acts in both the W-zone and the R-zone in terms of percentile proportions of N. Thus a supererogatory act that consists in giving 20 per cent of N will be described as an act which produces 20N, a wrong act that consist in giving 3.5 per cent of N will be described as an act which produces 3.5N, and so on. Finally, we will call the persuader or adviser in this case *Zeus*. Zeus has to decide how to expend his persuasive energies: in Norcross's example, he can either persuade Smith to give 10N rather than 9N of his income, which is a choice between an act in the W-zone and an act which coincides with the R-point, or he can persuade Jones to give 12N rather than 11N of his income, which is a choice between two acts in the R-zone.

The leading idea behind the Persuasion Argument can be put rather simply. Since Zeus has equally weighty reasons, by the lights of utilitarianism, for persuading Smith to give 10N (assuming that Jones gives 11N) as he does for persuading Jones to give 12N (assuming that

[7] Norcross, 'Reasons Without Demands', p. 41; original emphases; Norcross, 'Scalar Approach', p. 220, offers an identically worded argument.

https://doi.org/10.1017/S0953820812000295 Published online by Cambridge University Press

Smith gives 9N), it supposedly follows that there can be no special significance in the distances on the utility spectrum that take us from the W-zone into the R-zone; they are worth no more in utility than equal-sized utility gaps lying exclusively in the W-zone or in the R-zone. But that should not be so, Norcross claims, if rightness and wrongness enjoy irreducible importance in a theory of utilitarianism. As he puts it: 'If the difference between right and wrong is at all significant, it must be possible for it to offset at least some differences in goodness.'[8]

So if rightness is at all significant, then we would expect Zeus to conceive his role in a bifurcated way: he would conceive of himself as having a responsibility to expend his persuasive energies in such a way that overall utility is maximized, but he would also conceive of himself as having a responsibility to ensure that the acts performed by the particular agents concerned, Smith and Jones, fall in the R-zone. But as a reliable utilitarian persuader, Zeus experiences no such dual conception of his role: his only aim is to ensure that utility is maximized. So the rightness of individuals' acts carries no significance in itself; the only thing that matters is that these acts jointly maximize utility. And the rightness of individual acts does not matter in itself *because* the only thing that matters is the maximization of utility. That is the intended lesson of the Persuasion Argument.

III. A DILEMMA FOR THE PERSUASION ARGUMENT

The Persuasion Argument plainly relies upon a certain strategy. It proposes that we appeal to Zeus's indifference to, specifically, the rightness and wrongness of Smith's and Jones's individual acts in order to support the claim that utilitarianism has no deep interest in rightness and wrongness *simpliciter*. The basic strategy, more completely stated, must be this: if Zeus pays no attention to a putatively relevant feature of utilitarianism in his role of a persuader, then that feature is not an indispensable feature of utilitarianism, and utilitarians ought to pay no attention to it.

The basic problem with this strategy, I shall argue, is that it proves too much: facts about Zeus's indifference *over-generate* features which are supposedly irrelevant or inessential to utilitarianism. But this means, in turn, that the facts which are relevant to Zeus as he goes about his business of being a persuader cannot be deemed to be a reliable guide to the essential or indispensable features of utilitarianism. The only way of escaping this conclusion, I will suggest, will be equally unpalatable to utilitarians. It will thus turn out that the Persuasion Argument is exposed to a dilemma.

[8] Norcross, 'Scalar Approach', p. 220.

https://doi.org/10.1017/S0953820812000295 Published online by Cambridge University Press

My analysis will revolve around the following simple case: Zeus can either persuade Jones to give 20N, where Smith gives nothing (call this combination of acts *S1*), or he can persuade Jones to give 10N, where Smith gives 10N (call this combination of acts *S2*). By assumption, Zeus will be indifferent between S1 and S2, because both of them produce 20N overall.

The two following (related) facts will not matter to Zeus:

(A) Jones's act of giving 10N in S2 falls in the R-zone, whereas Jones's act of giving nothing in S1 falls in the W-zone.
(B) Jones's act of giving 10N and Smith's act of giving 10N in S2 both fall in the R-zone, but only Jones's act of giving 20N in S1 falls in the R-zone.

The fact that (A) and (B) do not matter to Zeus – the fact that they do nothing to alter his indifference between S1 and S2 – is supposed to indicate, in turn, that the rightness or wrongness of individual acts should not matter to utilitarians. But here are some more facts about S1 and S2 which will also not matter to Zeus:

(C) Jones's act of giving 20N in S1 is better than Jones's act of giving 10N in S2.
(D) Smith's act of giving 10N in S2 is better than Smith's act of giving nothing in S1.
(E) Jones's act of giving 20N in S1 is better than Smith's act of giving nothing in S1.
(F) Jones's act of giving 10N in S2 is just as good as Smith's act of giving 10N in S2.

Now (C) to (F) compare the individual acts of giving mentioned in S1 and S2 in purely evaluative rather than deontic terms. The relevant terms are 'better than' or 'as good as'. (Further claims could, of course, be easily generated for the 'worse than' relation. Note that I am not aiming to provide an exhaustive catalogue of the relevant evaluative facts.) These comparisons are perfectly legitimate, and any form of utilitarianism, whether scalar or non-scalar, is committed to accepting their truth. The fact that some of these acts are *better* than others follows from commitments that lie at the very heart of the utilitarian creed. However, Zeus's role as a persuader renders it unnecessary for him to be sensitive to the particular evaluative facts recorded by (C) to (F). He does not *care* about these facts, because he does not *need* to care about them. They need not impinge on his deliberations. This is because Zeus's task, as a persuader, is effectively to manage an interpersonal utility portfolio – it is to maximize the sum of benefits jointly produced by Smith and Jones. Relative to *this* role, Zeus will be indifferent to

https://doi.org/10.1017/S0953820812000295 Published online by Cambridge University Press

whether the utility Smith and Jones *individually* produce falls in the R-zone.

These considerations can now, I believe, be deployed to generate a dilemma for the Persuasion Argument. On one horn of the dilemma, we consult the pattern of Zeus's deliberations in order to dismiss certain features of traditional utilitarianism – the features of traditional utilitarianism to which Zeus displays indifference in his role as a persuader. Since Zeus is systematically insensitive or indifferent to the rightness or wrongness of Smith's and Jones's individual acts, it follows, on this line of thought, that we should dismiss the significance of rightness and wrongness. But this exercise proves too much, since Zeus is *also* insensitive or indifferent to the 'better than' and 'as good as' evaluative facts recorded by (C) to (F). And even scalar utilitarianism does not deny the general significance of the fact that some acts are *better* than others!

On this first horn of the dilemma, then, utilitarians can rid themselves of rightness and wrongness just as long as utilitarians are prepared to abandon the significance of scalar, better-and-worse comparisons as well.

On the other horn of the dilemma, we allow the significance of better and worse acts to be reinstated, by denying that the patterns of systematic indifference or insensitivity manifested in Zeus's role as a persuader are indicative of the truly significant or indispensable conceptual features of utilitarianism. But if we do that, we are simply denying, in effect, that Zeus's role-guided decisions have *anything* necessarily to teach us about the central conceptual features of utilitarianism.

Norcross wishes to argue that Zeus's insensitivity to, specifically, the rightness or wrongness of Smith's or Jones's acts demonstrates that there is nothing special about rightness and wrongness. But the conclusion we should reach instead is that there is nothing special about Zeus's sensitivities and insensitivities. The scalar utilitarian case against the indispensability of rightness and wrongness has therefore not yet been demonstrated.

IV. SCALAR UTILITARIANISM AND DEONTIC ASSESSMENT

The argument so far has been negative. I have attempted to show that Norcross's Persuasion Argument fails to demonstrate that traditional deontic utilitarianism is ill-founded. Of course, the failure of the Persuasion Argument does not *vindicate* traditional non-scalar utilitarianism, nor does it show that scalar utilitarianism cannot be secured in some alternative way. To make further progress, we shall

https://doi.org/10.1017/S0953820812000295 Published online by Cambridge University Press

now examine utilitarianism's association with the deontic questions in a more direct manner. These considerations will allow us to amass a more positive case for the indispensability of the deontic categories to utilitarianism.

The deontic categories from which Norcross wishes scalar utilitarianism to escape association include the notions of, variously, demands, requirements, obligations, and duties. (I shall abstain from any detailed investigation of the conceptual relationships which obtain among these deontic categories, and will tend to refer, for ease of exposition, solely to obligations.) These further deontic categories are evidently dismissed as a corollary of his dismissal of the notions of rightness and wrongness. However, another deontic category, or a category which may appear deontic upon initial inspection – reason-providingness – is explicitly accommodated by Norcross's account. That is, Norcross concedes that scalar utilitarianism provides us with reasons for acting. As he clearly says: 'Consequentialist morality ... provides reasons for actions, without issuing demands (or permissions).'[9]

An ambiguity in the notion of a reason must be confronted at this point: are we referring to normative reasons or motivating reasons?[10] I think it obvious, upon closer inspection of his arguments, that Norcross himself is referring to motivating reasons, not normative reasons. In order to substantiate this allegation, I will turn to Norcross's reasoning later on, in section VIII, and show it to be largely beside the point.

Regardless of Norcross's own way of handling the problem, the more pressing task is to consider, from first principles, two questions. First, should scalar utilitarianism admit that utilitarian standards generate normative reasons for action? Second, if scalar utilitarianism does accommodate normative reasons for action, does it then follow that items belonging to other deontic categories – duties, obligations, and so on – will also be generated? My answers to these two questions, respectively, will be: yes, and yes. These answers are defended in the next two sections.

V. SCALAR UTILITARIANISM AND REASONS FOR ACTION

Let us consider the first question: should scalar utilitarians admit that utilitarianism generates normative reasons for action? A normative reason for action states a consideration which counts in favour of that

[9] Norcross, 'Reasons Without Demands', p. 38; see also Norcross, 'Scalar Approach', p. 230.

[10] See Michael Smith, *The Moral Problem* (Oxford, 1994), ch. 4, for a detailed discussion of this distinction.

https://doi.org/10.1017/S0953820812000295 Published online by Cambridge University Press

action. If scalar utilitarianism were reason-providing in a normative sense, it would be able to acknowledge that utilitarian considerations count in favour of agents' acting in certain ways and not others. Does that look like a desirable feature for scalar utilitarianism to possess? Yes, most certainly.

Imagine, to take an easy case – we need only an easy case to make the point – a morally conscientious agent, *Anna*, who is committed to utilitarianism, and who is prepared to respond appropriately to her judgements about the verdicts issued by utilitarianism. Anna surely *regards* herself, in judging that *F*–ing is the best thing to do from a utilitarian point of view, as *endorsing* those utilitarian considerations; she will surely affirm that those considerations count in favour of her *F*-ing, or that they give her reason to *F*. Let us suppose, furthermore, that *F*–ing is indeed the best option by utilitarianism's lights, and that utilitarianism is indeed the correct moral theory. Scalar utilitarians will now be able to avoid disabusing Anna of her belief that these utilitarian considerations count in favour of her *F*–ing if, but only if, they can acknowledge that she has a normative reason for *F*–ing. If that condition cannot be met, scalar utilitarians would have to tell Anna that she is mistaken. Even if the commitments endorsed in the content of her judgements are correct, they would not count in favour of action. And that breakage in the connection between moral theory and normative reasons for action would make it difficult to see how the supposedly evaluative commitments exhibited by scalar utilitarianism could be *genuinely* evaluative at all. Surely a conscientious moral agent such as Anna can treat evaluative considerations, if those considerations are true, as reason-giving!

We have every reason, then, for supposing that scalar utilitarians will affirm that utilitarianism provides agents with normative reasons for action.

VI. UTILITARIANISM, REASONS AND OBLIGATIONS

We now turn to the second question, which is whether the category of normative reason-providingness is qualitatively different from the particular deontic categories Norcross wishes to dismiss: the categories of obligations, duties, requirements, demands, and so on. We can pose the question in an alternative way: can seepage of deontic assessment from the category of normative reasons for action into the other deontic categories be prevented?

There is a strong prima facie reason for thinking that this seepage *cannot* be prevented. This is for the simple reason that, as for traditional forms of utilitarianism, reasons for action are going to be generated by the evaluative considerations recognized by

https://doi.org/10.1017/S0953820812000295 Published online by Cambridge University Press

utilitarianism, and that scalar utilitarianism does not recognize any evaluative considerations which compete with exercises in traditional utilitarian evaluation. Evaluatively, scalar utilitarianism does not pretend to be anything except *fully* utilitarian. Scalar utilitarianism is not a pluralist or revisionary theory, at least in comparison with traditional act-utilitarianism; it does not recognize options or constraints as possible counterweights to the generation of impersonal utility. To put the point in slogan form: scalar utilitarianism is evaluatively conservative, but deontically revisionary. But then it should follow that, if scalar utilitarianism is reason-providing, it will also generate obligations. It will *demand* from agents those very same actions which it gives those agents *conclusive reasons* for performing. Its evaluative conservativeness – its non-accommodation of non-utilitarian considerations – will see to that. The two parts of the scalar utilitarian slogan therefore stand in mutual tension. Evaluative standards give us reasons, and reasons lead in their turn to obligations.

One possible reply to this line of reasoning is that, although scalar utilitarianism might be fully utilitarian in an evaluative sense, Anna can have a reason to act in such a way as to produce, via action *A*, some small positive balance of utility, even if it were possible, via action *B*, for Anna to produce even more utility. Plausibly, the fact that *B* would produce even more utility than *A* does not *defeat* the claim that Anna has *a* reason to perform *A*. Both *A* and *B* lie in the utility-producing zone, so both are reason-giving. In this sense, then, the claim that Anna has a reason to perform an action which produces suboptimal amounts of utility remains resilient even if there is another action available to her which would produce more utility. If this line of argument is coherent, then utilitarianism may be able to issue normative reasons calling for the production of *some* utility, without having to be dragged further down the path of maximization and correspondingly sterner forms of deontic assessment.

This line of argument is unconvincing. First, it seems swiftly to unravel, and in the following way: if we are prepared to say that the agent has a reason to produce some minimal or modest amount of utility, then it cannot be reasonably denied that she is *obliged* or *required* not to produce any *less* than that modest amount of utility. Thus other deontic categories are going to be wheeled in more or less immediately, even though they will apply here only to modest levels of moral effort. Second, it is going to be difficult in any case to defend this kind of modesty for very long. If utilitarianism issues Anna with a reason to produce some utility, it surely cannot fail to issue her with a *weightier* reason to produce *even more* utility. But the weightier Anna's reasons become, the more likely it is that they will fall into the catchment area of more heavy-duty deontic assessment.

https://doi.org/10.1017/S0953820812000295 Published online by Cambridge University Press

VII. MCELWEE ON REASONS AND OBLIGATIONS

I have argued, first, that utilitarianism should attempt to provide for normative reasons for action, and second, that there is no plausible way of preventing a slide between reasons for action and sterner forms of deontic assessment, such as obligations. Yet this conclusion might be challenged. After all, the notion of an obligation may have a different structure from the notion of a reason for action, or even a conclusive reason for action. It all depends on what we have to add to a reason to end up with an obligation.[11]

In this section, I shall simply consider one interesting recent way of arresting deontic haemorrhage from reason-givingness to the other, sterner deontic categories, and show why I think it fails. Of course, even if my argument is deemed successful, that still leaves us with potential other challenges. I will therefore settle for the lesser claim that the onus of proof lies with scalar utilitarians, not with non-scalar or traditional utilitarians.

The particular proposal I wish to discuss comes from Brian McElwee.[12] McElwee contends that, while consequentialism successfully demonstrates that an agent's *strongest reason* – not just her strongest *moral* reason, but her strongest reason *simpliciter* – is to bring about the optimal state of affairs, it does not follow that she is *morally obliged* to bring about that state of affairs. To illustrate this point, let us return to conscientious Anna. McElwee's discussion implies that if we were to ascribe to Anna a moral obligation to *F*, we would be implying more than that Anna's strongest reason is to *F*. We would also be implying that, if Anna failed to *F*, she would merit feelings of blame and guilt. But it would be unreasonable, McElwee argues, to conclude that Anna merited feelings of blame and guilt just in case her conscientious efforts to live a decent moral life fell slightly short of the very strenuous levels of self-sacrifice required by unadulterated versions of consequentialism.

I am not convinced that McElwee has uncovered a stable option for would-be scalar consequentialists to occupy. Why would it be unreasonable to suppose that Anna merits feelings of blame and guilt? Because, McElwee tells us, it is unreasonable to expect her to comply with a moral theory which is as demanding as act-consequentialism.

[11] Norcross, 'Scalar Approach', pp. 224–7, also discusses the connections between wrongness and blame; Brad Hooker, 'Right, Wrong, and Rule-Consequentialism', *The Blackwell Guide to Mill's Utilitarianism*, ed. H. West (Oxford, 2006), pp. 233–48, at 239–41, offers a good reply.

[12] See Brian McElwee, 'Consequentialism, Demandingness and the Monism of Practical Reason', *Proceedings of the Aristotelian Society* 107 (2007), pp. 359–74. McElwee clearly has act-consequentialism in mind.

https://doi.org/10.1017/S0953820812000295 Published online by Cambridge University Press

But why would that expectation be unreasonable? It is far from clear what the verdict of unreasonableness is supposed to be tracking in this case. The division between reasonableness and unreasonableness cannot be related to the reasons generated by consequentialism itself, since, after all, and by McElwee's own admission, consequentialism *does* conceive of itself as generating conclusive normative reasons for Anna to engage in extreme levels of self-sacrifice.

Of course, it could be maintained that very conscientious individuals, such as Anna, who fall short of the highest standards of moral compliance merit *much less* blame and guilt than most of us, who fall considerably short of these standards. If we are going to extend blame to agents for being insufficiently benevolent, Anna will surely be placed at the back of the queue. (And that may mean, in actual practice, that Anna will not collect any real blame at all.) But this point establishes less than McElwee needs, for even moral extremists such as Shelly Kagan can say *that* much without being exposed to any pressure to withdraw their claims about the strenuousness of moral obligation.[13]

To conclude, it is simply unclear, in McElwee's picture, why the sanctions of merited blame and guilt are *not* triggered by Anna's failure to act on the reasons prescribed by consequentialism (except for the reasons which even moral extremists can accept). If one is going to refuse to allow the Demandingness Objection to demonstrate that consequentialism is erroneous, then some other explanation must be sought of why Anna does not end up with an obligation to do what consequentialism says she has most reason to do. What could that explanation be?

VIII. NORCROSS ON ACTION-GUIDINGNESS

As forewarned, I now return briefly to Norcross's own argument, in order to show why it fails. Norcross himself approaches these issues by considering the following claim: 'If a moral theory doesn't tell you what to do, it is no good.'[14] His reply is that scalar utilitarianism is not embarrassed by this test. Admittedly, scalar utilitarianism may not tell you what you *ought* to do, but Norcross suggests that scalar utilitarianism can still count as *action-guiding*, which means that it does manage to tell you *what to do*.

To demonstrate why, he distinguishes between internalism and externalism. As Norcross explains this distinction, if internalism is true, then the agent will be motivated to act in an appropriate way

[13] See Shelly Kagan, *The Limits of Morality* (Oxford, 1989).

[14] See Norcross, 'Reasons Without Demands', pp. 46–7; and Norcross, 'Scalar Approach', pp. 229–30.

https://doi.org/10.1017/S0953820812000295 Published online by Cambridge University Press

by the content of her judgement. If externalism is true, the agent will not be so motivated, unless she has an additional independent desire to perform the act in question. So either the agent will be motivated to act in accordance with her moral judgements, or she will not be so motivated: either way, the verdicts of scalar utilitarianism will not leave her high and dry. In this sense, Norcross believes, scalar utilitarianism succeeds in telling an agent what to do.

This argument, I think, misses the point. First, note that the options Norcross identifies are plainly internalism and externalism about *motivation*. But this yields little more than the point that, when Anna makes a moral judgement, *something or other will happen at some point after that*: either Anna will be motivated to act on that judgement, or she will not be. As a characterization of the claim that scalar utilitarianism tells you what to do, however, that is surely not ambitious enough. The initial worry was not directly concerned with the connection between Anna's judgement and her *motivation* – it was concerned with the connection between Anna's judgement and her *reasons for action*. The worry was that the verdicts of scalar utilitarianism cannot be construed as considerations which count in favour of anything Anna subsequently does, or fails to do. And, if this is the case, it does not matter whether Anna's judgement is followed, or not followed, by any actual motivation to act on it. For, in making the moral judgement, Anna will not be endorsing any considerations which can count in favour of her so acting. *That*, I believe, is the only problem which is genuinely fit to be considered under the heading of 'action-guidingness'. And Norcross does not address it.

IX. UTILITARIANISM AND REASONS INTERNALISM

I have argued that utilitarians will want to hang on to the claim that the evaluative claims supported by utilitarianism generate reasons for action. But is this really true? In this penultimate section, I will sketch a more general challenge which may suggest otherwise.

Consider consequentialists who are also internalists about reasons (or reasons internalists),[15] such as David Brink and Peter Railton.[16] Brink and Railton think that reasons for action are fixed, broadly speaking, by our desires and personal projects. Reasons internalism contrasts with reasons externalism, which denies any such dependence of truthful reason-ascriptions on facts about agents' desires and

[15] To be clear: the distinction between internalism and externalism about reasons differs from the distinction between internalism and externalism about motivation, which was under the spotlight in sect. VIII.

[16] See David Brink, *Moral Realism and the Foundation of Ethics* (Cambridge, 1989), esp. chs. 3, 8; and Peter Railton, *Facts, Values and Norms* (Cambridge, 2003).

https://doi.org/10.1017/S0953820812000295 Published online by Cambridge University Press

personal projects. Yet Brink and Railton espouse consequentialism as the correct moral theory. Here their nomination of reasons internalism as the correct account of the nature of reasons for action is not usually thought to be in conflict with their espousal of consequentialism – though of course either of these commitments may be open to attack on independent grounds.

Interestingly, this point does not apply only to utilitarianism, or consequentialism. A similar pattern of advocacy is also discernible among some non-consequentialist philosophers, such as Bernard Williams. In a well-known passage, Williams considers the types of criticism we may heap upon someone who behaves horribly towards his spouse:

> There are many things I can say about or to this man: that he is ungrateful, inconsiderate, hard, sexist, nasty, selfish, brutal, and many other disadvantageous things. I shall presumably say, whatever else I say, that it would be *better* if he were nicer to her . . . There is one specific thing the external reasons theorist wants me to say, that the man has a *reason* to be nicer . . . But if it is thought to be appropriate, what is supposed to make it appropriate, as opposed to (or in addition to) all those other things that may be said?[17]

An additional possible source of support for this line of thought is that the conditions applying to reasons for action do not seem to have much to do with normative moral theory. They are concerned instead with a delicate mixture of considerations to do with action-explanation and rationality.[18]

For handy reference, let us refer to this idea that evaluative verdicts and reasons for action can diverge as the *Divergence Thesis*. As we have seen, the Divergence Thesis is compatible with scalar utilitarianism, but enjoys broader scope, since it makes a more general claim about the gap between values and reasons in moral theory. Again, though there has been much opposition to reasons internalism, the very combination of reasons internalism and a set of robust standards of morality, whether consequentialist or otherwise, has not elicited much critical comment. This suggests that a theory of moral values can coexist comfortably enough with a theory of reasons which fails to coincide with it.

Or does it? I have my doubts. Scalar utilitarians who are excited by the Divergence Thesis are, I think, exposed to a dilemma. Either

[17] Bernard Williams, 'Internal Reasons and the Obscurity of Blame', in his *Making Sense of Humanity* (Cambridge, 1995), pp. 35–45, at 39 (emphases added); see also Williams, 'Internal and External Reasons', in his *Moral Luck* (Cambridge, 1981), pp. 101–13, at 110.

[18] See the references in n. 17. (The further literature on these issues is, of course, enormous.)

https://doi.org/10.1017/S0953820812000295 Published online by Cambridge University Press

the Divergence Thesis has no revisionary implications for the moral reasons which are generated by utilitarianism, or it does have revisionary implications. If there are no revisionary implications, the Divergence Thesis cannot prevent the further emergence of obligations; if there are revisionary implications, obligations will be avoided, but at the cost of placing the evaluative verdicts issued by utilitarianism under severe conceptual pressure.

Let us begin with the first horn of the dilemma. Suppose that reasons internalism is the correct theory about reasons for action. What would then follow is that our *all-things-considered* or *conclusive* reasons for action are settled by the nature of the psychological items found in what Williams calls the 'subjective motivational set' of each agent. But it is consistent with reasons internalism that the putative reasons – the considerations which putatively favour action – which are certified by the internalist test *as* reasons for action can still be truthfully described as having a distinctive source or character: the class of prudential reasons, of moral reasons, and so on. So, if internalism about reasons is true, we can still distinguish between *moral* reasons (which might be all and only considerations deriving from utility production) and our *all-things-considered* reasons (which will come from whatever the agent cares about, and will normally encompass much more than utility production). We have seen no convincing arguments for denying that scalar utilitarianism, *as* a form of utilitarianism, will sponsor the very same moral reasons as non-scalar utilitarianism. But if that is so, moral obligations will be generated in addition to, and in the wake of, those moral reasons for action.

But perhaps that argument will not work. Perhaps, because of their supposedly overriding nature, moral reasons need to be all-things-considered or conclusive reasons in order to count *as* moral reasons. This takes us to the second horn of the dilemma. If reasons internalism is true, we will not have the moral reasons, and therefore we will not have the moral obligations, which many non-scalar utilitarians suppose us to have. We will have only the reasons related to our subjective motivational sets. But this solution appears to throw out the baby with the bathwater. For in what sense can the verdicts of utilitarianism now be considered *action-guiding*? The concerns raised in sections V and VI will not have been answered. Go back to conscientious Anna. *Because* Anna is conscientious, *she* will most likely have internal reasons to do what utilitarianism tells her to do. But even in this favourable case, Anna will surely think that her moral reasons for action are generated by utilitarian considerations, not by the fact that she happens to care about utilitarian considerations.

In summary, the Divergence Thesis is a poisoned chalice for scalar utilitarians.

https://doi.org/10.1017/S0953820812000295 Published online by Cambridge University Press

X. CONCLUSION

In this article, I have attempted to repudiate Norcross's Persuasion Argument, and I have suggested, more positively, that utilitarianism cannot easily distance itself from a commitment to the existence of obligations. My argument was relatively simple. It had two steps. First, the evaluative assessments upheld by utilitarianism plausibly generate reasons for action. Second, these reasons for action plausibly generate obligations. In the penultimate section, I aired a worry arising from reasons internalism, and tried to show that it offers little succour to scalar utilitarianism. I conclude that traditional utilitarians need not fear defeat from their scalar rivals just yet.[19]

g.r.lang@leeds.ac.uk

[19] I thank an audience at Durham University, especially Richard Cookson and Nick Zangwill, and also Daniel Elstein, Ulrike Heuer, Rob Lawlor and Georgia Testa, for their comments on earlier versions of this article. I also thank an anonymous reviewer for useful comments on the penultimate version. I'm especially grateful to Rob Lawlor and Brian McElwee for many extremely helpful and instructive conversations about these issues.

https://doi.org/10.1017/S0953820812000295 Published online by Cambridge University Press