Transitivity, Moral Latitude, and Supererogation

DOUGLAS W. PORTMORE
Arizona State University

On the standard account of supererogation, an act is supererogatory if and only if it is morally optional and there is more moral reason to perform it than to perform some permissible alternative. And, on this account, an agent has more moral reason to perform one act than to perform another if and only if she morally ought to prefer its outcome to that of the other. I argue that this account has two serious problems. The first, *the latitude problem*, is that it has counterintuitive implications in cases where the duty to be exceeded is one that allows for significant latitude in how to comply with it. The second, *the transitivity problem*, is that it runs afoul of the plausible idea that the one-reason-morally-justifies-acting-against-another relation is transitive. I argue that both problems can be overcome by an alternative account: *the maximalist account*.

The term 'supererogatory' is a technical term that one can use however one wants. But I'll use it such that an act counts as supererogatory if and only if, in performing it, the agent exceeds some morally required minimum. As such, my use of the term coincides with such colloquial expressions as 'going the extra mile', 'doing more than one has to', and, most illuminatingly, 'going beyond the call of duty'.[1]

In this article, I'll consider two competing accounts of supererogatory acts so conceived: the standard account and the maximalist account. I'll argue that the maximalist account is the more plausible of the two, because it avoids two serious problems associated with the standard account: the latitude problem and the transitivity problem.

[1] On my conception, to perform a supererogatory act is morally to exceed the demands of morality – that is, to do something morally better than the minimum required by morality. Others conceive of supererogation differently. For instance, Dale Dorsey holds that to perform a supererogatory act is morally to exceed the demands of rationality – that is, to do something morally better than the minimum required by rationality. To illustrate the difference, suppose that morality demands that I give at least 10% of my surplus income to the poor, whereas rationality requires only that I give at least 5% of my surplus income to the poor. Assuming, then, that it is morally better to give 6% of one's surplus income to the poor than to give only 5%, giving 6% would count as supererogatory on Dorsey's conception but not on my own. Clearly, then, Dorsey and I are talking about different concepts even though we are using the same technical term to refer to them. And, of course, there's nothing wrong with using a technical term in whatever way one chooses. Nonetheless, I think that when ordinary people talk about going beyond the call of duty, they're talking about going beyond the call of *moral* duty, not going beyond the call of *rational* duty. See Dale Dorsey, 'The Supererogatory, and How to Accommodate It', *Utilitas* 25 (2013), pp. 355–82. And for this and other criticisms of Dorsey's view, see Alfred Archer, 'The Supererogatory and How Not To Accommodate It: A Reply to Dorsey', *Utilitas* 28 (2016), pp. 179–88.

© Cambridge University Press 2016
Utilitas Vol. 29, No. 3, September 2017
doi:10.1017/S0953820816000364

CrossMark
https://doi.org/10.1017/S0953820816000364 Published online by Cambridge University Press

I take the standard account to be as follows.

> *The Standard Account:* (1) For any subject S and any option X, S's X-ing is supererogatory if and only if there is an alternative option Y such that: (a) she is both morally permitted to X and morally permitted to Y, and (b) she has more moral reason to X than to Y. (2) For any subject S and any two alternative options X and Y, S has more moral reason to X than to Y if and only if she morally ought to prefer how things would be if she were to X to how things would be if she were to Y.[2]

One problem with the standard account (call it *the latitude problem*) is that it has counterintuitive implications in cases where the duty to be exceeded is one that allows for significant latitude in how to comply with it. To illustrate this sort of duty and the problem it generates for the standard account, consider the following two examples.

> *The Miser:* Mr Scrooge lived from 1920 to 1995. In all that time, he never did a thing for anyone except on Christmas Eve 1938, which is when he found a five-dollar bill on the ground and donated it to the Salvation Army. Moreover, had he not done this, he would have ended up never having done anything for anyone.

> *The Visit:* Carlos's grandmother was in the hospital last week. And although he had promised that he would visit her at least twice during the week, he ended up visiting her only once: on Tuesday. Moreover, had he not visited her then, he would not have visited her at all, for things only got busier for him as the week progressed (although it never got so busy that he was unable to visit his grandmother).

In both cases, the agent has significant latitude in how to comply with the relevant duty. In *The Miser*, the relevant duty is the imperfect duty of beneficence. This duty requires each of us 'to make the happiness of others a serious, major, continually relevant, life-shaping end'.[3] And this in turn requires each of us to act throughout our adult lives so as to exhibit a sufficient propensity to promote the happiness of others – sufficient, that is, to count as having adopted the happiness

[2] Note that this is substantively neutral. It does not, for instance, presuppose consequentialism. After all, it could be that an agent morally ought always to prefer how things would be if she were to refrain from violating Kant's categorical imperative to how things would be if she were to violate Kant's categorical imperative. For more on this as well as a defence of this account of moral reasons, see Douglas W. Portmore, *Commonsense Consequentialism: Wherein Morality Meets Rationality* (Oxford, 2011), ch. 3.

[3] Thomas E. Hill, Jr, *Human Welfare and Moral Worth* (Oxford, 2002), p. 206.

https://doi.org/10.1017/S0953820816000364 Published online by Cambridge University Press

of others as 'a serious, major, continually relevant, life-shaping end'. For having such a propensity is constitutive of having adopted the happiness of others as such an end. Thus, if someone with the abilities and opportunities that are typical of those living in the First World did little or nothing to promote the happiness of others, this would show that she had not truly adopted the happiness of others as such an end.[4]

But as this is an imperfect duty, it gives us a lot of leeway as to how and when to promote this end. For having a propensity to promote the happiness of others does not necessitate that we take advantage of every favourable opportunity to do so. It requires only that we do so to some sufficient extent.[5] Thus, Mr Scrooge was not obligated to donate that five-dollar bill to charity on Christmas Eve 1938. He was permitted instead to spend it on himself. After all, he was only eighteen years old at the time, and there were to be plenty of future opportunities for him to promote the happiness of others. That said, it was good that he donated it, as he did have more moral reason to donate it to charity than to spend it on himself. After all, he morally ought to prefer the way the things were as a result of his donating it to how they would have been had he spent it on himself and never did anything for anyone. Still, he was permitted to forego this opportunity to promote the happiness of others and to choose instead to do so to a greater extent in the future. So, in donating that five-dollar bill, Mr Scrooge performed a morally optional act that he had more moral reason to perform than to perform the morally permissible alternative of spending that five-dollar bill on himself. Thus, on the standard account, his donating that five-dollar bill counts as supererogatory.

Now, consider *The Visit*. Here, the relevant duty was for Carlos to visit his grandmother at least twice during the week. And this duty allowed for significant latitude in how he was to comply with it. For instance, he wasn't required to visit his grandmother on any particular day that week. He was just required to visit her at least twice. Thus, he was not obligated to visit her when he did – that is, on Tuesday. For he could have fulfilled his duty by visiting her on, say, Wednesday and Thursday, Thursday and Friday, or Wednesday and Friday. Nevertheless, it was good that he visited her on Tuesday. For he had more moral reason to visit her on Tuesday than not to do so. That is, he morally ought to have preferred the way things were as a result of his visiting her

[4] Hill, *Human Welfare*, p. 204.

[5] Perhaps, it also requires us to take advantage of certain golden opportunities – that is, opportunities that are especially favourable and not likely to arise again. For more on this, see Robert Noggle, 'Give 'Till It Hurts? Beneficence, Imperfect Duties, and a Moderate Response to the Aid Question', *Journal of Social Philosophy* 40 (2009), pp. 1–16.

https://doi.org/10.1017/S0953820816000364 Published online by Cambridge University Press

on Tuesday to the way things would have been had he not done so and, consequently, ended up not visiting her at all. So, in visiting his grandmother on Tuesday, Carlos performed a morally optional act that he had more moral reason to perform than to perform the morally permissible alternative of not visiting her on that day. And, so, the standard account implies that his visiting his grandmother that Tuesday was supererogatory.

The problem, then, is that this is implausible. The standard account implies both that (1) in donating that five-dollar bill to charity on Christmas Eve 1938, Mr Scrooge went beyond the call of duty and performed a supererogatory act and that (2) in visiting his grandmother on Tuesday, Carlos went beyond the call of duty and performed a supererogatory act. Yet, to the contrary, it seems that not only did they each fail to exceed the minimum that morality required of them, but that they each fell well short of that required minimum. In other words, the problem with the standard account is that one can meet its conditions 1a and 1b without exceeding the minimum that is demanded by the relevant duty.[6] Take the duty of beneficence. Since there is almost always more moral reason to act beneficently than to act self-interestedly, and since this duty almost never requires one to perform any specific beneficent act, an agent can easily meet conditions 1a and 1b without exceeding the minimum required by the duty of beneficence. To illustrate, suppose that for someone with Mr Scrooge's abilities and opportunities to exceed the minimum that the duty of beneficence requires, he would have to dedicate more than 20% of his time and resources to promoting the happiness of others over the course of his adult life. The problem, then, is that Mr Scrooge was able to meet conditions 1a and 1b in donating that five-dollar bill without coming anywhere close to meeting this moral minimum.

So if we're going to avoid the latitude problem, we'll need to adopt an account that, unlike the standard account, acknowledges that one can meet or exceed the call of such a duty only by performing a certain series of actions over time. Moreover, we'll need to adopt an account that acknowledges that in order for such a series of actions to count as supererogatory, it must more than just partially or minimally fulfil the relevant duty. Thus, we'll need to adopt what I call the *maximalist account of supererogation*.

On the maximalist account, we are to distinguish between two types of options: maximally specific options and non-maximally specific options. An option X is a maximally specific option if and only if there

[6] This sort of problem was first introduced by Gregory Mellema, 'Supererogation and the Fulfilment of Duty', *Journal of Value Inquiry* 25 (1991), pp. 167–75.

https://doi.org/10.1017/S0953820816000364 Published online by Cambridge University Press

is no *other* option that entails it – that is, if and only if there is no option Y such that Y-ing entails X-ing but not vice versa.[7] Otherwise, it's a non-maximally specific option. And an option X entails another Y if and only if X-ing without Y-ing isn't an option. To illustrate, I have both the option of walking and the option of walking in the park, and my walking in the park entails my walking, for I don't have the option of walking in the park without walking [8] But walking in the park at t_1 won't be a maximally specific option if walking fast in the park at t_1 is an option, because walking fast in the park at t_1 entails walking in the park at t_1. Likewise, walking fast in the park at t_1 won't be a maximally specific option if walking fast in the park at t_1 and then resting at t_2 is an option. And walking fast in the park at t_1 and then resting at t_2 won't be a maximally specific option if walking fast in the park at t_1 and then resting on a bench at t_2 is an option. And so on and so forth until we get to an option that is not entailed by any other option – that is, a maximally specific option.[9]

We must distinguish between these two types of options, because, on the maximalist account, they are to be assessed differently.

[7] Two options are not the same option (and, thus, are other options) if and only if it is not the case that each entails the other.

[8] Walking and walking in the park are both options; they're just not alternative options. X and Y are alternative options if and only if, although X-ing and Y-ing are each options, both X-ing and Y-ing is not an option.

[9] One may doubt that there will always be a maximally specific option – that is, an option that is not entailed by any other option. For instance, one may think that I have all the following options: (Opt_1) thinking of a number greater than one, (Opt_2) thinking of a number greater than two, (Opt_3) thinking of a number greater than three, and so on and so forth, *ad infinitum*. Thus, one may conclude that for any option Opt_n, there will always be another option Opt_{n+1} that entails it. But, in any case, if there are no maximally specific options, then I just need to replace talk of maximally specific options with talk of normatively maximal options, where an option is normatively maximal if and only if every option that entails it is normatively equivalent to it – that is, equivalent to it in terms of all normatively relevant factors. For more on this, see Campbell Brown, 'Consequentialism, Maximalism, and the Structure of Acts' (2016, unpublished manuscript). Still, one may further worry that there won't always be a normatively maximal option. Suppose, for instance, that there is no limit to how large a number I can think of and that for whatever number I do think of, God will give me and my loved ones precisely those many days in heaven. In that case, no matter what number n I think of, there will be an alternative option, Opt_{n+1}, that will be better in terms of the normatively relevant considerations. And so, in this case, there seems to be no normatively maximal option. But this is a problem for the maximalist account only if we think that I must have a permissible option in this sort of case. After all, if there is no permissible normatively maximal option that entails my opting for Opt_n, the maximalist will just claim that my opting for Opt_n is impermissible. And, to my mind, this seems exactly right. In this case, it seems that I will do something impermissible no matter what I do, because no matter what I do there will be a better and no costlier option that I could have just as easily opted for but didn't. Thanks to Ángel Pinillos and Brad Armendt for pressing me on this.

https://doi.org/10.1017/S0953820816000364 Published online by Cambridge University Press

The Maximalist Account: (1) For any subject S and maximally specific option M_x, S's M_x-ing is supererogatory if and only if there is an alternative maximally specific option M_y such that: (a) she is morally permitted both to M_x and to M_y, and (b) she has more moral reason to M_x than to M_y. (2) For any subject S and any non-maximally specific option N_x, S's N_x-ing is supererogatory if and only if both: (a) in N_x-ing, she does not merely minimally or merely partially fulfil some positive duty D; and (b) there is an alternative non-maximally specific option N_y such that: (i) she is morally permitted both to N_x and to N_y, and (ii) she has more moral reason to N_x than to N_y. (3) For any subject S and any two maximally specific options M_x and M_y, S has more moral reason to M_x than to M_y if and only if she morally ought to prefer how things would be if she were to M_x to how things would be if she were to M_y. (4) For any subject S and any two non-maximally specific options N_x and N_y, S has more moral reason to N_x than to N_y if and only if she morally ought to prefer how things would be if she were to perform an optimal maximally specific option that entails her N_x-ing to how things would be if she were to perform an optimal maximally specific option that entails her N_y-ing.[10]

Now, when it comes to maximally specific options, the maximalist account and the standard account converge on the same verdicts. The two diverge only when it comes to non-maximally specific options. To illustrate, take Carlos's non-maximally specific option of visiting his grandmother on Tuesday. Whereas the standard account holds that this option is supererogatory given that he had more moral reason to perform it than to perform the permissible alternative of not visiting her that Tuesday, the maximalist account holds that this was not supererogatory given that, in performing this option, he merely partially fulfilled his duty to visit her at least twice that week. Let me explain.

In X-ing, S merely partially fulfils a positive duty D if and only if her X-ing is a proper subset of some set of actions by which she merely minimally fulfils D. And, in X-ing, she merely minimally fulfils a positive duty D if and only if, although her X-ing fulfils D, there is no more moral reason for her to X than to perform any other option

[10] I defend (4) in Douglas W. Portmore, 'Perform Your Best Option', *The Journal of Philosophy* 110 (2013), pp. 436–59. But see Joshua Gert, 'Perform a Justified Option', *Utilitas* 26 (2014), pp. 206–17, for some criticisms. In *Commonsense Consequentialism*, I defend (1), (2) and (3). And I defend the general idea that the deontic status of an act (e.g. baking) that is entailed by another (e.g. baking a pie) is derivative of the deontic statuses of the acts that entail it in both Douglas W. Portmore, 'Maximalism and Moral Harmony', *Philosophy and Phenomenological Research* (forthcoming) and Douglas W. Portmore, 'Maximalism vs. Omnism about Permissibility', an unpublished manuscript available at <http://bit.ly/1Svw0Eh> (2016).

https://doi.org/10.1017/S0953820816000364 Published online by Cambridge University Press

that would fulfil D. Now, Carlos has several ways of merely minimally fulfilling his duty to visit his grandmother at least twice this week: (W1) visit her on Monday and Tuesday, (W2) visit her on Tuesday and Wednesday, (W3) visit her on Monday and Wednesday, (W4) visit her on Wednesday and Thursday, etc. These are all ways of merely minimally fulfilling his duty, because, although they each fulfil his duty, he has no more moral reason to perform any one of them than to perform any other alternative that would fulfil his duty. Thus, visiting his grandmother on Tuesday counts only as partially fulfilling his duty to visit her at least twice this week. For, in visiting her on Tuesday, he performs only a proper subset of a set of acts that would merely minimally fulfil his duty, such as the set consisting of his visiting her only on Tuesday and Wednesday.

To exceed the minimum required by this duty, Carlos would have to do something that he has more moral reason to do than to just visit her twice that week – e.g. visit her three or more times. Thus, on the maximalist account, visiting her only on Tuesday would not count as supererogatory, but visiting her on Tuesday, Wednesday and Thursday would. The latter would count as a supererogatory, because Carlos is morally permitted both to (N_1) visit her on Tuesday, Wednesday and Thursday and to (N_2) visit her on only Tuesday and Wednesday, and Carlos has more moral reason to N_1 than to N_2. He has more moral reason to N_1 than to N_2, because he morally ought to prefer how things would be if he were to perform an optimal maximally specific option that entails his N_1-ing to how things would be if he were to perform an optimal maximally specific option that entails her N_2-ing, where M_x is an optimal maximally specific option for S if and only if M_x is a maximally specific option for S and there is no alternative maximally specific option for her, M_y, such that she has more moral reason to M_y than to M_x. The idea, then, is that an optimal maximally specific option that entails his N_2-ing isn't going to be as morally good as an optimal maximally specific option that entails his N_1-ing. And, therefore, he has more moral reason to N_1, making his N_1-ing supererogatory.

So it seems that, unlike the standard account, the maximalist account is immune to the latitude problem. For it has plausible verdicts in cases such as *The Visit*, where the duty to be exceeded is one that allows for significant latitude in how to comply with it. This gives us at least one reason to prefer the maximalist account. But there is another reason to prefer the maximalist account: it is immune to another problem facing the standard account – viz. the transitivity problem.[11] This problem has been presented by both Frances Kamm and Dale Dorsey, but I'll

[11] I borrow the label 'transitivity problem' as well as the basic way of presenting the problem as a kind of paradox from Alfred Archer, 'Moral Obligation, Self-Interest and

https://doi.org/10.1017/S0953820816000364 Published online by Cambridge University Press

focus on Dorsey's formulation of the problem.[12] The problem arises with respect to the following sort of triad of cases.

> '*Gus:* Gus finds himself the recipient of an inheritance from a wealthy relative. This inheritance will allow Gus to buy a new car, which Gus desires to do, and which will allow Gus to see much more of his significant other, who lives in a distant town. Alternatively, Gus could [use this inheritance to] . . . save ten people from death'.[13]
>
> *Roy:* Roy could save ten people from death, but only by assaulting Jerry.
>
> '*Stan:* Stan knows that Jerry has just inherited a substantial amount of money from a wealthy relative. Were Stan to intimidate Jerry into giving him the money as a result of beating Jerry up, this would be a prudential benefit to Stan, given that this would allow him to buy a new car, from which he will derive pleasure, and which will allow Stan to see much more of his significant other, who lives in a distant town. (Assume that Stan would avoid punishment.)'[14]

The problem is that the following four claims concerning this triad each seem individually plausible but are jointly inconsistent.

> (C1) The self-interested reason there is to buy a new car for oneself morally justifies acting against the moral reason that there is to save ten people from death. Thus, Gus is permitted to buy a new car for himself, and saving the ten is supererogatory, not obligatory.
>
> (C2) The moral reason there is to save ten people from death morally justifies acting against the moral reason there is to refrain from assaulting someone. Thus, Roy is permitted to assault Jerry.

the Transitivity Problem', *Utilitas* (forthcoming). Admittedly, Dorsey doesn't present the problem as one involving transitivity, but Archer and I find it illuminating to think of it in this way. Moreover, Dorsey's triad seems relevantly similar to the triad presented by Frances Kamm, and she does present the problem as one involving transitivity. See Frances M. Kamm, 'Supererogation and Obligation', *The Journal of Philosophy* 82 (1985), pp. 118–38. I should also note that Archer offers a different (and ingenious) solution to the transitivity problem, one that involves denying that the moral justifying strength of a moral reason must be proportionate to its moral requiring strength. I don't find this denial plausible and so favour the solution that I offer below. In *Commonsense Consequentialism*, I argue that the moral justifying strength of a *non-moral* reason (e.g. a self-interested reason) can be greater than its moral requiring strength, but I believe that the moral justifying strength of a *moral* reason must be proportionate to its moral requiring strength.

[12] See Dorsey, 'The Supererogatory' and Kamm, 'Supererogation and Obligation'.

[13] See Dorsey, 'The Supererogatory', p. 365.

[14] See Dorsey, 'The Supererogatory', p. 366.

https://doi.org/10.1017/S0953820816000364 Published online by Cambridge University Press

(C3) The self-interested reason there is to buy a new car for oneself does not morally justify acting against the moral reason there is to refrain from assaulting someone. Thus, Stan is not permitted to assault Jerry.

(C4) The one-reason-morally-justifies-acting-against-another relation is transitive.

C4 is difficult to deny.[15] And yet each of C1–C3 also seems plausible, and they jointly entail that C4 is false. So we have a paradox. I believe the solution to it lies in realizing that, depending on how we spell out the details of the three cases, we should reject either C1 or C2, which we can do if, and only if, we accept the maximalist account of when there is more moral reason to perform one non-maximally specific option than to perform another. Thus, the solution lies both with further specifying the details of the three cases and with adopting the maximalist as opposed to the standard account of supererogation.

There are at least two ways of further specifying the details in the three cases. One way (I'll call it *the strangers way*) is to substitute 'save ten distant strangers from death' for 'save ten people from death' throughout the descriptions of the three cases. The other way (I'll call it *the relatives way*) is to substitute 'save ten close relatives from death' for 'save ten people from death' throughout the descriptions of the three cases. I'll consider them in turn.

To start, then, let's take the strangers way of further specifying the details. When we do so, we get the following versions of the first two cases. (The third case remains the same, since it makes no mention of saving ten people from death.)

'*Gus[v1]:* Gus finds himself the recipient of an inheritance from a wealthy relative. This inheritance will allow Gus to buy a new car, which Gus desires to do, and which will allow Gus to see much more of his significant other, who lives in a distant town. Alternatively, Gus could donate his inheritance to Oxfam International, which would save ten [distant strangers] . . . from death'.[16]

[15] As Josh Gert has explained, the strength value of a reason is a concise representation of the way it affects the deontic statuses of acts across a range of contexts. Thus, if one reason has greater moral justifying strength than another, it must make it morally permissible to do anything that the other reason would make it morally permissible to do. That is just what it means for one reason to have greater moral justifying strength than another. And, if that's right, the one-reason-morally-justifies-acting-against-another relation must be transitive. See Joshua Gert, 'Normative Strength and the Balance of Reasons', *Philosophical Review* 116 (2007), pp. 533–62.

[16] See Dorsey, 'The Supererogatory', p. 365.

https://doi.org/10.1017/S0953820816000364 Published online by Cambridge University Press

> *Roy(v1):* Roy knows that Jerry has just inherited a substantial amount of money from a wealthy relative. Were Roy to intimidate Jerry into giving him the money as a result of beating him up, Roy could donate the ill-gotten money to Oxfam International, which would save ten distant strangers from death. If Roy doesn't beat Jerry up, Jerry is going to use it to buy himself a new car, and those ten people that Roy could have saved will die. Everything else, though, will be the same. Thus, we are to assume that Roy would avoid punishment even if he were to assault Jerry.

If this is how we are to think of the three cases, then we should, I believe, reject C2, which says: 'The moral reason there is to save ten people from death morally justifies acting against the moral reason there is to refrain from assaulting someone. Thus, Roy is permitted to assault Jerry.'[17] Admittedly, if we were to adopt the standard account, it would be difficult, if not impossible, to deny C2. For, on the standard account, it would seem that Roy has more moral reason to assault Jerry than to refrain from assaulting Jerry. After all, it seems that he morally ought to prefer, other things being equal (and they are stipulated to be equal), the world in which he assaults Jerry and thereby saves ten distant strangers to the world in which he refrains from assaulting Jerry and fails to save those ten distant strangers.

But if we accept the maximalist account, we can, and should, deny that Roy has more moral reason to assault Jerry than to refrain from assaulting Jerry. To see this, compare how things would be if Roy were to perform an optimal maximally specific option that entails his assaulting Jerry to how things would be if he were to perform an optimal maximally specific option that entails his refraining from assaulting Jerry. And let's suppose (plausibly, I believe) that any optimal maximally specific option will strike a balance between the use of one's limited time and resources to promote the happiness of others and the use of one's limited time and resources to promote (or respect) various other moral goals, such as justice, beneficence, self-improvement, and caring for one's family. That is, let's assume that an optimal maximally specific option will not be one where one dedicates all of one's time and resources just to promoting the happiness of others and, so, at the expense of neglecting other important moral ends, such as justice, self-improvement, and the well-being of one's family. Given this, Roy will not be morally required to save as many distant

[17] We should also deny the part of C1 that says, 'Saving the ten is supererogatory'. At best, Gus merely partially fulfils the imperfect duty of beneficences in saving the ten distant strangers.

https://doi.org/10.1017/S0953820816000364 Published online by Cambridge University Press

strangers as he could save if he were to set aside these other important moral goals. After all, we could all save more distant strangers if we were to disregard the well-being of our families, disregard our own self-improvement, and disregard justice by, say, robbing the rich to give to the poor. Thus, it seems that Roy will be morally required to save the lives of no more than n distant strangers over the course of his life, where n is significantly less than the number of distant strangers that he could save were he to disregard all other moral goals.

Admittedly, then, Roy could save more than n strangers if he were to some extent to neglect these other moral goals. But it would not be morally better for him to do so. After all, to exceed n is, by definition, to fail to strike an acceptable balance between the goal of saving more distant strangers and the goal of doing more to promote (or respect) various other moral ends. Moreover, if Roy were to insist on saving, say, $n + 10$ strangers, the best way for him to do so would not be by his setting aside the moral goal of respecting justice and assaulting Jerry, but instead by his setting aside some prudential goal (e.g. the goal of taking more overseas vacations) or some less important moral goal (e.g. the goal of improving himself a bit more). It seems, then, that if Roy were to insist on saving $n + 10$, it would be better for him to do so by using less of his time and resources on either his self-improvement or his overseas vacations than by obtaining additional resources through the assault of Jerry.[18] And, if all this is right, we should not think that Roy morally ought to prefer how things would be if he were to perform an optimal maximally specific option that entails his assaulting Jerry (one in which he saves the lives of $n + 10$ strangers over the course of his life and assaults Jerry) to how things would be if he were to perform an optimal maximally specific option that entails his refraining from assaulting Jerry (one in which he saves the lives of only n strangers over the course of his life and

[18] In the case of *Roy*, we are to assume that everything besides whether these ten distant strangers live or die will be the same regardless of whether or not he assaults Jerry. Thus, we are to assume that Roy will, in the future, dedicate the same amount of time and resources to each of his moral and prudential goals regardless of whether he assaults Jerry and saves these ten at present. And this is why advocates of the standard account must hold that Roy has more reason to threaten Jerry than to refrain from doing so, for how things would be if he were to assault Jerry is to be preferred to how things would be if he were not to do so. But, on maximalism, we don't look at how things would be if he were, at present, to act this way or that way. Instead, we look at how things would be if he were to perform this or that maximally specific option over the course of his future. And we derive his reasons for performing this or that action at present in terms of what these maximally specific options entail his doing at present. And what we find is that all his optimal maximally specific options entail his saving not n plus these ten strangers, but only n strangers.

https://doi.org/10.1017/S0953820816000364 Published online by Cambridge University Press

assaults no one). And, so, we should think that Roy does not have more moral reason to assault Jerry than to refrain from assaulting Jerry. Thus, we should deny that the moral reason that Roy has to save these ten strangers morally justifies his acting against the moral reason he has to refrain from assaulting Jerry. And, so, we should deny C2 – or, at least, we should so long as we adopt the maximalist account and take the strangers way of further specifying the details in these three cases.

But what if we were instead to take the relatives way of further specifying the details? In that case, we would get the following different versions of the first two cases.

> '*Gus[v2]:* Gus finds himself the recipient of an inheritance from a wealthy relative. This inheritance will allow Gus to buy a new car, which Gus desires to do, and which will allow Gus to see much more of his significant other, who lives in a distant town'.[19] But ten of Gus's close relatives have just been infected with a rare virus. And, unfortunately, these ten cannot afford the antiviral drugs that they need to survive. Of course, if Gus were to act altruistically, Gus would use his inheritance to purchase the needed antiviral drugs and save his ten relatives. But if he doesn't, they'll die.
>
> *Roy(v2):* Ten of Roy's close relatives have just been infected with a rare virus. Unfortunately, neither Roy nor these ten can afford the antiviral drugs they need to survive. But Roy knows that Jerry has just inherited a substantial amount of money from a wealthy relative. Were Roy to intimidate Jerry into giving him the money as a result of beating Jerry up, Roy would use the money to purchase the needed antiviral drugs and save his ten relatives. If Roy doesn't beat Jerry up, Jerry is going to use it to buy himself a new car and Roy's ten relatives will die. Everything else, though, will be the same. Thus, we are to assume that Roy would avoid punishment even if he were to assault Jerry.

And if this is how we are to think of the three cases, then we should, I believe, reject C1, which says: 'The self-interested reason there is to buy a new car for oneself morally justifies acting against the moral reason that there is to save ten people [in this case, ten close relatives] from death. Thus, Gus is permitted to buy a new car for himself, and saving the ten is supererogatory, not obligatory.' It seems to me that, in *Gus[v2]*, the relevant duty is the duty to take any decent opportunity to save the life of a close relative. And there is no way

[19] See Dorsey, 'The Supererogatory', p. 365.

https://doi.org/10.1017/S0953820816000364 Published online by Cambridge University Press

for Gus to fulfil this duty without using his inheritance to save his ten relatives. Thus, Gus's saving the ten, in *Gus[v2]*, is not supererogatory, but obligatory. This contrasts with *Gus[v1]*, where the relevant duty seems to be the imperfect duty of beneficence – the duty to adopt the happiness of others as an end and, thus, to act throughout one's adult life so as to exhibit a sufficient propensity to promote this end. And Gus needn't use his inheritance to save these ten particular strangers to fulfil this duty. For it doesn't matter, morally speaking, whether he forgoes this opportunity to benefit himself and saves these ten strangers or takes this opportunity to benefit himself but forgoes some later opportunity to benefit himself so as to save some other ten strangers. Either way, he will exhibit the same propensity to promote the happiness of others. But because, in *Gus[v2]*, the relevant duty is not the imperfect duty of beneficence but the duty to take any decent opportunity to save the life a close relative, we must deny that Gus is permitted to use his inheritance to buy a new car for himself rather than use it to save his ten relatives. Thus, we should reject C1 if we take the relatives way of further specifying the details of these three cases.

Either way, then, we have a solution to the transitivity problem. For no matter how we further specify the details of these three cases, saving the ten will either be an instance of taking one of several alternative paths to partially fulfilling an imperfect duty or it won't. If it is, as in the case of the strangers way of filling in the details, it will be plausible to deny C2 – or, at least, it will so long as we adopt the maximalist account of supererogation. And, if it isn't, as in the case of the relatives way of filling in the details, it will be plausible to deny C1. So, either way, we resolve the transitivity problem if we adopt the maximalist account.

To sum up, then, we seem to have two good reasons for favouring the maximalist account over the standard account, for only it is immune to both the latitude problem and the transitivity problem.[20]

dwportmore@gmail.com

[20] I thank Dale Dorsey, Brad Hooker, and the participants at the 2015 Workshop on Supererogation (Concept and Context) in Basel, Switzerland, for helpful comments on earlier drafts of this article.

https://doi.org/10.1017/S0953820816000364 Published online by Cambridge University Press