Moral Obligation, Self-Interest and the Transitivity Problem

ALFRED ARCHER
University of Tilburg

Is the relation 'is a morally permissible alternative to' transitive? The answer seems to be a straightforward yes. If Act B is a morally permissible alternative to Act A and Act C is a morally permissible alternative to B then how could C fail to be a morally permissible alternative to A? However, there are cases where this transitivity appears problematic. My aim in this article is to provide a solution to this problem. I will then investigate two ways in which we might justify rejecting the transitivity of the 'is a permissible alternative to' relation. Next, I will look at Dorsey's solution, which involves a reinterpretation of the intuitions used to generate the problem. I will argue that none of these solutions are fully satisfying, before going on to provide a novel solution to the problem and to argue that it avoids the problems facing the alternative solutions.

INTRODUCTION

Is the relation 'is a morally permissible alternative to' transitive? The answer seems to be a straightforward yes. If Act B is a morally permissible alternative to Act A and Act C is a morally permissible alternative to B then how could C fail to be a morally permissible alternative to A? However, as both Dale Dorsey[1] and Frances Kamm[2] point out, there are cases where this transitivity appears problematic. I will call this problem the Transitivity Problem.

This problem raises important issues for moral philosophy. If it turns out that this relation is intransitive then this might raise concerns about our ability to give a systematic account of what makes an act morally permissible or impermissible. As a result, this discussion will have implications for the question of which normative ethical view fits best with common-sense morality. After all, it is particularly incumbent upon consequentialists to give a full ranking of acts from better to worse and use this ranking to determine which acts are morally permissible. If it turns out that 'is a morally permissible alternative to' is an intransitive relation then this raises serious doubts about the prospects

[1] Dale Dorsey, 'The Supererogatory, and How to Accommodate It', *Utilitas* 25 (2013), pp. 355–82.

[2] Frances M. Kamm, 'Supererogation and Obligation', *Journal of Philosophy*, 82 (1985), pp. 118–38, and Frances M. Kamm, *Morality Mortality Volume 2: Rights, Duties and Status* (Oxford, 2001), ch. 12.

© Cambridge University Press 2016
Utilitas Vol. 28, No. 4, December 2016
doi:10.1017/S0953820816000091

CrossMark
https://doi.org/10.1017/S0953820816000091 Published online by Cambridge University Press

for doing so.[3] Finally, this discussion about moral requirements and transitivity raises issues about practical requirements more generally. If it turns out that 'is a morally permissible alternative to' is an intransitive relation then we might wonder whether the same is true for the 'is a rationally permissible alternative to' relation.

My aim in this article is to provide a solution to this problem. I will start in §I by giving a precise outline of the problem. I will then, in §II, set out three desiderata that a satisfactory solution to the problem should satisfy. In the next two sections I will investigate three potential solutions to the problem that fail to satisfy these three desiderata. First in §III, I will explore two ways in which we might justify rejecting the transitivity of the 'is a permissible alternative to' relation. I will then, in §IV, investigate Dorsey's solution, which involves a reinterpretation of the intuitions used to generate the problem in §I. I will then, in §V, lay out a novel solution to the problem which explains why transitivity should be rejected by outlining a new theory of moral reasons that holds that there are three different moral roles that reasons can play: a moral justifying role, a moral requiring role and a moral favouring role. I will finish in §VI by considering and responding to an objection that could be raised against my solution.

I. THE PROBLEM

In this section I will explain what the Transitivity Problem is. To start I will present three examples:

> *Case One:* Suppose that Charlie, a student nurse, is leaving his house one morning to attend a final year exam. As he leaves his house he notices an injured woman, Jane, lying on the pavement. After a quick assessment, it is clear to Charlie that, though her condition is far from critical, Jane needs to go to hospital. Moreover, the more quickly she gets to hospital the less complicated her recovery will be. He is then faced with a choice. He could drive Jane to the hospital himself, this would get her there more quickly but it would also mean missing his exam, as it is being held on the other side of town from the hospital. Missing the exam will mean having to wait two months to take a re-sit. Alternatively, he could phone an ambulance and wait with her until it comes. This would mean that he could still attend the exam, though it would increase the time it takes for the woman to get to hospital and, as a result, increase the time it takes her to recover from her injury.

[3] A similar point is made in Kamm, 'Supererogation', p. 136.

https://doi.org/10.1017/S0953820816000091 Published online by Cambridge University Press

Case Two: Charlie from Case One decides that he will drive Jane to the hospital. Unfortunately, his car won't start. Charlie thinks through his options. He has the keys to his neighbour's house, as he has been watering her plants while she is on holiday. He also knows where she keeps her car keys. He could, then, borrow his neighbour's car. However, Charlie's neighbour has not given him permission to borrow her car.

Case Three: Charlie from Cases One and Two leaves his house to go to the exam but this time he does not find an injured woman on the street. Unfortunately, his car won't start. Charlie thinks about what to do. If he doesn't drive to the exam he will not get there in time. He realises he could borrow his neighbour's car, though he does not have permission to do so.

Having presented these three cases I will now make three claims about these cases.

Claim One: In Case One it is morally permissible for Charlie to go to the exam rather than taking Jane to hospital. It would, though, be morally better for Charlie to drive Jane to hospital. The reason for this is that Charlie's self-interested reasons morally justify the performance of the less morally favoured act.[4]

In Case One it is plausible to think that whilst it would be morally good for Charlie to drive Jane to the hospital, he is not morally required to do so. In other words, it looks as if doing so would be supererogatory or beyond the call of duty. Moreover, it seems reasonable to think that what prevents the act from being morally obligatory in this case are the self-interested reasons that count against driving Jane to the hospital. If, for example, the exam were being held at the hospital then it might be plausible to say that Charlie is morally required to drive her there. This example, then, is a case where the morally best act available, taking the injured woman to the hospital, is prevented from generating an obligation by the self-interested reasons that count against performing the act.

This claim is one that would be endorsed by many theorists working on the nature of supererogation. It is commonly held that what prevents supererogatory acts from being obligatory are the costs to the agent's self-interest that performing the act involves.[5] According to Douglas Portmore's recent defence of this position, reasons that do not

[4] This is not to say that there are no moral reasons that support going to the exam but just that taking Jane to hospital is the morally better act of the two.

[5] See Jonathan Dancy, *Moral Reasons* (Oxford: 1993), p. 138; Gregory Mellema, *Beyond the Call of Duty: Supererogation, Obligation and Offence* (New York, 1991), p.

https://doi.org/10.1017/S0953820816000091 Published online by Cambridge University Press

count for or against performing the act from the moral point of view may nevertheless play a morally justifying role and make otherwise forbidden acts permissible.[6] Even those who don't think that this is true for all acts of supererogation accept that it is true for many such acts.[7] If this is right and this act is an example of such a case, then it is the agent's self-interested reasons that serve to morally justify the non-performance of the morally better act.

> *Claim Two:* In Case Two it is permissible and morally better for Charlie to take the car in order to drive the injured woman to hospital. This is because the moral reasons that support taking the woman to hospital override the moral reasons that count against taking the car.

In normal circumstances, it would be morally wrong for Charlie to borrow his neighbour's car without permission. However, clearly there are cases where it would be permissible to do so. Suppose, for example, that Charlie would save someone's life by taking the car. In this case it seems permissible to do so. The importance of respecting the neighbour's property rights is surely outweighed by the fact that a life is at stake. Any plausible moral theory then, will accept that it can be morally permissible to violate one person's property rights in order to prevent some level of harm befalling some other person. While transporting someone to hospital more quickly is not quite as important as saving a life, it seems reasonable to think that this is sufficiently important to justify taking the car and that the reason this is so is that this act is morally better than not taking the car and waiting for the ambulance. Both Dorsey and Kamm draw similar conclusions from similar cases. Dorsey claims that it would be morally permissible to beat someone up if this would save ten people from death.[8] Beating someone up is normally morally wrong, but in the case where it is morally better to do so than to not do so it is morally permissible. Kamm claims that it is morally permissible to break a promise to meet a friend for lunch in order to save a life. Again, breaking a promise is

179; and Douglas Portmore, *Commonsense Consequentialism: Wherein Morality Meets Rationality* (Oxford, 2011), ch. 5.

[6] Portmore, *Commonsense Consequentialism*, ch. 5.

[7] E.g. Terence Horgan and Mark Timmons, 'Untying a Knot from the Inside Out: Reflections on the "Paradox" of Supererogation', *Social Philosophy and Policy* 27 (2010), pp. 29–63. Elsewhere, I have also argued against the view that all acts of supererogation can be explained in this way. See Alfred Archer, 'Supererogation, Sacrifice and the Limits of Duty', *Southern Journal of Philosophy* (forthcoming). However, I do think that this will explain why some acts of supererogation are not obligatory.

[8] Dorsey, 'The Supererogatory', p. 366.

https://doi.org/10.1017/S0953820816000091 Published online by Cambridge University Press

usually morally wrong, but in the case where doing so can save a life it seems morally permissible.[9]

> *Claim Three:* In Case Three it is not permissible for Charlie to take his neighbour's car in order to get to the exam on time. This is because the moral reasons that count against taking the neighbour's car override the self-interested reasons that count in favour of doing so.

In Case Three, it does not seem plausible to think that it is morally permissible for Charlie to borrow his neighbour's car without permission. The fact that Charlie will miss his exam if he does not borrow the car does not seem to justify taking the car without permission. Dorsey draws a similar conclusion from the case he considers. Dorsey says that while it is permissible to beat someone up to save ten lives, it is not permissible to do so in order to get the money to buy a new car.[10]

Claims One, Two and Three are independently plausible yet together they seem puzzling. If we accept Claims One, Two and Three then this shows that 'is a morally permissible alternative to' is not a transitive relation. After all, the result of these three claims is that going to the exam is a morally permissible alternative to driving Jane to hospital. This in turn is a morally permissible alternative to not taking a neighbour's car without permission. However, going to the exam is not a morally permissible alternative to not taking a neighbour's car without permission.[11]

The basic thought that makes this puzzling is that it seems reasonable to think that if some reasons can play a moral justifying role when they conflict with moral reasons of a certain strength then they should also be capable of doing so when they conflict with weaker moral reasons. We can formalize this claim in the following way:

> *Claim Four:* If reasons R are capable of playing a moral justifying role against moral reasons of strength s then, for any positive value of x, R should be capable of playing that role against moral reasons of strength $s - x$.[12]

[9] Kamm, *Morality*, p. 314.

[10] Dorsey, 'The Supererogatory', p. 366.

[11] This issue is raised by Kamm, *Morality*, pp. 312–13. Kamm, though, says that it is the relation 'may take precedence over' that fails to be transitive.

[12] This is the way the problem is raised in Dorsey, 'The Supererogatory', p. 367. It is worth noting that Claim Four is necessary but not sufficient for transitivity. Even if we reject the claim that the 'is a morally permissible alternative relation' is transitive, then, we might still want to accept Claim Four. Thanks to Alexander Bird for helpful discussion here.

https://doi.org/10.1017/S0953820816000091 Published online by Cambridge University Press

This claim seems plausible, yet it is incompatible with the other three. To see why, let's stipulate that the moral reasons that count in favour of taking the woman to the hospital have strength *s*. Claim Two says that the moral reasons that count against taking the neighbour's car without permission are less strong than those that count in favour of driving the woman to hospital. Let's call this difference in strength *x*. Now we can see that if we accept Claim Four then if the costs to Charlie's self-interest can morally justify not performing the morally better act in Case One then they should also be able to in Case Three. Claims One, Two, Three and Four, then, are independently plausible but jointly inconsistent.

Before proceeding it is worth briefly mentioning that this problem is distinct from a related issue of transitivity that has been explored in depth by Larry Temkin and Stuart Rachels.[13] According to Temkin and Rachels a number of problems arise if we accept that the 'is better than all things considered' relation is transitive. The problem I will consider here, though, concerns the transitivity of the 'is a morally permissible alternative to' relation. Of course, we might think that if the first form of transitivity is problematic then it is unsurprising that the second is as well. Nevertheless, the problem discussed above is not dependent on any problems with the form of intransitivity that Temkin and Rachels discuss. After all, the 'is better than all things considered' does appear to be operating as if it were transitive in these three cases. Driving Jane to hospital is morally better than refusing to take the neighbour's car, which in turn is better than taking the car in order to go to the exam. Transitivity is preserved because it is also the case that driving Jane to hospital is better than driving to the exam. This problem is worth considering then, even if we reject the transitivity of the 'is better than all things considered' relation.

II. THREE DESIDERATA

In this section I will briefly outline three desiderata that a successful solution to the problem should satisfy.

First, a satisfying solution will respect the intuitions captured by Claims One, Two, Three and Four. We saw in §I good reason to accept these claims and, all else being equal, a solution that respects these

[13] See Larry, S. Temkin, *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning* (Oxford, 2012), and Stuart Rachels, 'Counterexamples to the Transitivity of Better Than', *Australasian Journal of Philosophy* 76 (1998), pp. 71–83. Ingmar Persson also denies the transitivity of the permissibility relation because he denies the transitivity of the betterness relation (though his reasons for denying this relation differ from Temkin's). See Ingmar Persson, *From Morality to the End of Reason: An Essay on Rights, Reasons, and Responsibility* (Oxford, 2013), p. 39.

https://doi.org/10.1017/S0953820816000091 Published online by Cambridge University Press

intuitions will be preferable to one that does not. Of course, there is more than one way to respect intuitions. We can do so by providing a vindicating explanation or we can respect the force of these intuitions by offering a debunking explanation. Either way, though, we should be looking for a solution to provide some kind of explanation for these intuitions. I will call this desideratum RESPECT THE INTUITIONS.

Second, we should seek a solution to the problem that does not lead us to abandon the project that led us to consider The Problem of Transitivity in the first place. This problem arises out of attempts to provide a unified explanation of the connection between moral permissibility and reasons for action. One way of responding to the problem, then, would be simply to abandon the project of looking for such a connection. However, this response should only be adopted as a last resort. Given that this project is one of great philosophical interest, both in its own right and for the potentially interesting implications for moral philosophy, we should look for a solution to this problem that does not require us to abandon this project. I will call this desideratum SAVE THE PROJECT.

Finally, we should look for a solution that provides a full solution to the problem. This desideratum is fairly self-explanatory but it is worth briefly stating what a full solution to the problem would be. A full solution to the problem will be one that does not leave any puzzling aspects of the problem unaddressed. This requires not only addressing the problem as outlined in §I but also any problems that arise from the proposed solution. I will call this desideratum SOLVE THE PROBLEM.

III. REJECTING TRANSITIVITY

In this section I will investigate whether it is plausible to solve the problem by rejecting the transitivity of the 'is a morally permissible alternative to' relation. There are two ways to justify rejecting this transitivity and I will examine each in turn.

III.1. Kamm's solution

Kamm's response to this problem is to provide an explanation as to why this relation is not transitive. According to Kamm, there is no single scale upon which we can decide whether one act may take priority over another. The reason Kamm gives for this is that the acts in the three cases are being evaluated from different points of view: the point of view of overall good, the point of view of personal interests and goals, and the point of view of the minimal standards of morality. We are, says Kamm, entitled to pursue our personal goals rather than performing a supererogatory act because of a personal prerogative to pursue our own projects and goals that is derived from the value of

https://doi.org/10.1017/S0953820816000091 Published online by Cambridge University Press

people as ends-in-themselves.[14] We are also, says Kamm entitled to perform a supererogatory act instead of a duty because this is best from the point of overall good from the impartial perspective.[15] However, we are not entitled to pursue our personal projects rather than performing our duties because of the moral importance that everyone meets the minimal standards of morality.[16]

The reason why there is no transitivity here is because some acts (minimal standards of duty) are requirable at high personal cost while others are not. So we can be required to sacrifice our personal preference in order to meet the minimal standards of duty even though we would not be required to sacrifice our personal preference in order to perform a supererogatory act. This explains why the 'is a permissible alternative relation' is intransitive. As these three acts are being evaluated from three different points of view and no point of view dominates over both of the others, we should not expect there to be a transitive relation between them.

When we apply this solution to our three cases we can see that it does a good job of respecting our intuitions. Claim One said that in Case One it is morally permissible for Charlie to go to the exam rather than taking Jane to hospital. Kamm's account can justify this intuition by saying that this is a case of a personal preference making it morally permissible not to perform a supererogatory act. Claim Two says that in Case Two it is permissible and morally better for Charlie to take the car in order to drive the injured woman to hospital. Kamm's account can justify this intuition by pointing out that in this case the fact that taking the woman to hospital will be better from the point of view of overall good means that it is permissible to perform this act instead of respecting the neighbour's property rights. After all, the costs facing Jane if she does not get to the hospital will be greater than those facing Charlie if he does not get to his exam on time. Finally, Kamm can also capture Claim Three that in Case Three it is not permissible for Charlie to take his neighbour's car in order to get to the exam on time. Kamm's account can capture this claim by saying that the minimal standards of duty require respecting property rights over the pursuit of personal projects in this case. Kamm's account, then, is able to satisfy RESPECT THE INTUITIONS.

Kamm's account also seems to satisfy SOLVE THE PROBLEM. Kamm's account provides a vindicating explanation for the apparent intransitivity of the morally permissible alternative to relation.

[14] Kamm, *Morality*, pp. 336–7 and Frances M. Kamm, *Intricate Ethics: Rights, Responsibilities, and Permissible Harm* (Oxford, 2007), p. 31.
[15] Kamm, *Morality*, p. 336, and Kamm, *Intricate Ethics*, p. 31.
[16] Kamm *Morality*, p. 336, and Kamm, *Intricate Ethics*, p. 31.

https://doi.org/10.1017/S0953820816000091 Published online by Cambridge University Press

According to Kamm, the reason this relation is not transitive is that these acts are being assessed from different points of view and that none of these points of view dominates the other two. This explains why the 'is a morally permissible alternative to' relation is not transitive. Though Kamm does not make this point explicit, her account also points towards a reasonable debunking explanation for the intuitive oddness of denying transitivity here. Kamm can say that the reason that the intuitive plausibility of the transitivity depends upon the false belief is that we can order these different points of view in order of normative importance. Kamm's account, then, not only explains why we should reject the transitivity of this relation but also explains why it seemed plausible to think that this relation would be transitive.

However, Kamm's account does less well with respect to the final desideratum, SAVE THE PROJECT. As I explained in §II, ideally we want a solution to the problem to allow us to explain the connection between moral permissibility and reasons for action. Kamm's account, though, leaves us unable to make any general claims about the connection between moral reasons and moral permissibility. On Kamm's account, we can say that in some situations it will be permissible to act in line with what is best from the point of view of overall good rather than what would meet the minimal standards of duty. In other cases it will be permissible to pursue our self-interest rather than overall good. However, given the importance of maintaining the minimal standards of duty, in many cases it will not be permissible to fail to meet the minimum standards of duty in order to pursue one's self-interest. Kamm's solution, then, abandons the project of trying to provide a unified account of the connection between moral reasons, self-interested reasons and moral permissibility.

Of course, this problem seems unlikely to worry Kamm much. As a non-consequentialist, Kamm is not committed to there being any unified way of explaining the connection between moral reasons and moral permissibility. Kamm's approach to moral philosophy is one that seeks to develop intricate principles that capture the complexity of our moral intuitions.[17] Given the complexity of our moral intuitions in these cases, it seems reasonable for Kamm to respond by claiming that we should not expect to be able to give a unified account of the connection between moral reasons and moral permissibility. Moreover, to demand this from Kamm might be thought to beg the question against her non-consequentialist approach. Starting from the assumption that we should be able to give a unified account of the connection between

[17] See Kamm, *Intricate Ethics*, pp. 14–15.

https://doi.org/10.1017/S0953820816000091 Published online by Cambridge University Press

moral reasons and moral permissibility might reasonably be thought to prejudice the debate against a deontologist's solution to the problem.

For these reasons, I think it would be wrong to say that Kamm's failure to satisfy SAVE THE PROJECT gives us reason to reject her solution. However, given that the project of finding a unified account of the connection between moral reasons and moral permissibility is of great importance to consequentialists, I think it does give us reason to ask whether an alternative solution can be offered that does not give up on this project. After all, if such a solution to this problem cannot be found then this seems likely to count as a mark against consequentialism. This, I think, gives us sufficient reason to ask whether we can give a solution that allows us to give a more unified account of the connection between moral reasons and moral permissibility.

III.2. A holist solution

Another way in which we might be tempted to reject Claim Four is to endorse some form of Holism about reasons. Reasons Holism is defined by Jonathan Dancy as the view that, 'a feature that is a reason in one case may be no reason at all, or an opposite reason, in another'.[18] This view contrasts with Reasons Atomism, which Dancy defines as the view that, 'a feature that is a reason in one case must remain a reason, and retain the same polarity, in any other'.[19] Accepting Holism opens up a new way to justify rejecting Claim Four. The reason why the 'morally permissible alternative to' relation is not transitive is that reasons that play one role in one case may play no role or an opposite role in one or more of the other cases.

However, the problem with this solution is that if this were the case then we should expect this to show up in our evaluations of Cases One, Two and Three. If the reasons in one of the cases were not reasons with the same polarity in one or more of the other cases then we should be capable of appreciating this fact. Unfortunately for this response, though, this is not the case.[20] The self-interested reasons that count in favour of going to the exam in Case One seem to count in favour of doing so in Cases Two and Three as well. Similarly, the moral reasons

[18] Jonathan Dancy, *Ethics without Principles* (Oxford, 2004), pp. 73–4. It is worth noting that Dancy's account of reasons offers additional resources for responding to this problem. In particular, his distinction between enticing and peremptory reasons might be thought to provide a possible solution to this problem. See Jonathan Dancy, 'Enticing Reasons', *Reason and Value: Themes From the Moral Philosophy of Joseph Raz*, ed. R. Jay Wallace, Philip Pettit, Samuel Scheffler and Michael Smith (Oxford 2006), pp. 91–118. See §VI for a discussion of this distinction.

[19] Dancy, *Ethics*, pp. 73–4.

[20] A similar response is made by Dorsey, 'The Supererogatory', p. 367 n. 16.

https://doi.org/10.1017/S0953820816000091 Published online by Cambridge University Press

that count in favour of taking the woman to hospital seems to count in favour of doing so in all of the cases. Finally, the moral reasons that count against taking the car in Case Two seem to count against doing so in Case Three as well.

Perhaps, though, Holism might be thought to undermine the motivation for thinking that this relation would ever be transitive. After all, if reasons for action are context sensitive then there is no reason to expect this relation to be transitive. However, while it is true that a holist should have no expectations about the transitivity of the permissible alternative relation and so no reason to accept Claim Four, they do have reason to accept the following claim:

> *Claim Four*:* If reasons R are capable of playing a moral justificatory role against moral reasons of strength s then, *assuming that reasons R retain their strength and polarity*, for any x, R should be capable of playing that role against moral reasons of strength $s - x$.

While holists about reasons may have no reason to accept Claim Four they do have reason to accept Claim Four*. After all, if we are comparing cases where the reasons in play retain their strength and polarity then surely the 'is a permissible alternative to' should be transitive. Simply accepting Holism, then, may give us good reason to reject Claim Four but it does not provide a solution to the problem.

However, Claim Four* points towards an alternative solution that is available to the holist. Rather than claiming that it is the existence or polarity of the reasons that alters between the three cases, a holist could instead claim that it is the *strength* of the reasons that changes. As Sean McKeever and Mike Ridge point out, a reason may retain its polarity in different contexts but vary in strength as the result of the presence or absence of intensifiers or diminishers.[21] If the strength of a reason is context sensitive then it should be no surprise that the 'is a permissible alternative to' relation is not transitive.

Perhaps, then, the holist could argue that there are features of the context in Cases One, Two or Three that are playing an intensifying or diminishing role. There are three options here. First, the self-interested reason that supports going to the exam could be being intensified in Case One or diminished in Case Three. Second, the moral reason that counts in favour of taking the woman to hospital could be diminished in Case One or intensified in Case Two. Finally, the moral reason that counts against taking the neighbour's car could be diminished in Case Two or intensified in Case Three. Unfortunately for this solution, it

[21] Sean D. McKeever and Mike Ridge, *Principled Ethics: Generalism as a Regulative Ideal* (Oxford, 2006), p. 28.

https://doi.org/10.1017/S0953820816000091 Published online by Cambridge University Press

does not seem as if any of these options are plausible. Most obviously, the strength of the moral reason that supports taking the woman to hospital appears completely unchanged between Cases One and Two. Similarly, the strength of the self-interested reason to go to the exam seems just as strong as in Cases One and Three. Finally, in Cases Two and Three, it does not seem to be the strength of the moral reasons against taking the neighbour's car that changes. In all of these cases, it appears as if these reasons are overridden rather than diminished.

Of course, there is a limit to how much weight we should give these intuitions. After all, while the strength of these reasons may not appear to change between these three cases this does lead to the counterintuitive result that the permissibility relation is not transitive. It is open then for the advocate of this approach to make a case for the existence of intensifiers or diminishers in one or more of these cases. However, in order for this solution to work the holist will need to point out what feature of the contexts in these three cases is playing an intensifying or diminishing role. Without such an explanation this solution will fail to satisfy SOLVE THE PROBLEM. This solution, then, will only be a live option if such an account could be given. In the absence of such an account we appear to have good reason to look elsewhere for a solution to the Transitivity Problem. Moreover, even if an account of the relevant intensifiers and diminishers were on the table there seems to be good reason to seek out an alternative account, for the simple reason that if no alternative account can be found then this might well be thought to provide an argument in favour of the existence of intensifiers and diminishers. We should not, then, rule out the possibility of some form of holist solution. We do, though, have sufficient reason to look elsewhere for a solution to the problem.

IV. DORSEY'S SOLUTION

Dorsey suggests an alternative response in his attempt to provide a solution to 'The Problem of Supererogation'.[22] Dorsey argues that it is a mistake to view supererogatory acts as acts that are morally better than the acts we are *morally* required to perform. Instead we should view these acts as morally better than the acts we are *rationally* required to perform.[23] According to Dorsey, we are always morally required to perform the morally best act available. However, not all acts that we

[22] Elsewhere I argue that Dorsey's attempt to solve the Paradox of Supererogation fails. See Alfred Archer, 'The Supererogatory and How Not to Accommodate It: A Reply to Dorsey', *Utilitas* (forthcoming). For the purposes of this article, though, I only aim to show that Dorsey's solution to the Transitivity Problem fails.

[23] Dorsey, 'The Supererogatory', p. 371. Dorsey expands his discussion in his forthcoming book *The Limits of Moral Authority* (Oxford, 2016), but his criticism of

https://doi.org/10.1017/S0953820816000091 Published online by Cambridge University Press

are *morally* required to perform are also acts that we are *rationally* required to perform.[24] If we accept Dorsey's view then it seems like we should reject Claims One and Two. Given that we are morally required to do what is best, in both Case One and Case Two we are morally required to drive to the hospital. Claims One and Two then should be revised as follows:

> *Claim One**: In Case One it is *rationally* permissible for Charlie to go to the exam rather than taking Jane to hospital. In other words, Charlie's self-interested reasons *rationally* justify the performance of the less morally favoured act.

> *Claim Two**: In Case Two it is *rationally* permissible and morally better for Charlie to take the car in order to drive the injured woman to hospital. This is because the moral reasons that support taking the woman to hospital override the moral reasons that count against taking the car.

Dorsey can, though, accept Claim Three as it stands. It remains the case on his view that Charlie is morally required not to take his neighbour's car in Case Three.

However, given that Dorsey is committed to accepting claims One* and Two* it looks as if he is also committed to accepting that it is rationally permissible for Charlie to take the neighbour's car without permission in order to get to the exam on time. After all, if it is rationally permissible to choose going to the exam over driving the woman to hospital and it is rationally permissible to choose driving the woman to hospital even if this means taking the neighbour's car without permission then surely it should also be rationally permissible to choose to go to the exam even if this means taking the neighbour's car without permission. The problem does not stop here, as if we accept Dorsey's account then not only is it rationally permissible for Charlie to take the car in Case Three, it is also supererogatory not to do so. This seems to be the wrong result. If you are running late for an exam it is not supererogatory to decide not to take a neighbour's car without permission, it is obligatory. Dorsey's solution then, fails to satisfy RESPECT THE INTUITIONS.

While acknowledging the force of these objections, Dorsey claims that these problems are far less serious than those we are faced with if

Portmore and his own account remain in all important respects the same. For ease of reference I will focus on his published paper.

[24] Dorsey, 'The Supererogatory', pp. 372–3. See also Dale Dorsey, 'Weak Anti-Rationalism and the Demands of Morality', *Noûs* 46 (2012), pp. 1–23.

https://doi.org/10.1017/S0953820816000091 Published online by Cambridge University Press

we accept the traditional view of supererogation. Dorsey supports this claim in three ways.[25]

First, there are additional resources available for responding to this objection that are not available to someone seeking to respond to the original problem. As the problem facing Dorsey's view is one of rational permission rather than moral permission he can appeal to additional non-moral reasons that might count against taking the car in Case Three. Dorsey says that reasons of neighbourliness, association and etiquette may provide additional non-moral reasons not to take the car in this case. If we accept that these additional non-moral reasons are in play in this case then this allows Dorsey to say that Charlie is not rationally permitted to take the car and, as a result, it would not be supererogatory for him not to do so.

This response, though, rests upon the controversial claim that these reasons are non-moral rather than moral reasons. Much more would need to be said before we can accept that this is a satisfactory response to the problems facing Dorsey's view. Moreover, if we accept this response in Case Three then it looks like it should also work in the other cases. In Case One for example, there appear to be a number of similar reasons that might be in play. Reasons of community and etiquette might well count in favour of driving to the hospital in this case. This, though, does not seem sufficient to prevent the act from being supererogatory.

Second, we could restrict the supererogatory to a subset of those acts that are morally better than what is rationally required. Dorsey suggests that if we restrict the supererogatory to beneficent acts then we avoid having to class not taking the car in Case Three as supererogatory.[26]

There are, though, two problems with this view. First, this is a view that several authors writing on supererogation argue against.[27] In order for this response to work Dorsey must respond to these arguments. More worryingly, though, even if we restrict the term 'supererogatory' to acts that are beneficent, the act of not taking the car in Case Three would remain beyond the call of duty on Dorsey's terms and this seems no less counter-intuitive than describing the act as supererogatory. This response fails to provide a full solution to the problem and so does not satisfy SOLVE THE PROBLEM.

[25] Dorsey, 'The Supererogatory', pp. 380–1.

[26] Dorsey, 'The Supererogatory', p. 381. Cf. David Heyd, *Supererogation: Its Status in Ethical Theory* (Cambridge, 1982), pp. 115, 136, 137.

[27] See Alfred Archer, 'Supererogation and Intentions of the Agent', *Philosophia* 41 (2013), pp. 447–62, and Mellema, *Beyond the Call of Duty*, pp. 19–20.

https://doi.org/10.1017/S0953820816000091 Published online by Cambridge University Press

Finally, Dorsey claims that even if we accept that these are genuine problems for his view these worries are significantly less problematic than the problems facing those who accept the standard picture of the supererogatory. The problem facing the standard view is that it appears to be committed to the view that it is morally permissible to take the car in Case Three. Dorsey claims that this is far more problematic than a commitment to saying that not taking the car is supererogatory. At best, though, this only gives us reason to think that Dorsey's solution is preferable to biting the bullet and accepting that it would be morally permissible for Charlie to take the car in Case Three. Given that we have already rejected this option, Dorsey's comparative point gives us little reason to accept his solution.

As it stands, then, none of these replies offers a satisfactory response to the problems facing this view. Furthermore, redefining supererogation in the way Dorsey suggests is problematic in its own right. The reason this is problematic is that on Dorsey's view we are unable to say everything we want to say about these three cases. Take Case One. We wanted to say that it is not morally wrong to go to the exam rather than to drive to the hospital. On Dorsey's view this is false. More generally, we want to say that the acts that saints and heroes perform are morally better than what they morally have to do; not that they are better than what they are *rationally* required to do. Again on Dorsey's view this would be wrong. The fact that Dorsey's account is revisionary in this way is far from a decisive objection. Nevertheless, given that his account also faces the problem of wrongly characterizing not taking the neighbour's car in Case Three I take this to be sufficient reason to look elsewhere for a solution.

V. THREE ROLES FOR MORAL REASONS

So far I have looked at two groups of solutions to the Transitivity Problem. One involved reinterpreting Claims One, Two and Three and the other involved rejecting Claim Four. I will now propose a novel solution that allows us to hold on to Claims One, Two and Three and endorse a slightly modified version of Claim Four.

To start let's look more closely at Portmore's claim, mentioned briefly in §I, about the two moral roles that reasons can play. According to Portmore we can distinguish between two forms of strength that moral reasons can possess:

> *Moral Requiring Strength:* A reason has moral requiring strength to the extent that it can make it morally impermissible to refrain from performing acts that it would otherwise be morally permissible to refrain from performing.

https://doi.org/10.1017/S0953820816000091 Published online by Cambridge University Press

> *Moral Justifying Strength:* A reason has moral justifying strength to the extent that it can make it morally permissible to perform acts that it would otherwise be morally impermissible to perform.[28]

Portmore claims that non-moral reasons, reasons that do not count for or against performing an act from the moral point of view, may nevertheless play a moral justificatory role. It is important to note though that once we have accepted that moral justifying strength is to some extent independent from moral requiring strength then an interesting possibility emerges. It might be the case that moral reasons can also possess moral justifying strength without possessing the same level of moral requiring strength.[29]

Once we accept the possibility that the moral requiring strength of a moral reason might be different from the moral justifying strength then an interesting solution to the Transitivity Problem emerges. Perhaps, the lack of transitivity between Cases One, Two and Three can be explained by the fact that the moral reasons in play possess moral justifying strengths that are not identical to their moral requiring strengths. On this view, the moral justifying strength of a moral reason can be different from its moral requiring strength. In fact this solution appears very attractive when we consider the three cases. In Case Two driving the woman to hospital is morally better than opting not to take the neighbour's car without permission. Note that though this act is morally better, it is not morally required. This suggests that an act's moral requiring strength may not be directly tied to how morally good the act is. This would also explain why the act that is less favoured in Case Two may nevertheless have more requiring force. We can see how this might work by assigning numbers to represent the moral requiring and justifying force of the reasons that support each act in the way shown in Table 1.

[28] Portmore, *Commonsense Consequentialism*, p. 121. Compare Joshua Gert's criteria for *rational* requiring strength and *rational* justifying strength in Joshua Gert, *Brute Rationality: Normativity and Human Action* (Cambridge, 2004). A related distinction is Dancy's distinction between peremptory and enticing reasons. According to Dancy, peremptory reasons are those that 'are in the business of telling us what to do' ('Enticing Reasons', p. 91). Enticing reasons, on the other hand, 'are more to do with making an option attractive rather than demanded, required or right' (Dancy, 'Enticing Reasons', p. 91). We might think that this distinction would also provide a solution to the Transitivity Problem. However, it seems plausible to think that all of the possible acts in the three cases are supported by peremptory reasons. Given that Dancy claims that requiring reasons will always defeat peremptory reasons ('Enticing Reasons.' p. 99), it doesn't look like this distinction will have sufficient resources to reconcile the four claims made in §I.

[29] We might think that Portmore is implicitly committed to this view, as he claims that some moral reasons may possess no moral requiring force (*Commonsense Consequentialism*, p. 122 n. 6 and p. 135 n. 22).

https://doi.org/10.1017/S0953820816000091 Published online by Cambridge University Press

Table 1

	Moral Requiring Strength	Moral Justifying Strength
Go To Exam	0	5
Take Woman to Hospital	5	10
Don't Take Neighbour's Car	8	8

Of course, we might quibble with the way in which the numbers are assigned in each case. We might think, for example, that going to the exam is supported by moral reasons with some moral requiring strength. The important point, however, is not about how these numbers have been assigned but in the possibilities that assigning numbers in these ways opens up. Let's assume that an act is morally permissible if and only if it is supported by reasons that possess moral justifying strength that is at least equal to the moral requiring strength of the reasons supporting any of the alternative acts. We can now see why it is that the 'is a permissible alternative to' relation is not transitive. Going to the exam is supported by reasons with sufficient moral justificatory force to make it a permissible alternative to taking the woman to hospital. This in turn is supported by reasons with sufficient moral justificatory force to make it permissible for Charlie to take the woman to hospital even though it means taking his neighbour's car without permission. However, going to the exam is not supported by reasons with sufficient moral justificatory force to make it permissible for Charlie to take his neighbour's car without permission. Note that this view allows us to retain Claim Four with one slight modification:

> *Claim Four****:* If reasons R are capable of playing a moral justificatory role against moral reasons of *requiring* strength s then, for any x, R should be capable of playing that role against moral reasons of *requiring* strength $s - x$.

The original version of Claim Four did not stipulate which form of strength (requiring or justifying) was being referred to. Once we make clear that we are discussing requiring strength then we can hold on to this claim. Strictly speaking, then, if we accept this solution then 'is a morally permissible alternative to' is not a transitive relation. However, this solution can provide a perfect explanation as to why Claim Four appeared plausible. The reason is that it is very close to a relation that does hold, namely that picked out by Claim Four**. This solution, then, satisfies RESPECT THE INTUITIONS by offering a vindicating explanation for the intuitions behind Claims One, Two and Three and a debunking explanation for the intuitions that motivated Claim Four.

https://doi.org/10.1017/S0953820816000091 Published online by Cambridge University Press

Moreover, this solution is able to accommodate these intuitions in a way that pushes forward the project of attempting to explain the connection between moral permissibility and reasons for action. As a result, this solution also satisfies SAVE THE PROJECT.

As it stands, though, an important problem remains. We might wonder how we can talk about one act being better than another if reasons can possess two different kinds of moral strength. In Case Two, for example, we want to say more than just that both acts are morally permissible. We also want to be able to say that taking the woman to hospital is the morally better of the two. At the moment though, it seems that all we can say is that this act possesses greater moral justificatory force but less moral requiring force.

We can solve this problem by positing a third kind of moral strength that reasons can possess. In addition to possessing moral requiring and moral justifying strength, reasons can also possess *moral favouring* strength. We can define this in the following way:

Moral Favouring Strength: A reason has moral favouring strength to the extent that it increases how favourably the act should be judged from the moral point of view.[30]

A more explicit endorsement of a similar position is offered by Terence Horgan and Mark Timmons, who argue that in addition to playing moral requiring and moral justifying roles moral reasons can also play a moral merit-conferring role.[31] Horgan and Timmons support their position by offering the following example:

Olivia's offer. Olivia and her husband Stan have recently moved to St. Louis, each having accepted an academic appointment at one of the local universities. During their first week in their new home, Olivia attends a block party organized by one of their new neighbors where she meets a recently widowed woman, Mary, a neighbor who lives a few doors down from Olivia and Stan. In conversation, Olivia learns that Mary lost her husband to cancer after forty-eight years of marriage. She also learns that Mary is an avid baseball fan and that she and

[30] It is worth noting that Portmore appears committed to the existence of this third moral role that reasons can play. In describing moral reasons as playing 'at least two normative roles', Portmore actually commits himself to the view that there are at least *three* different normative roles that such reasons play. This is because he has defined moral reasons as reasons that count for or against performing an act from the moral point of view. This role of counting for or against from the moral point of view must be distinct from a reason having moral requiring force, as Portmore accepts the possibility that a moral reason could have no moral requiring force. Portmore, *Commonsense Consequentialism*, p. 122 n. 6 and p. 135 n. 22.

[31] Horgan and Timmons, 'Untying a Knot', p. 54. This position is not identical to mine as their 'merit-conferring role' looks not only at the reasons favouring the act but also at the reasons on which the agent acted. I take this merit-conferring role to depend on the existence of a moral favouring role.

https://doi.org/10.1017/S0953820816000091 Published online by Cambridge University Press

Table 2

	Moral Requiring Strength	Moral Justifying Strength	Moral Favouring Strength
Go To Exam	0	5	0
Take Woman to Hospital	5	10	10
Don't Take Neighbour's Car	8	8	8

her husband used to regularly attend Cardinals games. But without anyone to go with, she doesn't go anymore. The next day, it occurs to Olivia that it would be a nice gesture to offer to go to a Cardinals game with Mary, although she herself has no particular interest in the game. But she thinks: 'Here is a chance to do something nice for someone, and the fall semester doesn't begin for another couple of weeks. Why not?' She calls Mary, who is delighted by the invitation, and they end up going to a game.[32]

Horgan and Timmons claim that in this example there are moral reasons that count in favour of Olivia making this invitation but that these reasons lack any requiring force.[33] In my terminology, then, this act is supported by reasons that are playing a moral favouring role (and also a moral justifying role) but not a moral requiring role.

Accepting that reasons can play a moral favouring role allows us to capture everything we wanted to say about our three cases, as Table 2 demonstrates.

With the introduction of the moral favouring role we are able to see that taking the woman to hospital is the morally best of the three acts while not taking the neighbour's car is morally better than going to the exam. As a result, this solution succeeds in satisfying SOLVE THE PROBLEM.

I have proposed a solution to the Transitivity Problem that appeals to a novel account of the way in which moral reasons operate. In the current literature on moral reasons there exist those who argue in favour of the need to distinguish between the three kinds of role that moral reasons can play, in the way I have done. Some of those who defend this view think that reasons can play more than one role at the same time.[34] However, no one has yet defended the view that reasons

[32] Horgan and Timmons, 'Untying a Knot', p. 47.

[33] Horgan and Timmons, 'Untying a Knot', p. 48.

[34] It is clear from Portmore's discussion that he views reasons as capable of playing a moral justifying and a moral requiring role at the same time. See Portmore, *Commonsense Consequentialism*, pp. 121–3.

https://doi.org/10.1017/S0953820816000091 Published online by Cambridge University Press

that are playing more than one role can vary in strength in regard to the different roles that they are playing.

We might wonder, though, why it is that reasons can vary in strength in regard to the different roles that they are playing. One plausible reason why these different kinds of strength might come apart can be found in the existence of different kinds of moral reason. In Jamie Dreier's discussion of the paradox of supererogation he argues that moral reasons can be divided into those concerned with justice and those concerned with beneficence.[35] While reasons of justice are capable of generating moral requirements, reasons of beneficence are not. As it stands this view is too strong. There are times when reasons of beneficence do generate moral obligations. In Peter Singer's example of a child drowning in a pond that can easily be saved at the expense of muddying a new suit it appears both that the reasons that support saving the child are reasons of beneficence and that they generate a moral obligation in this case.[36] Nevertheless, a weaker version of Dreier's view will prove fruitful here. Consider the following claim made by John Stuart Mill: 'Justice is a name for certain classes of moral rules, which concern the essentials of human well-being more nearly, and are therefore of more absolute obligation, than any other rules for the guidance of life.'[37] Mill's point is that because of the special importance of reasons of justice, it is generally more important that we follow these reasons than other forms of moral reasons, though there will be occasions where this is not the case.[38] If we accept this then it seems reasonable to conclude that reasons of justice have *greater requiring force* than other moral reasons, such as reasons of beneficence. This allows us to explain why the reasons that support taking the woman to hospital possess less requiring force than the reasons that speak against taking the neighbour's car without permission, despite the fact that the reasons that support taking the woman to hospital possess more favouring strength than the reasons against taking the car. After all, it seems plausible to suggest that the reasons that support taking the woman to hospital are reasons of beneficence while the reasons against taking the car are reasons of justice. By distinguishing between the different kinds of moral role a moral reason can play, then, we are able to explain why the moral requiring force of a reason can vary from its justifying or favouring force.

[35] Jamie Dreier, 'Why Ethical Satisficing Makes Sense and Rational Satisficing Does Not', *Satisficing and Maximising*, ed. Michael Byron (Cambridge, 2004), pp. 131–54, at 149.

[36] Peter Singer, 'Famine, Affluence, and Morality', *Philosophy & Public Affairs* 1 (1972), pp. 229–43.

[37] John Stuart Mill, *Utilitarianism* (Indianapolis, 2001; first published 1863), p. 59.

[38] Kamm, *Morality*, p. 64.

https://doi.org/10.1017/S0953820816000091 Published online by Cambridge University Press

In this section I have provided a new solution to the Transitivity Problem by proposing a novel account of the way in which moral reasons operate. I have argued that by allowing for three different moral roles that reasons can play and that the moral justifying strength of a reason can be different from its moral requiring strength we are able to say everything we wanted to say about our three cases. I have shown that, unlike the other solutions I have investigated, this solution succeeds in satisfying all three of the desiderata proposed in §II.

VI. OBJECTIONS AND RESPONSES

One objection that might be raised against the solution I outlined in the previous section is that it is incompatible with popular conceptions of what reasons are. One way to make this objection would be to claim that all reasons are deontic by nature. John Broome, for example, has claimed that, 'if you have a reason to q and no reason not to q, then you ought to q.'[39] We might think that the claim that reasons are capable of playing three different moral roles is incompatible with this claim, as it opens up the possibility that a reason could play a favouring or justifying role without playing a requiring role. Moreover, in the previous table it was claimed that going to the exam is an act that has no requiring force. If we accept this then it looks like the reasons in favour of this act do not fit with Broome's claim about reasons.

One way to respond to this objection would be to claim that Broome's view is one that ought to be rejected. I think there is good reason to think that this is right.[40] However, there is no need to commit myself to this claim in order to respond to this objection, as my view is compatible with Broome's claim. After all, Broome's claim is about the connection between reasons and *normative* requirements. My claim in the previous table that an act can have *moral* justifying force without possessing any *moral* requiring force is compatible with Broome's claim about *normative* requirements. In fact it seems very plausible to think that there is a *normative* requirement to go to the exam in the absence of any reason to the contrary.

Another claim about reasons that my claims in the previous section appear to be incompatible with is T. M. Scanlon's claim that what it is for x to be a reason for something is for x to be, 'a consideration that counts in favour of it'.[41] This view appears incompatible with the

[39] John Broome, 'Normative Requirements', *Ratio* 12 (1999), pp. 398–419.

[40] For arguments against this conception of reasons see Dancy, 'Enticing Reasons' and Margaret Olivia Little, 'In Defence of Non-Deontic Reasons', *Thinking about Reasons: Themes from the Moral Philosophy of Jonathan Dancy*, ed. David Bakhurst, Brad Hooker and Margaret Olivia Little (Oxford, 2013), pp. 112–36.

[41] T. M. Scanlon, *What We Owe to Each Other* (Cambridge, MA, 1998), p. 17.

https://doi.org/10.1017/S0953820816000091 Published online by Cambridge University Press

thought that the reasons that make it morally justified to go to the exam rather than to drive the injured woman to hospital do not morally count in favour of performing this act. After all, these are reasons but do not count in favour of performing the act.

On closer inspection, though, this worry is misplaced. While these reasons may not count in favour of performing the act from the point of view of morality they do count in favour of doing so from the *prudential* point of view and from the all-things-considered *normative* point of view.

One response that might be made here is to claim that this prevents these reasons from counting as moral reasons. Just as to be a reason in general for something is to count in favour of it, perhaps to be a *moral* reason for something is to *morally* count in favour of it. If we accept this then it seems we are committed to saying that reasons that play a moral justifying role without playing a moral favouring role are not moral reasons.

However, there doesn't seem to be any reason to think that it would be problematic to be committed to this position. As Portmore has pointed out, we can say that while a reason that does not count in favour of performing an act from the moral point of view may not be a moral reason, it can still be a *morally relevant* reason.[42] In other words, we can say that while the reasons that support going to the exam are not ones that count in favour of performing this act from the moral point of view, they do play a role in determining what is morally permissible. Once we appreciate that we can understand these reasons in this way there no longer appears to be any problem with saying that this act is not supported by moral reasons, as they can still be morally relevant reasons.

Another way in which the solution offered in the previous section might be thought to rest on a flawed conception of reasons is that it appears to make it unclear how we should understand the strength of a reason. Is the strength of a reason determined by its justificatory strength, its requiring strength or its favouring strength?[43] It isn't immediately clear what my account should say in response to this question, as it would seem arbitrary to pick one over the others. Perhaps it might instead be suggested that this phrase is simply ambiguous between the three terms. This though will appear unsatisfactory unless we can give an explanation for this ambiguity. Note that the simple favouring account faces no such problems here, as the strength of a

[42] Portmore, *Commonsense Consequentialism*, ch. 5.
[43] Thanks to Mike Ridge for raising this objection.

https://doi.org/10.1017/S0953820816000091 Published online by Cambridge University Press

reason can be understood straightforwardly in terms of how strongly a reason counts in favour of performing some action.

However, an alternative view of reasons that has attracted a number of supporters in recent years can provide us with the resources to give a unified account of a reason's strength. According to a number of theorists, reasons are premises of successful reasoning. As Jonathan Way describes the view, 'What it is for some consideration to be a reason to ψ is for it to be a premise of good reasoning towards ψing.'[44] This alternative account of the nature of reasons allows us to say that the strength of a reason can be understood in terms of how much weight it would be appropriate to give it when reasoning well.[45] Of course, the effect this strength will have on what would count as good reasoning will depend on whether the strength is requiring, justifying or favouring. In all three cases, though, the strength of the reason will determine what weight it would be appropriate to give it in our reasoning.

CONCLUDING REMARKS

In this article I have investigated whether the 'is a morally permissible alternative to' is a transitive relation. I started by introducing a set of cases that appeared to cast doubt on the transitivity of this relation. This created a puzzle, as there seemed to be good reason to think that the relation is transitive. I then set out three desiderata that a successful solution to this puzzle should satisfy. I then looked at three attempts to solve the problem and argued that they all failed to satisfy all three desiderata. Finally, I proposed my own solution to the problem, which involves positing three different moral roles that reasons can play: a justifying role, a requiring role and a favouring role. If we accept that moral reasons can play these three distinct roles and that a reason's strength may vary across these three roles then this provides us with a solution to the puzzle that satisfies all three desiderata.

The existence of a satisfactory solution to this problem has wider implications for moral philosophy. If no satisfactory solution to this problem had been available then this would have cast serious doubt on our ability to give a systematic account of the relationship between moral reasons and moral permissibility. This in turn would have had important implications for normative ethics, as it is particularly incumbent upon consequentialists to give a full ranking of acts from

[44] Jonathan Way, 'Reasons as Premises of Good Reasoning', *Pacific Philosophical Quarterly* (forthcoming). A similar view is defended in Kieran Setiya, 'What is a Reason to Act?', *Philosophical Studies* 167 (2014), pp. 221–35.

[45] See Way, 'Reasons' for a detailed account of how to give an account of the weight of reasons using this approach.

https://doi.org/10.1017/S0953820816000091 Published online by Cambridge University Press

better to worse in order to use this ranking to determine which acts are morally permissible. My solution allows us to avoid these doubts.

The precise nature of my solution is also important for wider issues concerning the nature of practical reasons. I have argued that reasons are capable of playing more than one moral role at the same time and that the strength of the reason can vary across these different roles. An interesting question for future research is whether the same can be said for practical reasons more generally. Are all practical reasons capable of playing more than one role at the same time and varying in strength between these roles? If so, is the 'is a rationally permissible alternative to' relation also intransitive? These questions are ones that strike me as worthy of future investigation.[46]

atmarcher@gmail.com

[46] Thanks to audiences at the following events for helpful questions and discussion: the University of Bristol Departmental Seminar, the 2015 Society of Applied Philosophy Annual Conference at the University of Edinburgh, the 2015 Meetings in Ethics and Political Philosophy at the University of Minho, and the 2015 MANCEPT workshop on Duties of Beneficence. Special thanks to Dale Dorsey, Mike Ridge and an anonymous referee for this journal for written comments on a previous draft of this article.

https://doi.org/10.1017/S0953820816000091 Published online by Cambridge University Press