A Dilemma for Permissibility-based Solutions to the Paradox of Supererogation

Marina Uzunova[1]   Benjamin Ferguson[2]

[1]VU Amsterdam, 1081 HV Amsterdam, The Netherlands. m.g.uzunova@vu.nl. URL: https://muzunova.weebly.com/

[2]VU Amsterdam, 1081 HV Amsterdam, The Netherlands and University of Warwick, Coventry CV4 7AL, United Kingdom. benjamin.ferguson@warwick.ac.uk. URL: http://benjaminferguson.org/

Abstract

We argue that permissibility-based solutions to the paradox of supererogation encounter a nested dilemma. Such approaches solve the paradox by distinguishing moral and rational permissions. If they do not also include a bridge condition that relates these two permissions, then they violate a very plausible monotonicity condition. If they do include a bridge condition, then permissibility-based solutions either amount to rational satisficing or they collapse back into the classical account of supererogation and fail to resolve the paradox.

1 The Paradox

Consider the following case:

> (Alice) Alice has received a much deserved $1,000 bonus from her employer. Lately she has worked long hours and would like to spend the money on a holiday. However, she recently heard a compelling appeal for funds to support a women's literacy initiative in low income countries. Alice reasons that although it would be morally *better* to donate the bonus, it would not be morally impermissible to use it for a holiday.

Donating the bonus is an example of a *supererogatory* act, according to the 'standard account' (Portmore 2016: 287) of supererogation: an act that is both morally optional and morally better than another permissible act. Ordinary moral intuitions support the existence of supererogatory acts. As Dreier (2004: 145) puts it, 'we take it for granted that there are supererogatory acts, and it would be incredible if the very idea of supererogation turned out to be incoherent'. Yet, the ordinary intuition encounters a *theoretical* challenge, known as the paradox of supererogation (Grigore 2019, Dreier 2004, Dorsey 2013, Horgan and Timmons 2010, Heyd 2016, Archer 2013a).

Following Terry Horgan and Mark Timmons (2010: 36–37), we characterise the paradox of supererogation as a conflict between three conditions: *supererogation*, *moral continence*, and the *very weak classical view*.[1] To define these conditions, we need some preliminaries. Suppose that there is a (finite) collection of token acts potentially available to an agent. The menu of token acts that is *actually* available in a particular context is a subset of this collection. Further, suppose that it is possible to compare *all* potentially available options according to their moral goodness.[2] Finally, suppose that it is possible to compare (at least some) acts with respect to the strength of moral reasons that support them. Then the three conditions can be formulated as follows.

> (Classical Supererogation) There exists an act – called a 'supererogatory act' – in the agent's menu of available options that is morally optional (permissible to perform and permissible to not perform) and also morally better than another available morally permissible act.

> (Moral Continence) If there are strictly stronger moral reasons to perform act x than there are to perform act y, then y is morally impermissible.

---

[1] Horgan and Timmons's account uses these conditions, but they do not use these *labels*. Continence comes from Davidson (1980) and the very weak classical view is a weakened version of a condition found in Ferguson and Köhler (forthcoming: §2.2), which itself is a weaker version of the classical view of intentional action (Davidson 1963).

[2] We are aware of challenges to the completeness assumption, primarily on the basis of incomparability. Such an answer, however, cannot vindicate supererogation because, as Heyd (1982: 5) points out, supererogation requires 'a common and continuous scale of values shared by supererogation and duty'.

(The Very Weak Classical View) If act x is morally better than act y, then there is a stronger moral reason to perform x than there is to perform y.

The conflict between these three conditions is straightforward. The very weak classical view implies that if act x is morally better than act y, then x is also supported by stronger moral reasons, and hence, by moral continence, y is morally impermissible. This conflicts with classical supererogation, which claims that there can exist morally *optional* acts that are morally better than other *permissible* acts.

There are many ways to escape the paradox. Most contemporary resolutions argue that the *reasons* in moral continence and the very weak classical view differ in kind (Dreier 2004, Horgan and Timmons 2010, Ferry 2015). Daniel Muñoz (forthcoming) provides a nice overview of this strategy. A second resolution claims that the *permissions* in classical supererogation and moral continence differ in kind. This permissibility solution, as we will call it, is defended by Dale Dorsey (2013, 2016).

In what follows, we focus on the permissibility solution. We show that both Dorsey's solution and other permissibility-based solutions that distinguish all-things-considered permissibility from moral permissibility encounter a dilemma. We outline the structure of permissibility-based solutions and then introduce the dilemma.

2   Permissibility-based Solutions

Permissibility-based solutions claim that there are two sorts of permissions: one in the classical supererogation condition and a different sort in the moral continence condition. According to Dorsey's (2013: 369) account, we should understand supererogation as granting *all-things-considered* (rational) permissions which may be sensitive to prudential, moral, aesthetic, legal and other considerations. Moral continence concerns the narrower domain of *moral* permissibility and hence grants *pro tanto* permissions. Dorsey thus escapes the paradox by replacing the classical supererogation condition with:

(All-Things-Considered Supererogation) There exists an act – called a 'supererogatory act' – in the agent's menu of available options that is *all-things-considered* optional (permissible to perform and permissible to not perform) and also morally better than another available *all-things-considered* permissible act.

The escape is straightforward. There is no paradox because, as Dorsey puts it, we should 'treat supererogatory actions not as morally better than [what] is *morally* required, but as morally better than [what] is *rationally* [or *all-things-considered*] required' (2013: 373, emphasis added). Morality may be stringent and demand the morally best (as moral continence and the very weak classical view imply), but *rationality* is more permissive and can let us do what is

morally suboptimal (as all-things-considered supererogation implies). In other words, there is no paradox because 'immorality can be rationally [or all-things-considered] permitted' (Dorsey 2013: 370).

Dorsey points out that there are 'two potential ways this solution might go wrong': there may be no workable bridge condition between all-things-considered and moral permissions, or 'even if there is such an account...this proposal is just implausible' (Dorsey 2013: 373). So far Dorsey's critics have focused on the latter problem, arguing that Dorsey's 'redefinition' of supererogation simply moves the goalposts (Archer 2013b: 185–186; Portmore 2016: 286) without vindicating ordinary intuitions about supererogation. For example, though it entails that Alice is rationally permitted to refrain from donating, it cannot deliver the intuition that she has a *moral* permission to take her holiday.

Yet, defenders of permissibility solutions could argue that puzzles surrounding supererogation arise precisely because our intuitions confuse moral and rational permissions. In this paper we set aside debates about the external plausibility of Dorsey's account and instead focus on the *first* kind of objection. We consider whether the permissibility solution can succeed on its own terms, with or without a bridge condition that allows for supererogation.

We take a bridge condition that we consider most plausible, and which is implied by the conjunction of two conditions Dorsey accepts: his own 'reasons permission' (2013: 374) condition and the very weak classical view.

> (Goodness Permission) If act x is morally better than y and y is all-things-considered permissible, then x is also all-things-considered permissible.

The goodness permission condition allows ordering the all-things-considered permissibility of acts along a moral goodness axis. Of course, a permissibility-based solution need not accept goodness permission. We consider the merits of goodness permission as well as the possibility of doing without a bridge condition and conclude that neither option is plausible. Thus, the permissibility solution is not viable. First, we argue that the permission to refrain from performing supererogatory acts is a *conditional permission*. Then, we show that this leads to a dilemma for permissibility accounts.

3   The Conditional Permissibility of Omissions

Supererogatory acts are optional, in the sense that it is permissible to perform them and to refrain from performing them. However, there is an important asymmetry in these permissions because the permissibility of *refraining* from performing a supererogatory act depends also on the act the agent performs instead.[3] Though Alice may be permitted to refrain from donating to charity, this permission is not unconditional. For example, it is not permissible for

---

[3]Note that this asymmetry differs from the more common asymmetry between acts and omissions defended by non-consequentialists.

Alice to refrain from donating if instead she plans to use the bonus to promote *illiteracy*.

Acts are supererogatory *with respect to* a particular agenda or menu of options. In this sense we may say that the omission of supererogatory acts is *conditionally* permissible. This point holds both for the permission solution's all-things-considered conception of supererogation and for classical supererogation. In order to incorporate this menu-dependence, both accounts of supererogation must substitute the permissibility of not doing x by a conditional permissibility which we define as follows (for all-things-considered supererogation):

> (Conditional Permissibility of Omissions) Refraining from doing act x is all-things-considered permissible conditionally on doing another act y if and only if doing y is all-things-considered permissible and in the menu of options.

The conditional permissibility of omissions might seem obvious and yet this distinction has important implications for the success of a permissibility-based solution. Before we explain why conditional permissibility is important, we first present the following two properties that the conditional permissibility of omissions might satisfy:

> (Monotonicity) If refraining from doing x is all-things-considered permissible conditionally on doing y and z is morally better than y, then refraining from doing x is also all-things-considered permissible conditionally on doing z.

> (Symmetry) If refraining from doing x is all-things-considered permissible conditionally on doing y and z is just as morally good as y, then refraining from doing x is also all-things-considered permissible conditionally on doing z.

These properties capture the intuition that if you are permitted not to do what is morally best because you are permitted to do a morally inferior act, then you should still be permitted not to do what is morally best if you were to do as much (symmetry) or more (monotonicity) good instead. However, as we argue in the next section, permissibility-based solutions run into problems regardless of whether they satisfy or violate one or both of these conditions. And what is important to note here is that these problems become apparent when we are explicit about and pay attention to the *conditional* nature of the permissibility of omissions.

4   The Dilemma

Recall that permissibility-based solutions solve the paradox of supererogation by claiming that morally suboptimal (and hence, by the moral continence condition, *morally* impermissible) acts are nevertheless all-things-considered permissible. Let $y^{\min}$ be the least morally good act that is all-things-considered

permissible.[4] For example, in Alice's case, we'll simply stipulate that $y^{\min}$ is the act of spending the \$1,000 bonus on a holiday.[5]

By definition, anything morally worse than $y^{\min}$ is all-things-considered impermissible. But what about acts that are morally better than $y^{\min}$? There are two possibilities: first, if the goodness permission condition is satisfied, then moral goodness orders all-things-considered permissibility in the sense that anything morally better than the permissible $y^{\min}$ is also all-things-considered permissible (see Figure 1a); second, if goodness permission is violated, then there could be acts that are morally better than $y^{\min}$ and yet all-things-considered impermissible (see Figure 1b). We believe the second case is more problematic so we'll discuss it first.


Figure 1: (a) illustrates what accepting the goodness permission entails: it is possible to order all all-things-considered permissible acts from the morally worst ($y^{\min}$) to the morally best ($x^{\max}$). (b) illustrates what rejecting the condition entails: there can exist acts (z) which are better than the morally worst permissible act ($y^{\min}$) but are all-things-considered impermissible.

4.1 Rejecting goodness permission

If a permissibility-based solution rejects goodness permission, then it allows acts like z in Figure 1b, which are morally better than the morally worst permissible act ($y^{\min}$) and yet which are all-things-considered impermissible. Here is an example:

> (Alice (b)) As before, Alice has a choice between donating the full \$1,000 bonus ($x^{\max}$) and spending it all on a holiday ($y^{\min}$). However, she can now also choose to spend \$500 on a holiday and donate the other half (z). Suppose that donating everything and spending it all on the holiday are both all-things-considered permissible while splitting the sum in half is all-things-considered impermissible.

---

[4]In the language prevalent in the literature on supererogation, $y^{\min}$ is our 'minimal duty' (and hence an erogatory act): anything morally worse than $y^{\min}$ is all-things-considered impermissible.

[5]One could, of course, imagine worse acts that are still permissible. The broader point is that goodness permission requires that *however* moral and rational theories set $y^{\min}$, they should also entail that those acts *they* claim are morally better than $y^{\min}$ are also all-things-considered permissible.

The Alice (b) case violates goodness permission and there might be good reasons for such a violation. For example, in his later writings, Dorsey (2016: 172–206) invokes what he calls the 'normative significance of the self'. This view allows giving heavier weight to agent-centric reasons related to the promotion of one's life projects and commitments. For instance, if the immersion in distant cultures is part of Alice's life project and if splitting the amount (z) is insufficient for contributing to this project, then splitting the amount might turn out to be all-things-considered impermissible.

The trouble is that while a violation of goodness permission might be plausibly justified, it necessarily implies a violation of monotonicity and that is less plausible. Recall that goodness permission concerns the *performance* of acts, while monotonicity concerns the *omission* of acts. Alice (b) violates monotonicity: not donating is all-things-considered permissible conditionally on spending the entire bonus on a holiday, but not donating is all-things-considered *im*permissible conditionally on spending only half. Put differently, Alice has a license not to donate if she were to donate nothing but no such permission if she were to donate half. Such non-monotonicity of omissions is at odds with common intuitions about the *optionality* of supererogatory acts. Additionally, it allows us to see that views, such as the 'normative significance of the self', that reject goodness permission are importantly incomplete: they need to provide a plausible account not just for violations of goodness permission but also for the violations of monotonicity which necessarily follow. And to do that, they need to pay attention to the conditional nature of the permissibility of omissions.

4.2 Accepting goodness permission

If a permissibility-based solution wants to avoid the non-monotonicity of omissions, it can accept goodness permission and grant that any act that is morally better than what is minimally all-things-considered permissible ($y^{\min}$) is also permissible (Figure 1a). This move avoids violations of monotonicity but leads to other problems. Consider:

> (Alice (a)) As before Alice has a choice between donating \$1,000 ($x^{\max}$) and spending it all on a holiday with her partner ($y^{\min}$). Additionally, she can now spend the bonus on a holiday with a distant relative instead of her partner (w) which is just as morally good as $y^{\min}$. All three acts are all-things-considered permissible.

Consider the two minimally permissible acts that are also equally morally good: going on a holiday with a partner ($y^{\min}$) and with a distant relative (w). Now start increasing the *non*-moral considerations that speak against w. The question is: is there some threshold beyond which non-moral considerations tip the previously all-things-considered permissibility of w into impermissibility?

Suppose there is no threshold. If the answer is 'no', then there is *no* amount of non-moral considerations that can make the all-things-considered

permissible w all-things-considered impermissible. Such a move satisfies symmetry of omissions but should be resisted by permissibility-based solutions. The problem is that all-things-considered – as opposed to moral – permissibility is then a trivial addition without real normative force. The real permissibility work would be done not by considerations that go *beyond* morality, but by moral goodness: there is a minimal threshold of moral goodness that grants a permission to do certain acts. If a permissibility-based solution were to take this route, then it collapses into a moral-goodness satisficing account (and hence into classical supererogation). Instead of disaggregating permissions, it should be disaggregating *moral* goodness.

Suppose there is a threshold. Alternatively, suppose the answer is 'yes'. Then we end up with two equally morally good acts one of which is all-things-considered permissible ($y^{\min}$) while the other is all-things-considered impermissible (w). This is a violation of symmetry. We find violations of symmetry problematic in themselves. However, others may not share this intuition. Particularly, they, like Dorsey, might believe that certain non-moral considerations – such as agent-centric reasons to promote one's life projects – can plausibly lead to asymmetries (2016: 172). Even so, the broader point is that such accounts must also explain *who* determines the nature and strength of the non-moral considerations that can lead to asymmetries. And regardless of whether this is determined 'internally' by the agent herself or 'externally' a problem arises.

To see how, take the 'external case first. Suppose that Alice is *not* the one who determines the nature and strength of the non-moral considerations that make it all-things-considered permissible or impermissible not to donate conditionally on taking a holiday with a distant relative. There are two problems with this. First, it is unclear how such an odd combination of agent and non-agent-centric considerations can explain why Alice should include them in her calculus. And, more importantly, if it is not Alice who determines the nature and degree of these considerations, it is hard to understand how these are not ultimately driven by *moral* considerations.[6] But if that is so, then, again, such a permissibility-based solution will collapse into a moral satisficing account and hence into classical supererogation.

Alternatively, suppose that it is Alice who sets the threshold and decides on the kind of considerations that enter the calculus. Such a response collapses into egoism: if Alice herself chooses the determinants of her conditional permissibility not to donate, then it is up to *her* to expand the domain of the minimally permissible. And if such a permissibility solution collapses into egoism, then it makes the permission to not perform the optimal act granted by all-things-considered supererogation equivalent to *rational* satisficing. Yet, rational satis-

---

[6]Dorsey suggests that the socialised nature of our own projects, and the role of *moral* education in eventually 'increasing the convergence between moral requirements and the demands of practical rationality' (2016: 207) provides a possible answer. But this answer shifts responsibility away from the agent by allowing her to set the terms of her own responsibility (that is, by granting her the all-things-considered permission not to do the morally superior act), *even while* these terms are judged morally undesirable.

ficing, as authors like Dreier (2004) and Davidson (1980) have forcefully argued, faces serious conceptual problems that undermine its coherence.

We have argued that permissibility-based solutions encounter a dilemma. If they reject goodness permission, then they allow for non-monotonic omissions which are counterintuitive. If they accept goodness permission, then they collapse into *rational* satisficing if the threshold is determined by the agent herself. And if there is no threshold, or if the threshold is determined by considerations other than the agent's, they collapse into the classical account of supererogation. These problems provide strong reasons to abandon permissibility-based solutions in favour of alternatives, such as reason-based solutions that focus on disaggregating moral and non-moral reasons.[7]

Funding

References

Archer, A. 2013a. Supererogation and intentions of the agent. *Philosophia* 41: 447–62.

Archer, A. 2013b. The supererogatory and how not to accommodate it: a reply to Dorsey. *Utilitas* 28: 179–88.

Davidson, D. 1963. Actions, reasons, and causes. *The Journal of Philosophy* 60: 685–700.

Davidson, D. 1980. How is weakness of the will possible? In his *Essays on Actions and Events*, 21–42. Oxford: Clarendon Press.

Dorsey, D. 2013. The supererogatory, and how to accommodate it. *Utilitas* 25: 355–82.

Dorsey, D. 2016. *The Limits of Moral Authority*. Oxford: Oxford University Press.

Dreier, J. 2004. Why ethical satisficing makes sense and rational satisficing doesn't. In *Satisficing and Maximizing*, ed. M. Byron, 131–54. Cambridge: Cambridge University Press.

Ferguson, B. and Köhler, S. forthcoming. Betterness of permissibility. *Philosophical Studies*.

---

[7]This paper was presented at VU Amsterdam, the University of Amsterdam and the University of Warwick. We thank those present, anonymous referees at *Analysis* and Sebastian Köhler for helpful comments.

Ferry, M. 2015. Beyond obligation: reasons and supererogation. *Royal Institute of Philosophy Supplement* 77: 49–65.

Grigore, N. 2019. On why there is a problem of supererogation. *Philosophia* 47(4): 1141–63.

Heyd, D. 1982. *Supererogation: Its Status in Ethical Theory.* Cambridge: Cambridge University Press.

Heyd, D. 2016. Supererogation. In *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta.

Horgan, T. and M. Timmons. 2010. Untying a knot from the inside out: reflections on the 'paradox' of supererogation. *Social Philosophy and Policy* 27: 29–63.

Muñoz, D. forthcoming. Three paradoxes of supererogation. *Noûs.*

Portmore, D. 2016. Transitivity, moral latitude, and supererogation. *Utilitas* 29: 286–98.