Forthcoming in *Noûs*
Draft of October 2019

Three Paradoxes of Supererogation

Daniel Muñoz
Monash University

Supererogatory acts—good deeds "beyond the call of duty"—are a part of moral common sense, but conceptually puzzling. I propose a unified solution to three of the most infamous puzzles: the classic Paradox of Supererogation (if it's so good, why isn't it just obligatory?), Horton's All or Nothing Problem, and Kamm's Intransitivity Paradox. I conclude that supererogation makes sense if, and only if, the grounds of rightness are multi-dimensional and comparative.

1. Introduction

We don't always have to do what's best. Some good deeds are *supererogatory:* they are optional and yet better than other permissible options. Examples include enduring injury to save a life, doing a favor, and granting forgiveness.[1] Lovely stuff; hardly required.

But as innocent as the examples may sound, the theory of supererogation is fraught with paradoxes. Three in particular have been influential. First, the classic Paradox of Supererogation asks how supererogation could be possible even in a simple pairwise choice, like the choice between saving someone from a crumbling building or safely loitering outside (Raz 1975, Dreier 2004: 148, Heyd 2016: §3, Archer 2018: §4). If being the hero is really better, why isn't it just obligatory?

The other puzzles add a third option. In Horton's (2017) All or Nothing Problem, the agent has two ways to be the hero: either rescue the one from the building, or rescue *two*, at the same level of cost and effort—say, two crushed arms either way. It seems wrong to save just one person; letting the other die is pointless badness. And yet, saving one seems no worse than permissibly saving zero. But surely it's always worse to do wrong. How could it not be?

[1] When I say that supererogation is "better," I just mean that there is more reason to do it, not that it has nice effects or merits praise. Later I will discuss views on which supererogation isn't better all things considered, but only morally better. None of my arguments depend on how we define 'better' or 'supererogatory'.

In Kamm's (1985, 1996) Intransitivity Paradox, meanwhile, the agent has two ways *not* to be heroic. Imagine that you come upon the crumbling building while on your way to the cafe, where you have promised to meet a friend. You may break your promise to save a life, or you may keep the promise, but it would be wrong to just sit there. This leads to an infamous intransitivity: you may do nothing instead of saving the one (in a pairwise choice), and may save the one instead of keeping your promise, but you may *not* do nothing instead of keeping the promise. How is this possible?

These puzzles are usually discussed in isolation. It is not standard to check whether a solution to Kamm's conundrum could work for Horton's, or vice versa, nor do people discuss these two in light of the classic Paradox. But we need a unified solution. My goal is to find one.

Why do we need to take on the puzzles together? For one thing, the choices involved are undeniably linked. In the classic Paradox, we have:

Do Nothing ——— ~ ——— *Save 1*

Where '*Do Nothing* ~ *Save 1*' means that *Do Nothing* "ties" *Save 1*: both options are permissible in a pairwise choice. In Horton's case, we seem to have:

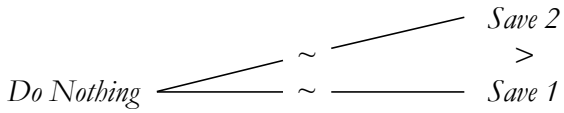

Where '*Save 2* > *Save 1*' means that *Save 2* "defeats" *Save 1*: only *Save 2* is permissible in a pairwise choice (cf. Dietrich and List 2017). Kamm's case is isomorphic:

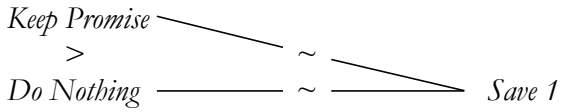

Indeed, the cases are built in the same way. We add an option that defeats *only one* of the classic pair. Kamm's addition beats *Do Nothing*; Horton's beats *Save 1*.

Intriguing links aside, the deeper reason to treat the puzzles together is that we need our solutions to cohere. To capture the intuitions in Kamm's or Horton's case, some philosophers

propose revisionary principles of reasons and wrongness. But these proposals, I argue, can't be extended to the combined case, where you have two ways to be heroic and two ways not to:

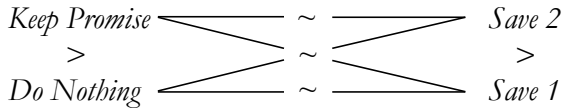

Other philosophers reject the intuitions in Kamm's and Horton's cases, citing a tension with certain familiar principles of reasons and wrongness, like the principle that wrong acts are worse than permissible ones. But these principles clash with, whereas the intuitions follow from, our best solutions to the classic Paradox. The best views about how supererogation is *possible* are, even by themselves, enough to debunk the principles that lead to paradox.

What makes supererogation possible is that wrongness doesn't just depend on worseness. Wrongness is determined by at least two factors that are independent—for example, moral reasons and non-moral reasons. This independence can lead to some wrong acts (like *Save 1*) being no worse than permissible ones (like *Do Nothing*). It also can make '≳' intransitive, though in Kamm's own case, I argue that the intransitivity arises from something more than the presence of independent factors; one factor must also be *comparative*. I will say more later about what this amounts to, and why I think comparativity is implicit in, or at least winked at by, solutions to the classic Paradox.

We begin with the classic Paradox (§§2–3), then Horton's (§§4–5) and Kamm's (§§6–7). If all goes well, the reward is a coherent picture of how morality works beyond and beneath the call of duty, and the happy conclusion (§§8–9) is that supererogation isn't really paradoxical—just more interesting than we expected.

2. The Classic Paradox: Against 1D Solutions

The classic Paradox of Supererogation asks: if these fabulous deeds are really so *good*, why aren't they *obligatory?* The question has bite if we are drawn to:

The "More Reason" Principle
An option is obligatory if there is more reason to choose it than any alternative.

Which sounds hard to resist. As Raz (1975: 165) asks: "How can one be permitted to refrain from action that is required by reason?" And yet the principle entails that anything less than best is wrong. Supererogation is impossible.[2]

Our task: replace the "More Reason" Principle and save the supererogatory. Now, one strategy is to stick with plain reasons, and pick a fancier, more forgiving principle of rightness. Why not just do that?

It's worth a shot, but the fancy rules turn out to have their own problems. The main rule:

The Baseline Principle
An option is obligatory just if it is better than the baseline. (Slote 1985, Hurka 1990)

Where "the baseline" is either an absolute standard (one's choice isn't *too* awful) or, like an average, relative to the set of one's options (one's choice isn't *too much worse* than it could have been). Either way, this view allows for supererogation in simple two-option cases; *Do Nothing* might be permissible, even though *Save 1* is better, if both are above the line. But the view fails to forbid acts that are above the line yet gratuitously bad. This problem predates Horton's puzzle (Bradley 2006), but his building case is a perfect example. It's wrong to save only one, given that saving two was no costlier, and yet *Save 1* must be above the line, since *Do Nothing* is. The Baseline Principle's permissions are too indiscriminate.

The principles we have just seen invite us to think of options as ranked from worst to best like numbers on a line. The "More Reason" Principle says anything below the top option is wrong; the Baseline Principle tells us that anything above a certain point is "high enough." But neither view

[2] Some discussions of the paradox also cite a principle linking reasons to some other notion, like praiseworthiness (Raz 1975: 164) or goodness (Heyd 2016: §3.1). Something like this: if x is better than y (or more praiseworthy), then there is more reason to do x. By contrast, I am not making substantive claims about how value links up to reasons; I use 'better' just as a snappy way to say 'more reason' (see n.1).

can rule out gratuitously worse options (like *Save 1*) while leaving in those that are justifiably worse (like *Do Nothing*).[3] To properly solve the classic Paradox, we need the resources to say *what it is* about the permissible suboptimal acts that *makes* them permissible, something beyond raw betterness.

Solving the Paradox of Supererogation thus requires us to think of an action's moral status as more than a number on a line. We need a second factor, something other than degrees of plain goodness—and this is exactly what we find in the two leading solutions to the Paradox of Supererogation.

3. The Classic Paradox: 2D Solutions

If supererogating is better, why is it permissible to do anything less?

The leading answer is that the supererogatory act is only better in one respect. Supererogation is favored by one kind of reason, but is opposed by weighty reasons of another kind, which are said to be "justifying" rather than requiring (Archer 2016, Lazar 2019, cf. Gert 2007), "partial" rather than impartial (Parfit 2011), or "non-moral" rather than moral (Slote 1991, Portmore 2011)—whatever they are, they counterbalance the reasons to supererogate, stopping them from grounding an obligation.[4] So we get something like:

---

[3] Another proposal is to assign each option an *interval* on the real line, not a point, and use this rule: an option is wrong just if its interval lies wholly below that of another option (cf. Gert 2004: 505). But this proposal can't handle the combined case (see §1, above), where the options are {*Keep Promise*, *Do Nothing*, *Save 2*, *Save 1*}. This case intuitively violates:

> *Interval Order Property*
> If A > B and C > D, then either A > D or C > B.

And no relation that violates this property can be represented with intervals and the proposed rule for determining wrongness (see Fishburn 1970: 20–23; Rabinowicz 2008: 33n.23).

[4] For some quite different 2D reasons-based views of supererogation, see Raz 1975, Dancy 1993a (critiqued in Postow 2005), Bedke 2011 (critiqued in Snedegar 2016), and Portmore 2019. Another notable view says that supererogatory beneficence is favored by "merely commendatory" reasons, which make options choiceworthy without tending to require them (Dreier 2004, Horgan and Timmons 2010, Little and McNamara 2017). The stock objection: this view wrongly implies that

> *The Non-Moral Reasons Principle*
> An option *x* is obligatory just if, for any alternative *y*, there is more moral reason to choose *x* than there is *combined* moral and non-moral reason to choose *y*.

Where moral and non-moral reasons are two independent factors, and only moral reasons favor supererogating.

Already, this is enough to solve the paradox. Heroic sacrifices are optional, even though they are morally better, because we have mighty non-moral reasons not to harm ourselves. (This view does, however, require us to gently revise our concept of supererogation, since it's no longer better overall—just morally better. See §5, below.)

My own preferred view admits just one kind of reason—which tends to favor, justify, and require—and posits something further: *prerogatives*, which justify without at all requiring or favoring.[5] Prerogatives are purely permissive. They have weights, like reasons, but their sole function is to justify acts that would otherwise be wrong, putting a targeted gap between "most reason" and "must." They are like non-moral reasons, except they don't affect what's best.

This view also lets us replace the "More Reason" Principle. We say:

> *The Prerogatives Principle*
> An option *x* is obligatory just if, for any alternative *y*, there is more reason to choose *x* than there is *combined* reason and prerogative to choose *y*.

And this gives us an easy way to explain the optionality of self-sacrifice. On any view of prerogatives, I have a prerogative not to give up my limbs, and so even though the reasons favor self-sacrifice, it lies well beyond the call of duty.[6]

---

beneficence as such is never required, even if one could save more lives at no cost (Dorsey 2013: 361, Archer 2016: 460). Horton's case is a special instance of this: *Save 1* seems wrong, but it would be permissible if the reasons to *Save 2* were merely commendatory.

[5] Hurka and Shubert (2012) prefer "prima facie permissions." ('Prerogative' is from Scheffler 1982, Slote 1984, and Kamm 1996). But just as nothing turns on how we define 'supererogation', nothing turns on how we define 'reason' and 'prerogative'; the substantive ideas here are justifying, favoring, requiring, and permissibility.

[6] Some say we have a basic prerogative to do what's prudent (Scheffler 1982, Slote 1984, Hurka and Shubert 2012); others derive prerogatives from rights (Benn 2017, Muñoz ms.b).

These are the leading solutions in the literature, and I can't stress enough that they both involve the idea that right and wrong depend on two dimensions. There isn't just a cutoff along a single scale from naughty to nice. Moreover, on either view, the two dimensions are independent, in the sense that an option's ranking along one doesn't tell us how it ranks with respect to the other. The weights of the moral reasons don't tell us how the non-moral reasons lie; the balance of reasons doesn't settle which prerogatives we have; and so on (cf. Lazar 2019: 246).

The crucial upshot of this independence is that an option can outrank its rival on one dimension while being outranked on the other. Indeed, this sort of moral mismatch is just what we would expect to find in our simple case of supererogation. Self-sacrifice is a wonderful thing, but a harm to oneself can have special significance. *Do Nothing* thus gets protected by the weightier prerogative (or non-moral reasons), even though *Save 1* is (morally) better.

We need two dimensions to make sense of supererogation, and having two dimensions can lead to moral mismatches. That is the main lesson of the Paradox of Supererogation. It's a simple point, but it is all we need to solve the All or Nothing Problem.

4. All or Nothing: Reasons and Prerogatives

We begin by stating the problem more carefully.

In Horton's case, I am faced with a collapsing building and three options. I can do nothing, which costs me nothing. Or I could sacrifice my arms to hold open a small escape route, saving one of the people inside. I could also, for the same price, open a bigger route and save both.

Intuitively, I do not have to be the hero.

May Save None
It is permissible for me to save no one.

After all, two arms would be a massive sacrifice. But it would be wrong to put that sacrifice to waste, saving only one person and abandoning the other:

Wrong to Save One
It is wrong for me to save just one.

There is no way to justify such gratuitous badness. Still, it is not as though saving one is worse than doing nothing.

No Worse to Save One
I don't have more reason, all things considered, to save no one than I have to save just one.

Morality wouldn't recommend saving fewer lives.

I assume that these are appealing intuitions. But they leave us with a wrong act (saving one) that is no worse than a permissible act (saving zero). And surely:

Worse to Do Wrong
If A is wrong and B is permissible, then I have more reason to do B.

Indeed, it is a "common assumption" in the literature that any permissible act is better than any wrong alternative (Pummer 2019: 284n.19).[7]

That is the All or Nothing Problem: if the only permissible options are saving all of the people and saving none, then we must revise our intuitions, or else deny the seemingly undeniable—Worse to Do Wrong.

Most solutions give up an intuition. Horton (2017) denies May Save None (he says it is false when the agent is willing to make the sacrifice);[8] Portmore (2019: §5.4), whose views are complex, denies that there is more moral reason to save one than none; and McMahan (2018: 99–100) suggests that saving one is wrong and yet somehow not impermissible. Although these concessions are costly, they would allow us to preserve Worse to Do Wrong.

---

[7] Worse to Do Wrong is endorsed in some form by Darwall 2013, Ferguson 2016: 959n.23, Horton 2017: 96, Ferguson and Köhler 2019 (as "reasons monotonicity of permissibility"), and Portmore 2019: 22. (Portmore's new view is subtle, but he thinks there must be more moral reason to do any morally permissible act over any alternative, because this link is needed for permissibility to depend on reasons.) Two important exceptions are Bader (forthcoming) and Pummer (2019). For a defense of the intuitions in Horton's case, see Muñoz ms.a. For discussion of cases like Horton's, see Parfit 1982; Kagan 1989: 16; Tadros 2011: 161–62; Portmore 2011: 147; and Snedegar 2015: 379.

[8] Sinclair (2018: 49) appears to share Horton's view of the building case.

But that principle is seriously undermined—and the intuitions are strongly predicted—by the idea of independent dimensions, which we got from working through the Paradox of Supererogation. Wrong acts *would* always be worse, if 'wrong' meant 'below the cutoff on the one true moral scale from bad to good'. Any permissible act would be above the line and therefore better than anything below. But what if there are two scales? Then we get counterexamples. Let's see why, starting with a view of reasons and prerogatives.

Recall the basic pair: *Do Nothing* ~ *Save 1*. (Where '$x \sim y$' means that x and y are both permissible in a pairwise choice.) Given that *Save 1* is better than *Do Nothing*, we might expect *Save 1* to be obligatory. But it isn't. We have prerogatives not to harm ourselves, so *Do Nothing* is justifiable despite being suboptimal. The reason to save a life doesn't outweigh the combined reasons and prerogatives to self-preserve.

Now the key question. What should we expect to happen when we add the option to *Save 2?* Consider how it fares in pairwise choices. Intuitively, *Save 2* is obligatory in a pairwise choice with *Save 1*: *Save 2* > *Save 1*. The cost is the same either way, so in effect, one is just choosing whether to pointlessly allow the second person to die. I assume that pointlessly letting die is wrong. Whatever prerogatives we have, they don't make a decisive difference in this kind of choice, where no additional harm to self (or weighty rights) are in play.

But our prerogatives do permit *Do Nothing* in a choice with *Save 2*. *Do Nothing* ~ *Save 2*. A massive harm to oneself can allow us to forgo a whole range of sacrifices; it is not as if anything beyond one life has got to tip the scales. I assume that two lives won't be the tipping point.[9]

So in the choice from {*Do Nothing*, *Save 1*, *Save 2*}, we should expect only one wrong option: *Save 1*. And there is a very natural way to get this from reasons and prerogatives:

[9] If you disagree with this judgment, feel free to bump up the size of the harm to self.

	Do Nothing	Save 1	Save 2
Reasons	1	5	10
Prerogatives	10	0	0

The result is that wrongly saving one is better than permissibly doing nothing. And this is possible because of reasons and prerogatives are independent. *Do Nothing* is worse than *Save 1*, but I have a mighty prerogative not to sacrifice, and only a negligible prerogative to sacrifice to save the one; I can't justify *Save 1* over purely better rescues. That is why the addition of *Save 2* is able to rule out only the better of two options. "Tied" options don't have to be morally equivalent, and they don't have to be defeated by the same things. Even the lesser option in a tie can be more resiliently permissible.[10]

So here is my main claim about the All or Nothing Problem. If we start with reasons, prerogatives, and pairwise judgments, we should expect that wrongly saving one will be better than permissibly saving no one. It is a prediction, not a paradox, that the three intuitions are true and that Worse to Do Wrong is false. We can solve the All or Nothing Problem without compromising on our intuitive judgments, all thanks to the insights we got from working through the Paradox of Supererogation.

5. All or Nothing: Parity of Reasons

But what about reasons-based views of supererogation? Do they make the same predictions about Horton's case? It's a bit complicated, but mostly: yes.

[10] This case also shows that we can't recover the facts about what's permissible from the *sum* of reasons and prerogatives in favor of each option. *Do Nothing* has a higher sum than *Save 2*, but only *Save 2* makes *Save 1* wrong. (Reasons can rule out alternatives; prerogatives can't!)

For simplicity, let's stick with the view that supererogation emerges from a clash of moral and non-moral reasons: an act is obligatory *iff* the moral reasons to do it outweigh the combined moral and non-moral reasons against. On this view, we would expect *Save 2* to be morally best, followed by *Save 1*, and we would expect both of these sacrifices to be non-morally worse than *Do Nothing*. So the prediction is that *Save 1* will be wrong and yet morally better than permissibly saving no one. This is violation of what we might call "Morally Worse to Do Wrong," which is the version of Worse to Do Wrong that is in fact endorsed by fans of non-moral reasons (Portmore 2019: 22).

But can this view accommodate our three intuitions? On the face of it, no. The trouble is with No Worse to Save One. Since *Do Nothing* is permissible, we know that it is at least as good overall as *Save 2*, which is strictly better than *Save 1*. Don't we have to conclude that *Do Nothing* is also better than *Save 1?*

This is a well-known problem.[11] But there is a solution, due to Derek Parfit: we say that *Do Nothing* is on a par with *Save 1* and *Save 2* (2011: 137–141). Two options are on a par when they are comparable, though neither is better, and they aren't exactly equal in goodness (Chang 2002). Using '~' to denote parity and '>' plain betterness, we get a familiar structure:

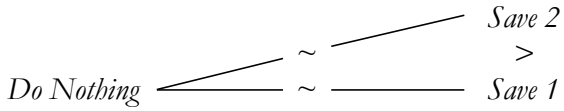

So *Save 1* is wrong because it's strictly worse than *Save 2*, but *Save 1* is still on a par with *Do Nothing*, which remains permissible because it is the best non-morally. With parity, I conclude, we can capture all three intuitions in the All or Nothing Problem and debunk Worse to Do Wrong.[12]

---

[11] For more on the problem of allowing for a range of permissible tradeoffs, on reasons-based views, see Kagan 1991: 927–28, Portmore 2011: Chapter 5, Hurka and Shubert 2012: n.7.

[12] Indeed, even non-moral parity makes trouble for Worse to Do Wrong. For example, start with *Tea* ~ *Coffee*, then add a "sweetened" third option: *Coffee*+ > *Coffee*. Plausibly, if *Coffee*+ is just a slight improvement, we have *Tea* ~ *Coffee*+. But that means a wrong option, *Coffee*, is on a par with (and therefore no worse than) a permissible option—*Tea*. In such cases, parity is "insensitive to mild sweetening" (Hare 2010). Sweetening is a recipe for counterexamples to Worse to Do Wrong.

Indeed, we can do all this even if we don't use moral and non-moral reasons in particular. We might instead say that the reasons to save lives are "other-regarding" while the reasons to do nothing are "prudential." Or we could call it a conflict between "impartial" reasons (like *it would save two lives*) and "partial" reasons (like *it would save* my *arms*). All good. So long as the reasons to sacrifice are on a par with the reasons to self-preserve, we can solve the All or Nothing Problem.[13]

6. The Intransitivity Paradox

Now the twist: we have *already* solved Kamm's Intransitivity Paradox. (And the bucket of cold water: it's not entirely obvious what "Kamm's Paradox" amounts to. We'll clear this up as we go.)

The core of the puzzle, as originally presented, is that some cases appear to feature a spooky intransitivity. In Kamm's case, I have three options: keep a promise to meet someone for lunch, save a life at great personal cost, or do nothing.[14] Because life-saving is costly, I am free to do nothing instead: *Do Nothing* ~ *Save 1*. (Where this "tie" means that I may do either option in a pairwise choice.) But since the sacrifice is for a splendid end—saving a life—it is permissible even when it entails breaking lunch obligations: *Save 1* ~ *Keep Promise*. And yet I may not skip lunch to lounge around: *Keep Promise* > *Do Nothing*. (This "defeat" means that I would have to keep the promise in a pairwise choice.) All together:

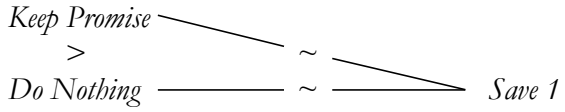

Here, we get an intransitivity in the relation of *weak defeat*, or '≳'.[15] Let '$x \gtrsim y$' mean that x is

[13] One limitation of the parity view, however, is that, since *Save 1* and *Do Nothing* are on a par, we can't say that *Save 1* is better. I think this is a point in favor of the Prerogatives Principle over the Non-Moral Reasons Principle (see Muñoz ms.a).

[14] Kamm's (1985) hero gives a kidney rather than losing arms. I ignore this cosmetic difference.

[15] Tying is also intransitive. Indeed, if weak defeat is intransitive, complete, and reflexive, then either trying or defeat must be intransitive, though not necessarily both (see n.22). Kamm (1985)

permissible in the pairwise choice. ('$x \gtrsim y$' means '$x > y$ or $x \sim y$'.) Kamm's intuitions about her case entail that '$\gtrsim$' is intransitive: *Do Nothing* $\gtrsim$ *Save 1* $\gtrsim$ *Keep Promise* > *Do Nothing*. This means that we have *Do Nothing* $\gtrsim$ *Save 1* $\gtrsim$ *Keep Promise*, but not *Do Nothing* $\gtrsim$ *Keep Promise*. That's Kamm's intransitivity.

The air of paradox creeps in when we ask: how is this possible? If x is tall enough to at least tie with y in a height contest, and y is tall enough to at least tie with z, then it's safe to say that x will at least tie with z, too. Why shouldn't the same hold for contests of permissibility? If x is at least good enough to choose over y, and y over z, doesn't that mean that x is good enough to choose over z, too?[16]

But we have *already* shown that this kind of question is confused, because being permissible, unlike being tallest, is a matter of two independent dimensions—reasons and prerogatives. Indeed, we have just seen how reasons and prerogatives lead to intransitivity in Horton's case:

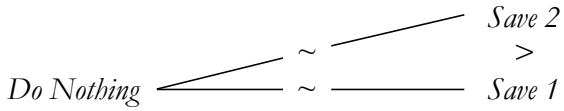

Save 1 $\gtrsim$ *Do Nothing* $\gtrsim$ *Save 2*, but not *Save 1* $\gtrsim$ *Save 2*. The key is a moral mismatch: my prerogatives protect *Do Nothing* more than they protect *Save 1*, even though I have more reason to save more lives. This flows naturally from the view that we have reasons to help others and prerogatives not to self-harm, which is precisely the view we needed to solve the original Paradox. (*Mutatis mutandis* for views with multi-flavored reasons.)

The fear of transitivity, like the pull of Worse to Do Wrong, depends on the assumption that right and wrong depend on only one dimension—a dogma that would rule out supererogation from

herself doesn't use 'tie', and her term for weak defeat is 'may permissibly take precedence over'.

[16] Archer (2016: 445) and Portmore (2017: 294) bring out the oddness in terms of reasons. If the reasons in favor of x can justify acting against the reasons for y, and same for y's reasons against z's, then why can't the reasons in favor of x justify acting against the reasons in favor of z?

the start. We should be glad to get rid of Worse to Do Wrong. For the same reason, we should enjoy our polite farewell to the transitivity of weak defeat.

7. Comparative Prerogatives

Now this is where things get tricky. Kamm's Paradox is kaput; there is nothing fishy about intransitivity per se. But we are still in a pickle insofar as we haven't figured out Kamm's *case*, which cannot be understood as easily as Horton's. Maybe we shouldn't count this as part of the paradox proper, since we are no longer just talking about transitivity (this is the unclarity alluded to earlier); still, the case demands attention.

Here's the problem. Because *Do Nothing* $\gtrsim$ *Save 1*, we have to say that there is a mighty prerogative to *Do Nothing*. But clearly, in the choice between *Do Nothing* and *Keep Promise*, there can't be such a prerogative. If there were, then breaking the promise would be permissible, which is precisely the verdict that we're trying to avoid. It's a datum that *Keep Promise* > *Do Nothing*.

An essential part of this problem is that we can't just say that *Keep Promise* rules out *Do Nothing* by virtue of its being even better than *Save 1*. Keeping the promise isn't better than giving up one's arms to save a life, as shown by the fact that saving the life is supererogatory. The problem is that the prerogative to *Do Nothing* seems to kick in against *Save 1*—the best option—but disappear up against *Keep Promise*.

This is really puzzling. What's going on here?

Kamm (1996: 313) proposes an answer: we are dealing with two independent dimensions. *Save 1* scores highest in one respect—it may "take precedence" over the biggest swatch of options, weakly defeating them—whereas *Keep Promise* ranks highest in another: it requires the greatest efforts. The idea seems to be that *Keep Promise* is urged by the bigger moral reasons ("requiring reasons"), but *Save 1* is overall more justified ("justifying reasons"—what I would call combined

reasons and prerogatives) (cf. Lazar 2019: 246). This view, which is analogous to my own view on Horton's case, has trouble accounting for the betterness of *Save 1*. But a deeper problem is that Kamm's view makes the wrong predictions about Horton's option to *Save 2*.

Recall the combined choice:

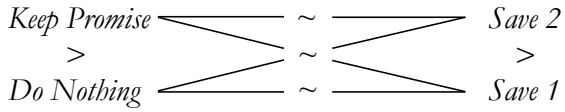

Now, if Kamm is right that we have a mighty prerogative to save lives, and only relatively weak ("requiring") moral reasons to do so, then the justification to *Save 1* should be enough to outweigh the reasons to *Save 2*, and *Save 1* will be permissible.

Here's why. We know that there is some justification for *Keep Promise*, since *Keep Promise* $\gtrsim$ *Save 1*. We also know that there is even *more* justification for *Save 1*; that is Kamm's main idea. But then adding *Save 2* can't make *Save 1* alone wrong. Since *Save 1* is even more justifiable than *Keep Promise*, anything that makes *Save 1* wrong will also rule out *Keep Promise*. But *Save 2* doesn't rule out *Keep Promise*. So it can't rule out *Save 1*, either. Kamm's view thus conflicts with the idea that it's wrong to save one in Horton's case. The view lets us be gratuitously nasty.[17]

The other leading take on Kamm's case, due to Alfred Archer (2016: 459), is that it involves *three* dimensions. Basically, Archer's view is just Kamm's plus the idea that an option can be "morally favored" independently of how strongly it is justified or required. This allows Archer to say that *Save 1* is more favored than *Keep Promise*. Now, perhaps it is odd that "favoring" should be independent of other factors (Portmore 2016: 293n.11). But the real problem for Archer isn't about favoring; it's

[17] Kamm (1996: 336) is sensitive to the worry that we may have to save lives, noting that the proposed view does not explain "why it is obligatory to save a life at small cost to ourselves rather than keep an important business appointment." She then suggests, roughly, that costless lifesaving is supported by stronger moral reasons than promise-keeping. I'm not sure that this will help in the four-option case, since *Save 2* isn't costless, and yet it is favored by strong moral reasons.

that he has to say the same things as Kamm about permissibility, including the claim that it's fine to *Save 1* even given the option to *Save 2*. Neither Kamm's view nor Archer's can be extended to plausibly cover Horton-style cases, where there are better and worse ways to save.

Now what? If we can't treat Kamm's case even with three dimensions, what on earth could do the trick? A leap to four dimensions? A formal *deus ex?*

I don't think we need anything so fancy. Just ask yourself: why am I free to stay out of the building, but not free to break the promise? The obvious answer is: keeping the promise wouldn't crush my arms! Keeping my word isn't any more harmful to me than doing nothing, so I don't have any harm-based prerogative to break the promise by doing nothing.

The big mistake was to think that we could always represent reasons and prerogatives with a single fixed weight. Sometimes, like in Kamm's case, it matters what we are comparing to what. I have a prerogative to *Do Nothing* rather than *Save 1*, because saving the life is more harmful to me; but I don't have any such prerogative to *Do Nothing* rather than *Keep Promise*, since the harm to me is the same either way. What matters is comparative harmfulness. That is why we can't just say that *Do Nothing* gets a prerogative of weight 5, or weight 10, or what have you. The weight depends on the alternative.

The upshot is that prerogatives are "contrastive" (Snedegar 2015), or as I like to put it: they are *comparative*. This means that the weight of a prerogative to do an option isn't fixed by that option's internal features—who dies, who lives, who is harmed how much. From the fact that *Do Nothing* lets me keep my arms, I can't tell whether I have a prerogative to do it, since I would need to know what happens to my arms (and legs!) in the relevant alternatives.

The key to Kamm's case is that prerogatives (or non-moral reasons) are comparative.[18]

[18] Why doesn't comparativity come up in Horton's case? Because the only option protected by a prerogative in that case—*Do Nothing*—is protected equally against both of the alternatives. *Save 1* and *Save 2* cost two arms apiece. (Also, let me be clear: I am not saying that prerogatives are

Stated in the abstract, this might sound like an ad hoc addition. But it's not. It is just common sense that our prerogatives depend on comparative harms, and there is nothing natural about the idea that the costs of heroism should justify one non-heroic option over another. This idea is just confused, as Kamm's case shows, and the solution is that harms matter in a comparative way.

Let me close with one more case, introduced by Dale Dorsey (2013), which has been treated as a variant on Kamm's.[19] Suppose I am able to beat up a stranger, Roy, and take his money, which I can use either to buy a car for myself, or to save ten lives via an amazing NGO. I have three options: *Get a Car*, *Respect Roy's Rights*, or *Save 10*. Intuitively, Dorsey thinks, I don't have to give up my own luxury for charity: *Get a Car* ~ *Save 10*. (A bit heartless, but let's grant it.) I may, however, beat up Roy when lives are on the line: *Respect Roy's Rights* ~ *Save 10*. And yet I would not be justified in assaulting someone just for the sake of new wheels: *Respect Roy's Rights* > *Get a Car*.

We get the same old structure:

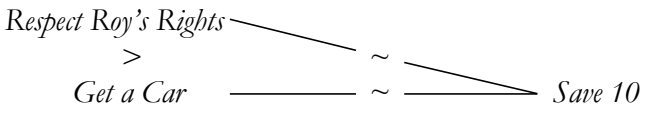

This time, however, we can't say that the two non-lifesaving acts involve the same null cost to self. Respecting Roy costs me a whole car! So why is my prerogative to *Get a Car* comparative? Why does it count for so much against the greater good and for so little against Roy's rights?

Again, I don't think the explanation here has to be ad hoc. On the contrary, the verdicts follow from two familiar views of prerogatives. First, on Hurka and Shubert's view, we have a prerogative to promote our own interests *except* if doing so would violate a right. They argue:

> Surely if it's wrong to kill one innocent person to save two others, as deontological moralities hold, it's also and equally wrong when one of the two is you: your agent-favoring

"essentially comparative" in Temkin's sense (2012: 371). That would mean that the weight of the prerogative to do *x* over *y* is not a function of the options' intrinsic properties. But for all I have said so far, the prerogative's weight might be a function of how much *x* and *y* harm the agent.)

[19] See Archer 2016, Portmore 2017. Dorsey doesn't himself mention transitivity or Kamm's paper. I have cleaned up his cases so that they fit into a single three-option choice.

> permission to pursue your own good must have more weight against the duty to promote the good than it does against deontological constraints…. (2012: 10)

On a rights-based view, meanwhile, it seems clear that a right to spend on oneself doesn't entail a right to steal from others, much less beat them up! (See Muñoz ms.b.)

Both views entail that we have no prerogative to violate rights, even for self-interest. We don't have to agree with this idea. But if we do, then we can explain the presence of a comparative prerogative in Dorsey's case, which turns out to be more complex than Kamm's. Instead of a difference in marginal costs, we find a difference in the significance of the same marginal cost of one shiny new car. The prospect of new wheels does little to justify stealing, but seems to do a fair bit to justify spending on oneself over strangers. This is the same comparativity as in Kamm's case—only with a different source.

8. The Challenge from Menu-Relativity

We have finally hacked our way through the paradoxes of supererogation. Let's recap.

The fundamental idea is that we build our way up to the outré many-option choices from judgments about pairs. The first obstacle is the classic Paradox, which asks: why may we choose a morally worse option over its superior? The answer is that there is a second factor, prerogatives (or non-moral reasons), that is independent of moral value and inherently permissive. That gives us our principle of permissibility: we may choose an option when the reasons and prerogatives in favor can outweigh the reasons against.

We then use this "Prerogatives Principle" to generate pairwise judgments about the options involved in Kamm's and Horton's cases, represented with "defeat" and "tying" relations:

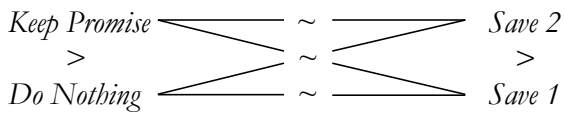

The final step is to say: there isn't a *sui generis* moral mutation when we put the options together,

because reasons and prerogatives are not *menu-relative*. By this, I mean that the reasons and prerogatives to do *x* rather than *y* don't change depending on the presence of any third option *z*.

In the absence of relativity, our pairwise judgments tell us everything we need to know about the balance of reasons and prerogatives, and so we can read off what's permissible and wrong in the full choice just by looking at pairs. The wrong options, like *Do Nothing* and *Save 1*, are those that are defeated by something. The permissible options are those that weakly defeat the alternatives. (If an option can take on its competitors one by one, then absent relativity, it can take them on together.)

It is a crucial assumption that our cases don't involve menu-relativity. But why, you might ask, is that safe to assume? I think this is the most serious challenge to a "pairs-first" approach to the paradoxes. Granted, I think it's a coherent idea that *z*'s presence could make a difference to how *x* and *y* compare. Perhaps this happens in some cases. And indeed, when we add *Save 2* (or *Keep Promise*), doesn't that seem to affect how *Do Nothing* compares to *Save 1?*

> No doubt there is some sort of effect here, as Pummer notes:
>
> > it is a familiar feature of nonconsequentialist ethics that the moral status of an act can depend on which alternative acts are available. In this case, the presence of [*Save 2*] alters the moral status of [*Save 1*], thereby altering the way that [*Save 1*] and [*Do Nothing*] compare morally. (2016: 86–7)

This can sound a bit like menu-relativity, as if adding *Save 2* suddenly makes *Save 1* worse than *Do Nothing*. Is that what's really going on in Horton's case?

I don't think so. We already have a simpler story. Adding *Save 2* makes *Save 1* wrong, while leaving *Do Nothing* permissible, precisely *because* of how the options compare in pairwise judgments:

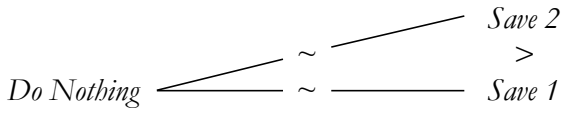

Instead of a change in reasons and prerogatives, we get a divergence in deontic status, because only one option (*Save 1*) is made wrong by the new addition (*Save 2*). The options are no longer both permissible. Indeed, this seems to be how Pummer himself sees the case:

> …with the full choice situation in view, it is clear that there is something to be said against [*Save 1*] that cannot be said against [*Do Nothing*] or [*Save 2*]: the performance of [*Save 1*] constitutes a deliberate refusal to do something much better at no extra cost. This is a serious moral failing. (2016: 86–7)

The failing consists in how *Save 1* compares to *Save 2*, not in how it directly compares to *Do Nothing*. Only *Save 1* is gratuitously worse than an alternative.[20]

So we don't need menu-relativity to explain the effects of adding *Save 2*. The key effect is that only *Save 1* becomes wrong, which can be explained by the fact that only *Do Nothing* is protected against *Save 2* by a prerogative. The same goes for adding *Keep Promise*, in Kamm's case, where the addition rules out *Do Nothing* without erasing the prerogative to *Do Nothing* rather than *Save 1*.

There is also a more formal clue that our cases don't involve menu-relativity. There are two kinds of cases where relativity is definitely the right explanation, cases where we can't explain what's permissible if we only look at the options pair-by-pair. In the first kind of case, subtracting an option can make another option wrong, violating:

> *Property* α
> If x is permissible to choose from a set of options S, then x is permissible to choose from any subset S* of S to which x belongs. (Sen 2017: Chapter 1*6)

In the other kind of case, an option is permissible to choose from any member of a set of sets of options, but not from their union, violating:

> *Property* γ
> Given some sets of options O_i, if x is permissible to choose from any O_i then x is permissible to choose from the union of all O_i. (Sen 1993: 500)

These are necessary and sufficient for us to derive the facts about what's permissible to choose from

[20] Of course, I agree that Horton's case involves the menu-relativity of *something*. It is menu-relative which things are permissible; *Save 1* is only wrong if *Save 2* is on the menu (Snedegar 2015: 479). Interestingly, it is also menu-relative whether two options have the same deontic status. For instance, *Do Nothing* and *Save 1* are both fine in a pairwise choice, but only *Save 1* is wrong if we add *Save 2*. I am not denying any of this. When I say I reject "menu-relativity," I just mean to deny that z's presence can change the reasons and prerogatives to do x over y. Our two factors are independent and comparative, but not themselves menu-relative.

S from pairwise judgments about S's members (including the trivial "pair" of an option and itself). More precisely: the properties hold *just if* the permissible options in S are exactly those that aren't defeated by anything in S.[21] If either property is violated, relativity follows.

But as we have described Horton's case, it obeys both Property *α* and Property *γ*. It only violates their less legendary cousin:

> *Property* β
> If *x* and *y* are both permissible to choose from a set of options S, then if *x* is permissible to choose from a superset S* of S, so is *y*. (Sen 2017: Chapter 1*6)

As shown by the fact that adding a *Save 2* makes only *Save 1* wrong. The essential upshot of violating *β*, it turns out, isn't relativity, but a certain kind of intransitivity.

This gives us an illuminating way to show that Horton's and Kamm's cases, even combined, violate only *β* while respecting *α* and *γ*. We can see that *α* and *γ* are obeyed because the permissible options are precisely those that weakly defeat everything; and we can see that *β* is violated because weak defeat is intransitive in a certain way:

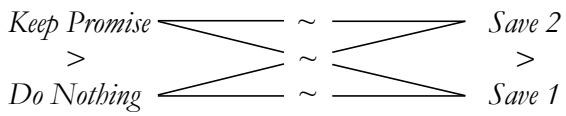

This violates what I call:

> *Transmission Over Ties*
> If $x > y \sim z$, then $x > z$.[22]

---

[21] This is true if we assume that, for any subset S* of our (finite!) set of options S, there is an option that would be permissible in a choice from S*; in such conditions, we say that there is a "choice function" *f* defined over S; α and γ are properties of *f* necessary and sufficient for *f* to be generated from a non-menu-relative relation '≿', such that $x \in f(S^*)$ *iff* $x \in S^* \,\&\, \forall y \in S^*, x \succsim y$ (Sen 1993: 499–500). We can't assume that a choice function is defined in a potential dilemma, where there might not be permissible options; but our case is clearly not a dilemma.

[22] Transmission Over Ties—which Sen (2017: 1*6) calls "PI-intransitivity"—is formulated to be intuitive. But this formulation might obscure the crucial fact that any violation of Transmission Over Ties implies an intransitive '≿'. If $x > y \sim z$, but $z \succsim x$, then we have $y \succsim z \succsim x$ without $y \succsim x$. The key is that, assuming '≿' is complete and reflexive, Transmission Over Ties is strictly weaker than the transitivity of '≿', which is equivalent to the conjunction of Transmission with the

Which, given *α* and *γ*, is equivalent to Property *β* (if a choice function is defined; see n.21, Sen 2017: 64). Transmission Over Ties says: if *x* defeats an option, it defeats anything that option ties with. But defeat doesn't transmit in our cases. For example: *Save 2* > *Save 1* ~ *Do Nothing* ~ *Save 2*. No surprise, either, since the case has two independent factors; there is a prerogative not to save anyone; more reason to keep the promise than do nothing, and more reason still to save more lives.

This is a rather nice result. We can see that the cases violate *β* alone in virtue of the exact kind of intransitivity that Kamm discovered: the failure of Transmission Over Ties.

I should also note that the cases do *not* violate the most plausible kinds of transitivity. In particular, as Temkin (2012: 196) notes of Kamm's case, our cases obey:

> *Acyclicity*
> If $x > y > \ldots > z$, then it's not true that $z > x$.

What a relief. The possibility of cycling is legendarily controversial (cf. Temkin 2012), as cycles leave the agent with no undefeated option. Transmission failures don't entail this moral doom.

That completes my defense against the challenge from menu-relativity. We don't need relativity, and so we are free to analyze Kamm's and Horton's cases using judgments about pairs.

Let me now switch to offense. Lazar and Barry (ms.) argue that Horton's case involves menu-relativity, and while I think their discussion is insightful, their appeal to menu-relativity may be open to objections.

Lazar and Barry think that, in a pairwise choice, *Save 1* isn't worse than *Do Nothing*. But they say that adding *Save 2* to the menu changes this. I have strong ("agent-relative") reasons not to disrespect people. Gratuitously leaving someone to die—as in the case where I save only one even though I could just as easily have saved two—is disrespectful. So, adding *Save 2* gives me new reasons not to *Save 1*, which are weighty enough to make *Save 1* worse than *Do Nothing*.

---

transitivity of '>', also known as the "quasi-transitivity" of '≳' (Sen 2017: 66).

There is something right about this idea. But I have three objections to it as a treatment of Horton-style cases. The first is dialectical: we don't *need* to posit changing reasons to account for what's permissible in Horton's case, as I have just tried to show.

Second, if we allow for changing reasons in this case, we open the door to more radical changes. After all, if reasons can be menu-relative, we should predict that there can be failures of Property *α*—like a case where *x* is fine in {*x*, *y*, *z*} but wrong in {*x*, *y*}. This is much stranger than a mere *β*-failure. As Vallentyne and Tungodden (2005: 143) ask: "If [*x*] is a winner against [*y*] and *z*, why would it not also be a winner against *y* alone?"[23]

I don't mean to say that we have some a priori guarantee that Property *α* never fails; there are putative counterexamples (Lazar and Barry cite one from Pettit 1991; see also Sen 1993, Kamm 1996: 343). But to the extent that these examples are controversial—and they are certainly more controversial than failures of *β*—a view that predicts *α*-failures is to that extent costly.[24]

Now the big challenge. Lazar and Barry's view of Horton's case does not generalize to Kamm's case. When we add *Keep Promise*, that makes *Do Nothing* wrong without changing the deontic status of *Save 1*. This is the same kind of effect that Lazar and Barry explain in Horton's case with relative reasons: *Save 2* makes *Save 1* alone wrong, so it must make *Save 1* worse. But this explanation doesn't fit Kamm's case. How could the option to *Keep Promise* alter the balance of reasons between

---

[23] There is also something funny about *α*-flouting preferences, as in Morgenbesser's joke:

BARTENDER: Would you like red or white wine?
PATRON: White, please.
BARTENDER: I forgot: we also have beer.
PATRON: In that case, I'll take the red.

(See Kamm 1996: n.31.)

[24] Property *α*, also known as "The Independence of Irrelevant Alternatives" (IIA), is not always kept distinct from nearby ideas. Kamm (1996: 344) mixes up IIA with its homophone in social choice theory (which says: society's preference over a pair depends only on how individuals rank that pair; see Arrow 1951). Rulli and Worsnip (2016) define IIA so that it entails Property *β*.

Do Nothing and *Save 1?* It adds reasons against doing nothing, since loafing would now break the promise. But so would being the hero! The promise doesn't exclusively count against *Do Nothing*; so, it can't make *Do Nothing* worse than *Save 1*. Presumably, then, Kamm's case needs some other explanation besides menu-relative reasons—such as my explanation. But then we might as well use that explanation for Horton's case, too, since it works there in much the same way.

Menu-relative reasons are unnecessary for treating Horton's case; they are costly to posit; and they don't help with Kamm's case. That is why I think we should explain both cases with static (but comparative!) reasons and prerogatives. This approach gives us a more unified picture of the paradoxes, since it covers both Kamm's and Horton's cases, and its core concept of independent dimensions is already motivated by the classic Paradox of Supererogation.

9. Conclusion

There are two kinds of objection to the possibility of supererogation (Dancy 1993b). The first says that the very concept of supererogation is confused or problematic. The second says that the concept, however coherent, is empty, because there is something morally objectionable in the idea that people may ever do less than best.

My concern here has been the conceptual objections—the paradoxes. I have argued that these shouldn't scare us, and that we can construct an elegant account of supererogation using nothing more than two dimensions—reasons and prerogatives, flavors of reasons—so long as we allow for comparativity. We don't need intervals, baselines, yet more dimensions, menu-relative reasons, or cycles of defeat. We just need the idea that you have to do what is best unless you have enough of a prerogative to do something else; you have to save lives, for example, unless it would harm you much more than the alternatives.

But solving paradoxes can also have an instrumental value. We are clearing the ground for a

more fruitful normative debate about which acts are really supererogatory. Does a non-moral reason to keep one's arms outweigh the moral reasons to save a life? Do we have any prerogative to eat the meat of animals? To close our borders to refugees, exhaust our natural resources? Or are we just plain required not to do these acts given the hefty reasons against? These questions are urgent and difficult. But we might see them in a new light, or at least with more clarity, if we could just get them out of the shadow of paradox.[25]

[25] This article was originally titled "Supererogation and Rational Choice: Incommensurability, Intransitivity, Independence." It has come a long way. I would like to thank Kieran Setiya, Caspar Hare, Tamar Schapiro, Jack Spencer, Anders Herlitz, Yael Loewenstein, Oli Rawle, Kerah Gordon-Solomon, and Theron Pummer for their invaluable comments on drafts of this paper. For helpful discussion, I am indebted to Joe Bowen, Kelly Gaus, Jocelyn Wang, Katie Steele, Sam Dishaw, Steve Yablo, Anni Räty, Frances Kamm, Quinn White, Justin Khoo, Nathaniel Baron-Schmitt, David Builes, Toby Handfield, David Barnett, and my audience at St Andrews's Centre for Ethics, Philosophy, and Public Affairs. I am also grateful to Seth Lazar for advice, and for sponsoring my visit in 2018 to the Australian National University, where I started this project. Finally, special thanks to the editors of *Noûs*, and to an anonymous referee for three rounds of brilliant comments, which decisively improved the paper's terms and structure.

REFERENCES

Archer, Alfred (2016). "Moral Obligation, Self-Interest, and the Transitivity Problem," in *Utilitas*, 28 (4): 441–464.

-------(2017). "Supererogation," in *Philosophy Compass* 13 (3).

Arrow, Kenneth (1951) [1963]. *Social Choice and Individual Values*. New York: Wiley. 2nd Edition 1963.

Bader, Ralf (forthcoming). "Agent-Relative Prerogatives and Suboptimal Beneficence," in *Oxford Studies in Normative Ethics*.

Bedke, Matthew (2011). "Passing the Deontic Buck," in R. Shafer-Landau (Ed.), *Oxford Studies in Metaethics*, volume 6, pp. 128–52. Oxford: Oxford University Press.

Benn, Claire (2017). "Supererogatory Spandrels," in *Ethics & Politics* 19 (1): 269–290.

Bradley, Ben (2016). "Against Satisficing Consequentialism," in *Utilitas* 18: 97–108.

Chang, Ruth (2002). "The Possibility of Parity," in *Ethics* 112 (4): 659–688.

Dancy, Jonathan (1993a). *Moral Reasons*. Oxford: Blackwell.

-------(1993b). "*Beyond the Call of Duty: Supererogation, Obligation, and Offence* By Gregory Mellema," in *Philosophical Books* 34 (1): 48–49.

Darwall, Stephen (2013). "'But It Would Be Wrong'," in *Morality, Authority, and Law: Essays in Second-Personal Ethics I*. Oxford: Oxford University Press: 52–71.

Dietrich, Franz and List, Christian (2017). "What Matters and How it Matters: A Choice-Theoretic Interpretation of Moral Theories," in *Philosophical Review* 126 (4): 421–479.

Dorsey, Dale (2013). "The Supererogatory, and How to Accommodate It," in *Utilitas* 25 (3): 355–382.

Dreier, James (2004). "Why Ethical Satisficing Makes Sense and Rational Satisficing Doesn't," in Michael Byron (Ed.), *Satisficing and Maximizing*. Cambridge: Cambridge University Press.

Ferguson, Benjamin (2016). "The Paradox of Exploitation," in *Erkenntnis* 81: 951–972.

Ferguson, Benjamin and Köhler, Sebastian (2019). "Betterness of Permissibility," in *Philosophical Studies*. Print version forthcoming.

Fishburn, Peter C. (1970). *Utility Theory for Decision Making*. New York: John Wiley & Sons.

Gert, Joshua (2004). "Value and Parity," in *Ethics* 114 (3): 492–510.

-------(2007). "Normative Strength and the Balance of Reasons," in *Philosophical Review* 116 (4): 533–562.

Hare, Caspar (2010). "Take the Sugar," in *Analysis* 70 (2): 237–247.

Heyd, David (2016). "Supererogation," in Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2016 Edition), URL = <https://plato.stanford.edu/archives/spr2016/entries/supererogation/>.

Horgan, Terence and Timmons, Mark (2010). "Untying a Knot from the Inside Out: Reflections on the 'Paradox' of Supererogation," in *Social Philosophy and Policy* 27(3): 29–63.

Horton, Joe (2017). "The All or Nothing Problem," in *Journal of Philosophy* 114 (2): 94–104.

Hurka, Thomas (1990). "Two Kinds of Satisficing," in *Philosophical Studies* 59 (1): 107–111.

Hurka, Thomas and Shubert, Esther (2012). "Permissions to Do Less Than Best: A Moving Band," in *Oxford Studies in Normative Ethics Volume 2*: 1–27.

Kamm, Frances (1985). "Supererogation and Obligation," in *Journal of Philosophy* 82 (3): 118–138.

-------(1996). *Morality, Mortality, Volume II: Rights, Duties, and Status*. New York: Oxford University Press.

Lazar, Seth (2019). "Accommodating Options," in *Pacific Philosophical Quarterly* 100 (1): 233–255.

Lazar, Seth and Barry, Christian (ms.). "Acting Beyond the Call of Duty: Supererogation and Optimization."

Little, Margaret and McNamara, Colleen (2017). "For Better or Worse: Commendatory Reasons and Latitude," in Mark Timmons (ed.), *Oxford Studies in Normative Ethics, Vol 7*: 138–160.

McMahan, Jeff (2018). "Doing Good and Doing the Best," in *The Ethics of Giving: Philosophers' Perspectives on Philanthropy*, Paul Woodruff (Ed.). New York: Oxford University Press: 78–102.

Muñoz, Daniel (ms.a). "Better to Do Wrong."

-------(ms.b). "Why Isn't Supererogation Wrong?"

Parfit, Derek (1982). "Future Generations: Further Problems," in *Philosophy and Public Affairs* 11 (2): 113–172.

-------(2011). *On What Matters, Volume One*. Oxford: Oxford University Press.

Pettit, Philip (1991). "Decision Theory and Folk Psychology," in Michael Bacharach and Susan Hurley (Eds.), *Essays in the Foundations of Decision Theory*. Basil Blackwell.

Portmore, Douglas (2011). *Commonsense Consequentialism: Wherein Morality Meets Rationality*. New York: Oxford University Press.

-------(2017). "Transitivity, Moral Latitude, and Supererogation," in *Utilitas* 29 (3): 286–298.

-------(2019). *Opting for the Best: Oughts and Options*. Oxford: Oxford University Press.

Postow, Betsy (2005). "Supererogation Again," in *Journal of Value Inquiry* 39: 245–253.

Pummer, Theron (2016). "Whether and Where to Give," in *Philosophy and Public Affairs* 44 (1): 77–95.

-------(2019). "All or Nothing, but If Not All, Next Best or Nothing," in *Journal of Philosophy* 116 (5): 278–291.

Rabinowicz, Wlodek (2008). "Value Relations," in *Theoria* 74: 18–49.

Raz, Joseph (1975). "Permissions and Supererogation," in *Philosophical Quarterly* 12 (2): 161–168.

Scheffler, Samuel (1982). *The Rejection of Consequentialism: A Philosophical Investigation of the Considerations Underlying Rival Moral Conceptions*. Oxford: Oxford University Press.

Sen, Amartya (1993). "Internal Consistency of Choice," in *Econometrica* 61 (3): 495–521.

-------(2017). *Collective Welfare and Social Choice: Expanded Edition*. Cambridge: Harvard University Press.

Sinclair, Thomas (2018). "Are We Conditionally Obligated to be Effective Altruists?" in *Philosophy and Public Affairs* 46 (1): 36–59.

Slote, Michael (1984). "Morality and Self-Other Asymmetry," in *The Journal of Philosophy* 81 (4): 179–192.

-------(1985). *Common-sense Morality and Consequentialism*. London: Routledge & Kegan-Paul.

-------(1991). "Shelly Kagan's *The Limits of Morality*," in *Philosophy and Phenomenological Research* 5: 915–917.

Snedegar, Justin (2015). "Contrastivism About Reasons and Ought," in *Philosophy Compass* 10 (6): 379–388.

-------(2016). "Reasons, Oughts, and Requirements," in R. Shafer-Landau (Ed.) *Oxford Studies in Metaethics*, XI. Oxford: Oxford University Press: 183–211.

Tadros, Victor (2011). *The Ends of Harm: The Moral Foundations of Criminal Law*. Oxford: Oxford University Press.

Temkin, Larry (2012). *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning*. Oxford: Oxford University Press.

Tungodden, Bertil and Vallentyne, Peter (2005). "On the Possibility of Paretian Egalitarianism," in *Journal of Philosophy* 102 (3): 126–154.