Forthcoming in *Oxford Studies in Metaethics*, Vol. 16.
Penultimate version. Please ask before quoting.

Reasons, Competition, and Latitude*

Justin Snedegar
University of St Andrews
js280@st-andrews.ac.uk

1 Reasons, Trade Offs, and Competition

Part of the task of moral theory is to explain why options have the overall moral statuses that they do. Familiar moral statuses include being *required*, *permitted*, *forbidden*, and what we really *should* do, even if we aren't required to. The overall moral status of a particular option depends on the features of that option, for example whether it is an act of promise-breaking, or how many people it will make better off. Since these kinds of morally significant features can often pull in opposite directions—sometimes breaking a promise can help lots of people, for example—we shouldn't try to explain the overall status of the options just in terms of simple, exceptionless rules like 'Don't break promises' or 'Help as many people as you can'. This would lead to many cases of moral dilemmas, in which you break a rule, and so act wrongly, no matter what you do.

Instead, many philosophers follow Ross (1930) in assigning the various morally relevant features an option has a *contributory* role. Each feature matters, but not by settling the question of the overall status of the option, as the exceptionless rules picture would imply. Rather, we have to take all the morally relevant features of the possible options into account, and then come to a conclusion about what the overall statuses of the options are. This is going to involve tradeoffs. To determine the overall moral status of the option, these contributory considerations that pull in different directions *compete* with one another. It is common to think of these contributory considerations as contributory or *pro tanto* reasons. So on this picture, the various contributory reasons bearing on our options compete with one another to determine the overall moral status of those options.[1]

I take this picture as my starting point: the key claim is that the overall normative status of our options is going to be explained by the competition between the contributory reasons bearing on those options. This paper aims to make progress on understanding how it is that reasons compete to determine overall moral status. In the next section I'll describe the most common account of how reasons compete. Then I'll present a familiar problem for this account, having to do with the fact that we typically have a significant degree of

* Thanks to audiences at St Andrews, West Virginia University, the 2019 Wisconsin Metaethics Workshop, Theron Pummer, and Doug Portmore.

[1] This kind of picture, at this level of abstraction, is near orthodoxy. See Lord and Maguire (2016) for a nice overview.

latitude or choice in deciding what to do. The task for a theory of how reasons compete is to explain how our reasons can compete in a way such that we could have this kind of latitude. A particularly important aspect of this challenge is to explain how some options could be supererogatory, where this means that it is permissible not to do them, even though they are better, in some sense, than our other options.

After illustrating the challenge by showing how it raises problems for the familiar account of how reasons compete, I'll argue against two ways of trying to explain latitude before offering an alternative. All three accounts aim to explain latitude by appealing to distinctions between reasons that compete in importantly different ways. The first two appeal to different kinds of *non-requiring* reasons—reasons which, even if they win the competition, do not ground requirements.[2] The third appeals to a sharp distinction between reasons for and reasons against.

Though the question of how reasons compete to determine overall status is an important one for all of normative philosophy—at least on the kind of reasons-based picture I am assuming here—I will be focused on morality. I will often try to formulate things in a general way, but ultimately I want to officially restrict any claims I make to the moral domain. I think that some extra complications arise outside of the moral domain; in particular, the phenomena of latitude and supererogation which are central to this paper are less plausibly features of, e.g., ordinary practical rationality or epistemology.[3]

2 Balancing

The most common way of thinking about how reasons compete is best introduced *via* a metaphor of weighing or balancing. Reasons compete by contributing some amount of normative weight in favor of the options that they're reasons for, and some reasons are weightier than others. That going to the movie would be mildly enjoyable is a fairly weak reason to go, and that skipping the movie to help someone in dire need would prevent a lot of suffering is a weightier reason to do that, instead. So if these are the only reasons in play, then the reason favoring helping the person in need wins the competition. On the most straightforward version of this account, this is what I'm *required* to do.

I call this a *Balancing Account* because the central idea is that the reasons compete in a way that's analogous to the way that differently weighted marbles on a scale or balance compete to let us see which are the heaviest. Then the option that wins is the one that the balance of reasons favors; this is how we get from the contributory reasons in a situation

[2] For very useful discussions of this general kind of strategy, see Little (2013), Little and Macnamara (2018, forthcoming).

[3] For discussion of the plausibility of supererogation, latitude, and satisficing in domains other than morality, see Slote (1989), the papers in Byron, ed. (2004), and McElwee (2017).

to the overall deontic status of the options. Here's the most common way to formulate this idea:

> *Balance Required*: You are required to ϕ if and only if the reasons for ϕ outweigh the reasons for any alternative to ϕ.[4]

> *Balance Permitted*: You are permitted to ϕ if and only if the reasons for ϕ are at least as weighty as the reasons for any alternative to ϕ.

If the reasons for helping the person in need outweigh the reasons for going to the movie, I am required to help, and going to the movie is not permissible. This view is intuitive and captures lots of what we want to say about how reasons compete. It's easy to imagine the deliberating agent weighing up pros and cons to determine where the balance of support lies. But it is also well known that this account, at least as stated so far, is oversimplified. In the next section I introduce a well-known problem for this account that is the focus of this paper.

3 Latitude

The problem that I'll focus on begins from the observation that on Balancing Account, typically, one option will win the competition and the others will lose. The option that wins will be the one that's most strongly supported by the reasons. Any option which is less strongly supported will lose. If you are required to take the option that wins this competition, then you are also required not to take the other options—they are impermissible. This seems *prima facie* very plausible: if one option is more strongly supported by the reasons than any other, then that's the one you have to take. Doing what is, or at least what you know to be[5], less strongly supported by reasons than some alternative would seem to be impermissible.

But there are many cases that suggest that this is not correct. Often, agents have *latitude*: they can permissibly select any of a range of options.[6] Some will of course be ruled out as impermissible, but at least often there will not be a single option that the agent is required to take. Suppose that you're trying to decide where to donate some money.

---

[4] See, e.g., Parfit (2011), p. 32, Schroeder (2015), and Lord (2018), pp. 10-11. I should note that Lord, at least, is explicitly concerned with rational requirement, rather than moral requirement, and so may well hold a different view about moral requirement. As mentioned above, I believe that the sort of maximizing requirement encoded here is more plausible in the rational and epistemic domains than in the moral domain, but will not address that here.

[5] I'll be ignoring this qualification here. So you can read this paper as primarily talking about objective reasons, requirement, etc., rather than subjective reasons, requirement, etc.

[6] See, e.g., Raz (1999).

Donating to the Neo-Nazis is impermissible, but plausibly there will be a range of options which you could permissibly choose. Effective Altruists, who think we're obligated to make sure our charitable donations are very effective, give you a list of several different charities on their website, suggesting that you can permissibly donate to any of them.[7] Perhaps you're even morally permitted not to donate at all.[8] There are many cases like this, where we have multiple permissible options rather than one single required option.

On the Balancing Account, an obvious way in which we could have latitude is if the reasons support multiple options equally—if we have ties amongst our options. In this case, none of the tied options will be required, and all will be permissible, assuming that the tied options are all better than any other options. None of them win the competition, but none of them lose, either (though of course some other options may lose). But if this is what is required to explain latitude, then latitude will be quite rare, since perfect ties amongst our options would seem to be quite rare. For example, even if there are slight differences between the effectiveness of possible charitable donations, several different ones might still be permissible.

A similar, but at least initially more plausible, way of explaining latitude on the Balancing Account is to hold that there is often *incommensurability* or *incomparability* amongst the reasons supporting our options.[9] Two options are incommensurable when neither is better supported by the reasons than the other, and they are not exactly equally as well-supported as one another. Rather, we can't draw these kinds of comparisons between them. Given that the reasons supporting our options are often quite different from one another—the pleasure of a lavish evening out seems quite different from the good of donating the money, for example—this is initially plausible.

The problem with this explanation of latitude becomes apparent when we think of one very important kind of case in which we seem to have latitude. These are cases in which we have the opportunity to perform a supererogatory action. An action is supererogatory when it goes "beyond the call" of morality—when it's above and beyond what morality demands of you. Paradigm cases are heroic actions, like running into burning buildings or diving on grenades to save others. What you consider supererogatory will of course depend on what you take to be morally required. But common sense morality also seems to recognize more everyday kinds of actions, like donating significant portions of your income to charity, or to borrow an example from Horgan and Timmons (2010), taking your elderly neighbor to a baseball game, which is something she used to enjoy doing with her husband.

---

[7] https://www.givewell.org/charities/top-charities. (Accessed on July 29, 2019.)

[8] See Pummer (2016).

[9] See, e.g., Raz (1999). See also Chang (2001); in particular, see the discussion of *parity*, which offers another theoretical option for explaining latitude on the Balancing Account.

These are actions that you undeniably have good reason to do, and that we'd praise you for doing, but which we would not blame you for not doing. So you have latitude: you can permissibly choose to do the supererogatory action, but you can also permissibly choose not to.

Now consider the idea that latitude is to be explained in terms of incommensurability. This would imply that in these cases in which you have the opportunity to do something supererogatory, the reasons supporting the supererogatory action and the reasons supporting omitting it—say, that you don't really feel like going to the baseball game, or have had your eye on some small luxury purchase, or that you don't want to risk your own life—are incommensurable. As noted, there is plausibly an important difference between these reasons: for example, the ones supporting the supererogatory option are other-involving, while the ones supporting the non-supererogatory option are often self-involving, in some sense. And there are important questions about how exactly these reasons compete with one another.[10] But to hold that they are incommensurable seems to be an overreaction. This would imply that neither option is more strongly supported by reasons than the other. But it is more plausible to hold that the supererogatory action genuinely is more strongly supported. Note that this also tells against the previous way of trying to explain latitude by positing frequent ties amongst our options, since it is (even more) implausible that the supererogatory action and the self-interested action are exactly equally well supported by reasons.

Neither of the obvious ways of explaining latitude on the Balancing Account look promising, especially once we remember that an important sort of latitude involves supererogation. One pessimistic reaction is to conclude that we're mistaken to think that there is such a thing as supererogation, and relatedly, that we're mistaken to think we really do have the kind of latitude that we seem to have. To reiterate what was *prima facie* attractive about the Balancing Account, how could it be permissible to do what you know you have less reason to do?[11] Those who accept very demanding accounts of morality, for example, may take the difficulty in explaining how we could be permitted to perform actions which we know are less strongly supported by reasons than others to show that in fact we are not so permitted, and that we're always required to do the best thing we can. We may add that we can nevertheless be excused for not doing the very demanding thing, though we still act wrongly. But I want to try to put off accepting this sort of demanding account of morality

[10] For an account of moral supererogation and latitude based on the distinction between moral reasons and non-moral reasons, though not on incommensurability, see Portmore (2011), especially chapters 5 and 6. I do not have the space here to properly engage with Portmore's account.

[11] For some important discussions of this "paradox of supererogation", see Kagan (1989), Gert (2004), Dreier (2004), and Horgan and Timmons (2010).

and instead consider ways to try to vindicate the common sense idea that we often do have latitude, including in cases of supererogation.

Before moving on to consider ways of improving upon the simple Balancing Account in order to explain latitude, I should make one qualification. One important kind of latitude involves cases in which there are several different *ways* of carrying out some relatively general action: I could pay you back the $20 you loaned me with two $10 bills today, or with a $20 bill tomorrow, or I could give you a $50 bill and let you give me $30 back, and so on. These cases raise complications that I want to try to avoid here, so I will focus on cases in which the options under consideration are not naturally described as ways of carrying out the same relatively general action.[12] Further, I will only be concerned with ways of addressing the challenge that appeal to distinctions between reasons that compete in importantly different ways.

4 Merely Justifying Reasons

The first account begins with a distinction between two different sorts of reasons, or different "dimensions of strength" (Gert 2004, 2007) that reasons might have. First are requiring reasons. These can require you to act or not act in some particular way, if they win the competition. Second are merely justifying reasons. These can justify you in acting in some way—that is, make it permissible for you to act in some way—but not require that you act in that way, even if they win the competition.

Merely justifying reasons are primarily defensive, on this account. There may be some requiring reasons that would require you not to ϕ, but as long as there are sufficiently good justifying reasons in favor of ϕ, it will be permissible. This kind of view is helpfully put in terms of criticism (Greenspan 2005, 2007, 2010): the requiring reasons provide criticisms of not following them, while the justifying reasons answer those criticisms, and so, if they're good enough, can justify you in acting against the requiring reasons.

For example, from the perspective of ordinary practical rationality, there are strong requiring reasons for you not to run out into traffic—there are good criticisms of doing so. But if there's a child stranded in the middle of the street, that is the kind of thing that can justify you in doing so by answering the criticism: "You could be killed!", "Yes, but there's

[12] In fact, on Portmore's (2011, 2019) view, these kinds of cases are really the same. That's because Portmore argues that the reasons we have to perform individual actions, like donating to this particular charity on this occasion, are derivative on the reasons we have to perform series of actions, or to live out our lives in certain ways (which are themselves derivative on our reasons to adopt certain life-shaping ends). So, for example, donating to charity A, charity B, or no charity at all on this particular occasion may all be permissible because each is compatible with a life that we have sufficiently good reason to adopt. More precisely, each is entailed by some larger series of actions that we have sufficiently good reason to perform. Thanks to Doug Portmore here. See also work on Kantian imperfect duties and supererogation, e.g., Baron (1987, 2015).

a kid out there!". But crucially, according to this view, answering a criticism of an action does not amount to giving a reason for it that will generate a requirement. Though the child being in the street has very strong justifying force—it can justify you in acting against strong requiring reasons, like that you might be killed—it does not have requiring force, and so you are not rationally required to run into the street. On a natural extension to the moral perspective, the dimensions of strength are reversed: that you can save the child's life by running out is a morally requiring reason to do so, but you can morally justify yourself in not running out into the street by citing considerations of personal safety.

The explanation of latitude, then, is this. An option is permissible as long as there is not unanswered requiring reason not to do it. This can happen because there is no requiring reason not to do it, or because the justifying reasons for that action are good enough to answer those criticisms. An action is impermissible—you're required not to do it—when there are unanswered requiring reasons not to do it. An action is required when it is the only permissible option—when all the others are ruled out. This leaves room for there to be multiple permissible options. They will all be supported by sufficiently strong justifying reasons. Put in terms of competition, there may be multiple options for which the justifying reasons win out over the requiring reasons not to do them, so each will be permissible. But since the reasons that win the competition are merely justifying reasons, none will be required.

As noted above, an important class of cases in which we have latitude are cases involving supererogation. In these cases, you can permissibly choose a supererogatory action, or an action that is not supererogatory, but just fulfills your obligations. So an account of how reasons determine overall normative status that is motivated largely by the existence of latitude should explain supererogation. And in fact, Joshua Gert (2004, 2007) argues for his view partly on the basis of its account of supererogation. In particular, he argues that he can solve the so-called "Paradox of Supererogation" of how a supererogatory action could be justified but not required by holding that the reasons supporting supererogatory actions are merely justifying. They answer the (often strong) criticisms of doing the supererogatory option, e.g. running into a burning building, and so make the supererogatory option permissible. But since they're merely justifying reasons, they can't generate a requirement to perform that option. On the other hand, the non-supererogatory but merely obligation-fulfilling option will also be justified, since there may not be any requiring reasons not to do it, or if there are you can justify yourself in acting against them by citing considerations of self-interest or other kinds of partial reasons.

But as I'll argue now, an appeal to merely justifying reasons does not give us a satisfying account of latitude, and in particular of supererogation. If we focus just on determining which actions are required and which are permitted, this account is attractive.

Some reasons tend to generate requirements, but others may justify us in acting against those reasons. If this happens for multiple options, then multiple options will be permitted, so the agent will have latitude. But there is at least one other important overall status that an option can have. There are sometimes options that we really *should* do, or *ought* to do, though we aren't *required* to do them.[13] Supererogatory actions are paradigm examples. Think of how you might describe the option of donating 20% of your income to charity, assuming you can afford it: "Well, you really *should*, though of course you don't have to". What we mean is something like that it would be good if you did, and that there is good reason to do so, but it goes beyond the call of duty, and so we won't blame you if you don't. Whether or not you really should perform some action is going to be determined by the reasons you have, at least on the reasons-based picture I'm assuming here. So our theory of how reasons determine overall normative status needs to account for this status. Moreover, cases in which there's some option that we should perform, but aren't required to, we have latitude. So a theory of how reasons compete to determine overall status that has as a central goal explaining how we could have latitude should definitely account for this status. Otherwise, all we can say about a supererogatory action is that it is permitted but not required; but of course this is true for the non-supererogatory but merely obligation-fulfilling option, as well.[14] At least some supererogatory actions have the special status of also being what you should do. But the appeal to merely justifying reasons does not account for this.

Whereas the simple Balancing Account focuses on the competition between reasons supporting different options, the justifying reasons account focuses just on the competition between the reasons bearing on a particular option. There are requiring reasons against running out into the street, like that it is a serious risk to your life. The reasons that compete with these are the justifying reasons in favor of running into the street, like that there is a child stranded there. These justifying reasons may be good enough to answer the criticisms given by the requiring reasons, in which case you'd be permitted to run out into the street. Importantly, reasons bearing on other options, like having an ice cream cone, don't enter this competition. Instead, there will be a separate competition between the requiring reasons to have or to not have an ice cream cone, and the justifying reasons that give answers to the relevant criticisms, which will determine whether having an ice cream cone is permissible.

---

[13] This kind of distinction is not always fully appreciated in ethics, but it is very old. See, for example, Hobbes on command vs. counsel in the *Leviathan*, Chapter 25, Sec. 1 and Grotius on law/right vs. counsel in *On the Law of War and Peace*, Book I, Chapter 1, Section 9. Recent discussions in connection with the issues in this paper include Bedke (2011) and Snedegar (2016). In deontic logic, see McNamara (1996). In linguistics, see, e.g., von Fintel and Iatridou (2008) and Portner (2009) on weak vs. strong deontic necessity modals.

[14] Bedke (2011) raises this kind of objection to Gert's view.

But to determine which option (if any) you *should* perform, in the relevant sense, we need to be able to draw comparisons amongst the options, in terms of how well supported they are by the reasons. This is because the relevant notion of *should* is comparative: to say that you should ϕ is to say that ϕ is the best, that it is the option that's most strongly supported by reasons (see Sloman (1970) and Finlay (2009)). If the only competition between reasons is the competition between the reasons bearing on particular options—the requiring reasons not to run into the street and the justifying reasons to do so, for example—then we can't make these kinds of comparative evaluations to determine which option is most strongly supported, and so what you should do. This overly restrictive view of how the competition between reasons determines overall normative status is thus what lies behind this account's inability to give a satisfying account of the overall moral status of supererogatory actions.

The basic problem for the appeal to merely justifying reasons is that we can't make comparisons amongst the options of how strongly supported those options are by the reasons, overall.[15] This is because the competition between the requiring reasons and the merely justifying reasons is a competition just between the reasons for and against an individual option.[16] It may be that there are connections between the requiring or justifying reasons bearing on one option and those bearing on the alternatives. For example, requiring reasons for one option may be requiring reasons against alternatives. But this would not solve the problem. These "transmitted" requiring reasons against the alternative would still only compete with the justifying reasons in favor of that same alternative to determine whether or not it is permissible. So we would still have no way of comparing how strongly the reasons support different options. Determining which option is *most* strongly supported is what we need to capture the overall status of being what you should do. Next I will consider a different way of explaining latitude that does allow us to draw these comparisons between different options.

5 Merely Commendatory Reasons

While the previous account aimed to explain latitude via an appeal to purely defensive justifying reasons, the account I'll consider now appeals to a special class of reasons that

[15] Gert (2004, 2007) is particularly clear that questions about how strong the reasons are, overall, are misguided.
[16] Tenenbaum (2007) also criticizes Gert's (2004, 2007) view for failing to take account of the alternatives in evaluating the overall status of an option. But in a way Tenenbaum's criticism is even more serious because he shows that as a result the view gives a mistaken account of permissibility and requirement, which is what the view was designed to explain. See also Portmore (2019), pp. 10-11, who points out that Gert (2012), p. 118, has revised the view to take account of the alternatives. But it is still not clear how to make the kinds of comparative claims about *how strongly* the reasons support different options, which is what would be needed for an account of *should*, as opposed to requirement.

positively support the options they are reasons for. I'll call these *merely commendatory reasons*, though different philosophers give them different names.[17] What these reasons share with merely justifying reasons is that even if they win the competition between reasons, they cannot generate requirements for the options they support. This distinguishes both justifying reasons and these commendatory reasons from the more familiar *requiring* reasons, which do generate requirements when they win the competition.

While merely commendatory reasons do genuinely support options, they are also such that you can ignore them without facing any criticism, even if they win the competition, and in fact even if they are unopposed. It is perhaps easiest to see the appeal here if we focus first on a non-moral case, which seems to involve what Dancy (2004) calls *enticing* reasons—the non-moral variety of what I'm calling merely commendatory reasons. Suppose it's a nice day out, and that a picnic would be pleasant. This fact favors going on a picnic. But if you decide not to go on a picnic, you don't need to justify yourself in not doing so by citing some reason in favor of spending a lazy day at home, instead. According to this view, you have the latitude to just ignore this reason in favor of going on a picnic, because the reason favoring it is merely commendatory. This is so even though we admit that this reason makes the option better than the alternatives.[18] More generally, if in many cases the reasons supporting our options are largely commendatory ones, then we will often have this kind of latitude.

But it may seem that this kind of account—whatever its merits in the non-moral domain—is not well-suited to morality, since moral reasons cannot just be ignored without criticism. But as Horgan and Timmons (2010) and others point out, this appeal to merely commendatory reasons in the moral domain would give us a neat explanation of supererogation and moral latitude. Consider what Horgan and Timmons think of as a kind of everyday case of supererogation. This is the case mentioned above, in which you have the chance to take your elderly, widowed neighbor to a baseball game, which you know she'd enjoy. Suppose that this would be supererogatory: it would surely be a nice thing to do, but you aren't morally required to do so—you won't be blamed for not doing it, just like we aren't all blameworthy for passing up many opportunities for small kindnesses everyday. If we agree that it would be better to take her than to stay home, then as we saw the Balancing Account struggles to explain this. But note also that we don't typically require justification for not doing these small kindnesses. So the justifying reasons account does not seem quite right, either. If the reasons supporting doing this small kindness are merely

[17] Little and Macnamara (2018) use 'commendatory'. See, e.g., Dancy (2004) on 'enticing' reasons, Little (2013) and Little and Macnamara (forthcoming) on 'non-requiring' or 'non-deontic' reasons, Horgan and Timmons (2010) on 'merit-conferring' reasons, and Scanlon (2014), p. 3 on 'optional' reasons. See also Dreier's (2004) discussion of reasons from the point of view of beneficence vs. reasons from the point of view of justice.

[18] See Little and Macnamara (2018) for a rich discussion of the "comparative lives" of commendatory reasons.

commendatory, then even though they can explain why the option is good, and even better, morally speaking, than the alternative of staying home, they are not the kind of reasons that can ground a moral requirement. But they can explain why we might think especially well of someone who performs these actions—there are good (commendatory) reasons for them, after all. So insofar as we think that there is supererogation, and competing accounts like the simple Balancing Account and the appeal to merely justifying reasons struggle to explain it, this supports the appeal to merely commendatory reasons. Moreover, there is something plausible about the idea that, at least in cases involving these kinds of everyday, small kindnesses, we can ignore the reasons supporting them without facing moral criticism, even if we don't really have anything better to do.

But when we turn to other cases of supererogation—in particular, cases of the more heroic variety—this account looks less plausible. On this account, we can say that the supererogatory action is most strongly favored by the reasons and is what the agent really *should* do. This is an advantage, compared to the merely justifying reasons account. We can also say that, since the reasons supporting the supererogatory action are merely commendatory, you are nevertheless not required to do it. This is an advantage over the simple Balancing Account. But note that the reasons supporting supererogatory actions, at least in these more heroic cases, are things like that the action would save a life, or prevent an immense amount of suffering. These are not merely commendatory reasons—reasons that make an option nice but which can be ignored. As Archer (2017) points out, they are exactly the kinds of reasons that could generate requirements.

Of course, in cases of supererogation, they do not generate requirements. But this is not because the reasons favoring the supererogatory actions are the sorts of reasons that can just be ignored. In many cases, the supererogatory action comes with an extremely high cost to the agent; in such cases, this cost should be part of the explanation for why the supererogatory action is not required. We can see this by considering a case of supererogation—saving someone from a burning building, say—and then tweaking the case so that the cost diminishes greatly. Suppose that you have a fire-retardant suit, and so are sure that you won't be hurt by the fire, for example. In this case, you are required to go and save the person, even though the reasons in favor of going in are the same between the two cases. Alternatively, in some cases the explanation for why the supererogatory action is not required may be better understood in terms of the strength of the reasons favoring the non-supererogatory alternative, rather than in terms of the cost of the supererogatory action. For example, suppose you are deciding between taking the bus four hours to surprise your sister on her birthday, or else calling her and sending a thoughtful gift. Visiting would be supererogatory, and this is arguably because the alternative is good enough. They key point

for now is just that the explanation for why the supererogatory action is not required is not that the reasons supporting it can be ignored with impunity.[19]

If we accept that there are both merely justifying reasons and merely commendatory reasons, then we can combine these two kinds of solution to get a more adequate explanation of latitude, especially the sort involved in supererogation. This is the strategy that Horgan and Timmons (2010) seem to endorse. The reasons favoring taking your neighbor to the baseball game are merely commendatory: they favor the action but you can ignore them without facing any criticism. The reasons favoring running into the burning building to save a life, on the other hand, are requiring, but there are merely justifying reasons that can justify you in not doing so, like the fact that doing so is a serious risk to your life.[20]

In the end we may need to accept that there are several different kinds of reasons which compete in importantly different ways. But I am interested in exploring an alternative, simpler explanation of latitude. In particular, I am interested in seeing how far we can get if we just rely on the sort of competition that is captured in the Balancing Account, which can be understood, at least metaphorically, in terms of straightforward comparisons of weight.[21] I will sketch an account that does this in the next section.[22]

6 Reasons Against

The discussion of the previous accounts of how reasons compete, with a focus on latitude and in particular on supererogation, has brought out two important points. First, the reasons favoring a supererogatory option are (at least typically) reasons that could ground a requirement, were the costs or alternatives different. This is where the merely commendatory reasons view faltered. Second, giving a satisfying explanation of latitude and supererogation requires being able to draw comparisons amongst the alternatives in terms of how strongly the reasons support them. This is where the purely justifying reasons view faltered.

[19] On this point see Archer (2016, 2017).

[20] Horgan and Timmons think of what I was calling heroic cases as not genuinely cases of supererogation (they use "quasi-supererogation"), so do not think that we need to give a univocal explanation. But I think that a univocal explanation would be preferable.

[21] Of course, we will also need to account for other ways reasons can interact, including attenuating, intensifying, and excluding; see Dancy (2004), Chapter 3, Bader (2016), Raz (1999). Attenuating and intensifying are typically thought of in terms of the weights of the attenuated or intensified reasons (and in particular modification of those weights); exclusion, if it is a distinct way that reasons interact, can be thought of just as having reason to take the excluded reasons out of consideration.

[22] For further discussion of this account, see Snedegar (forthcoming).

The account I sketch here, like the others, appeals to a distinction between different kinds of reasons. This is the distinction between reasons for and reasons against.[23] Everyone should recognize that there are reasons against our options, in addition to the reasons for them. So this account has at least the *prima facie* advantage of relying on a distinction between kinds of reasons that we should all already accept. But it is still controversial because many people have—usually implicitly—assumed that reasons against an option are really nothing more than reasons for alternatives to that option, so that we really just have reasons *for* competing options.[24] This is implicitly assumed by the simple Balancing Account, since according to that view we just have to figure out which option has the weightiest reasons in its favor—just like we determine which set of marbles is heaviest by seeing which has the greatest combined weight.

This assumption is mistaken. If the reasons against an option are really just reasons for alternatives to that option, then we need to say precisely *which* alternatives. I argue in Snedegar (2018) that there is no good, general answer to this question. For just one example, a reason against *A* cannot be identified with a reason for *each* alternative to *A*, since some alternatives might be even worse in the relevant respect: the fact that restaurant A is crowded is a reason against going there, but it isn't a reason *for* going to restaurant B instead, if restaurant B is even more crowded (and there are non-crowded alternatives). Instead, we do better to recognize a genuinely distinct category of reasons against an option that need not be identified with reasons for alternatives to that option. If that's right, then we must explain how reasons against an option—as distinct from reasons for alternatives to that option—factor into the competition between reasons. While this may initially seem like an extra complication to be avoided if possible, the account I sketch here aims to show that in fact this lets us give a more nuanced account of how reasons compete to determine overall normative status.

The easiest way to understand the account is by thinking about the competition between reasons as involving two stages.[25] First, there is the competition, for each option, between the reasons for it and the reasons against it. Second, we use the results of these individual competitions for each option to draw a comparison between the options.[26] This gives us accounts of how reasons determine what we (i) are permitted to do, (ii) are

[23] Euan Metz (forthcoming) has independently proposed a similar view. I do not have the space here to compare my view to Metz's.

[24] For more explicit endorsements of this idea, see Nagel (1970), p. 47, and Schroeder (2007), Chapter 7, and Schroeder (2015).

[25] I do not claim that practical reasoning actually has this kind of two-stage structure.

[26] This isn't a particularly novel way of thinking about how reasons compete; see, e.g., Berker (2007) and Sher (2019). What is novel here is making use of these two different aspects of the competition between reasons in giving a more nuanced picture of how the competition between reasons determines overall status.

required to do, and (iii) should do. First, whether an option is permitted or required is determined by the competition between the reasons for it and the reasons against it:

> *Permission*: You are permitted to ϕ if and only if the reasons for ϕ are at least as weighty as the reasons against ϕ.
>
> *Required*: You are required to ϕ if and only if ϕ is the only permissible option.[27]

The reasons for an option will either be weightier than the reasons against it by a certain amount, less weighty by a certain amount, or there will be a tie.[28] This difference in weight between the reasons for and against ϕ (which may be positive, negative, or zero) is the *overall* weight of reasons bearing on ϕ. Comparing the weights of reasons will be complicated in various ways. But as long as we can compare the weights of the reasons for an option to the reasons against it, such that for some options the reasons for will win or lose to a greater extent than for others, the account here can get up and running.

The overall weight of reasons bearing on each of the options can then be compared. As I said above, these comparative evaluations amongst the options are important in themselves, and also what we need to determine which option is *most* strongly supported, and so what you should do:

> *Should*: You should ϕ if and only if the overall weight of reasons bearing on ϕ is greater than the overall weight of reasons bearing on any of the alternatives.

So by recognizing reasons against in addition to reasons for and giving them a distinct role in the competition between reasons, we can give accounts of how this competition determines what you are permitted to, required to, and should do.

This account lets us explain why agents have latitude such that multiple options will often be permissible. An option is only impermissible when the reasons against that option outweigh the reasons for it. But since we are distinguishing between reasons against an option and reasons for the alternatives, strong reasons for an alternative to ϕ are not

---

[27] An important choice point in developing this theory involves the possibility of moral dilemmas. So far I have said nothing that would rule out a case in which every option is impermissible. We could rule this out in various ways. For example, we could insist that the fact that some option is the 'least bad' is always a sufficiently strong reason in favor of doing it, or alternatively just accept a principle that in such cases you are required to do the least bad alternative. In the latter case, there is then the further question of whether the option is still impermissible, in addition to being required. These are important issues for developing the account, but go beyond the goal here of illustrating how accepting the distinction between reasons for and reasons against help us explain latitude.

[28] I am ignoring incommensurability here to keep things simple.

necessarily strong reasons against ϕ. So it is possible for there to be multiple permissible options, rather than a single required option. It does not follow, at least from the structure of the account, that the option that's most strongly supported by reasons is the one that you are required to take, unlike on the Balancing Account. One option can be more strongly supported, overall, than another while both are permissible.

This account has certain similarities to the other views considered here. For example, like commendatory reasons, the reasons *for* options, on this view, cannot themselves ground requirements; in a sense all reasons for options are merely commendatory (cf. Greenspan (2005)). But the explanation for why supererogatory actions are not required is not that there is a distinctive kind of reason supporting them that is incapable of grounding requirements. It is rather that there are sufficiently good reasons in support of the non-supererogatory, merely obligation-fulfilling options—they are good enough. On some views, including the Balancing Account, sufficiency is understood to mean that the reasons for one option are at least as good as the reasons for the alternatives.[29] On this view, on the other hand, the reasons in favor of the non-supererogatory options are sufficiently good in the sense that they can stand up to the reasons against those same options. For example, that staying on the sidewalk rather than running into the burning building keeps you safe is a reason to do so that's good enough to make the option permissible, and so keep the supererogatory action of running into the building from being required. So even if the reasons supporting the supererogatory action—that you can save a three lives, say—outweigh the reasons for staying on the sidewalk, there can be sufficiently good reason to stay on the sidewalk, compared to the reasons against, and doing so may still be permissible.

That staying safely on the sidewalk means giving up a chance at saving three lives seems to be a strong reason *against* doing so. So if there were not strong reasons in favor of doing so, this option would be impermissible. In a case in which you have a fireproof suit, for example, staying on the sidewalk is impermissible. This is because in this case there are not sufficiently good reasons for this option. The reason to stay on the sidewalk in the case in which you do not have a suit is that doing so will keep you safe. But if you would *also* be safe if you ran into the building, as in the case in which you have the suit, then this is not a strong reason in favor of staying on the sidewalk.

The assumption here is that features of the alternatives to a given option can affect what reasons there are for that option, or at least how strong they are. In the case in which you have a fire-retardant suit, you will be safe whether you run in or not, and so considerations of personal safety are not (very) significant in determining the overall status

[29] See Schroeder (2015) and Lord (2018), p. 10; cf. also Portmore (2011), Chapter 5.

of the two options. The claim that considerations that do not distinguish between the two options are not (significant) reasons for either is one that I have defended elsewhere.[30] To see, briefly, why it is plausible, note that you could not justify your decision to stay on the sidewalk by pointing out that doing so keeps you safe, if you would also be safe if you ran into the building to help, since you have the suit. This is some evidence that this fact is not a reason to stay on the sidewalk, in this case, since one key thing reasons do is provide justification for the things they are reasons for.

Here is an important objection to this account. To hold that running into the burning building is permissible, it must be that the reason in favor of doing so—that you can save three lives—is at least as weighty as the reason against doing so—that it's dangerous. To say that you *should* (but aren't required to) run into the building, we must in fact hold that the reasons in favor of doing so are *weightier* than the reasons against doing so—perhaps significantly so. But to hold that it's not required, the alternative of staying on the sidewalk must be permissible, which means that the reason in favor of staying on the sidewalk—that it will keep you safe—must be at least as weighty as the reason against—that you can save three lives by running in, instead. This looks inconsistent: it looks like we have to hold both (i) the life-saving reason outweighs the personal safety reason, and (ii) the personal safety reason is at least as weighty as the life-saving reason.

First, this objection brings out that one and the same consideration can and often will be both a reason against one option and a reason for alternatives. For example, the reasons for one option, constituted by good features of that option, may often be reasons against alternatives, since performing the alternatives may mean missing out on these good features; this is one way to think about *opportunity cost*.[31] This is compatible with denying that a reason against one option is *nothing more than* a reason for alternatives. The crucial point is that once we recognize a distinction between reasons for and reasons against, we can deny that the weight of a consideration as a reason *for* one option will be the same as the weight of that consideration as a reason *against* a different option. If the reasons against one option were really nothing over and above reasons for alternatives, then the reasons against would have the same weight as the reasons for with which they are identified. But since the starting point for the account here is a distinction between reasons for and reasons against, we need not accept this. So considerations that strongly support one option may not count against competing alternatives with corresponding strength, and *vice versa*. This is the key feature of this account, which lets us give the explanation of how multiple options could be permissible, even when some of them are better supported than others.

[30] See especially Snedegar (2019).
[31] See Snedegar (2017), Chapter 4 and (2018).

Giving a full account of which reasons can outweigh which others is going to involve lots of first-order theorizing. The account here guides us in how to interpret judgments about these questions, and questions about which options are permissible, required, and so on. For example, the judgments in these cases suggest that the life-saving reason *against* staying safely on the sidewalk is less weighty than the personal safety reason in favor of doing so. But the reasons against staying on the sidewalk are weightier than a monetary reason in favor of doing so in a case in which you have a fireproof suit, but someone offers you $100 not to run in. The important thing for this paper is that both of these claims—in particular, the first one—are compatible with saying that running in is more strongly supported, overall, than staying out. Once we distinguish reasons against one option from reasons for alternatives, the weight of a consideration as a reason against an option need not be the same as the weight of the same consideration as a reason for an alternative.[32]

The distinctions between requiring reasons, on the one hand, and either merely justifying reasons or merely commendatory reasons, on the other, were offered, at least in part, in order to explain latitude. What I have argued here is that we can explain latitude by instead relying on a different distinction between kinds of reasons, and one which stays closer to the simple and intuitive Balancing Account. By recognizing a distinct category of reasons against, we can give a promising account of how reasons compete to determine the overall status of our options. There are many open questions about this account, but I think that it is worth pursuing. Though the competition between reasons is understood in terms of weight, like on the simple Balancing Account, recognizing a distinct category of reasons against gives us a rich enough structure to allow for latitude without making contentious claims about widespread ties or incommensurability.

References

Archer, A. (2017). Supererogation. *Philosophy Compass*, 13(3).

--- (2016). Supererogation, sacrifice, and the limits of duty. *Southern Journal of Philosophy*, 54(3): 333-354.

[32] That one consideration could bear differently on different alternatives is not controversial. That the chancellor of my university is going to be at my talk is both a reason for me to put on a tie and a reason for me to write an email thanking her for coming. But there's no reason to think that these reasons will have the same weight. There are of course different issues involved in the cases I am discussing, but again, once we distinguish reasons for and reasons against, treating them as having distinct weights does not seem ad hoc.

Bader, R. (2016). Conditions, modifiers, and holism. In Lord and Maguire, eds., *Weighing Reasons*. Oxford: Oxford University Press, pp. 27-55.

Baron, M. (2015). The supererogatory and Kant's imperfect duties. In *Reason, Value, and Respect: Kantian Themes from the Philosophy of Thomas E. Hill, Jr.*, eds. Timmons and Johnson. Oxford University Press, pp. 215-231.

--- (1987). Kantian ethics and supererogation. *Journal of Philosophy*, 84(5): 237-262.

Bedke, M. (2011). Passing the deontic buck. In Shafer-Landau, ed., *Oxford Studies in Metaethics*, Vol. 6. Oxford: Oxford University Press, pp. 128-152.

Berker, S. (2007). Particular reasons. *Ethics*, 118(1): 109-139.

Byron, M, ed. (2004). *Satisficing and Maximizing: Moral Theorists on Practical Reason*. New York: Cambridge University Press.

Chang, R. (2001). *Making Comparisons Count*. London: Routledge.

Dancy, J. (2004). Enticing reasons. In Wallace, Pettit, Scheffler, and Smith, eds., *Reasons and Value: Themes from the Moral Philosophy of Joseph Raz*. Oxford: Oxford University Press, pp. 91-118.

Dreier, J. (2004). Why ethical satisficing makes sense and rational satisficing doesn't. In Byron, ed., *Satisficing and Maximizing: Moral Theorists on Practical Reason*. Cambridge: Cambridge University Press, pp. 131-154.

Finlay, S. (2009). Oughts and ends. *Philosophical Studies* 143(3): 315-340.

von Fintel, K. and Iatridou, S. (2008). How to say *ought* in foreign: the composition of weak necessity modals. In Guéron and Lecarme, eds., *Time and Modality*. Springer, pp. 115-141.

Gert, J. (2012). *Normative Bedrock: Response-Dependence, Rationality, and Reasons*. Oxford: Oxford University Press.

--- (2007). Normative strength and the balance of reasons. *Philosophical Review*, 116(4): 533-562.

--- (2004). *Brute Rationality*. Cambridge: Cambridge University Press.

Greenspan, P. (2010). Making room for options: moral reasons, imperfect duties, and choice. *Social Philosophy and Policy*, 27(2): 181-205.

--- (2007). Practical reasons and moral *ought*. In Shafer-Landau, ed., *Oxford Studies in Metaethics*, Vol. 2. Oxford: Oxford University Press, pp. 172-194.

--- (2005). Asymmetrical practical reasons. In Reicher and Marek, eds., *Experience and Analysis: Proceedings of the 27th International Wittgenstein Symposium*. Oebv&Hpt, pp. 387-394.

Grotius, H. (1625/2012). *On the Law of War and Peace*. Neff, ed., Cambridge: Cambridge University Press.

Hobbes, T. (1651/1996). *Leviathan*. Tuck, ed., Cambridge: Cambridge University Press.

Horgan, T. and Timmons, M. (2010). Untying a knot from the inside out: reflections on the "paradox" of supererogation. *Social Philosophy and Policy*, 27(2): 29-63.

Kagan, S. (1989). *The Limits of Morality*. Oxford: Oxford University Press.

Little, M. (2013). In defense of non-deontic reasons. In Bakhurst, Hooker, and Little, eds., *Thinking about Reasons: Themes from the Philosophy of Jonathan Dancy*. Oxford: Oxford University Press, pp. 112-136.

Little, M. and Macnamara, C. (forthcoming). Non-requiring reasons. In Chang and Sylvan, eds., *The Routledge Handbook of Practical Reasons*. New York: Routledge.

--- (2018). For better or worse: commendatory reasons and latitude. In Timmons, ed., *Oxford Studies in Normative Ethics*, Vol. 7. New York: Oxford University Press, pp. 138-160.

Lord, E. (2018). *The Importance of Being Rational*. Oxford: Oxford University Press.

Lord, E. and Maguire, B. (2016). An opinionated guide to the weight of reasons. In Lord and Maguire, eds., *Weighing Reasons*. Oxford: Oxford University Press, pp. 3-24.

McElwee, B. (2017). Supererogation across normative domains. *Australasian Journal of Philosophy*, 95(3): 505-516.

McNamara, P. (1996). Must I do what I ought (or will the least I can do do)? In Brown and Carmo, eds., *Deontic Logic, Agency, and Normative Systems*. Berlin: Springer-Verlag, pp. 154-173.

Metz, E. (forthcoming). How to be a deontic buck-passer. Forthcoming in *Philosophical Studies*.

Nagel, T. (1970). *The Possibility of Altruism*. Princeton: Princeton University Press.

Parfit, D. (2011). *On What Matters*. Oxford: Oxford University Press.

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

--- (2011). *Commonsense Consequentialism*. Oxford: Oxford University Press.

Pummer, T. (2016). Whether and where to give. *Philosophy and Public Affairs*, 44(1): 77-95.

Raz, J. (1999). *Practical Reason and Norms*. Oxford: Oxford University Press.

Ross, W.D. (1930). *The Right and the Good*. Oxford: Oxford University Press.

Scanlon, T.M. (2014). *Being Realistic About Reasons*. Oxford: Oxford University Press.

Schroeder, M. (2015). What makes reasons sufficient? *American Philosophical Quarterly*, 52(2): 159-170.

Sher, I. (2019). Comparative value and the weight of reasons. *Economics and Philosophy*, 35(1): 103-158.

Sloman, A. (1970). 'Ought' and 'better'. *Mind*, 79(315): 385-394.

Slote, M. (1989). *Beyond Optimizing: A Study of Rational Choice*. Cambridge: Harvard University Press.

Snedegar, J. (forthcoming). Competing reasons. In Brown and Simion, eds., *Reasons, Justification, and Defeat*. Oxford: Oxford University Press.

--- (2019). Deliberation, reasons, and alternatives. *Pacific Philosophical Quarterly*, 100(3): 682-702.

--- (2018). Reasons for and reasons against. *Philosophical Studies*, 175(3): 725-743.

--- (2017). *Contrastive Reasons*. Oxford: Oxford University Press.

--- (2016). Reasons, oughts, and requirements. In Shafer-Landau, ed., *Oxford Studies in Metaethics*, Vol. 11. Oxford: Oxford University Press, pp. 155-181.

Tenenbaum, S. (2007). Brute Requirements: A critical notice of Joshua Gert's *Brute Rationality*. *Canadian Journal of Philosophy*, 37: 153-172.