

The All or Nothing Ranking Reversal and the Unity of Morality

Chris Tucker
William & Mary

Abstract: Supererogatory acts are, in some sense, morally better than their non-supererogatory alternatives. In this sense, what is it for one option A to be better than an alternative B? I argue for three main conclusions. First, relative rankings are a type of all-in action guidance. If A is better than B, then morality recommends that you A rather than B. Such all-in guidance is useful when acts have the same deontic status. Second, I argue that Right > Wrong: permissible acts are always better than their impermissible alternatives. If Right > Wrong were false, then morality's deontic verdicts would sometimes conflict with its relative rankings. Such conflict would undermine the thought that morality is a coherent, authoritative guide to action. Third, the All or Nothing Problem is not a counterexample to Right > Wrong as is commonly thought. Instead, it involves an interesting ranking reversal: whether A > B can depend on whether a certain alternative is added as a third option.

Keywords: supererogation; morally better; All or Nothing Problem; Right > Wrong; conditional deontic status; reasons; commending weight; choiceworthiness; ought

This paper concerns the All or Nothing Problem, which I take to be an interesting ranking reversal. There are two children in the burning building, and you can easily carry both at the same time. This gives us three potential options. Remain a bystander and avoid getting burned (**Bystander**), save one child and get severe burns (**Save1**), and save both children and get the same severe burns (**Save2**). The relative ranking of Save1 and Bystander depends on whether Save2 is added as a third option.

If Bystander and Save1 are the only two options, then Save1 is supererogatory and Bystander is permissible but not supererogatory. So, in the two-option case, *Save1 > Bystander*. The **All or Nothing Problem** gets its name because of what happens when Save2 is a third option: you have to save *all* the children (Save2) or *none* of the children (Bystander). More specifically, Save2 is supererogatory, Bystander is permissible but not supererogatory, and Save1 is *impermissible*. In the three-option case, then, *Bystander > Save1*. The addition of Save2 somehow causes a ranking reversal.

This apparent ranking reversal isn't due to some rogue intuitions about the All or Nothing case. These intuitions are supported by conceptual truths. In the two-option case, my intuition that *Save1 > Bystander* is supported by the conceptual truth **Supererogation is Better**, that supererogatory acts are morally better than their non-supererogatory alternatives. In the three-option case, my intuition that *Bystander > Save1* is supported by the conceptual truth **Right > Wrong**, that permissible options are morally better than their impermissible alternatives.

In treating the All or Nothing Problem as a ranking reversal, I depart from any existing treatment of the case. Horton endorses Right > Wrong (2017: 96). Yet he denies that Bystander is permissible if you are "willing" to enter the building (96-7). I disagree with him on this point, but his own account is bolstered by the arguments I provide for Right > Wrong.

To the extent that there is a standard story, it is that the All or Nothing Problem is a counterexample to Right > Wrong. The story begins with **Pairwise Precedence**: if *Save1 > Bystander* when there are two options, it remains better when you add Save2

as a third option. Given Pairwise Precedence, we can infer that Save1 > Bystander in the three-option case. Since Save2 makes only Save1 impermissible, we get a counterexample to Right > Wrong: an impermissible option (Save1) is morally better than a permissible alternative (Bystander).¹

I deny Pairwise Precedence and thus hold that the All or Nothing Problem is a ranking reversal. The standard story denies Right > Wrong and thus denies that the case is a ranking reversal. To resolve this debate, we need to go back to the basics.

Q1: What is the relation between an action's deontic status (im/permissible)² and its relative ranking?

Q2: How do reasons interact to determine an action's deontic status?

Q3: How do they interact to determine its relative ranking?

The relation between deontic status and relative ranking constrains how reasons determine those statuses. Thus, this paper provides an integrated answer to these three questions.

Let's start with Q1. Morality's deontic verdicts provide all-in action guidance. A verdict of impermissible is morality's way of saying *don't do it*. Less obviously, morality's relative rankings also provide all-in action guidance. If morality ranks A over B, then it recommends that you choose A over B. Bystander and Save2 are both permissible in the three-option case. Yet morality ranks Save2 as your best option, i.e., it recommends Save2 over all other options. We capture this recommendation by saying that you (morally) *ought* to choose Save2 even though you aren't required to do so. But now we face a question at the heart of this paper. Can morality's all-in deontic guidance conflict with morality's all-in relative ranking guidance?

Right > Wrong amounts to the claim that morality is single-minded, that its deontic verdicts and relative rankings speak in one voice. I give two arguments for this claim. The more important one highlights the tension involved in saying "Don't do what it is better to do". In this utterance, the deontic directive and relative ranking conflict. It's like saying don't do it, but I recommend you do it anyway. Morality can't say such things if it is to be the coherent, authoritative guide to action that it is widely taken to be. If morality doesn't say such things, then Right > Wrong.

In reply to Q2, the *deontic status* of ϕ is determined by *two* features of reasons: the **justifying weight for ϕ** (weight that pushes ϕ toward permissibility) and **requiring weight for $\sim\phi$** (weight that pushes ϕ toward impermissibility). In reply to Q3, the *relative ranking* of ϕ is determined by *three* features of reasons: justifying weight for ϕ , requiring weight for $\sim\phi$, and **commending weight** (that which explains why some actions with the same deontic status are better than others). Note the difference in functional roles. Justifying and requiring weight are defined by their functional roles in *fixing deontic status*. Commending weight is defined by its functional role in making one action better than another *when the actions have the same deontic status*.

These integrated answers to Q1-Q3 entail that the All or Nothing case is a ranking reversal. Save1 is morally better than Bystander when they are the only two options. We explain this by pointing out that Save1 is always more commended than Bystander.

¹ See, e.g., Pummer (2019), Ferguson & Köhler (2020), and Muñoz (2021: 703); cf. Rulli 2020: 379. Portmore's story is also standard insofar as he assumes Pairwise Precedence and that Save1 is better than Bystander in the three-option case (2019: 207-210).

² I don't want to pick fights about the meaning of *deontic status*, but you need to know what I mean by that term, which is just permissible, impermissible, and combinations thereof. Requirement to ϕ is a deontic status because it is just the combined verdict that ϕ is permissible and $\sim\phi$ is impermissible. In my terminology, supererogation is a hybrid of deontic and comparative evaluative status (cf. Horgan and Timmons 2010: 31-2).

More generally, altruistic benefits always commend more than self-interested benefits of the same size. This general truth about commending explains why paradigmatic supererogatory action tends to be *altruistic* rather than self-interested.

When we add Save2 as an option, Save1 is impermissible but more commended than Bystander. It does *not* follow that Save1 is morally better than Bystander. The functional role of commending kicks in—it is an active ingredient in the relative ranking—only when Save1 and Bystander have the same deontic status. But they don't have the same deontic status in the three-option case. The interaction of these three features of reasons gives us **Deontic Precedence**. When an impermissible option is more commended than a permissible one, the difference in deontic status takes precedence over the difference in commending. Thus, Bystander is better than Save1 in the three-option case.

In §1, I clarify the relevant sense of *morally better than*, as well as identify two conceptual truths about it: Supererogation is Better, and Right > Wrong. In §2, I clarify the concept of moral commending and distinguish it from justifying and requiring weight. In §3, I argue that the All or Nothing Problem is a counterexample to Pairwise Precedence, not Right > Wrong. In §4, I use Deontic Precedence to explain the All or Nothing ranking reversal. In §5, I give a second argument for Right > Wrong, the Normative Unity Argument. In §6, we consider some strong intuitions that seem to conflict with Right > Wrong. Yet, given Deontic Precedence, there is no conflict after all.

1. Morally Better

1.1. Supererogation is Better

In addition to issuing deontic verdicts, morality issues comparative evaluative verdicts: some acts are morally better than their alternatives. I take the following to be a definitional constraint on the relevant sense of morally better:

Supererogation is Better: A supererogatory act is morally better than its non-supererogatory, permissible alternatives. [definition of supererogation]

If supererogation exists, which I assume that it does, then some morally permissible act is morally better than some permissible alternative.³ Liv's jumping on the grenade to save 5 others is supererogatory because it is morally better than a permissible alternative, remaining in safety.

Supererogation is Better is a conceptual constraint only on how good a supererogatory act is relative to *its* non-supererogatory alternatives. It does not follow that any supererogatory act is morally better than any non-supererogatory act, no matter how different the alternatives and relevant reasons are.

1.2. Morally Better as Morally More Choiceworthy

I take the relevant sense of morally better to involve an act's being (*morally*) *more choiceworthy* than another. A verdict of 'more choiceworthy' is a type of all-in action guidance. It is morality's way of "recommending" that we choose an option over some alternative (cf. Little & Macnamara 2020: §4). Relative rankings guide us when alternatives have the same deontic status. It is supererogatory, so morally better, for

³ I focus on moral supererogation, but the structure of supererogation is the same for any normative perspective that allows for it. To make sense of *moral* supererogation, we focus on *moral* justifying weight, *moral* requiring weight, and *moral* commending (§2-3). To make sense of *rational* supererogation, we would focus instead on *rational* justifying weight, *rational* requiring weight, and *rational* commending. Kawall's (2003) example of the self-regarding prisoner provides a potential case of rational supererogation.

Liv to jump on the grenade to save five others rather than remain in safety. Morality is, therefore, recommending that Liv jump on the grenade rather than remain in safety.

Sometimes ethicists claim that you ought to do “what you have most reason to do” and, more generally, that A is better than B just when there is *more reason* for A than B.⁴ This way of talking is unclear at best if there is more than one weight of reasons (justifying and requiring weight). Thinking of morality’s relative ranking in terms of *more reason* and thinking of it in terms of *more choiceworthy* both acknowledge three things: that morality’s ranking of options is normative, providing all-in action-guidance; that relative rankings can guide us when alternatives have the same deontic status; and that morality’s relative rankings are determined solely by the interaction of reasons.⁵ The only difference is that *more choiceworthy* is neutral. It doesn’t smuggle in controversial assumptions about how many weights of reasons there are (just *the* reason’s weight) or how many weights determine morality’s relative rankings (just *requiring* weight). We can say, then, that an agent (**morally**) **ought to ϕ** if and only if ϕ is the agent’s most (morally) choiceworthy option.⁶ Morality ranks A as better than B if and only if A is more choiceworthy than B.

1.3. *Right > Wrong*

Our second conceptual truth is:

Right > Wrong: Permissible acts are more choiceworthy than their impermissible alternatives.

It is permissible to take a nap after an exhausting day. It is *impermissible* to murder the irritating boss who made your day exhausting. Right > Wrong entails that the former is better, or more choiceworthy, than the latter. Right > Wrong is a conceptual constraint only on how good a permissible act is relative to *its* impermissible alternatives. It does not follow that any permissible act is more choiceworthy than any impermissible act, no matter how different the alternatives and relevant reasons are.

The *Deontic Evaluation Argument* is my first argument that Right > Wrong is a conceptual truth. It holds that ‘permissible’ and ‘impermissible’ aren’t merely action guidance; they are also positive and negative evaluations of choiceworthiness, respectively. The argument has three premises:

1. It is a conceptual truth that, if ϕ is permissible, then it is okay (alternatively: ϕ is choiceworthy or neutral).⁷
2. It is a conceptual truth that, if ϕ is *impermissible* to ϕ , then ϕ is not okay (alternatively: ϕ is *unchoiceworthy*, it is morally bad).
3. It is a conceptual truth that any choiceworthy or neutral option is more choiceworthy than any unchoiceworthy alternative.

The conjunction of 1-3 entails that it is a conceptual truth that Right > Wrong, that permissible options are more choiceworthy than their impermissible alternatives.

The third premise seems straightforward. An action’s being *neutral* with respect to choiceworthiness is neither the absence of moral evaluation nor a verdict of undefined.

⁴ See, e.g., Muñoz (2021: 703, 712, nt. 1). Portmore (2019: 185-9) also endorses the claim when it is restricted to some privileged way of individuating options.

⁵ Consequently, an act can be choiceworthy without being praiseworthy. Choiceworthiness is determined *solely* by the interaction of reasons. Praiseworthiness is determined by the interaction of reasons, the agent’s motivations, and maybe more besides (Massoud 2016).

⁶ Snedegar (2016: §7.3) and Portmore (2019: 14-16) would also require that ϕ be permissible. Feel free to add that condition, if you like.

⁷ Some folks think that permissibly pointless actions, like counting blades of grass, aren’t worth choosing. That’s compatible with 1, as long as counting blades of grass is neutral with respect to choiceworthiness.

It is a substantive moral evaluation. Acts that are okay (at least neutral) are always better than those that are bad with respect to choiceworthiness.

The first two premises are more controversial, but they still seem plausible. Kid wants to do something fun.

Kid: Can I go to the pool?

Dad: That's permissible.

Mom: Going to the pool is *not* okay.

What Dad giveth, Mom taketh away. Dad's deontic directive conflicts with Mom's evaluation. The first premise explains: Dad's directive *that's permissible* entails that going to the pool is okay. 'Permissible' is not only a directive but also a positive evaluation. That's why it conflicts with Mom's negative evaluation (it's not okay).

This deontic evaluation dialogue is a bit unnatural since "that's permissible" sounds so pedantic. Yet the attempt to make it less pedantic is just grist for my mill. What is the most natural, non-pedantic substitute for 'that's permissible'? I bet you think it is 'that's okay', but 'that's okay' is an evaluation! For example, you evaluate the broccoli and rice when you say they were okay. Why is it that the most natural, non-pedantic substitute for 'permissible' is an evaluation? Because 'permissible' is itself an evaluation.⁸

If 'permissible' is a positive evaluation of that's okay, then presumably 'impermissible' is a negative evaluation of that's not okay or bad. The second deontic evaluation dialogue confirms this presumption. Yet again Kid's one-track mind is stuck on fun.

Kid: Can I go to the movies?

Mom: That's okay.

Dad: I prohibit it.

What Mom giveth, Dad taketh away. Why does Dad's deontic directive conflict with Mom's evaluation? The second premise explains: Dad's directive *I prohibit it* (i.e., going to the movies is impermissible) implies that it's *not* okay to go to the movies. 'Going is impermissible' is not only a directive but also a negative evaluation. That why Dad's deontic directive conflicts with Mom's positive evaluation (it's okay).

Dad comes across as stiffly formal in the second dialogue. Yet again, though, making him less formal is just grist for my mill. The everyday way to communicate that something is impermissible is to say that *it's not okay*, which is used to communicate a negative evaluation. (Semantically, *it's not okay* is merely the denial of a positive evaluation; pragmatically, denials of positive evaluations generally communicate negative evaluations. If someone says, "the movie is not good", they generally mean that it is bad.)

If you still reject 1 and 2, don't worry about it. I give a second argument for Right > Wrong in §5. Even if both arguments fail, I will still show that Right > Wrong is compatible with the intuitions that allegedly threaten it (§6).

My overall argument in this chapter requires only that Supererogation is Better and Right > Wrong are true, not that they are conceptually true. If you want to drop the talk of conceptual truths, feel free. I went with the bolder claim that they are conceptual truths, because my arguments for Right > Wrong arguably do support that bold thesis. The Deontic Evaluation Argument appeals to the idea that 'permissible' functions as a positive evaluation, and the forthcoming Normative Unity Argument appeals to the idea that 'morally better' functions as a recommendation. Truths about how concepts

⁸ I don't assume that 'okay' means exactly the same thing when applied to both broccoli and actions. The two uses may be polysemous and so have different but related meanings. If so, what unifies those two uses of 'okay' under polysemy is that they are both positive evaluations of neutral or good.

function in a conceptual scheme (or how terms function in a language) often are conceptual truths.

2. Moral Commending Weight

In my view, relative choiceworthiness is determined by the interaction of *three* weights of reasons: justifying weight, requiring weight, and commending weight. I quickly review the concepts of justifying and requiring weight, but the section's main goal is to clarify commending weight.

2.1. The Basic Concept

Justifying and requiring weight are defined by their roles in fixing whether ϕ is permissible or impermissible. Justifying weight for ϕ is weight that pushes ϕ toward permissibility. Requiring weight for $\sim\phi$ is weight that pushes ϕ toward impermissibility. It is this competition that determines whether ϕ is permissible:

Pluralist Permissibility: ϕ is permissible if and only if the justifying weight for ϕ is at least as great as the requiring weight for each alternative; otherwise, ϕ is impermissible.⁹

This assignment for permission doubles as an assignment for requirement: ϕ is required just when you apply Pluralist Permissibility to every alternative and ϕ is the only permissible option. I use the label “*Pluralist Permissibility*” because the assignment makes permissibility a function of more than one weight value (namely, justifying and requiring weight).

Justifying and requiring weight fix relative choiceworthiness only to the extent that an act's deontic status determines its relative choiceworthiness. Justifying and requiring weight do not (as a conceptual truth) settle which acts are better than others when those acts have the same deontic status.

Given that supererogation exists, some permissible acts are better than their permissible alternatives. We assumed that relative choiceworthiness in such cases is explained by the interaction of reasons (§1.2). Let us say that **moral commending weight**¹⁰ is that gradable feature of reasons that makes an act morally better than at least some alternatives with the same deontic status (im/permissible). For simplicity, I assume that all justifying, requiring, and commending weight values are commensurable. I think that the incommensurability of weight and choiceworthiness is most plausible given what I call ‘Deontic Precedence’; however, that additional wrinkle requires more space than I can give it here.

Our intuitions about supererogation suggest that, other things equal, altruistic action is morally better than self-interested action. Even if Liv had to choose between her life and saving only one fellow soldier, it would be morally permissible to do either but morally better to jump on the grenade. Commending weight is that feature of reasons that explains why some permissible acts are morally better than some permissible alternative. Hence, altruistic action is generally better than self-interested action because:

⁹ This assignment for permissibility is closest to Tucker (2023: §6; cf. 2022a: 381), but you can find similar ideas in Portmore (2011: ch 5), Massoud (2016), Archer (2016a), and Muñoz (2021).

¹⁰ Another term for commending is ‘favoring’ (e.g., Dancy 2004: 29, 33; Portmore 2011: 92, nt 26; and Muñoz 2021: 713, nt. 5). I prefer ‘commending’, because one might say that there is justifying or requiring weight *in favor of* an option without meaning to talk about the sense of commending/favoring at issue. Other accounts of commending weight, or something like it, include Horgan and Timmons (2010: §VI.B) and Little & Macnamara (2017, 2020).

Altruism Commends More: altruistic benefits always commend more than self-interested benefits of the same size.

Of course, altruism can fail to be better than self-interested action when other things aren't equal. One way things may not be equal is for the self-interested benefits to be much greater than the altruistic benefits. Suppose that Bernie can remain in safety or he can jump on a grenade to spare someone else a mild burn. Arguably, it is morally better for him to remain in safety than to sacrifice his life. I would add that his life morally commends remaining in safety more than preventing the mild burn commends jumping on the grenade.¹¹ The reason why his self-interested reasons commend more than the altruistic ones is that the self-interested benefits are *much greater* than the altruistic ones.

I will argue that the All or Nothing Problem reveals a second way things may not be equal, namely there is a difference in deontic status. Suppose A is more commended than B. The greater commending is a conditional guarantee: on the condition that A and B have the same deontic status, then A is morally better than B. If A is impermissible and B is permissible, then all bets are off.¹²

2.2. *Commending vs Justifying and Requiring Weight*

As I've defined commending for ϕ , it is *conceptually distinct* from justifying weight for ϕ , requiring weight for ϕ , and the sum of justifying and requiring weight for ϕ .¹³ Yet conceptually distinct concepts can still refer to the same thing. *Biden* is conceptually distinct from *the President of the US*, but Biden is (identical to) the President of the US. I take no stand on whether commending weight is identical to requiring weight or the sum of justifying and requiring weight. I do assume, however, that justifying and commending weight are *independent variables*. My main opponents generally grant this assumption.¹⁴ We generally agree that the reason with the most justifying weight need not be the reason with the most commending weight. My solution to the All or Nothing Problem capitalizes on this independence. Altruism Commends More tells us that Save1 is more commended than Bystander. Yet the standard account of supererogation tells us that there is more justifying weight in favor of Bystander than Save1 (§4.2). We'll see that this difference triggers the ranking reversal.

¹¹ See my 2022b for a discussion of cases like Bernie's, cases in which one could go *too far* beyond the call of duty.

¹² Reasons might *discommend* as well, but I largely ignore discommending. **Moral discommending weight** is that gradable feature of reasons that makes an act *morally worse* than at least some alternatives with the same deontic status in that context. I focus on cases that involve supererogation. And I assume that the supererogation involves moral commending (a reason making the supererogatory action better than some permissible alternative) rather than moral discommending (a reason making the permissible, non-supererogatory alternative worse than the supererogatory action).

¹³ Justifying and requiring weight are defined by their opposing functional roles in fixing the (im)permissibility of an act. Commending is likewise defined by its functional role in fixing some status, namely choiceworthiness. Yet that status is conceptually distinct from (im)permissibility. Permissibility is a non-gradable concept. No act is more permissible than any another. Choiceworthiness is a gradable concept. Some (permissible) acts are more choiceworthy than others. Since justifying/requiring weight and moral commending are defined in terms of conceptually distinct statuses, they too are conceptually distinct.

¹⁴ Muñoz (2021: 702, 712-3, nts. 1 and 5) endorses the assumption and then holds that commending is identical to requiring weight. Portmore (2019: 188, nt 4) and Little and Macnamara (2020) would agree with me that commending is distinct from justifying weight, requiring weight, and their sum.

3. The All or Nothing Problem for Pairwise Precedence

We now understand how reasons interact to determine deontic status and relative choiceworthiness. It is time to apply that understanding to the All or Nothing Problem.

3.1. *The All or Nothing Verdicts*

There are two children in the burning building, and you can easily carry both at the same time. This gives us three potential options. Remain a bystander (**Bystander**), save one child and get severe burns (**Save1**), and save both children and get the same severe burns (**Save2**). My main opponents—the proponents of the standard story¹⁵—and I agree completely on what the deontic and relevant rankings are for the two-option case (and we even agree about the explanation I will provide in §4.2). That is, we agree:

Two Option Verdict: If you had to choose between Save1 and Bystander, Save1 is supererogatory and Bystander is permissible but not supererogatory. (Save1 > Bystander)

The conceptual truth Supererogation is Better confirms the relative ranking in this verdict: if Save1 is supererogatory and Bystander isn't, then Save1 is morally better than Bystander.

In my view, what happens in the three-option case is this:

Three Option Verdict: when you have all three options, it is supererogatory to choose Save2, permissible to choose Bystander, and *impermissible* to choose Save1. (Bystander > Save1)

The proponents of the standard story agree with this verdict insofar as it concerns the deontic status of each option. They even agree about how these deontic verdicts are to be explained. I'll give a more technical explanation in §4.3, but the basic idea is as follows. It is permissible to not save lives only if there is a countervailing consideration for some alternative. Since Bystander is your only way to avoid severe burns, there *is* a countervailing consideration for choosing it over the best option, Save2. Yet there *isn't* a countervailing consideration for choosing Save1 over Save2. For Save2 involves saving an additional person over Save1 at no additional cost.

Where the proponents of the standard story and I part ways is over the relative ranking in the Three Option Verdict. In the three-option case, I endorse the verdict and say that Bystander > Save1, which reverses the relative ranking in the two-option case. The proponents of the standard story reject the relative ranking in the Three Option Verdict and say that Save1 > Bystander. If the standard story is correct, then the All or Nothing Problem is a counterexample to Right > Wrong: impermissible Save1 would be better than permissible Bystander. Is there a reason to endorse the standard story over the ranking reversal interpretation?

3.2. *Pairwise Precedence*

There would be a reason to prefer the standard story if the Two and Three Option Verdicts were to conflict, if there were something incoherent about treating the All or Nothing Problem as a ranking reversal. These verdicts will seem to conflict if we implicitly assume:

Pairwise Precedence: if A is morally better than B when A and B are the only options, then A is morally better than B when C is added as an alternative.¹⁶

¹⁵ Recall the references from note 1.

¹⁶ My opponents and I already agree that the All or Nothing Problem is a counterexample to *Deontic Pairwise Precedence* (aka β , aka expansion consistency): if A and B have the same deontic status when they are the only two options, then they have the same deontic status when C is added as an alternative (Pummer 2019: 288; Muñoz 2021: 711). This principle fails because the addition of Save2 makes Save1

If Two and Three Option Verdicts are both true, then Pairwise Precedence is false. Just let A = Save1, B = Bystander, and C = Save2. The Two Option Verdict says that Save1 is morally better than Bystander when they are the only options. Given Pairwise Precedence, it follows that Save1 is morally better than Bystander when Save2 is added as an option. But that conflicts with the Three Option Verdict, which holds that Bystander is morally better than Save1 in that case.

Two Option Verdict, Three Option Verdict, and Pairwise Precedence form an inconsistent triad. We must reject at least one of them. Pairwise Precedence seems to be the least plausible claim. The proponents of the standard story and I agree that the Two Option Verdict is true, and it is also supported by an apparent conceptual truth, Supererogation is Better (§1.1). Three Option Verdict (i.e., that permissible Bystander is better than impermissible Save1) is supported by another apparent conceptual truth, Right > Wrong (§1.3). Even proponents of the standard story find Right > Wrong plausible.¹⁷ By the end of the paper, I will have given two arguments that it is (conceptually) true. The proponents of the standard story don't provide any considerations to support Pairwise Precedence. These facts make for a strong initial case against Pairwise Precedence. At least, it's strong *as long as* we have a plausible way to explain both Two and Three Option Verdicts. We have such an explanation.

4. All or Nothing Ranking Reversal Explained

4.1. Deontic Precedence

Pairwise Precedence claims that the choiceworthiness of A and B when A and B are the only options constrains their relative choiceworthiness when you add new alternatives. I suggest that we replace Pairwise Precedence with Deontic Precedence, the idea that the deontic statuses of A and B constrain their relative choiceworthiness. In other words, I suggest that we endorse:

Deontic Precedence: the conjunction of:

Pluralist Permissibility: ϕ is permissible if and only if the justifying weight for ϕ is at least as great as the requiring weight for each alternative; otherwise, ϕ is impermissible. (§2.1)

Right > Wrong: Every permissible option is more choiceworthy than its impermissible alternatives. (§1.3, §5)

Definition of Commending Weight: Relative amounts of moral commending weight determine which options are morally better than their alternatives with the same deontic status. (§2)

Given Deontic Precedence, justifying and requiring weight arguably make a more important contribution to an act's choiceworthiness than moral commending weight. Only justifying and requiring weight determine whether acts are unchoiceworthy or not. No amount of moral commending weight can override the relative choiceworthiness determined by deontic status.

Moral commending weight plays only a secondary role in explaining why some (im)permissible acts are more choiceworthy than others. It kicks in only insofar as two alternatives have the same deontic status. This secondary role is all that's needed to explain Supererogation is Better, to explain how a supererogatory act can be morally

impermissible while leaving Bystander permissible. Since Deontic Pairwise Precedence fails, it shouldn't be surprising if the relative choiceworthiness version does too. See nt 21 for a discussion of how my ranking reversal interpretation of the All or Nothing Problem bears on α , or IIA.

¹⁷ Ferguson and Köhler note that Right > Wrong is widely accepted and intuitive. To reject it is to "incur a high theoretical cost" (2020: 2452). Also see Muñoz (2021: §4).

better than its non-supererogatory, permissible alternatives. Now that we understand Deontic Precedence, let's see how it resolves the All or Nothing Problem.

4.2. Two Option Verdict Explained

In order to determine the deontic status of Bystander, Save1, and Save2, we need to work with some account of how to weight the relevant reasons. I work with the standard account of supererogation. This account holds that self-interested benefits can have more justifying weight than even greater altruistic benefits. More specifically, it holds that the life of another person has an equal amount of justifying and requiring weight (say, 100 justifying and requiring weight per life) and that self-interested benefits have lots of justifying weight but no requiring weight (say, 500 justifying weight for Bystander).¹⁸ These assumptions give us:

All or Nothing Weights¹⁹

| | JW | RW |
|------------------|-----------------------------|---------------------------|
| Bystander | 500 $JW_{\text{Bystander}}$ | 0 $RW_{\text{Bystander}}$ |
| Save1 | 100 JW_{Save1} | 100 RW_{Save1} |
| Save2 | 200 JW_{Save2} | 200 RW_{Save2} |

In the two-option case, the only two alternatives are Bystander and Save1. ϕ is permissible just when the justifying weight for ϕ is at least as great as the requiring weight for each alternative. Hence:

- Save1 is permissible, because 100 justifying weight for Save 1 is greater than 0 requiring weight for Bystander.
- Bystander is permissible, because 500 justifying weight for Bystander is greater than 100 requiring weight for Save1.

Since they are both permissible, they are both choiceworthy. But neither deontic status nor justifying/requiring weight tell us whether either option is better than the other. That's where commending weight comes in. Since Save1 and Bystander are alternatives of each other with the same deontic status, their relative degrees of choiceworthiness are determined by the difference between how commended they are. Altruism Commends More than self-interested action (§2.1). Hence, Save1 is supererogatory, and $\text{Save1} > \text{Bystander}$.

4.2. Three Option Verdict Explained

Now add Save2 as a third option. The addition of Save2 makes the second child's life relevant to the choice. This addition raises the bar for permissibility for both Save1 and Bystander. For it raises how much requiring weight they must defeat in order to remain permissible, namely 200 requiring weight for Save2. 500 justifying weight for Bystander meets this higher bar, but 100 justifying weight for Save1 does not. Thus, we get the intuitive deontic verdicts:

¹⁸ See, e.g., Portmore (2011: ch 5), Massoud (2016). Archer (2016a), and cf. Muñoz (2021).

¹⁹ If you endorse comparativism—that reasons for ϕ are ways that ϕ is better than some alternative—then you will deny that there is justifying weight for Save1 when it competes with Save2. For there is no way that Save1 is better than Save2. Such disagreements with All or Nothing Weights don't matter. What determines whether ϕ is permissible is *net* justifying weight for ϕ (justifying weight for ϕ – requiring weight for the alternative). These net values will remain the same if you prefer to say that, in the pairwise competition between Save1 and Save2, the second child has 0 JW and RW for Save1 and only 100 JW and RW for Save2. For detailed discussion of how comparativism affects the assignment of weights, see my 2023: §4-6.

- Save2 is permissible, because 200 justifying weight for Save 2 is greater than both 100 requiring weight for Save1 and 0 requiring weight for Bystander.
- Bystander is permissible because 500 justifying weight for Bystander is greater than both 200 requiring weight for Save2 and 100 requiring weight for Save1.
- Save1 is *impermissible*. While 100 justifying weight for Save1 is greater than 0 requiring weight for Bystander, it is *less than* 200 requiring weight for Save2.

We also get (what I take to be) the intuitive relative ranking verdicts. Bystander is permissible and Save1 is not. Given Right > Wrong, it follows that Bystander is better than Save1 when Save2 is added as a third option. Save2 is supererogatory and best of all, because it is more commended than the only other permissible option, Bystander.

Deontic Precedence, therefore, explains both Two Option and Three Option Verdicts. In the two-option case, Save1 is morally better than Bystander because they are both permissible but Save1 is more commended and thus supererogatory. In the three-option case, Bystander is morally better than Save1, because the addition of Save2 makes Save1 impermissible while leaving Bystander permissible. Save1 remains more commended in the three-option case; however, the difference in deontic status takes precedence over the difference in commending.^{20, 21}

5. The Unity of Morality

My account of the All or Nothing Problem relies heavily on Right > Wrong, so let's consider a second argument for it. A miracle occurs. Kid wants to be helpful. He considers doing the dishes, but then speaketh Dad:

Dad: Don't do the dishes. Clean your room.

Kid starts to walk to his room, but then speaketh Mom:

Mom: It is better to do the dishes than to clean your room.

What is Kid to do? His parents are divided. Dad's directives conflict with Mom's rankings. There is no way to satisfy them both. The parents are *insatiably divided*.

Perhaps Kid will be relieved to hear that sometimes it is possible to satisfy both parents even when they are divided. Again, Kid wants to help.

Dad: Don't do the dishes. Wash my car or your Mom's car.

As he is walking to the garage, Mom calls out:

²⁰ Contra Portmore (2019: 208, especially his defense of 6.12), this section shows that the following conjunction is coherent: the relative weights of reasons determine both deontic status and relative choiceworthiness, and yet his 6.12 can be false, i.e., an impermissible option (Save1) can have more commending weight in its favor than a permissible option (Bystander).

²¹ Some readers may wonder whether Deontic Precedence is compatible with the Independence of Irrelevant Alternatives (aka IIA, aka α , aka contraction consistency). It *is* compatible with the standard version, namely *Deontic IIA*: if A is permissible when the alternatives are B and C, then A is permissible when B is the only alternative. Yet it does conflict *Relative Choiceworthiness IIA*: if A is better than B when C is an alternative, then A is better than B when they are the only two options. This principle holds that the subtraction of a third option cannot cause a ranking reversal. The All or Nothing Problem is a plausible counterexample to this principle. Deontic Precedence provides a plausible explanation of why the principle fails when it does: subtracting C/Save2 can affect the deontic status of A/Save1 without affecting the deontic status of B/Bystander. The intuitively problematic violations of IIA do not involve the third option differentially affecting the deontic status of the first two options (see, e.g., Rulli and Worsnip 2016: 219-20). So Deontic Precedence is committed to violating Relative Choiceworthiness IIA only when it is plausible to do so. Furthermore, the standard motivation for IIA comes from maximizing versions of decision theory (Rulli and Worsnip 2016: 206-7). Any such maximizing view is poorly suited to explain our intuitions in the All or Nothing Problem. Those intuitions are decidedly non-maximizing. Save2 makes only Save1 impermissible, even though you do more good in Save1 than Bystander. I submit, then, that IIA principles pose no problem for Deontic Precedence or my ranking reversal interpretation of the All or Nothing Problem.

Mom: It is better to do the dishes than to wash your Dad's car, but it is best of all to wash my car.

He can satisfy both parents by washing Mom's car. Yet Mom and Dad are still divided because of what they say about doing the dishes vs washing Dad's car. Mom says it is better to do what Dad says don't do. That's division, even if there is a way of satisfying them both. The parents are *satiably divided*.

Those who reject Right > Wrong to solve the All or Nothing Problem hold that morality can be double-minded. They hold that morality can be divided against itself in the same way that Mom and Dad are divided at least in the most recent parental division case. The Dad in morality says: Don't choose Save 1. Choose Bystander or Save2 instead. The Mom in morality says Save1 is morally better than Bystander but Save2 is best of all. You can satisfy the Mom and Dad in morality by choosing Save2. Yet morality is still divided against itself because of what it says about Bystander vs Save1. It says don't do what it is better to do. That's double-minded, that's morality being divided against itself.

If you don't feel the doublemindedness yet, recall that morality's relative rankings are recommendations, a type of all-in guidance. Save1's being impermissible is morality's way of saying don't choose it. Ranking Save1 over Bystander is morality's way of recommending that you choose Save1 over Bystander. Put the two together, and morality says, "don't choose Save1 over Bystander, but I recommend that you choose it over Bystander anyway". Such guidance is incoherent.

These reflections give us the *Normative Unity Argument*:

P1: It is a conceptual truth that, if morality isn't divided against itself, then Right > Wrong.

P2: It is a conceptual truth that morality isn't divided against itself.

C: Therefore, it is a conceptual truth that Right > Wrong.

P1 is true, because all it takes to get the relevant kind of incoherence is to deny Right > Wrong. If you deny Right > Wrong, you are committed to there being some case in which: A is morally wrong, B is morally right, but A is morally better than B. And such a case illustrates the double-minded conflict between morality's deontic verdicts and relative rankings. In contrast, Right > Wrong allows morality to be single-minded. It ensures that morality's deontic verdicts never conflict with its relative rankings.

P2 is true, because morality is supposed to be a coherent, authoritative guide to action. Yet you can't be coherently guided by morality insofar as morality's guidance is itself incoherent, insofar as it says, "Don't do what it is better to do". And no perspective that is divided against itself can have the authority morality is supposed to have. This point is especially important if morality is the normative perspective that has *final* normative authority, the perspective that settles what you are to do and aspire to. Conflicted guidance is itself unsettled. Unsettled guidance can't settle for you what to do or aspire to.

I have provided two arguments that Right > Wrong is a conceptual truth, the Deontic Evaluation and Normative Unity Arguments. The former capitalizes on plausible conceptual connections between deontic concepts (im/permissible) and evaluative ones (okay to choose vs not okay/unchoiceworthy). The latter capitalizes on our intuitions about when deontic verdicts and relative rankings compete. Together they show that Right > Wrong is rooted in a relatively robust range of intuitions.

Yet the arguments are not entirely independent of one another. Why do the relevant deontic verdicts (A is *impermissible*, B is *permissible*) seem in tension with the relevant ranking verdict (A is better than B)? I submit that the premises of the Deontic

Evaluation Argument explain our intuitions in the Normative Unity Argument. Deontic verdicts can conflict with relative choiceworthiness verdicts because deontic verdicts are choiceworthiness verdicts in disguise.

While the two arguments are not entirely independent of each other, they reveal a coherent network of intuitions that speak in favor of Right > Wrong as a conceptual truth. Right > Wrong is itself intuitively plausible, as even its critics grant. If you give it up to resolve the All or Nothing Problem, you better have a good reason for doing so.

6. Conflicting Intuitions?

Consider Monima who, out of hatred, refuses to save the second child (Portmore 2019: 207-9). In discussions of the All or Nothing Problem, it is generally taken for granted that the following sorts of claims are true:

AorN Conditional Ought (CO_{AorN}): if Monima isn't going to Save2, then she ought to Save1 (cf. McMahan 2018: 94-5)

AorN Contrastive Ought: Monima ought to choose Save1 rather than Bystander. (Horton 97)

Such intuitions seem to provide a devastating objection to Deontic Precedence, as it may seem that Deontic Precedence predicts that, if Monima isn't going to Save2, then she ought to remain a bystander. After all, Right > Wrong guarantees that Bystander is morally better than Save1.

I think the conditional and contrastive ought claims are true, but *only if* we understand their semantics in a very particular way. Take the conditional ought. If the conditional is a material conditional, and we are concerned with what Monima ought to do all things considered, the claim is clearly false. You can't get out of an ought simply because you refuse to do it. All things considered, Monima ought to Save2 even if she isn't going to do it. And since she ought to Save2, it is false that she ought to Save 1 (when she isn't going to Save2).

A similar point holds for the contrastive ought. "Rather than" usually means something like "and not" (Snedegar 2013: 234). The contrastive ought claim, then, might be understood as saying something like this: Monima ought to choose Save1 and it is false that she ought to choose Bystander. But this claim is clearly false. Monima ought to choose Save2, not Save1.

If these conditional ought and contrastive ought claims are true, we cannot interpret them as being concerned with what Monima all-in ought to do. I take it that the proponents of the standard story agree with me on this point. So, how do we understand such claims?

So far in this chapter we have been concerned with morality's *unconditional* guidance, its unconditional deontic verdicts and relative rankings. Yet it is common to think that morality offers conditional guidance to those who will ignore some aspect of morality's unconditional guidance. Take wrongdoers, for instance. Morality tells Melvin that he (unconditionally) ought not murder anyone. But it also tells him that:

- If Melvin is going to murder, then he ought to murder *gently* (Pummer 2019: 281-2; Muñoz & Pummer 2022: §1).
- Melvin ought to murder gently rather than brutally.

The correct semantics for such statements does not yield the claim that Melvin ought to murder gently (full stop). The reason why the AorN Conditional and Contrastive Oughts seem plausible is that they are true conditional guidance to be understood on analogy with this conditional guidance to Melvin. At least, that's what I and at least some proponents of the standard story assume. In the remainder of the section, I focus

on the relevant conditional ought, CO_{AorN} , but the same points apply to the AorN Contrastive Ought as well.

When CO_{AorN} is understood as morality's conditional guidance to Monima, Pummer holds that it is a counterexample to Right > Wrong.²² For he assumes that if CO_{AorN} is true, then Save1 must be "the next best thing" behind Save2 (284). If he's right, then Save1 would be better than Bystander and thus Right > Wrong would be false. Yet the assumption is false. The truth of CO_{AorN} does not require that Save1 be second best behind Bystander. It just gives us a new kind of ranking reversal. Although Bystander is *unconditionally* better than Save1, Save1 is *conditionally* better than Bystander. Deontic Precedence predicts and explains this ranking reversal.

According to Muñoz & Pummer, there is only a single difference between morality's conditional and unconditional guidance. The condition "If the agent is not going to ϕ " *excludes* the option of ϕ -ing, such that pairwise comparisons with excluded options are ignored in the conditional moral evaluation (Pummer 283-5; Muñoz & Pummer 2022: §5).²³

Pummer, Muñoz, and I all accept Pluralist Permissibility, which holds that ϕ is permissible just when the justifying weight for ϕ never loses a pairwise competition against the requiring weight for *any* alternative. When combined with Muñoz & Pummer's account of conditional guidance, Pluralist Permissibility tells us that ϕ is *conditionally* permissible just when it never loses a pairwise competition with any *unexcluded* alternative.

Consider the case in which Monima refuses to save the second child. The pairwise comparison with Save2 is what makes Save1 *unconditionally* impermissible (§4.2 above; Muñoz & Pummer 2022: §5). The 100 justifying weight for Save 1 loses to the 200 requiring weight for Save2. The condition "If Monima isn't going to Save2" excludes Save2 and now the 100 justifying weight for Save1 is greater than the requiring weight of all *unexcluded* alternatives, namely 0 requiring weight for Bystander. Hence, if Monima won't Save2, it is *conditionally* permissible for her to Save1.

It remains conditionally permissible for Monima to choose Bystander, as the 500 justifying weight for Bystander continues to outweigh the 100 requiring weight for the only unexcluded alternative, Save1. Hence, Bystander and Save1 have the same *conditional* deontic status. Muñoz & Pummer agree (2022: §5; Pummer 2019: 284-5). Since Bystander and Save1 have the same conditional deontic status, Deontic Precedence entails that Save1's greater commending weight kicks in and Save1 is conditionally better than Bystander. Muñoz and Pummer miss the way that the condition itself changes the relative ranking of Bystander and Save1. And, thus,

²² Pummer (2019: 283) also gives another apparent counterexample to Right > Wrong:

Hot Death: You can either (I) do nothing, (II) save one hundred children by pressing a button, or (III) save these very same hundred children, and prevent a separate child from losing a foot, by pressing a different button. Pressing either button will also cause you to drop into a pit of red-hot coals, where you will die an excruciatingly painful death. (283)

Pummer assumes that, unconditionally, it is permissible to choose I, impermissible to choose II, and supererogatory to choose III. III makes II impermissible, as it ensures that II would allow significant suffering (the loss of the foot) that could be costlessly prevented. III doesn't make I impermissible, because I is your only way to avoid death. Yet Pummer insists that the intuitive conditional deontic verdict is: ($CO_{HotDeath}$) if you aren't going to choose III, then you ought to choose II. Thus, he considers Hot Death to be another counterexample to Right > Wrong. But $CO_{HotDeath}$ is compatible with Right > Wrong in the same way as CO_{AorN} .

²³ Excluding a pairwise comparison with a given alternative is *not* tantamount to treating that alternative as if it weren't a genuine option. See Muñoz & Pummer (2022: 1438-9, including n28) for what the difference amounts to and why the difference matters.

Deontic Precedence explains the intuitive judgment that, if Monima isn't going to Save2, then she should Save1.

But wait! If Bystander is *unconditionally* better than Save1 and Save1 is *conditionally* better than Bystander, isn't that conflicting guidance? No. Conditional guidance doesn't overrule or compete with unconditional guidance. Conditional guidance is determined by excluding morally relevant considerations from consideration. Consequently, if you act in accordance with morality's conditional but not its *unconditional* guidance, then you are making a moral mistake. If Monima does what she conditionally ought (Save1), then she ignores morality's unconditional guidance (don't choose Save1) and makes a moral mistake (costlessly allows the second child to die).

Conditional guidance is for those who ignore some aspect of morality's unconditional guidance. This aim requires that conditional statuses sometimes diverge from their unconditional counterparts. It is this divergence that allows morality to say things like, "if Melvin is going to murder, then he ought to murder gently" and "if Monima isn't going to Save2, then it is permissible for her to Save1" (cf. Pummer 285). The idea that conditional relative rankings might diverge from their unconditional counterparts is just another instance of this same phenomena.

It is worth stressing just how wonderfully generous Muñoz and Pummer are. They object to Right > Wrong but then generously provide half of what we need to refute their objection (thanks, guys!). It may seem obvious that, if Monima refuses to choose Save2, then she ought to choose Save1. Such intuitions seem like damning counterexamples to Right > Wrong. But they're not. When you combine my Deontic Precedence with their account of morality's conditional guidance, the intuitive conditional oughts are compatible with Right > Wrong.

7. The Unity Explained

Morality is single-minded. Right > Wrong is what guarantees this single-mindedness, what guarantees that morality's deontic verdicts never oppose its relative rankings. Deontic Precedence is an account of how morality ranks alternatives while remaining single-minded. First, the competition between justifying and requiring weight sorts the alternatives into (i) the okay (neutral or choiceworthy) and (ii) the unchoiceworthy. This initial sorting is coarse-grained. Morality then relies on relative amounts of commending weight to sort the choiceworthy actions into those more or less choiceworthy, and to sort unchoiceworthy actions into those more or less unchoiceworthy. This sorting is fine-grained. It is also secondary and can't override the initial sorting. If an impermissible action is more commended than a permissible alternative, it is less choiceworthy than that permissible alternative.

The All or Nothing Problem involves an interesting ranking reversal. In the two-option case, Save1 > Bystander. In the three-option case, Bystander > Save1. The reversal occurs because Save2 adds new reasons that make Save1 impermissible. It is interesting, because these new reasons aren't relevant to the pairwise competition between Save1 and Bystander. In that pairwise competition, the justifying, requiring, and commending weight in favor of each option remains the same even after Save2 is added.

The All or Nothing ranking reversal is possible because of how commending weight behaves. First, the most commended alternative need not be the alternative with the most justifying weight in its favor. Save1 is more commended than Bystander, but only Bystander has enough justifying weight to beat Save2's requiring weight. Second, commending weight affects relative ranking only insofar as alternatives have the same

deontic status. Save1 is more commended than Bystander in both the two- and three-option cases. Yet this commending weight makes Save1 better than Bystander only in the two-option case, the case in which they have the same deontic status.²⁴

References

- Archer, Alfred. 2016a. "Moral Obligation, Self-Interest, and the Transitivity Problem." *Utilitas* 28: 441-464.
- _____. 2016b. "Are Acts of Supererogation Always Praiseworthy?" *Theoria* 82: 238-55.
- Benn, Claire. 2019. "Intentions, Motives, and Supererogation." *The Journal of Value Inquiry* 53: 107-23.
- Dancy, Jonathan. 2004. *Ethics Without Principles*. Oxford University Press.
- Ferguson, Benjamin and Sebastian Köhler. 2020. "Betterness of Permissibility." *Philosophical Studies* 177: 2451-2469.
- Gert, Josh. 2007. "Normative Strength and the Balance of Reasons." *Philosophical Review* 116: 533-562.
- Horgan, Terry and Mark Timmons. 2010. "Untying a Knot from the Inside Out: Reflections on the 'Paradox' of Supererogation." *Social Philosophy and Policy* 27:29 – 63.
- Horton, Joe. 2017. "The All or Nothing Problem." *Journal of Philosophy* 114: 94-104.
- Kawall, Jason. 2003. "Self-Regarding Supererogatory Actions." *Journal of Social Philosophy* 34: 487-98.
- Little, Margaret Olivia and Coleen Macnamara. 2020. "Non-Requiring Reasons." *The Routledge Handbook of Practical Reason*. Ruth Chang and Kurt Sylvan (eds.). Routledge, 393-404.
- _____. 2017. "For Better or Worse: Commendatory Reasons and Latitude." *Oxford Studies in Normative Ethics* 7: 138-60.
- Massoud, Amy. 2016. "Moral Worth and Supererogation." *Ethics* 126: 690-710.
- McMahan, Jeff. 2018. "Doing Good and Doing the Best." *The Ethics of Giving: Philosophers' Perspectives on Philanthropy*. Paul Woodruff (ed.). Oxford, 78-102.
- McNamara, Paul. 2011. "Supererogation, Inside and Out: Toward an Adequate Scheme for Common Sense Morality." *Oxford Studies in Normative Ethics, Volume I*: 202-35.
- Muñoz, Daniel. 2021. "Three Paradoxes of Supererogation." *Nous* 55: 699–716. <https://doi.org/10.1111/nous.12326>
- Muñoz, Daniel & Theron Pummer. 2022. "Supererogation and Conditional Obligation." *Philosophical Studies* 179: 1429-43. <https://doi.org/10.1007/s11098-021-01724-y>
- Portmore, Douglas. 2019. *Opting for the Best: Ought and Options*. Oxford University Press.
- _____. 2011. *Commonsense Consequentialism: Wherein Morality Meets Rationality*. Oxford University Press.

²⁴ This paper has benefited from the helpful people of many people. This includes the audiences of the *Arizona Workshop in Normative Ethics*, *APA Pacific Division, 14th Annual RoME*, *Virtual Metaethics Colloquium*, and the *Winter Ethics Workshop Miniseries*. It also includes Daniel Fogal, Doug Portmore, Theron Pummer, Kerah Gordon-Solomon, Joe Horton, and especially Tina Rulli and Daniel Muñoz who commented on multiple drafts of the paper.

- Pummer, Theron. 2019. "All or Nothing, but if Not All, Next Best or Nothing". *The Journal of Philosophy* 116: 278-291.
- Rulli, Tina. 2020. "Conditional Obligations." *Social Theory and Practice* 46: 365-90. doi: <https://doi.org/10.5840/soctheorpract20204189>
- Snedegar, Justin. 2016. "Reasons, Oughts, and Requirements." *Oxford Studies in Metaethics, Volume 11*. Oxford University Press, 155-181.
- Tucker, Chris. 2023. "Parity, Moral Options, and the Weights of Reasons". *Nous* 57: 454-80. doi: doi.org/10.1111/nous.12410
- _____. 2022a. "The Dual Scale Model of Weighing Reasons." *Nous* 56: 366-92. doi: <https://doi.org/10.1111/nous.12361>
- _____. 2022b. "Too Far Beyond the Call of Duty: Moral Rationalism and Weighing Reasons." *Philosophical Studies* 179: 2029-52. doi: <https://doi.org/10.1007/s11098-01-01747-5>