

Parity, Pluralism, and Permissible Partiality

Chris Tucker 5/15/23
William & Mary cstucker@wm.edu

We don't always have to perform the morally best action. It is permissible to go out to dinner even if you could have donated that money to feed five strangers for a single meal instead. It is better to donate and yet still permissible to eat out. But there are limits. It would be *impermissible* to go out to dinner tonight if you could have otherwise permanently ended world hunger. In the latter case, you would be required to forgo eating out for the sake of ending world hunger.

Whether you could otherwise feed five people or permanently end world hunger, going out to dinner involves *partiality* to yourself. That is, it involves giving a certain priority to your own wellbeing, or treating your own wellbeing as though it were more significant than that of strangers. The first case is an example of *permissible* partiality. It is okay to give yourself a benefit when you could have provided a comparable benefit to each of five strangers. The second case is an example of *impermissible* partiality. It is not okay to give yourself a small benefit when you could have provided a much bigger benefit to millions of other people.

Choosing to eat out over feeding five involves permissible partiality to *yourself*. We also can be permissibly partial *toward loved ones*. It is permissible to send my three kids to college even if I could have done more good by sending three needier, smarter, morally better children to college instead. And we can be permissibly partial *toward projects*. Even if you could do more good by spending your Saturdays volunteering at the local homeless shelter, it is permissible to watch college football instead.

With this variety acknowledged, I focus on cases of permissible partiality toward yourself, especially insofar as they involve choosing a worse self-interested action over a better altruistic action. This kind of partiality has a special connection to supererogation, or going above and beyond the call of moral duty. Permissible partiality and supererogation tend to be two sides of the same choice. If you eat out instead of feeding five, you are being permissibly partial toward yourself. If you feed five instead of eating out, then you are doing something supererogatory (or at least something that is a candidate for supererogation¹). By understanding how to explain permissible partiality, we thereby better understand how to explain supererogation.

The title of the volume refers to *virtuous* permissibility, which goes beyond permissible partiality by involving, at the very least, a disposition that is responsive to opposing reasons. You and I go out to eat instead of feeding five people, and so we both do something permissibly partial toward ourselves. Yet suppose that only you care about other people, that only you have a disposition to respond to altruistic reasons that oppose self-interest. For example, you would not have eaten out if you could have otherwise permanently ended world hunger, and I would have. Your partiality might be virtuous, but mine clearly isn't. This chapter's focus is on permissible partiality. Since virtuous partiality always involves permissible partiality, we can clarify virtuous partiality by clarifying permissible partiality.

A natural thought is that the im/permissibility of an act ϕ is determined by weighing the reasons for ϕ against the reasons for the alternatives. An *alternative* to ϕ is a different action, incompatible with ϕ , that you can still perform in the circumstances. Spending

money to eat out and donating that money are alternatives insofar as there is no way to do both. When they are alternatives, the reasons to eat out compete with the reasons to donate. Intuitively, the benefit to you of eating out is a weighty enough reason to make it permissible to eat out when the only competing reason is the altruistic benefit of feeding five strangers for a single meal; however, it is completely overwhelmed when the only competing reason is the altruistic benefit of permanently ending world hunger.

I assume that this “weighing reasons approach” to ethics is fundamentally correct. Yet, to be compatible with the full range of permissible actions, we must make an important and surprising claim about the weights of reasons. How many weight values does each reason have? **Weight Monism** is the idea that reasons have a single weight value. There is just *the* weight of reasons. Weight Monism will strike many as the intuitively obvious view, but it struggles to explain the full range of permissible actions. In particular, no extant version of Weight Monism explains both permissible partiality and a certain puzzle that I call *The Normative Significance of Small Improvements*. On the other hand, I will show that **Weight Pluralism**—the idea reasons have more than one weight value—easily explains both issues.

In §1, I explain the Normative Significance of Small Improvements. In §2, I identify the simplest versions of Weight Monism and show that they cannot resolve this puzzle. We must pay the price of complication. But not all complications are equal. We should go for the complication that gives us the most bang for our buck. I focus on two potential complications: appealing to a new comparative, parity, or appealing to a second weight value *a la* Weight Pluralism.

In §3, I explain how to revise Weight Monism to resolve this puzzle. This version—Parity Monism—complicates weight comparisons by appealing to the idea that reasons can be incommensurable (i.e., not precisely comparable) and, more specifically, the idea that they can be “on a par”. Unfortunately, in §4, we’ll see that this version of Weight Monism is incompatible with permissible partiality.

In §5, I explain how Weight Pluralism can explain permissible partiality. In §6, I show that it can resolve the Normative Significance of Small Improvements. Since Weight Pluralism can explain both issues and parity can explain only one, we have better reason to accept Weight Pluralism than parity. (We can accept both complications; we would just need to find something for parity to do that Weight Pluralism doesn’t already do.) More generally, this paper provides a reason for Weight Pluralism, because it provides the only known mechanism for explaining both permissible partiality and the normative small improvement puzzle.²

1. The Normative Significance of Small Improvements

The **normative significance of small improvements** is the question, how can *pro tanto* maximizing reasons generate stable moral options? To understand this question, we need to understand the two key terms, *pro tanto maximizing reasons* and *stable moral options*. Let’s start by explaining the former.

Pro tanto maximization concerns what happens when there is a single morally relevant consideration at play, i.e., when we don’t have to trade off different morally relevant considerations against each other. Consider *altruistic reasons* to save the lives of strangers. Such reasons are *pro tanto maximizing* in that, in the absence of other morally relevant considerations, you are required to save as many lives as you can.

Suppose that you have only two options: you have the altruistic option of saving 5 people (**A**) and a given small improvement of that option, namely saving those 5 people and an additional person (**A+**). If the number of lives saved is the only morally relevant consideration, you are required to choose **A+**, to save all six lives. If you can *costlessly* save an additional life, then you are required to save it.

Reasons to respect rights are also pro tanto maximizing in that, when you aren't trading rights off against other morally relevant considerations, you must respect as many rights as you can. Suppose that your only two options are respecting Jerry's right not to be beaten up (**R**) and a given small improvement of that option, namely respecting both Jerry's right not to be beaten up *and* his right not to be insulted (**R+**). If the number of rights that you respect is the only morally relevant consideration, then you are required to choose **R+**, to respect both rights. If you can *costlessly* respect an additional right, then you are required to do it.

Now that we better understand pro tanto maximization, let's consider stable moral options. A **moral option** is a case in which you have more than one permissible option. We often have more than one permissible option when we have to trade off different morally relevant considerations against each other. Cases of permissible partiality are one sort of example. When you have to choose between worse self-interested and better altruistic benefits, it is often permissible to choose either one. Or consider a choice between saving lives and respecting rights. If the only way to save five lives is to beat up Jerry, then you arguably have the moral option to choose **A** or **R** (i.e., it is permissible to save five and it is permissible to respect Jerry's right not to be beaten up).

A moral option is **stable** to the extent that the moral option persists despite increases/decreases to one of the opposing reasons. Moral options between altruism and respect are often stable in this sense. If it is a moral option to beat up Jerry to save 5 lives, then it is also a moral option to beat up Jerry to save 6 lives. That is, if the choice between **A** and **R** is a moral option, then so is the choice between **A+** and **R**. The moral option between altruism and respect is also stable when we improve the reasons of respect. If we have the moral option to beat up Jerry to save 5, then we also have the moral option to beat up *and insult* Jerry to save 5 lives. That is, if the choice between **A** and **R** is a moral option, then the choice between **A** and **R+** is also a moral option.

There are typically limits to how stable moral options are. For example, the moral option to beat up Jerry to save 5 is converted into a requirement if we increase the altruistic reason enough. We are required to beat up Jerry if doing so will save a million lives.

Small improvements, such as **A+** and **R+**, pose a puzzle because their liability to generate requirements depends on whether we are trading off morally relevant considerations. When we aren't trading off different considerations, then we are required to choose the small improvement. It is *impermissible* to choose **A** over **A+** and *impermissible* to choose **R** over **R+**. On the other hand, when we are trading off morally relevant considerations against each other, moral options are stable and the small improvement isn't liable to generate a requirement. If the choice between **A** and **R** is a moral option, then the choice between **A** and **R+** is a moral option too.

In general, the normative significance of small improvements is just the question of how pro tanto maximizing reasons, such as altruistic reasons to save lives and reasons to respect rights, can generate stable moral options. When we apply it to the cases we've

been considering, it boils down to this question: how do we explain all four deontic verdicts (all four verdicts about which options are im/permissible)?

SI1: It is *impermissible* to choose altruistic A over A+.

SI2: It is *impermissible* to choose respectful R over R+.

SI3: It is permissible to choose A over R and R+ (and permissible to choose both R and R+ over A).

SI4: It is permissible to choose R over A and A+ (and permissible to choose A and A+ over R).

SI1 and SI2 exemplify the pro tanto maximization of altruistic and respect reasons, respectively. SI3 and SI4 exemplify the stability of moral options when we are trading off altruistic and respect related reasons. We resolve the puzzle by providing a good explanation of how all four claims can be true. This is easier said than done.

2. Against Simple Weight Monism

2.1. Simple Weight Monism and Pro Tanto Maximization

Recall that Weight Monism is the idea that reasons have a single weight value. There is just *the* weight of reasons. The simplest version of Weight Monism is:

Simple Weight Monism: Weight Monism plus these three claims:

Three Comparatives: Every reason is either weightier than, less weighty than, or equally weighty as every other reason;

Monist Permissibility: ϕ is permissible if and only if the reason for ϕ is not less weighty than the reason for $\sim\phi$.

Monist Requirement: ϕ is required if and only if the reason for ϕ is weightier than the reason for $\sim\phi$.

Three Comparatives makes weight comparisons work like numerical comparisons. Just as every number is either greater than, less than, or equal to every other number, every weight value is either weightier than, less weighty than, or equally weighty as every other reason. Monist Permissibility and Monist Requirement are assignments of deontic status in terms of the single weight of reasons.

Simple Weight Monism's simplicity may seem attractive; however, it is too simple to explain the Normative Significance of Small Improvements. While it can explain pro tanto maximization, it is incompatible with stable moral options.

Simple Weight Monism's commitment to Monist Requirement does the heavy lifting in its explanation of pro tanto maximization. That is, it does the heavy lifting in explaining why:

SI1: It is *impermissible* to choose altruistic A over A+.

SI2: It is *impermissible* to choose respectful R over R+.

Monist Requirement says that you are required to do whatever you have weightier reason to do. The reason to save all six strangers (A+) is weightier than the reason to save only five of them (A), so Monist Requirement entails that you are required to choose A+ over A. SI1 is thereby explained. The reason to respect both Jerry's right not to be beaten up and his right not to be insulted (R+) is weightier than the reason to respect only his right not to be beaten up (R), so Monist Requirement entails that you are required to choose R+ over R. SI2 is thereby explained.

2.2. Simple Weight Monism and Stable Moral Options

Recall that, intuitively, trade-offs between altruism and respect tend to generate stable moral options. Moral options are cases in which more than one option is permissible, and they are stable insofar as we can increase the reason to choose one of the options and the moral option persists (rather than being converted into a requirement to choose one option over the other). The choice between A (save five) and R (respect Jerry’s right not to be beaten up) is a moral option. When we strengthen the respect reasons, the choice between A and R+ is also a moral option.

Simple Weight Monism is incompatible with stable moral options. It makes moral options *fragile*. If we begin by assuming that the choice between A and R is a moral option, the small improvement R+ converts the moral option into a requirement to choose R+ over A. Likewise, if we begin by assuming that the choice between A and R is a moral option, the small improvement A+ converts the moral option into a requirement to choose A+ over R. The problem emerges from its implicit assignment for moral options.

Simple Weight Monism’s commitment to Three Comparatives tells us that there are only three possibilities concerning the relative weight of the reasons for A and R: the reason for A is weightier, the reason for R is weightier, or they are equally weighty. Its commitment to Monist Requirement tells us that you are required to do whatever you have weightier reason to do, and so the first two possibilities give us requirements, not moral options. Simple Weight Monism entails, then, that the choice between A and R is a moral option exactly when the reasons for A and R are *equally weighty*. This gives us the first premise in the below argument, an argument which shows that Simple Weight Monism is incompatible with *stable* moral options (alternatively, it is committed to the *fragility* of moral options).

Improvements Would Make Requirements (for Simple Weight Monism)

P1: The reasons for A and R are equally weighty.

P2: The reason for R+ is weightier than the reason for R.

P3: Any reason that is weightier than R is weightier than A. [P1, substitutability principle to be introduced below]

C1: Therefore, the reason for R+ is weightier than the reason for A. [P1, P2, P3]

P4: If the reason for R+ is weightier than the reason for A, then you are required to choose R+. [Monist Requirement]

C2: Therefore, you are required to choose R+ over A. [C1, P4]

Again, P1 follows from Simple Weight Monism, given our assumption that the choice between A and R is a moral option. P2 follows from the commonsense observation that we have weightier reason to respect both of Jerry’s rights rather than only one of them. P4 just is an application of Monist Requirement. That leaves us with P3. Why think that premise is true?

P3 is just an application of this general principle:

Substitutability of Equality: “if two items are equally good with respect to V, one can always be substituted for the other in comparisons with respect to V” (Chang 2017: 4).

When V is weight of reasons, the Substitutability of Equality says that equally weighty things always compare equally to the weights of other things. If the reason for R is equally weighty as the reason for A, then no reason can be weightier than R without being weightier than A. Analogously, since the square root of 25 equals 5, 10 can’t be greater than 5 without being greater than the square root of 25. P1 tells us that the reasons for A

and R really are equally weighty. Hence, when we apply the Substitutability of Equality to P1, it follows that P3 is true, that any reason weightier than R is also weightier than A.

The argument, Improvements Would Make Requirements, has four premises: P1, P2, P3, and P4. I claim that P2 and the Substitutability of Equality (the general principle behind P3) are true. I don't claim that P1, P3, and P4 are true; I claim only that Simple Weight Monism is stuck with them. Hence, what the argument shows us is that Simple Weight Monism is incompatible with stable moral options. If the choice between A and R is a moral option, then Simple Weight Monism entails that we are required to choose R+ over A.

To explain the stability of moral options, we need a way to explain moral options when the opposing reasons *aren't* equally weighty. One way to do this is to reject Simple Weight Comparativism's commitment to Three Comparatives and then find a fourth comparative to explain *stable* moral options. This is Ruth Chang's strategy.

3. Parity Monism

Parity (Weight) Monism is nearly identical to Simple Weight Monism. Like *Simple Weight Monism*, it claims that reasons have a single weight value (Weight Monism), and that deontic status is to be assigned in terms of this single weight value. ϕ is permissible just when the reasons for ϕ aren't outweighed (Monist Permissibility), and ϕ is required just when the reasons for ϕ outweigh the reasons for the alternatives (Monist Requirement). The only difference is that it replaces Simple Weight Monism's assumption that there are only three comparatives with the claim that there are four:

Four Comparatives: every reason is either weightier than, less weighty than, equal to, or *on a par* with every other reason;

In other words, Parity Monism is the conjunction of Weight Monism, Monist Permissibility, Monist Requirement, and Four Comparatives.

Parity is like an imprecise version of equal to. The altruistic reasons to save lives and the reasons to respect rights apparently ensure that the choice between A and R is a moral option. But how do those reasons compare? Perhaps neither seems weightier than the other and they don't seem *exactly* equal either. The Parity Monist says that there is a fourth way that they might compare. They might be on a par.

If parity is a genuine comparative, then it is a *distinct* way that reasons might compare than the standard three comparatives (weightier than, less than, and equal to). If a reason R1 is weightier than another reason R2, then it can't be less weighty than, equally weighty as, or on a par with R2. Likewise, if the reasons for A and R are genuinely on a par, then neither is weightier than the other and they aren't equally weighty either.

Why would the reasons for A and R be on a par rather than being related by one of the other three comparatives? Our altruistic reasons to save lives are qualitatively very different than our reasons to respect rights. The proponent of parity will insist that qualitative difference can limit the extent to which those two reasons can be precisely compared. In contrast, the difference between the reasons for A and A+ is just a quantitative difference. It is just a difference in the total number of lives saved. Mere *quantitative* differences tend not to generate parity, because it is *qualitative* difference that limits the precision of comparisons. It is qualitative difference that makes parity possible (assuming that parity is a genuine comparative).

We are assuming that the reasons for A and R generate a moral option. Monist Permissibility says that both A and R are permissible just when neither the reason for A nor the reason for R is weightier than the other. Single Weight Monism's commitment to Three Comparatives entailed there is only one way that neither reason can outweigh the other, namely that the reasons for A and R are *equally weighty* (as noted by P1 above). Parity Monism's commitment to Four Comparatives allows a second way those reasons may not be weightier than each other, and so a second way they may generate a moral option. The reasons for A and R might be *on a par*. Parity Monism says that parity *always* generates moral options. When reasons for two options are on a par, neither is weightier than the other and thus Monist Permissibility says that they are both permissible. And, unlike equality, parity is well-suited to explain stable moral options.

Recall that equality is so precise of a comparative that equivalent values can always be correctly substituted. Since the square root of 25 is equal to 5, any number that is greater than 5 is also greater than the square root of 25. If A and R are equally weighty, then any reason weightier than R must be weightier than A (as noted by P3 above). Yet parity's imprecision yields:

Failures of Substitution: If the reasons for A and R are on a par and the reason for R+ is weightier than the reason for R, it does *not* follow that the reason for R+ is weightier than the reason for A. The reason for A might be on a par with the reason for R+ too.

These failures of substitution can generate *stable* moral options. A moral option is stable to the extent that you can de/increase the weight of one of the reasons and the moral option persists. Small improvements always break ties (exact equalities), but they don't always break parities. If the reasons for A and R are on a par, the reasons for A and R+ might still be on a par and, thus, the choice between A and R+ would still be a moral option.

We are now in position to see how Parity Monism explains the normative significance of small improvements, how it can explain all of SI1-SI4. Parity Monism explains pro tanto maximization (SI1-SI2) in the same way as Simple Weight Monism, by appealing to Monist Requirement (weightier reasons give you requirements).

SI1 Explained: Since the reasons for A+ are weightier than the reasons for A, you are required to choose A+ over A.

SI2 Explained: Since the reasons for R+ are weightier than the reasons for R, you are required to choose R+ over R.

Parity Monism explains the stability of moral options (SI3-SI4) by appealing to the fact that parity always generates moral options.

SI3 Explained: You have the moral option to choose A (save five) over both R (respect right not to be beaten up) and R+ (respect that right and right not to be insulted), because the reason for A is on a par with the reasons for both R and R+.

SI4 Explained: You have the moral option to choose R (respect right not to be beaten up) over both A (save five) and A+ (save six), because the reason for R is on a par with both the reason for A and the reason for A+.

We should pause to appreciate how Parity Monism's explanation of pro tanto maximization coheres with its explanation of stable moral options. The standard three comparatives are all transitive. For example, if R1 is weightier than R2 and R2 is

weightier than R3, then R1 is also weightier than R3. Since parity yields failures of substitution, it is *intransitive* (Chang 2017: 15). This intransitivity is what allows Parity Monism to combine pro tanto maximization and stable moral options. Consider, for example, Parity Monism's explanations of SI2 and SI3. Its explanation of SI3 entails both that the reason for R+ is on a par with the reason for A and that the reason for A is on a par with the reason for R. But, per its explanation of SI2, it does not follow that the reason for R+ is on a par with the reason for R. The reason for R+ is clearly weightier than the reason for R.

4. Against Parity Monism

Recall that, intuitively, there are cases of permissible partiality in which you have the moral option to choose a *worse* self-interested action over a *better* altruistic action. These cases of permissible partiality are also paradigmatic opportunities for supererogation. When you perform a supererogatory option, you go beyond the call of duty. You choose a permissible (altruistic) option that is better than another permissible (self-interested) option. If you forgo eating out to feed five hungry children for a single meal, then you have done something supererogatory.³

There's not a lot riding on a choice between dinner out and donating to feed five. Other choices between permissible partiality and supererogation are more significant. If you live in a city with adequate public transportation, you might find yourself in:

Car Case: It is permissible to buy a \$17k car for self-interested reasons (allows you to take a preferable job and enjoy preferable leisure activities) rather than donate the money to an effective malaria relief charity that will save the lives of five people.⁴ It would be supererogatory, so better, to donate the money.

Choices between a permissible worse self-interested option and a better altruistic option are ubiquitous. Given how intuitive it is that such self-interested choices can be permissible, it is a strike against a moral theory if it entails that, after all, we are required to do the better altruistic action. Unfortunately, this is precisely what Parity Monism entails.

In order to accommodate these cases of permissible partiality, we need to combine two ideas: the self-interested option is (i) permissible but (ii) worse than some altruistic alternative. If the self-interested option is *worse*, then the reason for the self-interested option is *less weighty* than the reason for the altruistic option. Recall that Parity Monism is committed to Monist Requirement, which says that we are always required to do what we have weightier reason to do (and so it is impermissible to do what we have less weighty reason to do). Hence, Parity Monism entails that the relevant cases of permissible partiality are impossible. Once it is acknowledged that the self-interested option is worse than the altruistic alternative, Parity Monism entails that the self-interested option is *impermissible*.

The proponent of Parity Monism might try to redescribe apparent cases of permissible partiality as cases of parity. Donating to feed five is not really better than eating out; rather, the two options are on a par. Donating to save five lives is not really better than your moderately improved career and leisure activities; rather, the two options are on a par. Such a redescription has a genuine advantage. Parity Monism holds that parities generate moral options. Given the redescription, then, Parity Monism captures the intuition that apparent cases of permissible partiality are moral options, and, as such, we

are not required to do the altruistic option. Nonetheless, there are two problems. First, the redescription strains all plausibility. I really do think you are great. You should be proud of yourself. But it vastly overstates your importance to say, e.g., that the improvement to your career and leisure are on a par with the lives of five people.

The second problem is that the redescription does nothing to make Parity Monism compatible with supererogation. Like permissible partiality, supererogation also combines two ideas: the altruistic option is (i) morally optional (so not required) *and* (ii) better than the self-interested alternative. The redescription preserves the former but loses the latter. For it holds that saving five lives is, not better, but on a par with getting the car.

Intuitively, there are many instances of permissible partiality that involve choosing a worse self-interested option over a better altruistic alternative, and these cases are also cases in which one might supererogatorily choose the better altruistic option over the worse self-interested option. Parity Monism's commitment to Monist Requirement is incompatible with permissible partiality and supererogation. For it guarantees that the weightier reasons for the better option require you to do that better option.

Simple Weight Monism entailed that moral options are possible only when the reasons for the options are tied. To resolve the Normative Significance of Small Improvements, we saw that we need a way to accommodate moral options when the weights of reasons are *not* tied. Parity Monism retained Simple Weight Monism's idea that permissibility and requirement are determined by a *single* weight of reasons. It accommodated moral options without ties by appealing to a fourth comparative, parity. But maybe that was a wrong turn. It took us some distance insofar as it resolved the Normative Significance of Small Improvements, but it ultimately faced a dead end and could not explain permissible partiality or supererogation.⁵ Perhaps we should revisit the idea that permissibility and requirement are determined by a single weight of reasons. Perhaps a second weight of reasons take us everywhere we want to go without the need for a fourth comparative.

5. Weight Pluralism, Permissible Partiality, and Supererogation

Recall that Weight Pluralism is the idea that reasons have more than one weight value. In this section, I show that it can explain permissible partiality and supererogation. In the next section, I show that it can explain the Normative Significance of Small Improvements.

5.1. Pluralist Permissibility and Stable Moral Options

Following Josh Gert (2004: ch 4), my version of Weight Monism holds that deontic status is determined by the interaction of two weight values, justifying and requiring weight. A reason's **justifying weight** for ϕ is how good it is at making ϕ permissible, i.e., how hard it "pushes" ϕ toward permissibility. A reason's **requiring weight** for ϕ is how good the reason is at making ϕ required, when ϕ is also permissible. If ϕ is permissible, it is required just when the alternative, $\sim\phi$, is *impermissible*. Hence, a reason's requiring weight for ϕ is how hard it "pushes" $\sim\phi$ toward *impermissibility*.

Notice that justifying weight for ϕ and requiring weight *for* $\sim\phi$ are in direct competition. Justifying weight for ϕ pushes ϕ toward permissibility. Requiring weight

for $\sim\phi$ pushes the alternative to $\sim\phi$, namely ϕ , toward *impermissibility*. It is this competition that determines whether ϕ is permissible:

Pluralist Permissibility: ϕ is permissible just when the justifying weight for ϕ is not outweighed by the requiring weight for any alternative. Otherwise, ϕ is impermissible.

Monist Permissibility and Requirement were assignments of deontic status in terms of a *single* weight value of reasons. Pluralist Permissibility is intended as a replacement. Most directly, it tells us which options are permissible. Indirectly, it tells us which options are required. We simply apply Pluralist Permissibility to ϕ and each alternative. If it tells us that ϕ is permissible and all alternatives are impermissible, it thereby tells us that ϕ is required.

Recall that moral options are stable to the extent that you can de/increase the reasons for one of the options and the moral option remains (rather than being converted into a requirement to choose one option over the other). To whatever extent reasons have more justifying than requiring weight—to whatever extent reasons are better at making an act permissible than making it required—to that extent Pluralist Permissibility generates stable moral options.

Intuitively, self-interested reasons have more justifying weight than they have requiring weight. Consider the self-interested reason *that you enjoy eating out*. This reason is good at making eating out permissible but bad at making it required. You would be permitted to eat out if you could otherwise feed five people, but you are hardly required to do so. Indeed, you are not required to eat out even if you could otherwise feed only a single child. It seems, then, that self-interested reasons are *at least* five times as good at making an action permissible than they are at making that action required. For illustration, suppose that the self-interested reason *that you enjoy eating out* has 10 justifying weight and 1 requiring weight.

If you could costlessly feed a hungry child, you would be required to do so. Altruistic reasons to feed children, then, must have at least some requiring weight. For illustration, assume that the altruistic reason to feed hungry children has 2 justifying weight and 1 requiring weight per child that you can feed. When we plug these weight values into Pluralist Permissibility, we get stable moral options involving self-interested reasons:

- Eating out vs feeding *one* child: This choice is a moral option. Eating out is permissible, because 10 justifying weight outweighs 1 requiring weight to feed the one child. Donating is also permissible, because 2 justifying weight to feed a single child outweighs 1 requiring weight to eat out.
- Eating out vs feeding *five* children: The reason to feed five children is, we are assuming, five times weightier than the reason to feed a single child. Nonetheless, Pluralist Permissibility entails that this choice remains a moral option. It is permissible to eat out rather than feed five, because 10 justifying weight to eat out outweighs 5 requiring weight to feed five (five children x 1 requiring weight per child). It is also permissible to feed five, because 10 justifying weight to feed five (five children x 2 justifying weight per child) outweighs 1 requiring weight to eat out.

Intuitively, choices between self-interested and altruistic benefits are often stable. The choice to eat out or feed one is a moral option. It remains a moral option even when we strengthen the altruistic reasons so that we could otherwise feed five. We explained

the stability of this moral option by appealing to Pluralist Permissibility and the intuition that self-interested reasons have more justifying than requiring weight. The more that their justifying weight outstrips their requiring weight, the more stable the moral option will be.

5.2. Explaining Permissible Partiality and Supererogation

The choice to eat out tonight or feed five children is essentially a choice between permissible partiality and supererogation. If you self-interestedly eat out tonight, then you are being permissibly partial to yourself. If you altruistically feed five, then you perform a supererogatory action. So understood, the choice must have at least two features: (i) it is a moral option (both options are permissible, neither is required), and (ii) the self-interested option is *worse* than the *altruistic* option. Pluralist Permissibility, as explained in the previous sub-section, explains how this case is a moral option. But how can the Weight Pluralist explain why the altruistic action is better?

Different Weight Pluralists can give different accounts of when one action is better than another. They can appeal to axiological rankings of impersonal goodness; they can appeal to rankings that involve a third weight value, such as commending weight; or they can appeal to rankings that involve just requiring weight. These differences don't matter for our purposes. Let's work with:

Simple Betterness: ϕ is (morally) better than some alternative just when the requiring weight for ϕ is greater than the requiring weight for that alternative.⁶

Given our assumptions about how the self-interested and altruistic reasons are weighted, Simple Betterness entails that it is better to feed five than to eat out. There is 5 requiring weight for feeding five children but only 1 requiring weight for eating out.

Simple Weight Monism and Parity Weight Monism could not explain permissible partiality and supererogation, because of their commitment to Monist Requirement, the idea that weightier reasons generate requirements. Since we have weightier reasons to do what is better, Monist Requirement doesn't let options be better without being required over what is worse. To explain permissible partiality and supererogation, we need to let what's permissible and required float free of what's better. That's what the combination of Pluralist Permissibility and Simple Betterness does.

Pluralist Permissibility holds that whether ϕ is permissible is determined by one competition: the *justifying* weight for ϕ vs the requiring weight for $\sim\phi$. Simple Betterness holds that what's better is determined by a distinct competition: *requiring* weight for ϕ vs requiring weight for $\sim\phi$. Self-interested reasons can make a worse action permissible because they have more justifying weight than requiring weight. Eating out is worse than feeding five, because the self-interested reason's 1 requiring weight is *less than* the 5 requiring weight to feed five. It is nonetheless permissible to eat out, because the self-interested reason's 10 justifying weight is *greater than* the 5 requiring weight to feed five. When we plug reasons that have more justifying weight than requiring weight into Pluralist Permissibility and Simple Betterness, we will tend to get cases in which a worse option is permissible over a better one. And that's just what we need to explain permissible partiality and supererogation.

6. Weight Pluralism and the Normative Significance of Small Improvements

6.1. Altruistic Reasons and Stability

Parity Monism could explain the Normative Significance of Small Improvements, but it couldn't explain permissible partiality or supererogation. Weight Pluralism can explain permissible partiality and supererogation. But can it explain the Normative Significance of Small Improvements? Yes.

The Normative Significance of Small Improvements, in general, is the question of how pro tanto maximizing reasons can generate stable moral options. We focused on a particular instance of this puzzle, which concerned two examples of pro tanto maximizing reasons, altruistic reasons to save lives and reasons to respect rights. We can resolve this puzzle if we can explain all four of:

SI1: It is *impermissible* to choose altruistic A over A+.

SI2: It is *impermissible* to choose respectful R over R+.

SI3: It is permissible to choose A over R and R+.

SI4: It is permissible to choose R over A and A+.

To explain SI1 and SI2 is to explain how altruistic and respect reasons are pro tanto maximizing. To explain SI3 and SI4 is to explain how these reasons can nonetheless give rise to stable moral options.

We already have the machinery to explain SI3 and SI4, so let's start there. Moral options are stable to the extent that we can de/increase one of the options and the moral options persists (rather than being converted into a requirement). Pluralist Permissibility generates stable moral options to the extent that the opposing reasons have more justifying than requiring weight. We saw that tradeoffs between self-interest and altruism generate stable moral options, because at least self-interested reasons have more justifying than requiring weight (§5.1). Tradeoffs between altruism and respect generate stable moral options, because at least altruistic reasons have more justifying than requiring weight (cf. Archer 2016: 455-61).

In §5.1, I assumed that self-interested reasons have *ten times* more justifying than requiring weight whereas altruistic reasons have only *twice* as much justifying as requiring weight. That's because I take self-interested reasons to be better at generating moral options than altruistic reasons, i.e., moral options involving self-interest tend to be more stable than moral options involving altruistic reasons. In any event, let's continue assuming that altruistic reasons have twice as much justifying weight as requiring weight and see where that takes us. For illustration, suppose that there are 100 justifying weight and 50 requiring weight for each additional life you can save.

Also assume that reasons to respect rights are just as good at making an act required as they are at making it permissible. That is, assume that these reasons have just as much requiring weight as justifying weight, roughly in proportion to the significance of the right. The more significant the right, the more justifying and requiring weight there is to respect it. These assumptions about altruistic and respectful reasons give us:

The Weight of the Altruistic and Respect Reasons

5 Lives: 500 justifying and 250 requiring weight for A.

6 Lives: 600 justifying and 300 requiring weight for A+.

Right not to be beaten up: 400 justifying and 400 requiring weight for R
Right not to be beaten up or insulted: 405 justifying and 405 requiring weight for R+

When we plug these weight values into Pluralist Permissibility, we explain SI3 and SI4:

SI3 is Explained: A is permissible over R, because 500 justifying weight for A outweighs 400 requiring weight for R. A is also permissible over R+, because 500 justifying weight for A also outweighs 405 requiring weight to respect both rights.

SI4 Explained: R is permissible over A, because 400 justifying weight for R outweighs 250 requiring weight for A. R is permissible over A+, because 400 justifying weight for R outweighs 300 requiring weight for A+.

Weight Pluralism, then, has no problem explaining the stability of moral options involving altruistic and respectful reasons. Alas, there is trouble brewing.

6.2. *Comparativism and Pro Tanto Maximization*

At first glance, my explanation of stable moral options seems to rule out pro tanto maximization. When we face a choice between saving five and saving six, we are required to save six and so it is *impermissible* to save five. Yet Pluralist Permissibility seems to entail that it is permissible to save five, because 500 justifying weight to save five is greater than 300 requiring weight to save six (cf. Muñoz 2021: 707 and Cullity 2018: 431n 9). I guess I should give up and go home.

Or we can accept the following plausible claim and solve all our problems.

Comparativism about Reasons: a reason for φ is a way that φ is better than some specified alternative.

Comparativism requires us to think carefully about what the reasons are in the cases that we've been considering. Your altruistic reasons to A+ are not determined by the total number of lives you save in A+; they are determined by the *additional* lives you save in A+ that you don't save in the alternative. The reason to choose A+ over A is not that you would save six people in A+. Rather, it is the way that A+ is better than A: that you would save one additional person in A+ that you don't save in A. Given our assumptions about how to weight altruistic reasons (§6.1), the reason to choose A+ over A has 100 justifying weight and 50 requiring weight.

Crucially, since *the ways that φ is better* can vary as you vary the alternative, comparativism implies that *your reason for φ* can vary as you vary the alternative. In the choice between A and R, you have a very weighty reason to choose A, namely that you save five additional lives in A that you don't save in R. We assumed that this reason has 500 justifying weight and 250 requiring weight (§6.1). Yet there is *no* reason to choose A over A+, because there is no way that A is better than A+. You don't save anyone in A that you don't also save in A+. Hence, there is 0 justifying weight and 0 requiring weight to choose A over A+.

According to comparativism, your respectful reasons to R+ are not determined by the total number and significance of the rights you respect in R+; they are determined by the number and significance of the *additional* rights you respect in R+ that you don't respect in the alternative. The reason to choose R+ over R is not that you would respect *both* the important right and the small additional right. It is the way that R+ is better than R: that

you would respect a small additional right in R+ that you don't respect in R. This reason, we can assume, has 5 justifying weight and 5 requiring weight to choose R+ over R.

Reasons of respect also can vary as you vary the alternative. You have a very weighty reason to choose R over A, namely the important right you respect in R that you don't respect in A. We assumed that this reason has 400 justifying weight and 400 requiring weight (§6.1). Yet there is *no* reason to choose R over R+. You don't respect any right in R that you don't also respect in R+. Hence, there is 0 justifying weight and 0 requiring weight to choose R over R+.

We now have everything we need to give a Weight Pluralist explanation of pro tanto maximization.

SI1 Explained: It is impermissible to choose A over A+, because no reason's 0 justifying weight for A *is outweighed by* one additional life's 50 requiring weight for A+.

SI2 Explained: It is impermissible to choose R over R+, because no reason's 0 justifying weight for A *is outweighed by* the additional small right's 5 requiring weight for A+.

Let's recap the Weight Pluralist explanation of the Normative Significance of Small Improvements. Comparativism did the heavy lifting in the explanation of pro tanto maximization. When we aren't trading morally relevant considerations off against one another, then you are always required to choose the small improvement (A+ over A, and R+ over R). That's because there is a way that the small improvement is better and no way that the unimproved option is better. This relationship will generate a requirement to choose the small improvement as long as the reason to choose the small improvement has at least some requiring weight.

Pluralist Permissibility did the heavy lifting in our explanation of how these pro tanto maximizing reasons can generate stable moral options. When we have a moral option to trade off two morally relevant considerations (e.g., A vs R), we can improve one of the reasons and the moral option remains. Pluralist Permissibility tells us that it will be permissible to choose, say, altruistic A as long as its justifying weight is not outweighed by the requiring weight of the respectful alternative. Since altruistic reasons have more justifying than requiring weight, you can improve the respectful reason quite a bit before you convert the moral option into a requirement to respect rights.

Comparativism also does the heavy lifting in explaining why my explanation of pro tanto maximization is compatible with my explanation of stable moral options. Comparativism implies that the reason for the unimproved options (A and R) can vary as you vary the alternative. This variance allows A to be permissible when it competes with R, but not when it competes with A+.

Conclusion

Simple Weight Monism was too simple to explain the Normative Significance of Small Improvements. We must pay the price of complication. I focused on two potential complications: follow Parity Monism and add a new comparative, parity, or follow Weight Pluralism and a second weight value. We saw that Parity Monism was able to explain the small improvement puzzle, but it could not explain supererogation or permissible partiality.⁷ Weight Pluralism was able to explain all three, and it provides the only known mechanism for explaining all three. Parity is no substitute for pluralism.⁸

Works Cited

- Archer, Alfred. 2016. "Moral Obligation, Self-Interest, and the Transitivity Problem." *Utilitas* 28: 441-464. doi: doi.org/10.1017/S0953820816000091
- Chang, R. (2017). Hard Choices. *Journal of the American Philosophical Association*, 3, 1-21. doi: doi.org/10.1017/apa.2017.7
- Cullity, G. (2018). Weighing Reasons. Daniel Star (Ed.), *The Oxford Handbook of Reasons and Normativity* (pp. 423-42). Oxford: Oxford University Press.
- Dorr, Cian, Jacob M. Nebel and Jake Zuehl. Forthcoming. "The Case for Comparability." *Nous*. doi: doi.org/10.1111/nous.12407
- Gert, Josh. 2004. *Brute Rationality: Normativity and Human Action*. Oxford University Press.
- Givewell. (2020). GiveWell cost-effectiveness analysis — version 1. https://docs.google.com/spreadsheets/d/1BmFwVYeGMkpys6hG0tnfHyq_ZFZf-bmXYLSHODGpLY/edit#gid=1034883018
- 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. doi: <https://doi.org/10.1017/S026505250999015X>
- 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.
- Muñoz, Daniel. (2021). Three Paradoxes of Supererogation. *Noûs*, 55, 699–716. doi: doi.org/10.1111/nous.12326
- Schoenfield, Miriam. "Decision Making in the Face of Parity." *Philosophical Perspectives* 2014: 263-77.
- Tucker, Chris. Manuscript. *The Weight of Reasons: A Framework for Ethics*.
_____. Forthcoming. "Parity, Moral Options, and the Weights of Reasons". *Nous*. doi: doi.org/10.1111/nous.12410

¹ The parenthetical caveat is needed in case there are other conditions on supererogation that aren't discussed in this paper. For example, some argue that feeding five is supererogatory only if you perform that action for the altruistic reasons that make it better and not merely because, say, you'll get your name in the paper.

² With the exception of §5.2, the ideas in this chapter were first published in my very long and technical "Parity, Moral Options, and the Weights of Reasons", forthcoming in *Nous*. This chapter's contribution is that it is more accessible and focuses more on partiality.

³ At least, it is supererogatory provided that any other conditions on supererogation are also satisfied. (See note 1 for an example of an alleged further condition.)

⁴ \$17k was roughly the going rate to save five people from malaria, as of 2020 (Givewell).

⁵ In my ms: §6.3, I show that a version of Weight Monism that appeals to vagueness runs into the same problems as the version that appeals to parity.

⁶ This is basically the view in Daniel Muñoz, "Three Paradoxes of Supererogation" *Nous* (2021): 55, 699–716, especially 702, 712-3ns 1 and 5. My official view appeals to a third weight value. See my manuscript *The Weight of Reasons: A Framework for Ethics*, ch. 9. Other three value views include Terry Horgan and Mark Timmons, "Untying a Knot from the Inside Out: Reflections on the 'Paradox' of Supererogation," *Social Philosophy and Policy* 27 (2010): 29 – 63 (especially §IV.B); Alfred Archer. "Moral Obligation, Self-Interest, and the Transitivity Problem," *Utilitas* 28 (2016): 441-64; and Margaret Olivia Little & Coleen Macnamara, "Non-Requiring Reasons," in *The Routledge Handbook of Practical Reason*, eds. Ruth Chang and Kurt Sylvan, (Routledge, 2020), 393-404.

⁷ The existence of parity also faces additional difficulties that the distinction between justifying and requiring weight does not face. Cian Dorr, Jacob M. Nebel & Jake Zuehl, “The Case for Comparability,” forthcoming in *Nous* (see especially §4) show that parity conflicts with apparent conceptual truths about comparatives. Miriam Schoenfield, “Decision Making in the Face of Parity,” *Philosophical Perspectives* (2014): 263-77 shows that parity conflicts with standard decision theory. See my *The Weight of Reasons*, §4.4.2, for an explanation of how to make decision theory compatible with the two weight values.

⁸ I received helpful feedback on this chapter from Preston Budd, Eric Silverman, and the audience of the 2022 *Virtuous and Vicious Partiality* conference that led to this volume. Thanks!