

The Dual Scale Model of Weighing Reasons

Chris Tucker 12/17/20
William & Mary cstucker@wm.edu

(Forthcoming in *Nous*)

0. Introduction

Ethical theory is steeped in two related metaphors: *reasons have weight* and *reasons are weighed on a balance scale to determine an act's deontic status* (e.g., whether an act is permissible, impermissible, or required). Such metaphors are implicit in the ubiquitous talk of *outweighing* and the *balance of reasons*. The ubiquity of these metaphors gives them importance. If we are going to use *or reject* them, we should understand the best way of cashing them out. We may find, as I argue in this paper, that we can construct a promising model of how reasons interact to determine deontic status simply by cashing them out carefully.

I assume that the metaphors of weight and weighing are apt for morality and practical rationality. So understood, substantive moral (rational) theories can make two kinds of mistake. First, a theory might incorrectly *weight* reasons. Perhaps, for example, it overestimates how much weight morality gives to self-interest. Second, a substantive theory might *weigh* reasons incorrectly. In principle, a theory might make one kind of mistake without making the other.

This paper is after **the fundamental model of weighing reasons** to determine a deontic status, the model that determines what it is to weigh reasons correctly. Such a model can represent any normative theory that correctly weighs reasons, even if the theory incorrectly weighs them. It can't represent any theory that incorrectly weighs reasons. By identifying what it is to weigh reasons correctly, the fundamental model identifies structural constraints on any substantive moral or rational theory.

The metaphor of weighing reasons brings to mind a *single (double-pan balance) scale*. The reasons for φ (R_φ) go in one pan and the reasons for $\sim\varphi$ ($R_{\sim\varphi}$) go in the other. The relative weights, as indicated by the relative heights of the two sides of the scale, determine the deontic status of the act.




fig. 1

This model, *Single Scale*, is not by itself a complete normative theory. Among other things, it must be combined with a function that assigns relative weights to deontic status. A natural assignment is:

Max: you are *permitted* to ϕ iff your reasons for ϕ are at least as weighty as your reasons for $\sim\phi$; you are *required* to ϕ iff your reasons for ϕ are weightier than your reasons for $\sim\phi$; otherwise, ϕ is *impermissible*.

Alternatively, one might allow an act to be permissible as long as the reasons for ϕ are “good enough,” even if they are, within limits, outweighed by the reasons for $\sim\phi$.

Another common metaphor is the single vector sum, where the deontic status of an act is determined by comparing the total force (aka pressure) in one direction (toward permissibility) with the total force in the opposite direction (toward impermissibility). There is no competition between these metaphors. For normative purposes, talk of weight, force, and pressure are rightly taken to be interchangeable. You put an object on the ϕ -pan. Its weight just is the downward force it exerts on the ϕ -pan and the upward force that it exerts on the $\sim\phi$ -pan. In the physical world, forces can be exerted on an object in any direction: up, down, left, right, backwards, forwards, and everything in between. The balance scale is just a convenient way to focus our attention on two opposing directions: down (toward permissibility) and up (toward impermissibility).¹

When understood as an attempt to capture the fundamental model of weighing reasons, *Single Scale* is simple and intuitive. It can also represent a wider range of normative theories than you might think. Nonetheless, *Single Scale* has a serious representational limitation. As a reason pushes ϕ down toward permissibility (has justifying weight for ϕ), it pushes the alternative up toward impermissibility (has requiring weight for ϕ). In other words, *Single Scale* assumes *Single Proportion* (first pass): all reasons have the same proportion of justifying and requiring weight. We’ll see that this assumption prevents *Single Scale* from being able to represent certain normative theories, such as the existence of dilemmas, the standard account of supererogation, and certain satisficing theories.

You might think it is a plus for *Single Scale* that it can’t represent such theories; however, it is a problem insofar as *Single Scale* is an attempt to capture the fundamental model of weighing reasons. The problem is not (merely) that one or more of the theories is true. Even if they are all false, *Single Scale* still misconstrues why they are false. It claims that they misunderstand what it is to weigh reasons in order to determine a deontic verdict. Yet if these theories are false, their mistake is about the *weights* that certain reasons have, not how to *weigh* them. If *Single Scale* isn’t the fundamental model, then its use must be justified. It must be shown that, given its relation to the fundamental model, *Single Scale* won’t distort the normative phenomena it is used to model.

Single Scale fails to capture the fundamental model of weighing reasons, not because it appeals to the image of a scale, but because it appeals to the image of a *single* scale. We need *two* scales. The first scale determines whether an act is permissible by

¹ ‘Strength’ is also used interchangeably with ‘weight’ and ‘force’. Strength, though, does have one connotation that I want to ignore in this paper. I can have the strength to lift the box (put enough upward force on the box to lift it off the floor) even if I don’t exert that strength on the box. This connotation is useful in contexts where we are trying to distinguish between reasons that you possess and reasons that you don’t possess (forthcoming, §3). When the weight of a reason is on a pan, there is no question as to whether it exerts force on the pan or whether it is relevant to an action’s deontic status.

comparing the justifying weight for an act and the requiring weight for the alternative. The second scale determines whether an act is a commitment by comparing the requiring weight for the act and the justifying weight for the alternative. (ϕ is a **commitment** iff $\sim\phi$ is *impermissible*.) The two scales working together determine whether ϕ is a requirement or a dilemma. (ϕ is a **requirement** iff ϕ is both permissible and a commitment. ϕ is a **dilemma** iff ϕ is both *impermissible* and a commitment.)

This model, *Dual Scale*, can represent justifying and requiring weight as dependent variables. This is just another way of saying that it can represent normative theories which assume that Single Proportion is true. Values that always come in the same proportion don't need to be tracked independently: if you know the value of one, then you are able to calculate the value of the other. Thus, Dual Scale can represent everything Single Scale can represent—and arguably represent it better. It can also represent justifying and requiring weight as *independent* variables. This is just another way of saying that it can represent normative theories that reject Single Proportion. When values don't always come in the same proportion, you must track them independently insofar as you can't calculate the value of one from the value of the other. Thus, it can represent what Single Scale can't.

As I present it here, Dual Scale inherits three limitations from Single Scale. First, it focuses on a simple choice context in which there are just two options, ϕ and $\sim\phi$. Second, it assumes that every reason against ϕ is some sort of reason for $\sim\phi$, which some regard as a problematic assumption (Snedegar 2018). Finally, it is susceptible to certain holist challenges. I will identify these limitations as they become relevant and, usually in footnotes, briefly explain how Dual Scale needs to be refined to resolve them. These refinements would vindicate scale-based models from the most pressing objections; however, full defense of these refinements, the metaphor of a reason's weight, and the metaphor of weighing reasons on scales must wait for another occasion.

The simple version of Dual Scale presented here nonetheless *approximates* the fundamental model of weighing reasons. As an approximation, it helps us to better understand the structural constraints on any theory of morality, rationality, and any other normative perspective that is structured by the metaphors of weight and weighing.

In §1, I present Single Scale. In §2, I refine Gert's justifying/requiring weight distinction and show that it is unavoidable as long as we assume that reasons have weight. Dual Scale and my criticism of Single Scale, in a way, just systematically develop and apply this refined distinction.² In §3, I explain *what* Single Scale can represent and *how* it can represent it. This explanation will clarify Single Proportion, the component that prevents Single Scale from representing even more. In §4, I show that the commitment to Single Proportion prevents Single Scale from representing a range of normative theories. In §5, I explain why these representational failures reveal that Single Scale isn't the fundamental model. In §6, I present Dual Scale. In §7, I highlight Dual Scale's advantages and show that it can represent the three theories that Single Scale can't represent. It has a lot in common with Single Scale, but it has a much better claim to being the fundamental model of weighing reasons.

² Gert generally uses the distinction to make trouble for views he rejects, including something like Single Proportion (2007: 536-7). Yet he never systematically develops the distinction into a model of weighing reasons (see 2004: 73-79 for his most explicit discussion), and some of his claims about the distinction are incompatible with my model (§2.1). For other differences between our views, see nt 26, and nt 43.

When our best image/metaphor for weighing reasons is Single Scale, Single Proportion seems like a default assumption of normative inquiry, if not a constraint on the coherent weighing of reasons. In §8, I argue that unjustified reliance on Single Scale has made it difficult for philosophers to even make sense of theories that reject Single Proportion. Dual Scale shows decisively that we *can* make sense of such theories.³

1. Single Scale

Single Scale is the conjunction of four theses. These theses are implications of taking the single scale metaphor seriously. The first implication begins with the observation that the options, ϕ and $\sim\phi$, are represented by the scale itself, more specifically, by the two pans. In fig. 1 above, the left pan is ϕ and the right pan is $\sim\phi$. Since one's reasons don't influence the shape or structure of the pans, the scale (the individuated options) is prior to, and so independent of, what goes on the scale (reasons):

Independent Individuation of Options: an agent's options are individuated independently of what reasons the agent has and what the weights of those reasons are.⁴

Independent Individuation of Options is a constraint on how options are to be individuated. This constraint has important implications,⁵ but a complete theory of option individuation is itself not part of Single Scale.

Any plausible theory of option individuation will allow one to be in a situation with three or more options. I focus, however, on the case in which we have two options, ϕ and $\sim\phi$. Generalizing Single (and Dual) Scale to handle any number of options raises complications that I won't be able to address here.⁶

The second implication restricts what goes on the scale. To whatever extent something counts in favor of ϕ (pushes down the ϕ -pan), it also counts against $\sim\phi$ (pushes up the $\sim\phi$ -pan). In §2.2, we'll discuss this implication under the label *Single Proportion*. The model is otherwise largely noncommittal about what goes on the scale. A *natural assumption* is that all and only reasons for ϕ go on the ϕ -pan and all and only reasons for $\sim\phi$ go on the $\sim\phi$ -pan. We'll refine this assumption in light of a certain alternative to Max (§4), but it'll do for now.

The third implication restricts which factors make a difference to an act's deontic status. Deontic status is a function of the relative heights of the pans, and so:

Weight Comparativism: An action's deontic status is determined solely by the relative weights of all reasons for and against it.⁷

³ Despite what Portmore (2011: 137-143) may think, his dual ranking theory and principle META do not show that we can make sense of the standard account of supererogation, at least not without the help of Dual Scale. They are meta-principles linking two normative perspectives, morality and rationality. Those principles *presuppose* that, within the single normative perspective of morality, there is a model of weighing reasons that can make sense of the standard account of supererogation. It is good news for Portmore that he is unwittingly committed to Dual Scale (see nt 38 below).

⁴ Cf. Snedegar (2018: 726).

⁵ For example, it blocks Rubio's (2018) Satan's Apple Paradox, which depends on using which reasons a divine being has as a constraint on what can be an option for a divine agent.

⁶ See nt 39 for a brief discussion.

⁷ Weight Comparativism is similar to what Chang (2016: 215) calls "Comparativism (reasons version)." Cf. Portmore (2019: 8-11).

Weight Comparativism claims that the weights of the relevant reasons *and only these weights* determine the deontic status of a particular action. And these weights don't determine the deontic status in any old way: it is the *relative weight* of the reasons for and against an action that determine its deontic status. Different versions of Single Scale will endorse different assignments of relative weights to deontic status. Max is a particularly demanding assignment, but we'll soon consider some less demanding alternatives.

The fourth implication concerns the analogy with *weight*. Things with weight exhibit counterfactual stability across contexts. If my phone is heavier than my pen, then whatever scale you put them on, in whatever room you're in, the phone will exert more downward force than my pen. If reasons have counterfactual stability, then it may seem that they have the same weight—they exert the same force in the same valence—across all contexts. These fixed weights are determined prior to determining the deontic status of any particular action, because reasons have a given weight *before* you put them on the scale. The fourth implication, then, is:

Prior Fixed Weight: the same reason has the same weight in all contexts, and what particular weight it has is prior to the deontic status of any particular action. If something is a reason, then it has some weight or another. What particular weight it has includes its valence (weight toward permissibility vs weight toward *impermissibility*) as well as just how weighty it is for (against) an act. Thus, Prior Fixed Weight entails that if something is a reason with a specific weight, it is always a reason with that weight. If it doesn't have a weight, then it isn't a reason in any context.⁸

If Prior Fixed Weight is true, it holds only for some privileged way of individuating reasons. Single Scale does not itself take a stand on exactly what this privileged way is, but it likely involves at least a restriction to fundamental, non-derived, or basic reasons (Gert 2004: 77-79; Nair 2016: 94-5).⁹

There is a fifth implication of the weighing metaphor, but I don't build it into Single Scale. The weight of everything on a pan just is the sum total weight of each individual thing on the pan. By analogy, we get additive aggregation: the weight of all the reasons for ϕ just is the sum total weight of the individual reasons for ϕ .¹⁰ There are many worries about additive aggregation,¹¹ and I don't want these worries to distract you. The

⁸ Dancy (2004: 105, 190); Gert (2004: 73-77, 2007); cf. Berker (2007, especially 122-4). Elsewhere I argue that if holism is compatible with Single (Dual) Scale, the counterfactual stability of weight must be compatible with the idea that a reason's weight can vary according to context (cf. Cullity 2018: 425-7). This compatibility would require some refinement to Prior Fixed Weight.

⁹ That a given restaurant exclusively serves Thai food might be a reason for me to go and a reason for you to refrain from going. At first glance, this is a violation of Prior Fixed Weights since the valence of *exclusively serving Thai food* seems to vary according to who the diner is. Yet *exclusively serving Thai food* probably won't make our list of fundamental or basic reasons. It is a derived reason, i.e., it is a reason only insofar as it bears a certain relation to a fundamental reason, such as its promoting someone's wellbeing or pleasure. Even if Prior Fixed Weights applies to basic reasons, we should not expect it to hold of derived reasons precisely because the same thing (Thai food) can bear different causal relations (promote/prevent) to a given basic reason (pleasure).

¹⁰ See, e.g., Dancy 2004: 9-10, 190 and Snedegar 2018: 725-743; cf. Berker 2007: 122-3.

¹¹ A common worry about additive aggregation is that it will double-count some reasons. Consider, for example, the reasons *it would cause pain* and *it would cause severe pain* (cf. Kearns' 2016: 186; Maguire and Snedegar forthcoming: 6 EV). This seems to be a counterexample to additive aggregation, but it's not. Elsewhere I argue that it's a counterexample to the claim that the correct aggregation principle—whether additive or not—takes all reasons as inputs. Although they don't put the point this way, I take it that this

other four implications can survive even if additive aggregation does not. Just keep in mind that Single Scale needs to be combined with whatever the true aggregation principle happens to be.

Single Scale is the conjunction of Independent Individuation of Options, Single Proportion, Weight Comparativism, and Prior Fixed Weights. There are well-known holist challenges to Single Scale. If conditions (enablers, disablers) or modifiers (amplifiers, attenuators) exist, then a reason's weight is partly determined by its context. Single Scale denies this context-sensitivity, as will Dual Scale. Single (Dual) Scale can and should be refined to allow for this context-sensitivity, but I save these refinements for another occasion.¹² (A holist version of Single Scale is still committed to Single Proportion, roughly, that justifying and requiring weight always come in the same proportion. So it too suffers from the problem I develop for the non-holist version.)

2. Justifying and Requiring Weight

To understand Single Scale's problem and Dual Scale's promise, we need a nuanced understanding of the justifying/requiring weight distinction.

2.1. Justifying and Requiring Weight as Opposing Forces

A consideration has **justifying weight** for ϕ iff the consideration makes it permissible to ϕ in the absence of sufficiently weighty countervailing considerations (more precisely: sufficiently weighty requiring weight for $\sim\phi$). The fact that eating this cookie will make me happy has justifying weight insofar as it is a consideration that pushes *eating the cookie* toward permissibility. A consideration has **requiring weight** for ϕ iff it makes $\sim\phi$ impermissible in the absence of sufficiently weighty countervailing considerations (more precisely: sufficiently weighty justifying weight for $\sim\phi$). If stopping the car is the only way to avoid killing one of the children playing in the street, then you have a reason for stopping that has requiring weight, a reason that pushes the alternative (not stopping) toward impermissibility.¹³

You'll notice that 'justifying weight' and 'requiring weight' are inter-defined in typical functionalist fashion (cf. Dorsey 2016: 8-12 and especially Gert 2004: ch 4). When concepts are a package deal, there is no problem with inter-defining them. Consider necessity and possibility. Conceptually, they seem to be a package deal and so

is the upshot of Gert (2004: 77-79); Bader (2016: 29); Kearns (2016: 186); Nair (2016: 63-70, 94-5); Cullity (2018: 424-5); and Maguire and Snedegar (forthcoming).

¹² These refinements tweak Prior Fixed Weight and then clarify that whenever Single (Dual) Scale uses the term "weight", it means the *embedded* weight of a reason. A reason's (better: a ground's) *default* weight is the weight it has abstracted from all contexts. A reason's (better: a ground's) *embedded* weight is the weight it has when conjoined with a specific set of conditions and modifiers (cf. Bader 2016: 40). Deontic statuses are determined most directly by the embedded weights of reasons, so that's the weight value that we consider when we put reasons on the scale. Elsewhere I argue that these refinements are enough to resolve the most serious holist challenges to Single (Dual) Scale.

¹³ Justifying/requiring weight *against* ϕ is conceptually equivalent to justifying/requiring weight for $\sim\phi$. To go back and forth between conceptually equivalent categories of weight, hold fixed whether the weight is justifying/requiring and then vary *both* the for/against and the $\phi/\sim\phi$. For example, justifying weight against ϕ is conceptually equivalent to justifying weight for $\sim\phi$: they both push ϕ toward permissibility. See my manuscript for an account of the for/against distinction and an explanation of how Dual Scale handles it. If you are inclined to reject the justifying/requiring weight distinction because you think the reason for/against distinction is all we need, see nt 36 below.

they can be inter-defined. Some proposition P is necessary iff it is not possible that $\sim P$, and P is possible iff it is not necessary that $\sim P$. I am claiming that justifying and requiring weight are likewise a package deal. There is a formal methodology for inter-defining terms without vicious circularity, and it is straightforward to restate my definitions using that methodology.¹⁴

To have justifying weight for ϕ is just to play a certain role in fixing deontic status, i.e., pushing ϕ toward permissibility. To have requiring weight for $\sim\phi$ is to play an *opposing* role in fixing deontic status, i.e., pushing ϕ toward *impermissibility*. Hence, contrary to what Gert himself asserts (2007: 536), justifying and requiring weight are opposing weights/forces to be understood on analogy with downward and upward force. Just as the direction of the particle's movement (or the height of a balance scale pan) is determined by a competition between the downward and upward forces exerted on it, the permissibility of a particular action is determined by a competition between justifying weight for ϕ and requiring weight for $\sim\phi$. Something's *amount* (quantity, magnitude) of justifying and/or requiring weight is a way of specifying which role wins the competition in a given case.

Gert sometimes assumes that, if justifying and requiring weight were independent variables, then *amounts* of justifying and requiring weight would be incomparable (2007: 537). This is a mistake. Downward and upward forces are distinct but comparable: the downward force can be greater than, less than, or equal to upward force. Likewise, justifying weight and requiring weight are distinct but comparable: justifying weight for ϕ can be greater than, less than, or equal to requiring weight for $\sim\phi$. This comparability is essential for any view of weighing reasons that assigns deontic status in terms of *relative* weights. I'll have more to say about amounts in §5 and the competition between justifying and requiring weight in §6.

Following Portmore (2011: 88-9) and Dorsey (2016: 166), I use a thick notion of requirement: ϕ is a **requirement** iff ϕ is both permissible and $\sim\phi$ is impermissible. I use a different term for Gert's thinner, purely negative notion of requirement: ϕ is a **commitment** iff $\sim\phi$ is *impermissible*. Commitment to ϕ is what you add to permission to ϕ in order to make ϕ required. Nonetheless—and in contrast to Dorsey (2016: 166)—I follow Gert in using a thin sense of 'requiring weight' that does *not* entail justifying weight. This lack of entailment is necessary to represent the conceptual possibility of dilemmas when an agent has finitely many options (§4.1). So understood, an undefeated requiring reason to ϕ does not, by definition, entail a requirement to ϕ (both ϕ is permissible and $\sim\phi$ is impermissible) but only a commitment to ϕ ($\sim\phi$ is impermissible). Were Gert's terminology not ingrained in the literature, I would have avoided this infelicity by replacing the term "requiring weight" with "committing weight".

2.2. *The Inescapability of the Justifying/Requiring Distinction*

¹⁴ Lewis (1970) explains how to (inter)define terms by their functional relations to each other. If you are familiar with his framework, you can let 'permissible' (ok to do) and 'impermissible' (not ok to do) be the "Old" terms, defined independently of the theory/functional role. Justifying weight and requiring weight are the "Theoretical" terms defined in terms of their relations to the Old terms and each other. (Were I making Dual Scale compatible with holism, I would also add 'enabler', 'disabler', and other holist notions as Theoretical terms.)

Given the assumption that reasons have weight, the conceptual distinction between justifying and requiring weight is inescapable and should be uncontroversial.¹⁵ We all agree—or at least I assume—that some actions are permissible and that some are required. We also all agree—or at least I assume—that it isn't brute that an action is permissible or required. Actions are made permissible (required) because they are supported by reasons that make them permissible (required). Yet, once you admit the existence of reasons that make an act permissible, you thereby admit that those reasons have justifying weight (push the act toward permissibility). When you admit the existence of reasons that make an act required (both permissible and a commitment), you thereby admit that those reasons have requiring weight (push the act's alternatives toward impermissibility).

If the conceptual distinction should be as uncontroversial as I make it out to be, then you may wonder why the distinction is not more widely understood and appreciated. Well, many normative theorists *begin* normative theorizing by assuming that justifying and requiring weight come in the same proportion. The assumption is so deep that many don't even recognize it as such. (I ultimately blame Single Scale for this predicament in §8.) If justifying and requiring weight always come in the same proportion—if they are dependent variables—then we don't need to track them independently. Perhaps, for example, a single factor plays both roles: the greater the single factor (e.g., the greater the aggregate net wellbeing), the greater the justifying and requiring weight. As long as we track this single factor, we don't need to think of ourselves as tracking justifying and requiring weight to describe normative reality. We can think of ourselves as tracking something more generic, such as the single factor itself, the (unqualified) pressure for/against an option, or the (unqualified) weight of the reasons for/against it.

In other words, insofar as we assume that justifying and requiring weight always come in the same proportion, we thereby assume that the distinction is insignificant for substantive normative theorizing. Yet we do not thereby assume it out of existence. Nor do we thereby make the distinction irrelevant to the fundamental model of weighing reasons. *Polygon with exactly three sides* and *polygon with exactly three internal angles* pick out conceptually distinct features of all and only triangles. Likewise, justifying and requiring weight pick out conceptually distinct roles that a reason can play—even if all and only reasons play these roles and even if the two roles always co-vary. As I illustrate in §5, there are conceptual truths about how these roles interact, and these conceptual truths provide structural constraints on substantive normative theorizing.

While the justifying/requiring distinction itself should be uncontroversial, it is controversial *which* considerations have some sort of weight. The maximizing act utilitarian says that a stranger's wellbeing has both justifying and requiring weight and the ethical egoist says that it has neither. A different controversial issue—and one at the heart of this paper—is which proportions of justifying and requiring weight that a given reason can have. In principle, two normative theories can agree about which things have

¹⁵ Some of the more fine-grained details are, however, up for negotiation. For example, my definitions entail that any amount, no matter how little, of undefeated requiring weight for ϕ makes ϕ a commitment. One might opt instead for a notion of requiring weight that demands a certain amount n of undefeated requiring weight for ϕ before ϕ becomes a commitment. But you'll still need the functionalism to ensure that the concepts—'(amounts of) justifying weight', '(amounts of) requiring weight'—are properly coordinated and neutral between rival (conceptually possible) accounts of what plays the role picked out by each concept.

weight and disagree about what proportion of justifying and requiring weight each reason has.

3. The Limits of What Single Scale can Represent

In this section, we consider the full range of *what* Single Scale can represent and *how* it can represent what it can represent. This consideration will clarify Single Proportion, the component of Single Scale that prevents it from representing even more.

Single Scale is limited in which proportions of justifying and requiring weight that it can represent, because the ϕ -pan goes down (toward permissibility) only insofar as the $\sim\phi$ -pan goes up (toward *impermissibility*). Single Scale may seem to entail, then, that a reason must have *exactly* the same amount of justifying weight (weight toward ϕ 's permissibility) and requiring weight (weight toward $\sim\phi$'s *impermissibility*). In other words, it may seem that Single Scale can represent only **balanced requiring reasons**, i.e., reasons that have exactly as much requiring weight as they have justifying weight. It isn't that limited, however.

Some elements of Single Scale are represented visually or "on the image of the scale". The pans represent the options, for example. Yet some elements are "off-image". The assignment of relative weights to deontic status is not represented visually, though the visual elements do rule out some assignments (see §5 below).¹⁶ Max is one assignment function compatible with Single Scale. It holds that an act is permissible iff it is an act for which one has the greatest (or tied for the greatest) balance of reasons. Non-maximizing, or satisficing, versions of Single Scale hold that an act is permissible as long as there is a *good enough* balance of reasons in its favor.

For Single Scale, the assignment of relative weights to deontic status is tightly correlated with which kind(s) of reasons are represented. (The parallel point won't hold for Dual Scale, which has richer representational resources.) For example, Max requires that all reasons that go on the pans be balanced requiring reasons, that all such reasons push an act toward permissibility exactly to the extent that they push the alternatives toward *impermissibility*. It is this feature of Max that ensures that one is not permitted to ϕ whenever the reasons against it outweigh (even by the smallest margin) the reasons for it. Single Scale can also represent two other kinds of reasons and is, therefore, compatible with two importantly different kinds of satisficing.

Consider a view that holds that reasons are always **justifying heavy requiring reasons**, i.e., reasons with both justifying and requiring weight but more of the former than the latter. For illustration, let's say they have twice as much justifying weight as requiring weight. Single Scale can represent such a view using:

Relative Weight Satisficing: ϕ is permissible iff the reasons against ϕ are no more than twice as heavy as the reasons for ϕ .

On the scale, there will be some height differential that represents the reasons against ϕ as being twice as heavy as the reasons for ϕ . Relative Weight Satisficing holds that this "twice as heavy point", not the counterbalanced point, is the threshold for permissibility. So Single Scale can represent justifying heavy requiring reasons if they *always* have the same proportion of justifying and requiring weight.

¹⁶ Conditions and modifiers will also be represented off-image, if the response to holist challenges I suggest in note 12 is on the right track.

Single Scale can even model a *special* case of reasons that have different proportions of justifying and requiring weight. Perhaps an act comes with a “default cushion of permissibility”. That is, perhaps there is a default reason (a reason to choose any action in any situation) and this default reason is **merely justifying**—it has justifying but not requiring weight. Assume the default justifying weight is five units.

Recall that the image of the scale doesn’t itself say that *all* reasons go on the scale (§1) and that some elements of a complete model must be represented off-image. We might suggest, then, that the *default* reason to choose an action is represented off-image in the assignment of relative weights to deontic status. What goes on the scale are all and only relevant *non*-default reasons, which we can suppose are balanced requiring reasons. We can represent this combination of merely justifying default reasons and non-default balanced requiring reasons using:

Default Reason Satisficing: it is permissible to ϕ iff the reasons for ϕ are not outweighed by at least 5 units.

Thus, Single Scale can allow default reasons to have a different proportion of justifying and requiring weight than non-default reasons.

Ethicists don’t normally talk about default reasons or default biases toward permissibility, but one way to motivate prerogatives might be recast in these terms. Some motivate prerogatives by appealing to partiality or relationships (I have the prerogative to give *my* wellbeing extra weight). Set that idea aside. Others motivate it by appealing to agential autonomy: all moral agents are autonomous and the value of their autonomy explains why they can self-direct even if their preferred direction is, within limits, outweighed by opposing reasons. The “within limits” can be given a default reason reading. If you think of reasons as facts, the default reason of five units might be *that the agent is autonomous* or that *ϕ is an available act*.¹⁷

Relative Weight Satisficing and Default Reason Satisficing are less demanding than Max because they allow for a cushion of permissibility, i.e., they allow the total justifying weight to outstrip the total requiring weight in favor of an act. The bigger this cushion, the more the balance of reasons can be against the act and it still be permissible. Relative Weight Satisficing represents a *proportional cushion*: the more total justifying weight, the bigger the cushion. If the total justifying weight is 10, then the total requiring weight is 5 (a difference of 5). If the total justifying weight is 100, the total requiring weight is 50 (a difference of 50). Default Reason Satisficing represents a *fixed cushion*: the difference is always 5. No matter what the total justifying weight is, the total requiring

¹⁷ It is not crucial that we think of default biases as default *reasons*. I think of them that way because I hold that: (i) a default bias toward permissibility makes an act permissible in the absence of sufficiently weighty requiring weight for the alternative; (ii) by the definition of justifying weight for an act (§2.1), it follows that default biases have justifying weight for an act; and (iii) anything that has justifying weight for an act is a reason for that act. All of (i)-(iii) are negotiable (recall, for example, nt 15). What is non-negotiable for the fundamental model is that it takes default biases into account; that it is able to represent disagreements about the size and valence of the bias; and that the size of the default bias is measurable by how much opposing weight it takes to overcome it.

weight is 5 units less.¹⁸ Thus, Single Scale can represent two importantly different kinds of satisficing (and three if you count their combination).¹⁹

If we can represent default reasons “off-image”, one might wonder whether any sort of reason can be represented off-image and so wonder whether Single Proportion has any representational limitations at all. In reply, default reasons really are a special case.

Consider two features of combining Single Scale with Max. The first feature is that *all variation is on-image*. Regardless of which action and situation we consider, the only variation is which reasons go on the pans. Default Reason Satisficing captures this feature too. The only reasons that are represented off-image are default reasons which apply to every act in every situation. Consequently, all variation between different actions and situations is captured on-image by considering which (non-default) reasons go on the scales. Once you start to move variations in the reasons off-image, the less the image of a single scale is being used as a model for weighing reasons.

Second, Max is *biased toward permissibility*. When the reasons for and against are counterbalanced or there are no reasons either way, the act is permissible. Not impermissible, not undefined. The default reason picture agrees that there is a bias toward permissibility, but it objects that the bias is 5 units weightier than Max recognizes.

While default reasons might sound exotic, it is to Single Scale’s credit that it can represent them. We can intelligibly disagree about both whether there is a default bias toward permissibility²⁰ and how weighty that bias is. The fundamental model of weighing reasons should be able to represent such disagreements.

As one pan goes down, the other pan goes up. At first glance, this feature suggests that all reasons must be balanced requiring reasons, reasons that have justifying weight for ϕ exactly to the extent that they have requiring weight for ϕ . We’ve seen, however, that Single Scale can represent justifying heavy requiring reasons (reasons with both kinds of weight but more justifying weight), as long as reasons always have the same proportion. We’ve also seen that Single Scale can represent views that give default reasons a different proportion than non-default reasons. Here, then, is the representational limitation of the image of a single scale:

Single Proportion: all non-default reasons have the same proportion of justifying and requiring weight.

In other words, Single Scale cannot allow the justifying and requiring weight of non-default reasons to be independent variables. This limitation prevents Single Scale from being the fundamental model of how reasons are to be weighed.

4. What Single Scale Can’t Represent

In this section, I show that Single Scale’s commitment to Single Proportion prevents it from representing the existence of dilemmas, the standard account of supererogation, and certain versions of satisficing. In the next section, I explain why these

¹⁸ If an agent’s autonomy gives them a fixed cushion of permissibility, then that is all the more reason to think of autonomy-grounded prerogatives as default reasons.

¹⁹ If you combine both forms of satisficing, you’ll get a more complex equation: an act is permissible as long as it is at least $[\frac{1}{2}(\text{greatest balance}) - 5]$.

²⁰ While it is common to assume that *actions* are permissible by default, the standard view in epistemology is that *(dis)beliefs* are impermissible by default. This standard view is revealed by how controversial doxastic (aka: epistemic) conservatism is, where doxastic conservatism is just the view that beliefs are rational/justified/reasonable by default.

representational failures are a problem for Single Scale insofar as it is an attempt to capture the fundamental model of weighing reasons.²¹

4.1. Failure 1: Can't Properly Represent Dilemmas

Dilemmas are cases in which neither ϕ nor $\sim\phi$ is permissible. Dilemmas can't be represented properly in Single Scale. Comparative Weight requires that we assign deontic status in terms of relative weights.²² Max divides relative weights into three categories: counterbalanced, reasons to ϕ are weightier, and reasons to $\sim\phi$ are weightier. There are, however, four conceptually possible deontic combinations: both are permissible (deontic option), only ϕ is permissible (ϕ is required), only $\sim\phi$ is permissible ($\sim\phi$ is required), and neither are permissible (dilemma).²³ Max assigns *ϕ is required* to weightier reasons for ϕ , *$\sim\phi$ is required* to weightier reasons for $\sim\phi$, and *both are permissible* to counterbalanced weights; therefore, dilemma is left out in the cold.

We could divide the relative weights into four categories: exactly counterbalanced, within 10 units of counterbalanced, reasons to ϕ are weightier by more than 10 units, and reasons for $\sim\phi$ are weightier by more than 10 units. This makes conceptual room for dilemmas at the cost of normative absurdity. The difference between whether one is in a dilemma or in a deontic option isn't a difference in whether the reasons are exactly counterbalanced or almost counterbalanced. Whether someone is in a dilemma is a function of the particular reasons he has, and perhaps the particular failures that he's made. Consider the case in which a first time cheater gets his one night stand pregnant and must decide between continuing a relationship with the one night stand or his fiancé. Whether this case counts as a dilemma doesn't seem to depend on whether one's reasons are exactly counterbalanced or almost counterbalanced. This is the wrong way to try to settle whether it is a dilemma or a deontic option.

Single Scale does capture something important: whether this case counts as dilemma depends on the relative weight of reasons. Dilemmas are cases in which the reasons make the alternative, $\sim\phi$, impermissible without making ϕ permissible. Recall that the above two versions of satisficing appeal to **justifying heavy reasons**, reasons with more

²¹ When I say 'Single Scale can't represent normative theory T', you may prefer to say 'Single Scale represents normative theory T *as impossible*'. The two expressions amount to the same thing for my purposes. Insofar as Single Scale is an attempt to capture the fundamental model of weighing reasons, either expression is tantamount to the charge that T is confused about how to weigh reasons.

²² Note that Comparative Weight rules out allowing trivial cases of counterbalancing (no reasons on either pan) to have a different deontic status than non-trivial cases (there are reasons on both pans that are equally weighty). Even if we ignore this implication of Comparative Weight, the distinction between trivial and non-trivial counterbalancing won't provide a normatively satisfactory account of dilemmas. The most promising examples of dilemmas are cases in which an agent seems to have reasons that push in opposite directions (like the one time cheater case in the main text). But some deontic option cases are also cases in which different reasons push in opposite directions (e.g., when we choose between altruism and self-interest).

²³ Some theorists think of supererogation as a deontic status. I think of it as a deontic/evaluative hybrid: it involves a deontic option in which one option is morally better than the other.

justifying than requiring weight.²⁴ In contrast, the existence of dilemmas requires **requiring heavy reasons**, i.e. reasons that have more requiring than justifying weight.²⁵

One seems to have requiring heavy reasons, if at all, only in very specific circumstances, so they can't be default reasons. Unless we risk overgeneralizing which cases are dilemmas, such reasons better be few and far between: the dilemma generating non-default reasons will be requiring heavy, whereas most other (non-default) reasons will be balanced or justifying heavy. Thus, to represent dilemmas in a normatively promising way, we need to be able to represent *non*-default reasons with different proportions of justifying and requiring weight. Yet it is precisely such a possibility that Single Proportion rules out.²⁶

4.2. Failure 2: Can't Represent Standard Account of Supererogation

Suppose that Liv jumps on a grenade to save the lives of five other soldiers. The standard account of supererogation holds that Liv's act is supererogatory. In a nutshell:

The lives of the five soldiers have requiring (and justifying) weight, because Liv would be morally required to save the five soldiers if she could do so at no cost to herself. Morality is other-centered, so Liv's own wellbeing lacks requiring weight. She is nonetheless permitted to remain in safety because her life is a very weighty merely justifying reason. Since both alternatives are permissible and it is morally better to save the others, she goes "beyond the call of duty" in jumping on the grenade.²⁷

Altruistic and self-interested reasons are ones that we don't have by default, because the costs and benefits to you and others varies from case to case. Thus, Single Scale can't represent the standard account of supererogation, because it depends on certain non-default reasons (altruistic ones) having a different proportion of justifying and requiring weight than others (self-interested ones). Altruistic reasons have both justifying and requiring weight whereas self-interested reasons have only justifying weight.

²⁴ Justifying heavy reasons can be justifying heavy *requiring* reasons (reasons with both justifying and requiring weight but more justifying weight) or merely justifying reasons (reasons that have justifying weight but no requiring weight).

²⁵ More precisely: requiring heavy reasons are needed to properly represent the conceptual possibility of dilemmas *when there are only finitely many options*. See note 39 for some discussion of dilemmas when there are infinitely many options. Also note that the sort of dilemma that I'm talking about is *not* due different perspectives (morality vs rationality or truth vs epistemic rationality) issuing conflicting all-in requirements. It is due to a single perspective issuing a single combined verdict of impermissible commitment.

²⁶ Irrational aims and unjustified beliefs are another potential source of dilemmas. Such dilemmas are instances of *structural* irrationality, which is often distinguished from *substantive* irrationality. While structural irrationality is often explained in terms of the violations of requirements, I agree with Fogal (2019) that both structural and substantive irrationality should be explained in terms of pressure, or weight. (Here I disagree with Gert, who explains structural irrationality in a different way (e.g., 2004: 69-72).) Yet Fogal underestimates the case for treating structural issues as a proper subset of substantive ones, perhaps because he is seduced by Single Scale and its commitment to Single Proportion (2019: 1054).

²⁷ Portmore (2008; 2011: ch 5) is perhaps the clearest proponent of this account. His account is assumed by Massoud (2016) and Archer (2016). Similar accounts are given by Clark (1978), Curtis (1981), Muñoz (forthcoming), and Raz (1990: 94). Raz's response utilizes "exclusionary permissions" rather than merely justifying reasons. Yet "exclusionary permission" seems to be an awkward way of referring to an undefeated merely justifying reason (Gert 2004: 107-110).

4.3. Failure 3: Good-Centered Satisficing

Recall that Single Scale can represent Relative Weight and Default Reason Satisficing, versions of *weight-centered satisficing*. Such theories understand the ‘good enough’ as a threshold in the relative weight of all relevant reasons. Single Scale cannot model what I take to be the most promising form of satisficing, *good-centered satisficing*. This form of satisficing lets the good enough be a threshold, not in the overall relative weights, but in an intrinsic good, such as wellbeing.

The standard account of supererogation focuses on the supererogatory relief of suffering. It correctly supposes that one has requiring reason to make (a part of) someone’s life better when it is (or otherwise would be) poor. Yet what if someone else’s life is already very good? It is less clear whether there is requiring reason to make their life even better. A proponent of good-centered satisficing might say, then, that we have requiring reason to make a life better to the point that it is good enough and only merely justifying reason to make it better beyond that point. The good enough can be given absolute (at least 100 units of wellbeing) and relative readings (at least 80% of the most wellbeing one can give), though I prefer the absolute reading. Such a view supposes that Single Proportion is false. For it supposes that the (non-default) reasons provided by someone’s wellbeing are sometimes requiring (have requiring strength) and sometimes merely justifying (cf. Tucker 2017).

One benefit of this kind of satisficing is that it makes possible a kind of supererogation that is different from the supererogatory relief of suffering, the sort of supererogation the standard account focuses on. Perhaps there is a sort of supererogation that involves a “gratuitous promotion of wellbeing”, a promotion of someone’s wellbeing when such promotion isn’t even pro tanto required (cf. Horgan and Timmons 2010).²⁸ The standard account holds that the altruistic prevention of five deaths generates a pro tanto requirement to jump on the grenade, which implies that such reasons have requiring weight. We need merely justifying altruistic reasons to explain how we can have cases of supererogation in which altruism is gratuitous, or not even pro tanto required.

5. Are these Representational Failures Really a Problem for Single Scale?

Single Scale can’t (properly) represent three normative theories: dilemmas exist, the standard account of supererogation, and good-centered satisficing. All that follows so far is that there is a problem with Single Scale *or* there is a problem with these three theories. Why blame Single Scale?

We are considering Single Scale insofar as it is an attempt to capture the fundamental model of weighing reasons. This model is analogous to the physical equations that specify how forces interact to determine which direction a particle is moving. One can use the correct equations (analogue: understand how reasons are to be weighed) but be mistaken about whether a given object exerts a certain force in a certain direction (analogue: be mistaken about the weight of a given reason). Physical equations specify structural relationships that abstract away from particular judgments about which forces a given object has. Likewise, the fundamental model of weighing reasons specifies

²⁸ This qualifies my earlier contention that the existence of supererogation does not require good-centered satisficing (2017: 1371-2). Good-centered satisficing isn’t needed for the sort of supererogation that folks ordinarily focus on (supererogatory relief of suffering); however, it is needed to explain the sort of supererogation that intrigues Horgan and Timmons (supererogatory gratuitous promotion of wellbeing).

structural relationships that abstract away from particular judgments about what weight(s) a given reason has.

Single Scale can't represent the relevant three theories because of Single Scale's commitment to Single Proportion. Single Proportion is a claim about the *weights* of reasons rather than the *weighing* of reasons. The main problem with Single Scale, then, is *not* (merely) that it can't represent a normative theory that is true. Assume that the relevant three theories are false if you like. The main problem is that one can reject Single Proportion without being confused about how to weigh reasons. Thus, Single Scale is not the fundamental model of weighing reasons.

It will be instructive to consider a foil, a view that Single Scale can't represent where we should blame the view rather than Single Scale. To preserve the image of weighing one thing against another, deontic verdicts must be a function of *relative* weights of the reasons for and against. To preserve the image, in other words, we must endorse something like Weight Comparativism. Not every normative theory has this structural feature. Consider:

Absolute Weight Satisficing (AWS): ϕ is permissible iff the reasons for ϕ have a weight of at least 100 (no matter how much weight the reasons for $\sim\phi$ have); otherwise, ϕ is impermissible.

The proponent of Single Scale must reject AWS, because AWS assigns deontic status only in virtue of the *total* weight for ϕ , not the *relative* weights of the reasons for ϕ and for $\sim\phi$. Thus, we are forced to choose between Single Scale—more specifically, Weight Comparativism—and AWS. In this case, we should blame AWS.

In contrast to Single Proportion, Weight Comparativism is not a substantive normative judgment about the weights of reasons. Indeed, something like it is regularly taken to be a structural constraint on substantive moral and rational theorizing.²⁹ It is no surprise to find, then, that the proponent of AWS lacks a coherent conception of weighing reasons.

Recall that the conceptual distinction between requiring and justifying weight is uncontroversial. What's controversial is what proportion of justifying and requiring weight a given reason can have. Let's consider a version of AWS which holds that all reasons are balanced requiring reasons. We discover that the total justifying weight for ϕ is 100. AWS entails that ϕ is permissible *no matter how much requiring weight there is for $\sim\phi$* . This entailment is incoherent.

The basic problem can be illustrated with a vector analogue to AWS, one which says that a particle goes down iff at least 100 units of downward force are exerted on it. Suppose that a total of 100 downward units are exerted. The analogue absurdly entails that the particle goes down regardless of how much upward force is exerted on it. Numerical units of force make it easy for us to compare the downward and upward forces and determine which of the two opposing forces wins. 1,000 units of upward force beats 100 units of downward force every single time.

The same sort of competition holds between (amounts of) justifying weight for ϕ and (amounts of) requiring weight for $\sim\phi$. A reason's *amount of justifying weight* for ϕ is a measure of how much total requiring weight for $\sim\phi$ it can successfully override, i.e., how much total requiring weight for $\sim\phi$ there can be and the reason still make ϕ permissible. A reason's *amount of requiring weight* for ϕ is a measure of how much total justifying

²⁹ See, e.g., Chang (2016: 213-5) and Portmore (2019: 8-11).

weight for $\sim\phi$ it can successfully override, i.e., how much justifying weight for $\sim\phi$ there can be and the reason still make $\sim\phi$ impermissible.³⁰ Numerical values are assigned to justifying and requiring weight because it is an easy way to track which sort of weight will win when they compete. (I inter-defined of amounts of justifying and requiring weight, but this is no problem as I explained in §2.1, including nt. 14.)

Suppose that the total justifying weight for ϕ is 100 (which includes any default justifying weight) and the total requiring weight for $\sim\phi$ is 1,000. Justifying weight for ϕ pushes ϕ down toward permissibility. Requiring weight for $\sim\phi$ pushes ϕ up toward impermissibility. These weights are in direct competition and 1,000 units of weight toward impermissibility beats 100 units of weight toward permissibility every single time. In other words, the very concepts of justifying and requiring weight rule out AWS. Their conceptual relationship entails that ϕ would be impermissible in this case (100 JW $_{\phi}$ vs 1,000 RW $_{\sim\phi}$), and AWS falsely entails that it would be permissible. Perhaps a given version of AWS can be normatively adequate if, e.g., it assigns the correct weight to all reasons. Yet no version of AWS can be structurally sound, for it misunderstands how reasons are to be weighed to determine a deontic status.³¹

On the surface, AWS might look no different than an absolute threshold satisficing consequentialism (e.g., an action is permissible iff the net aggregate wellbeing is at least n). But such a satisficing consequentialism is coherent if we understand it as good-centered satisficing (good enough is a threshold in the good) rather than weight-centered satisficing (good enough is a threshold in weights). At least, this is an option if we can sensibly reject Single Proportion. If satisficing consequentialism has an unavoidable problem, the problem is normative: it assigns weights to reasons incorrectly.³²

We are considering Single Scale as the putative fundamental model of how reasons are to be weighed. The fundamental model represents all and only those normative theories that weigh reasons correctly, regardless of whether they weight reasons correctly. Someone who endorses Absolute Weight Satisficing (an act is permissible iff there is at least n weight in its favor) misunderstands how reasons are to be weighed. Consequently, the conflict between Single Scale (more precisely: Weight Comparativism) and AWS is a problem for AWS.

In contrast, someone who endorses good centered satisficing (e.g., one has requiring reason to make a life good enough but merely justifying reason to make it even better) need not make any mistake about how reasons are weighed. If they make a mistake, it may be only a normative mistake about the weights of reasons. Consequently, the conflict between Single Scale (more precisely: Single Proportion) and good-centered satisficing is a problem for Single Scale. And the same point applies to Single Scale's conflict with the existence of dilemmas and the standard account of supererogation.

³⁰ Cf. Gert (2007: 537-539) and Portmore (2011: 120-121).

³¹ The conceptual incoherence of AWS makes it especially clear that its mistake is structural and not a substantive normative judgment about the weights of reasons; however, not every non-normative, structural mistake suffices for conceptual incoherence (e.g., Rubio's mistake I allude to in nt. 5).

³² Portmore (2019: 9) conflates AWS and absolute threshold satisficing consequentialism. He accuses the latter of holding that, once the absolute threshold is reached, "the alternatives are irrelevant in assessing its moral permissibility" (9). But that's AWS. On a good-centered understanding of the threshold, satisficing consequentialism can say that, beyond the threshold, the reasons for the alternatives get taken into account; it's just that all such reasons are merely justifying.

Single Scale is, therefore, not the fundamental model of weighing reasons. Yet it is such a simple model, you might wonder whether we can use it anyway. In principle, yes, as long as you justify its use in reference to the fundamental model, the model that sets the standard for what it is to weigh reasons correctly. The goal of this justification is to show that Single Scale doesn't distort the normative phenomena it is used to model. For example, you might argue that Single Proportion is true and that Single Scale follows from the conjunction of Single Proportion and the fundamental model. At present, however, the case for Single Proportion is rather thin. It is ordinarily accepted by faith, not argument. It conflicts with many otherwise plausible normative theories and intuitions.³³ Furthermore, it might seem plausible only because Single Scale was our best metaphor for weighing reasons (§8). Perhaps you are more optimistic that Single Proportion can be adequately defended. Fine. The main point is that the use of any non-fundamental model *needs* justification in reference to the fundamental one, justification that has never been provided for Single Scale.

6. Dual Scale

The problem with Single Scale isn't the image of the scale, but the image of a *single* scale. We need *two* scales to represent views that allow the requiring weight of non-default reasons to vary independently of their justifying weight. **Permission Scale** (for ϕ) determines whether ϕ is permissible by comparing the justifying weight for ϕ (JW_ϕ) with the requiring weight for $\sim\phi$ ($RW_{\sim\phi}$). **Commitment Scale** (for ϕ) determines whether ϕ is a commitment (i.e., $\sim\phi$ is impermissible) by comparing the requiring weight for ϕ (RW_ϕ) with the justifying weight for $\sim\phi$ ($JW_{\sim\phi}$).³⁴ The two scales working together reveal whether ϕ is a requirement (permissible commitment) or a dilemma (*impermissible* commitment).³⁵

³³ In addition to the theories mentioned in §4, there are many other normative theories that cannot be represented in Single Scale due to its commitment to Single Proportion. The most obvious will be any theory that, like the standard account of supererogation, appeals to merely justifying reasons or terminological variants, such as some uses of “prerogatives” (Scheffler 1982) or “mere permissions” (Hurka and Shubert 2012). In this paper, I'm ignoring the possibility that we might have three or more options. If I weren't ignoring that possibility, I'd point out that Single Scale cannot represent what I regard as the most promising solutions to Kamm's Intransitivity Paradox (Archer 2016; Muñoz forthcoming) and Horton's All or Nothing Problem (Muñoz forthcoming).

³⁴ The Commitment Scale for ϕ is just the mirror image of the Permission Scale for $\sim\phi$. This shouldn't be surprising. Commitment to ϕ is just another way to talk about the *impermissibility* of $\sim\phi$.

³⁵ The vector sum analogy to weighing reasons can also be given a dual reading: it takes two vector sums to determine the deontic status of an act, one for permission and one for commitment.




fig. 2

This model, **Dual Scale**, has much in common with Single Scale. The pans still represent ϕ and $\sim\phi$, and so we still get:

Independent Individuation of Options: an agent's options are individuated independently of what reasons the agent has and what the weights of those reasons are.

The counterfactual stability of weight still applies:

Prior Fixed Weights: The same reason has the same weight in all contexts, and what particular weight it has is prior to the deontic status of any particular action.

So far the model is exactly the same as Single Scale. It is also like Single Scale insofar as it has a comparative element. Yet Weight Comparativism needs refinement. Since it is conceptually possible that justifying and requiring weight are independent variables, the comparison isn't as simple as the reasons for ϕ get weighed against the reasons against ϕ .³⁶

We weigh one weight value against another only to the extent that there is *competition* between those two weight values. Justifying and requiring weight are defined by their contributions to fixing deontic status. Thus, we can clarify when justifying and requiring weight compete by clarifying which deontic verdicts are jointly impossible. There is no competition between justifying weight for ϕ and justifying weight for $\sim\phi$, because deontic options are conceptually possible, i.e., you can be permitted to ϕ and permitted to $\sim\phi$. In contrast, you can't be permitted to ϕ and committed to $\sim\phi$. To be committed to $\sim\phi$ just is for ϕ to be impermissible. Hence, justifying weight for ϕ competes with *requiring* weight for $\sim\phi$. It is this competition that determines whether ϕ is permissible *a la* the Permission Scale in fig. 2.

Requiring weight for ϕ does *not* compete with requiring weight for $\sim\phi$. It is at least a conceptual possibility that one be in a dilemma, a case in which neither ϕ nor $\sim\phi$ is

³⁶ In discussion, some have tried to resist Dual Scale by saying that all we need is reasons for and reasons against, as long as we deny that reasons against ϕ are always reasons for $\sim\phi$. This view just is Dual Scale in worse terminology. These folks use 'reasons for ϕ ' as code for justifying weight for ϕ (reasons that push an act toward permissibility) and 'reasons against $\sim\phi$ ' as code for requiring weight against $\sim\phi$ (reasons that push $\sim\phi$ toward impermissibility), and they assume that reasons for ϕ and reasons against $\sim\phi$ need not always come in the same proportion. That's Dual Scale stated in a way that obscures that there are two reason for/against distinctions (justifying weight for/against and requiring weight for/against), and so it obscures the conceptually necessary connections between reasons for and reasons against (every reason for is some kind of reason against—see note 13).

permissible. In contrast, you can't be committed to ϕ and permitted to $\sim\phi$. To be committed to ϕ just is for $\sim\phi$ to be impermissible. Hence, requiring weight for ϕ competes with *justifying* weight for $\sim\phi$. It is this competition that determines whether ϕ is a commitment *a la* Commitment Scale in fig. 2.

Now that we better understand what competes with what, we can replace Weight Comparativism with:

Dual Weight Comparativism:

- i. Whether ϕ is permissible is determined by comparing the *justifying* weight of all reasons for ϕ with the *requiring* weight of all reasons for $\sim\phi$.
- ii. Whether ϕ is a commitment is determined by comparing the *requiring* weight of all reasons for ϕ with the *justifying* weight of all reasons for $\sim\phi$.

Like Single Scale, Dual Scale is officially neutral on which particular relative weights are assigned to a given deontic status. For illustrative purposes, I work with:

Permission Assignment: ϕ is permissible iff the justifying weight of all reasons for ϕ is *at least as weighty* as the requiring weight of all reasons for $\sim\phi$.

Commitment Assignment: ϕ is a commitment iff the requiring weight of all reasons for ϕ is *weightier than* the justifying weight of all reasons for $\sim\phi$.

Whether you opt for these assignment functions or not, you must coordinate the assignments for permissions and commitments, because the two assignment functions interact. If you are permitted to ϕ , then $\sim\phi$ isn't a commitment. If ϕ is a commitment, then you aren't permitted to $\sim\phi$. Once I endorsed the "at least as weighty" language in Permission Assignment, I had to endorse the "weightier than" language of Commitment Assignment. This coordination ensures that there are no gaps in the deontic assignments (every act is permissible or impermissible) and there are no overlap (no act is both permissible and impermissible).^{37,38}

Like Single Scale, Dual Scale will need to be combined with whatever the correct aggregation principle happens to be. The scale metaphor suggests that the correct principle will be additive (the justifying weight of all reasons is the sum total justifying weight of the individual reasons, and the requiring weight of all reasons is the sum total requiring weight of the individual reasons). Since there are many worries about additive aggregation and the other aspects of Dual Scale can survive even if additive aggregation doesn't, I don't build additive aggregation into Dual Scale.

³⁷ The interaction is due to a specific mathematical relationship. Permission Assignment holds that ϕ is permissible iff the **net justifying weight** for it (justifying weight for ϕ – requiring weight for $\sim\phi$) is at least 0. Commitment Assignment holds that $\sim\phi$ is a commitment iff the **net requiring weight** for it (requiring weight for $\sim\phi$ – justifying weight for ϕ) is greater than 0. The net justifying weight for ϕ and the net requiring weight for $\sim\phi$ subtract the same two things in different order; therefore, *they are additive inverses of each other*. If the net justifying weight for ϕ is 1, the net requiring weight for $\sim\phi$ is ~ 1 . Given that a net justifying weight of *at least 0* is the cut-off for permission: (i) to ensure that every action is assigned a deontic status, I had to assume that any net requiring weight for ϕ of *greater than 0* (net justifying weight for $\sim\phi$ of less than 0) is *sufficient* for commitment to ϕ , and (ii) to ensure that no action is both permissible and impermissible, I had to assume that a net requiring weight for ϕ of greater than 0 (justifying weight for $\sim\phi$ of less than 0) is *necessary* for commitment to ϕ .

³⁸ Thanks to the mathematical relationship in the previous footnote, the standard account of supererogation entails Dual Scale. Elsewhere I argue for this at length, but here's a summary: (i) the standard account of supererogation entails Permission Assignment, and (ii) Permission Assignment entails Commitment Assignment.

Single Scale is the conjunction of *four* theses: Independent Individuation of Options, Prior Fixed Weights, Weight Comparativism, and Single Proportion. Dual Scale is the conjunction of *three* theses: Independent Individuation of Options, Prior Fixed Weights, and Dual Weight Comparativism. Dual Scale has no analogue for Single Proportion. That’s because Single Proportion is a substantive normative judgment and has no business in an attempt to capture the fundamental model of weighing reasons.

As I present it here, Dual Scale needs refinement and/or development in at least three ways. First, it must be generalized to handle cases in which there are three or more options. The basic idea of the generalization is simple enough. The above simple Permission Assignment holds that φ is permissible iff it wins a single competition with a single alternative ($JW_{\varphi} \geq RW_{\sim\varphi}$). The generalized Permission Assignment holds that φ is permissible iff it wins a *tournament*, a pairwise competition with each alternative (for each alternative A, $JW_{\varphi} \geq RW_A$). The same point applies to the generalized Commitment Assignment: φ is a commitment iff φ wins a tournament, a pairwise competition with each alternative (for each alternative A, $RW_{\varphi} > JW_A$). But there are some complications that I can’t address here.³⁹ Second, it assumes that every reason against φ is some sort of reason for $\sim\varphi$, which some regard as a problematic assumption (Snedegar 2018).⁴⁰ Third, Dual Scale can and should be refined to allow for conditions and modifiers.⁴¹ These refinements must wait for another occasion. Now that we know what Dual Scale is, let’s see what it can do.

7. The Advantages of Dual Scale

7.1. Anything Single Scale Does, Dual Scale Does Better

Recall that Single Scale can represent maximizing views using Max and two distinct non-maximizing views: the justifying heavy requiring reason view (Relative Weight Satisficing) and the default merely justifying reason view (Default Reason Satisficing). Max holds (i) that all non-default reasons are balanced and (ii) that the (default) bias toward permissibility holds only when reasons are perfectly counterbalanced. Relative Weight Satisficing rejects (i) and Default Reason Satisficing rejects (ii). The first disagreement looks like a disagreement about the weights of *non-default* reasons and the second looks like a disagreement about the weights of *default* reasons. Single Scale obscures these differences, because it can represent normative disagreements in only one way, namely by modifying the function that assigns deontic status to relative weights. Thus, the representation of normative disagreements is taken “off-image”.

³⁹ One complication is that this basic idea can’t represent the conceptual possibility that ever better cases (for every option, there is another better) are dilemmas. If ever better cases are dilemmas, then every option is trivially a commitment: you make a mistake if you don’t choose that option, because you make a mistake no matter which option you choose. The generalized Commitment Assignment *incorrectly* entails that no option is a commitment: for any φ , RW_{φ} is less than the justifying weight of some alternative. Another complication is that justifying and requiring weight are relative to pairwise competitions between φ and each alternative. Consequently, JW_{φ} can vary for non-holist reasons as you vary the alternative (manuscript §6). This is because every reason *against* an alternative is functionally equivalent to a reason *for* φ (recall nt 13). If there is more justifying weight *against* alternative A1 than alternative A2, then there is more JW_{φ} in φ ’s pairwise competition with A1 than in φ ’s pairwise competition with A2.

⁴⁰ The alleged problem depends on a false view of the relationship between reasons for and against. I have a separate paper addressing these issues. Yet again recall footnote 13.

⁴¹ See note 12 for a brief explanation of the sort of refinement I have in mind.

Dual Scale can represent normative disagreements in different ways, and it can represent normative disagreements “on-image”. It represents the disagreement about the weights of *non-default* reasons by changing what goes on the scales. Maximizing views hold that only balanced requiring reasons go on the scales. So whatever justifying weight you assign to the reason for φ on Permission Scale, you assign the same requiring weight to that reason on Commitment Scale. If $JW_{\varphi} = 10$, then $RW_{\varphi} = 10$.



fig. 3

(A requiring reason for $\sim\varphi$ would be represented by assigning the same weight to $RW_{\sim\varphi}$ and $JW_{\sim\varphi}$.) Relative Weight Satisficing, in contrast, holds that whatever justifying weight goes on Permission Scale, you assign *half* to the requiring weight on Commitment Scale. If $JW_{\varphi} = 10$, then $RW_{\varphi} = 5$.



fig. 4

This way of representing the disagreement illustrates that it is about the weights of *non-default* reasons. (A justifying heavy requiring reason to $\sim\varphi$ would be represented by assigning *half* the weight to $RW_{\sim\varphi}$ that you assign to $JW_{\sim\varphi}$.)

Single Scale couldn’t represent disagreements about non-default reasons on-image because the φ side goes down *exactly* to the extent that the $\sim\varphi$ side goes up. It had to represent the disagreement between Max and Relative Weight Satisficing by modifying the assignment function. Since Dual Scale uses one scale for permission and another for commitment, it allows the justifying weight of non-default reasons to differ from their requiring weight. And that’s what we need to represent *on the image of the scale(s)* the difference between balanced and justifying heavy requiring reasons.

Dual Scale has two different ways of representing disagreements about the *default* bias toward permissibility. It can represent it off-image like Single Scale does, by modifying the assignment function. It can also represent it on-image by having the scales begin in a state of bias toward permissibility for both ϕ and $\sim\phi$ (remember default reasons are reasons for every act in every situation). Maximizing views have both Permission and Commitment Scales begin in the counterbalanced state before you add non-default reasons to the pans. A default merely justifying reason can be represented by biasing Permission Scale 5 units in ϕ 's favor and Commitment Scale 5 units in $\sim\phi$'s favor *before* non-default reasons are added. (Remember that the two scales must be coordinated. If you bias Permission Scale toward ϕ without equally biasing Commitment Scale toward $\sim\phi$, it will be possible for an action's deontic status to change simply by changing whether it counts as ϕ or $\sim\phi$.) The starting positions of the two biased scales will then resemble fig. 5.



fig. 5

You couldn't represent the default reason strategy on-image in Single Scale because it represented both permissibility and commitment on the same scale. If you biased the scale toward ϕ 's being permissible, you thereby biased the scale toward $\sim\phi$'s being impermissible. Since Dual Scale uses one scale for permission and another for commitment, it allows the default justifying weight in favor of ϕ to vary independently of the default requiring weight for ϕ . And this is exactly what we need to represent *on the image of the scales* disagreements about how weighty the default bias toward permissibility is.

7.2. Dual Scale Represents What Single Scale Can't

To illustrate how Dual Scale represents what Single Scale can't, it will be helpful to convert the dual scale image into the following matrix:

Dual Scale Matrix

Permission Scale		Commitment Scale	
JW_{ϕ}	$RW_{\sim\phi}$	RW_{ϕ}	$JW_{\sim\phi}$
Verdict: permissible/impermissible		Verdict: commitment/not a commitment	
Combined verdict: (im)permissible (but not a) commitment			

The first row identifies the two scales. The second row tracks the weights of the relevant non-default reasons. The third tracks the verdicts of each individual scale. The fourth combines those verdicts to yield the combined deontic status of an act, which depends on both whether the act is permissible and whether it is a commitment. (This matrix assumes that the default bias toward permissibility holds only in the case of counterbalanced reasons. If we wanted to represent alternative accounts of default reasons, we could add a new row between the first and second to track disagreements about default weight.)

Dilemmas. I stupidly make it my deepest aim to punch every brick wall I find as hard as I can. I find a brick wall. *If* this case is a dilemma, how could we explain why it is a dilemma? My two options are to Punch or \sim Punch. The harm of punching the brick wall is, say, a balanced requiring reason not to punch the wall ($50 JW_{\sim\text{Punch}}$ and $50 RW_{\sim\text{Punch}}$). Stupid aims have no capacity to justify, i.e., they lack justifying weight, but even stupid aims can commit you to an action, i.e., they have requiring weight. In other words, the stupid aim is a **merely requiring reason**, a reason that has requiring weight but no justifying weight. Since it is your deepest aim, it has more requiring weight than the harm has justifying weight not to ($0 JW_{\text{Punch}}$ and $100 RW_{\text{Punch}}$).

Dual Scale represents this dilemma as follows. Permission Scale tells us that punching the wall isn't permissible. My reason to punch the wall has no justifying weight and my reason not to has considerable requiring weight ($0 JW_{\text{Punch}} < 50 RW_{\sim\text{Punch}}$). Yet Commitment Scale says that punching the wall is a commitment. My reason to punch the wall has more requiring weight than my reason not to has justifying weight ($100 RW_{\text{Punch}} > 50 JW_{\sim\text{Punch}}$). So punching the wall is an impermissible commitment and my stupid aim has put me in a dilemma.

(when assessing the deontic status of Punch)

Permission Scale		Commitment Scale	
$0 JW_{\text{Punch}}$	$50 RW_{\sim\text{Punch}}$	$100 RW_{\text{Punch}}$	$50 JW_{\sim\text{Punch}}$
Verdict: Punch is <i>impermissible</i>		Verdict: Punch is a commitment	
Combined verdict: Punch is an impermissible commitment (dilemma)			

Sometimes ethicists distinguish between *requirement dilemmas* (both ϕ and $\sim\phi$ are required) and *prohibition dilemmas* (both ϕ and $\sim\phi$ are impermissible). I have shown that Dual Scale can represent the latter, but not the former. Is that a problem? No. Given the sense of 'requirement' in the paper, requirement dilemmas are conceptually impossible. In §2.1, following Dorsey and Portmore, we defined 'requirement to ϕ ' as a compound deontic verdict of permissible to ϕ and impermissible to $\sim\phi$. Thus, a requirement to ϕ entails that $\sim\phi$ is not permissible and so not required.

Standard Account of Supererogation. Liv must choose whether to sacrifice her life (Sacrifice) or remain in safety (Safety). The standard account of supererogation treats altruistic reasons, such as the lives of the five soldiers, as requiring reasons (say, $500 JW_{\text{Sacrifice}}$ and $500 RW_{\text{Sacrifice}}$). The account treats self-interested reasons, such as the value of her life, as very weighty merely justifying reasons (say, $1,000 JW_{\text{Safety}}$ and $0 RW_{\text{Safety}}$).

Permission Scale entails that the sacrificial act is permissible, because the justifying weight of Liv's altruistic reason is weightier than the non-existent requiring weight of her self-interested reason ($500 JW_{\text{Sacrifice}} > 0 RW_{\text{Safety}}$). Commitment Scale entails that Sacrifice isn't a commitment, because the requiring weight of the altruistic reason is outweighed by the justifying weight of her self-interested reason ($500 RW_{\text{Sacrifice}} < 1,000 JW_{\text{Safety}}$).

(when assessing deontic status of Sacrifice)

Permission Scale		Commitment Scale	
500 JW _{Sacrifice}	0 RW _{Safety}	500 RW _{Sacrifice}	1,000 JW _{Safety}
Verdict: Sacrifice is permissible		Verdict: Sacrifice is not a commitment	
Combined verdict: Sacrifice is permissible but not required			

Since the altruistic act is permissible but not a commitment (and so not required), one goes beyond the call of duty in performing it. At least, one goes beyond the call on the plausible assumption that it is morally better to save the five soldiers than to save one's own life.

Good-Centered Satisficing. Suppose that 100 units of wellbeing is the threshold for the good enough. Good-centered satisficers generally hold that if you have to choose between Good Enough (100 units of wellbeing) and Not Good Enough (95 units of wellbeing), you are required to choose Good Enough. That's because this kind of satisficer generally assumes that one has *balanced* requiring reason to promote wellbeing all the way up to the good enough. This assumption gives us something like this: there is 95 JW and RW for Not Good Enough and 100 JW and RW for Good Enough.

When we plug these values into Dual Scale, Permission Scale tells us that choosing Good Enough is permissible, because the justifying weight for Good Enough is greater than requiring weight for Not Good Enough ($100 JW_{GE} > 95 RW_{NGE}$). Commitment Scale tells us that Good Enough is also a commitment, because the requiring weight for Good Enough is greater than the justifying weight for Not Good Enough ($100 RW_{GE} > 95 JW_{NGE}$).

(when assessing deontic status of GE relative to NGE)

Permission Scale		Commitment Scale	
100 JW _{GE}	95 RW _{NGE}	100 RW _{GE}	95 JW _{NGE}
Verdict: GE is permissible		Verdict: GE is a commitment	
Combined verdict: GE is required (a permissible commitment)			

On the other hand, if you are choosing between Good Enough (100 units of wellbeing) and Even better (105 units of wellbeing), you are permitted but not required to choose Good Enough. Even Better is better than Good Enough, so you have more justifying weight to choose it ($105 JW_{EB}$ vs $100 JW_{GE}$). Yet the good-centered satisficer's view is that requiring weight stops at the good enough amount of wellbeing (beyond the good enough threshold any additional reason is merely justifying). Consequently, there is no more requiring weight to choose Even Better than there is to choose Good Enough ($100 RW_{EB}$ vs $100 RW_{GE}$). Let's plug these numbers into Dual Scale.

Permission Scale tells us that choosing Good Enough is permissible because the justifying weight for Good Enough is equal to the requiring weight for Not Good Enough ($100 JW_{GE} = 100 RW_{EB}$). Commitment Scale tells us that Good Enough is not a commitment, because the requiring weight for Good Enough is less than the justifying weight for Even Better ($100 RW_{GE} < 105 JW_{EB}$).

(when assessing deontic status of GE relative to Even Better)

Permission Scale	Commitment Scale
------------------	------------------

100 JW _{GE}	100 RW _{EB}	100 RW _{GE}	105 JW _{EB}
Verdict: GE is permissible		Verdict: GE is not a commitment	
Combined verdict: GE is permissible but not required			

8. Counteracting Single Scale's Distorting Influence

Inapt but salient metaphors can distort our theorizing, even in the face of compelling counterarguments.⁴² When our best metaphor for weighing reasons is Single Scale (and the related single vector sum), Single Proportion seems like a default assumption of normative inquiry, if not a constraint on the coherent weighing of reasons.

Some actions are required, so there must be requiring reasons. The standard account of supererogation and good-centered satisficing claim that there are also merely justifying reasons. This is clearly a coherent normative claim, and we've seen how easily it can be represented in Dual Scale. When relying on Single Scale, however, it is easy for intelligent philosophers to conclude that merely justifying reasons don't even make sense:

I simply do not understand how a reason *could* be [merely justifying]. To be a reason is to support a particular course of action. For a reason to be [merely justifying], however, it would have to be rationally justifiable for the agent to neglect that support. But if that support is genuine, I cannot see how such neglect could be justified. It would not be implausible, then, to suspect that we cannot give a coherent account of [merely justifying] reasons. (Kagan 1989: 381)

Kagan is thinking with Single Scale: to the extent that a reason supports a course of action (pushes the ϕ -pan down), it also counts against not doing that action (pushes the $\sim\phi$ -pan up). Single Scale rules out (non-default) merely justifying reasons. So it's no surprise that he can't see how they are coherent.

Single Scale has also obscured good-centered satisficing. When you are working with Single Scale, the only way to understand the *good enough* is to think about it as a threshold within relative weight. To understand the good enough as a threshold within, say, wellbeing, you need to make sense of the possibility that wellbeing can provide both requiring and merely justifying reasons. But you can't make sense of such a possibility when you reason with Single Scale and its commitment to Single Proportion. It is no surprise, then, that philosophers tend to ignore good-centered satisficing when they discuss weighing reasons.^{43,44}

I have emphasized Dual Scale's handy visualizations because they counteract the distorting influence of Single Scale. They give us a new, more accurate image/metaphor to think with. In any event, Dual Scale has a much better claim to being the fundamental

⁴² At least, this is suggested by the psychological research on the availability heuristic.

⁴³ In discussions of weighing reasons, satisficing comes up almost exclusively as an assignment of deontic status to relative weight, one that is a non-maximizing alternative to Max (e.g., Berker 2007: 136; Dorsey 2016: 10-11). That's weight-centered satisficing. This focus leads Gert to take the failure of Single Scale to be a failure of satisficing theory (2007: 535). Even Slote's (1989) defense of satisficing shies away from the merely justifying reasons of good-centered satisficing (see my 2017: 1372, nt. 10 for an explanation).

⁴⁴ When you reason with Single Scale, an absolute threshold version of good-centered satisficing might look especially suspicious. The discussion of Absolute Weight Satisficing revealed that an absolute threshold in the weights of reasons was incoherent, because such a threshold is incompatible with anything in the neighborhood of Weight Comparativism. Since Single Scale can represent only weight-centered satisficing, it is easy to infer that no version of absolute threshold satisficing is coherent. (Cf. Portmore's conflation discussed in nt 32 above.)

model of weighing reasons. It can easily weigh the merely justifying reasons needed for the standard account of supererogation and good-centered satisficing. It can easily weigh the requiring heavy reasons needed for dilemmas. Perhaps there are problems with these normative theories and any other rejection of Single Proportion. Yet Dual Scale shows that these theories make sense and that they can weigh reasons correctly.

Conclusion

The fundamental model of weighing reasons determines what it is to weigh reasons correctly. By endorsing Single Proportion, Single Scale takes a substantive normative stand on which *weights* a reason can have. One can reject this stand without being confused about how reasons are to be weighed. Dual Scale distills and refines the promising structural content of Single Scale, and it can represent a much broader range of normative theories. Dual Scale is a crucial step in our search for the fundamental model of weighing reasons. It needs refinement to deal with holist challenges and choice contexts with any number of options, but I submit that it approximates the fundamental model. It is commonly assumed that the metaphors of weight and weighing are apt for morality and rationality. If so, then Dual Scale helps us better understand what is and is not a structural constraint on substantive moral and rational theorizing.⁴⁵

References

- Archer, Alfred. 2016. "Moral Obligation, Self-Interest, and the Transitivity Problem." *Utilitas* 28: 441-464.
- Bader, Ralf. 2016. "Conditions, Modifiers, and Holism." In Lord, Errol and Barry Maguire (eds.), 27-55.
- Berker, Selim. 2007. "Particular Reasons." *Ethics* 118: 109-139.
- Chang, Ruth. 2016. "Comparativism: The Grounds of Rational Choice." In Lord, Errol and Barry Maguire (eds.), 213-240.
- Clark, Michael. 1978. "The Meritorious and the Mandatory". *Proceedings of the Aristotelian Society* 79: 22-33.
- Cullity, Garrett. 2018. "Weighing Reasons." *The Oxford Handbook of Reasons and Normativity*. Daniel Star (ed.). Oxford University Press, 423-442.
- Curtis, Barry. 1981. "The Supererogatory, the Foolish, and the Morally Required." *Journal of Value Inquiry* 15: 311-318.
- Dancy, Jonathan. 2004. *Ethics Without Principles*. Oxford University Press.
- Dorsey, Dale. 2016. *The Limits of Moral Authority*. Oxford University Press.
- Fogal, Daniel. 2019. "Rational Requirements and the Primacy of Pressure." *Mind* 129: 1033-70.
- Gert, Josh. 2016. "The Distinction between Justifying and Requiring: Nothing to Fear." In Lord, Errol and Barry Maguire (eds.), 27-55.

⁴⁵ It is the work of a village to produce a single-authored paper. This paper was possible because of the generous research grant and sabbatical provided by William & Mary. Over the years, conversations with Josh Gert and Philip Swenson clarified my thinking about the justifying/requiring weight distinction. Very helpful comments on the paper were provided by Jeremy Fantl, Daniel Fogal, Lanny Goldman, Daniel Muñoz, Wes Siscoe, Justin Snedegar, anonymous referees, and audience at the 5th Annual Theistic Ethics Workshop. I am extremely grateful for all this assistance.

- _____. 2007. "Normative Strength and the Balance of Reasons." *Philosophical Review* 116: 533-562.
- _____. 2004. *Brute Rationality: Normativity and Human Action*. Oxford University Press.
- 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.
- Hurka, Thomas and Shubert, Esther. 2012. "Permissions to Do Less than the Best: A Moving Band." *Oxford Studies in Normative Ethics* 2: 1-27.
- Kagan, Shelly. 1985. *The Limits of Morality*. Oxford University Press.
- Kearns, Stephen. 2016. "Bearing the Weight of Reasons." In Lord, Errol and Barry Maguire (eds.), 173-90.
- Lewis, David. 1970. "How to Define Theoretical Terms." *Journal of Philosophy*. 67(13): 427–446.
- Lord, Errol and Barry Maguire (eds.). 2016. *Weighing Reasons*. Oxford University Press.
- Maguire, Barry and Justin Snedegar. Forthcoming. "Normative Metaphysics for Accountants." *Philosophical Studies*. Early view available at: https://link.springer.com/article/10.1007/s11098-020-01435-w?wt_mc=Internal.Event.1.SEM.ArticleAuthorOnlineFirst
- Massoud, Amy. 2016. "Moral Worth and Supererogation." *Ethics* 126: 690-710.
- Muñoz, Daniel. Forthcoming. "Three Paradoxes of Supererogation." *Nous*. Manuscript version available at: <https://philpapers.org/rec/MUOTPO-3>.
- Nair, Shyam. 2016. "How Reasons Accrue." In Lord, Errol and Barry Maguire (eds.) 2016: 56-73.
- Portmore, Douglas. 2019. *Opting for the Best: Ought and Options*. Oxford University Press.
- _____. 2011. *Commonsense Consequentialism: Wherein Morality Meets Rationality*. Oxford University Press.
- _____. 2008. "Are Moral Reasons Morally Overriding?" *Ethical Theory and Moral Practice* 11: 369-388.
- Raz, Joseph. 1990. *Practical Reason and Norms*. 2nd ed. Oxford University Press.
- Rubio, Daniel. 2018. "God meets Satan's Apple: the Paradox of Creation." *Philosophical Studies* 175: 2987-3004.
- Slote, Michael. 1989. *Beyond Optimizing*. Harvard University Press.
- Snedegar, Justin. 2018. "Reasons for and Reasons Against." *Philosophical Studies* 175: 725-743.
- Tucker, Chris. Manuscript. "Weighing Reasons Against."
- _____. Forthcoming. "Experience as Evidence." *Routledge Handbook of Philosophy of Evidence*. Maria Lasonen-Aarnio and Clayton Littlejohn (eds). Routledge. Manuscript version available at: <https://philpapers.org/rec/TUCEAE-3>.
- _____. 2017. "How to Think About Satisficing." *Philosophical Studies* 174: 1365-1384.