

## A Holist Balance Scale

Chris Tucker 3/14/22  
William & Mary [cstucker@wm.edu](mailto:cstucker@wm.edu)

(forthcoming in the *Journal of the American Philosophical Association*)


**Abstract:** Scale-based models of weighing reasons face challenges concerning the context-sensitivity of weight, the aggregation of weight, and the methodology for determining what the weights of reasons are. I resolve these challenges.

**Keywords:** weighing reasons, holism, additive aggregation, partiality, diminishing marginal utility

### 0. Introduction

Ethicists regularly talk of *the balance of reasons* and of *one reason outweighing another*. Such talk invites the metaphor of weighing reasons on a scale to determine an act's deontic status. Some philosophers explicitly rely on this metaphor when developing their accounts of how reasons interact to determine deontic status (e.g., Broome 2013: 46-55). Other philosophers object that the metaphor distorts normative theorizing in some way or another. They include Dancy (2004), Drai (2018), Snedegar (2018), Hawthorne & Magidor (2018), and Titelbaum (2019). I side with those who find the metaphor useful. I even have a pet model (Tucker 2022, forthcoming, manuscript). But today I speak in defense of all scale-based models.

I focus on a simple model, Single Scale, so that we won't be distracted by the idiosyncratic details of more elaborate alternatives.<sup>1</sup> According to Single Scale, the reasons for  $\varphi$  ( $R_\varphi$ ) go in one pan and the reasons for  $\sim\varphi$  ( $R_{\sim\varphi}$ ) go in the other. For simplicity, I assume that every reason against  $\varphi$  is equivalent to some reason for  $\sim\varphi$ .<sup>2</sup> The relative weights, as indicated by the relative heights of the two pans, determine the deontic status of the act. This paper develops Single Scale—and by extension all scale-based models—to address challenges concerning the context-sensitivity of weight, the aggregation of weight, and the methodology for determining what the weights of reasons are.



Dancy argues that the kitchen scale model—his term for Single Scale—is:

far too atomistic...On the kitchen scale model, each consideration has a practical weight, which it keeps irrespective of what it is combined with—just

<sup>1</sup> More elaborate alternatives include Sher's (2019) probabilistic model, Wedgwood's (forthcoming) expected utility model, and my "Dual Scale" model. I borrow the diagram of Single Scale from my forthcoming.

<sup>2</sup> See Wedgwood (forthcoming, especially §5.iii) for a version of Single Scale cashed out entirely in terms of reasons against rather than reasons for. For my official account of reasons against, see my 2022.

as a kilogram of butter weighs a kilogram whatever it is combined with. A more holistic picture...has it that the presence of one feature can affect the weight of another. (2004: 105)

This *Holist Challenge* boils down to two claims. The first claim is that Single Scale assumes **atomism**: a feature that is a reason in one context is a reason with the same polarity and weight in every other context. Suppose that your promise to  $\phi$  is a reason to  $\phi$  in some context. Atomism holds that it is a reason and, more specifically, an *equally weighty reason for  $\phi$*  in any choice context. The second is that atomism is false, i.e., **holism** is true: a feature that is a reason in one context may have a different magnitude (from weak to strong) or polarity (from reason for to reason against), or may not be a reason at all, in some other context. For example, while one's having promised to  $\phi$  is generally a reason for  $\phi$ , perhaps it isn't a reason at all (for or against  $\phi$ ) when the promise was coerced (38). The dispute between atomism and holism concerns whether a reason's weight is context-sensitive, whether a reason's weight can differ in different contexts. Atomists say no; holists say yes.

Hawthorne and Magidor (2018: 133-5) raise a similar challenge insofar as they claim that putting weights on a balance scale is monotonic (placing more reasons on the scale can only increase the total weight) whereas reasons are non-monotonic and context-sensitive. I focus on Dancy's version of the challenge and address Hawthorne and Magidor's arguments along the way.

I endorse holism, the second premise of the Holist Challenge. The challenge nonetheless fails because the first premise is false: Single Scale can allow context to help fix the weight of a given consideration. It can reject atomism and endorse holism.

In §1, I clarify Single Scale and explain why it might seem incompatible with holism. In §2, I present what I call the 'Standard Holist Framework'. The basic idea of this framework is, roughly, that a reason's *embedded weight*—its practical significance in a context—is a function of the reason's *varying context* and *invariant default weight*. In §3, I quickly rehearse what I take to be the most compelling argument for holism. This argument appeals to permissible partiality. While the argument won't settle whether holism is true, it will give those with no antecedent commitment to holism some reason to take it seriously. In §4, I explain how to make Single Scale compatible with the Standard Holist Framework. You first use context to fix the embedded weights of all reasons (in accordance with the Standard Holist Framework). Then the deontic status of  $\phi$  is determined by comparing the embedded weights of the reasons for and against  $\phi$  (in accordance with Single Scale).

Single Scale is committed to additive aggregation, and holists worry that no additive aggregation principle can accommodate certain "peculiar and irregular ways" that reasons combine given holism. If they are right, then holism is a threat to additive aggregation, and thus also a threat to Single Scale, after all.<sup>3</sup> To properly assess this threat, we need to clarify exactly when Single Scale is committed to additive aggregation. Putative counterexamples generally attack additive aggregation when Single Scale is not committed to it. Physical weight itself is not additive *regardless of*

---

<sup>3</sup> See, e.g., Dancy (190-1) and the references in nt 6. Berker (2007: 118-22, including nt 24) also understands Dancy and another holist, Little (2000: 280), as rejecting additive aggregation. Lechler (2012: 771-2) interprets Dancy differently, such that Dancy is open to additive aggregation but only in a way that fits poorly with scale-based weighing models.

*how objects are individuated.* The total weight of my phone and its processor is *not* their sum. In §5.1, we'll see that Single Scale is committed to additive aggregation only when the inputs to aggregation are the non-redundant embedded weights of basic reasons. In §5.2, I show that holism is not a threat, but a shield. It allows us to reinterpret putative counterexamples to additive aggregation as further instances of weight's context-sensitivity.<sup>4</sup>

Dalia Drai (2018: 67-68) anticipated that one day someone sneaky like me would come along, proposing that holism resolves putative counterexamples to Single Scale and additive aggregation. She's having none of it. She argues that a holist Single Scale cannot assign default weights in a principled way. The combination of holism and Single Scale is a methodological dead end. In §6, I argue that we can *reverse engineer* the default weights of reasons, i.e., we can identify which acts are permissible and work backwards to what the default weights of reasons must be. This methodology may seem suspicious for various reasons. I show that it is nothing more than a common, well-understood aspect of scientific (and philosophical) methodology.

In sum, I have three main aims. The first is to turn the Holist Challenge to Single Scale on its head, showing that holism is not a threat but a shield. The second is to identify the additive aggregation principle that Single Scale is committed to, thereby avoiding many putative counterexamples. The third is to show that the first two aims are compatible with a principled methodology for determining the (default) weights of reasons.

My strategy is easily generalized to any model that assigns deontic status in terms of weights, forces, pressures, or strengths. For example, Gert argues that deontic status is a function of justifying and requiring strength and that, consequently, Single Scale is false (2004, 2007). Yet he could insist (i) that *embedded* justifying and requiring strengths are a function of varying context and invariant default justifying and requiring strengths, while also insisting (ii) that deontic status is a function of *embedded* justifying and requiring strength. This simple move would make his (currently atomistic) alternative to Single Scale compatible with holism.

Sher (2019: 143-5) denies that there is a holist challenge to his model, but that's because he takes the atomism/holism distinction to ignore the issues discussed in §5.1. A holist challenge arises for his view if the following is possible: we hold fixed the reasons already taken into account during aggregation and yet we can still vary some feature F's magnitude, polarity, or whether it counts as a reason at all. It is a good thing, then, that my strategy can be extended to his model too.

### 1. Atomistic Single Scale

In this section, I present what I call *Atomistic Single Scale*. It has three components. Each component emerges from taking the scale metaphor seriously, at least when that metaphor is interpreted atomistically. (See my manuscript if you wonder why we should bother taking such metaphors seriously.) This discussion will help us understand in exactly which ways Single Scale is supposed to be incompatible

---

<sup>4</sup> Bader (2016) and Cullity (2018: 426) allow that holism and Single Scale are compatible, but neither claims that holism can be used in defense of additive aggregation. Indeed, Bader (manuscript: cf. 2016: 52-4) objects to the way I turn holism into a shield. See note 11 for discussion.

with holism. Throughout the paper we'll see that objections to Single Scale often result from *ignoring* salient aspects of the metaphors of weight and weighing.

Atomistic Single Scale's first component restricts which factors determine an act's deontic status. It makes deontic status a function of the relative heights of the pans. In other words:

**(Atomistic) Weight Comparativism:** An action's deontic status is determined solely by the relative weights of all reasons for and against it. (cf. Chang 2016: 213-5 and Portmore 2019: 8-11)

I intend 'determined solely' as an explanatory supervenience relation. There is no difference in deontic status without a difference in relative weights *and* relative weights explain why each particular action has the deontic status it does. Weight Comparativism takes no stand on what determines the weights of reasons. It constrains only the relation between (relative) weights and deontic status.

Different versions of Single Scale will endorse different assignments of relative weights to deontic status. A popular assignment holds that an act  $\phi$  is permissible iff the reasons against  $\phi$  are not weightier than the reasons for  $\phi$ . Less demanding assignments allow one to perform  $\phi$  even if one, within limits, has weightier reason to do something else.

Single Scale's second component concerns the analogy with *weight*. Things with weight exhibit counterfactual stability across contexts. The atomistic interpretation of this stability holds that things have the same weight—they exert the same force in the same polarity—across all contexts.<sup>5</sup> 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. 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. Hence, the second component is:

**(Atomistic) 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.

Prior Fixed Weight entails atomism. If something is a reason, then it has some weight or another. What particular weight it has includes its polarity (for vs against). Thus, Prior Fixed Weight entails that if something is a reason for (against)  $\phi$ , it is always a reason for (against)  $\phi$ . If it doesn't have a weight, then it isn't a reason in any context.

The third component concerns the aggregation of weight, or the relationship between the weights of individual reasons and the total weight of all reasons for  $\phi$ . The weight of everything on a pan just is the summed weight of each individual thing on the pan. Likewise:

**(Atomistic) Additivity:** the weight of all reasons for  $\phi$  is the summed weight of each reason for  $\phi$ . (cf. Dancy 2004: 9-10, 190 and Snedegar 2018: 726)

Dancy (2004: 190-1) claims that holism rules out Additivity. Additivity predicts that the weights of reasons combine in some regular way (addition). Holism predicts that the weights of reasons "combine in peculiar and irregular ways" (15). Suppose, for example, that diminishing marginal utility sometimes applies to a reason's weight: the more people that an action will benefit, the less weight that each benefit contributes

---

<sup>5</sup> Dancy (2004: 105, 190-1); Gert (2004: 73-77, 2007); Hawthorne and Magidor (2018: 133-5); and cf. Berker (2007, especially 122-4).

to the total. Let ‘Bert<sub>B</sub>’ and ‘Ernie<sub>B</sub>’ refer to Bert and Ernie’s receiving some benefit, respectively.

**Bert&Ernie:** In contexts where Bert<sub>B</sub> (Ernie<sub>B</sub>) is the only reason for  $\phi$  (only he can be benefited), the total weight for  $\phi$  is 10 units. In contexts where they are both reasons to  $\phi$  (both can be benefited), the total weight for  $\phi$  is 18 units.

Bert&Ernie, our illustration of diminishing marginal weight, seems to violate Additivity. Bert<sub>B</sub> and Ernie<sub>B</sub> each have a weight of 10 units when taken individually, which sums to 20. But, by hypothesis, the total weight is 18 when both are reasons for  $\phi$ . Thus, Bert&Ernie is a potential counterexample to Additivity.<sup>6</sup>

Atomistic Single Scale, then, entails the atomistic versions of Weight Comparativism, Prior Fixed Weights, and Additivity. To make Single Scale compatible with holism, we must refine Prior Fixed Weights so that it no longer entails atomism. This refinement to Prior Fixed Weights will force parallel refinements to Weight Comparativism and Additivity. These refinements will make Single Scale fit more tightly, not less, with the metaphor of weight.

## 2. A Standard Holist Framework

The holist asserts that a reason’s weight is sensitive to context in ways that the atomist denies. I make three assumptions which constrain how such context-sensitivity is to be explained. I call the conjunction of these three assumptions the **Standard Holist Framework**, because it isolates and polishes standard assumptions made by even extreme holists. If I can show that Single Scale is compatible with this framework, I will have shown that Single Scale is compatible with holism, thereby resolving the Holist Challenge to Single Scale. This framework will also underwrite the reverse engineering methodology that I defend at the end of the paper.

The *first assumption* is that there are reasons where, to be a reason, just is to have weight—or, in other terminology, force, pressure, favoring/disfavoring (cf. Dancy 2004: 37; Maguire and Snedegar 2021: 368). A reason’s weight includes both its polarity (for/against) and magnitude. (This account of weight, as well as Single Scale, ignores Gert’s 2004, 2007 distinction between justifying and requiring weight. We shouldn’t ignore Gert’s distinction, but that’s a lesson for another day.) *Magnitudes* specify how weighty a reason is and allow us to compare reasons for with reasons against. Suppose that a certain reason for  $\phi$  is equally weighty as a certain reason for  $\sim\phi$ . The two reasons have different polarities but the same magnitude.

Horty (2007) develops a sophisticated model of weighing reasons, partly to vindicate holism. His model may reject this first assumption (as does Draï 2018: 70-6). But my goal isn’t to show that holism entails Single Scale; it is to show that holism *doesn’t preclude* Single Scale.

To state the second assumption, we need a little more terminology. Let’s say that a **ground** is anything that is a reason in at least one context. Given the conceptual connection between reasons and weight (the first assumption), it follows that a ground

---

<sup>6</sup> The Bert&Ernie case is mine. Dancy (15-16), Draï (2018: 65-7), Horty (2012: 61), Hawthorne & Magidor (2018: 133-5), and Prakken (2005: 86) raise other cases which supposedly threaten Additivity. In §5, I discuss such cases, the now familiar response to them, and why the familiar response doesn’t address Bert&Ernie.

is anything that has weight for (against) some action in at least one context. A ground fails to be a reason in some context just when it fails to have weight in that context.

The Standard Framework's *second assumption* is that every ground has a default weight value.<sup>7</sup> The **default weight** of a ground R is the weight it has itself, the weight it has when abstracted from any given context (Bader 2016: 40). A given ground (one's promise to  $\phi$ ) can have default weight for  $\phi$  even though it fails to have weight in a given context (the context in which it was coerced). If so, the ground's default weight is distinct from its embedded weight. The **embedded weight** of ground R is just the weight of R taking into account its context. In effect, it is the ground's practical significance in that context. A ground's (embedded) weight is sensitive to context iff there is at least one case in which a ground's embedded weight differs from its default weight.

The default weight of a ground structures explanations for why a ground has the embedded weight it does. If a ground's embedded weight is equivalent to its default weight, "there is nothing to explain" (Dancy 113). Or, perhaps better, the default weight of a ground by itself explains why the ground has the embedded weight it does. On the other hand, if the embedded weight is *not* equivalent to the default weight, the ground's default weight does not, by itself, explain why a ground has the embedded weight it does.

The Standard Framework's *third assumption* is that deviations from the default require explanation of a specific sort. Changes in whether the ground has weight at all are to be explained by certain contextual features called *conditions*. There are two kinds, distinguished by their distinct functional roles. To be an **enabler** is to explain why a ground that has no default weight has embedded weight. To be a **disabler** is to explain why a ground with some default weight has no embedded weight.

These characterizations allow there to be higher order enablers and disablers. For example, suppose E1 enables G. Some fact E2 might be a second order enabler (roughly, an enabler of an enabler) insofar as it explains why E1 enables G when E1 doesn't enable G by default. The most extreme versions of holism tend to posit the existence of higher order enablers and disablers (e.g., Dancy 91).

A ground's **conditioned weight value** is the weight value it has taking into account just the ground and any condition (disabler or enabler) that may apply. Any further change to a ground's weight is explained by *modifiers*. There are two kinds, distinguished by their distinct functional roles. To be an **amplifier** is to explain why a ground has more embedded weight in a given context than it has conditioned weight. To be an **attenuator** is to explain why a ground has less embedded weight in a given context than it has conditioned weight.

In short, the *Standard Holist Framework* holds that (i) to be a reason is to have weight, (ii) reasons have default weight values, and (iii) any deviation from a reason's default weight value is to be explained by conditions and/or modifiers. This framework introduces a number of "tools", or functional roles: ground, default weight value, conditions (enablers, disablers), higher order conditions on conditions, conditioned weight value, modifiers (amplifier, attenuator), and embedded weight value. Let us refer to these roles collectively as the **Holist's Toolkit**. Dancy (52)

---

<sup>7</sup> Cf. Dancy (2004: 112-3, 117); Lance and Little (2007: §7); Cullity (2002: 188-9); and Bader (2016: 40).

claims that the holism (and particularism) of his *Ethics Without Principles* needs nothing more than the tools we've already mentioned.

This Standard Holist Framework is intended to be as normatively neutral as possible. It takes no stand on, say, which things have default weight or which contextual features count as conditions or modifiers. Our normative theories give us those details, filling in the structure that the framework provides. The framework is also *neutral* between atomism and holism. Even the atomist can allow that *if* embedded weight ever were distinct from default weight, it would have to be explained by conditions or modifiers. She just denies that there are any such deviations, and so she denies that there are any conditions or modifiers. This neutrality means that an argument for holism can rely on this framework without begging the question against the atomist.

### 3. A Motivation for Holism

The primary motivation for holism is case-based. The holist puts forward (contextual) contrast cases: cases that share some alleged ground R but R's context and alleged weight value differ. These cases support holism to the extent that there clearly is a contrast in R's weight, i.e., that exactly the same ground has different weight values in different contexts. Many of the earliest contrast cases are now widely rejected, even by holists.<sup>8</sup> There is increasing agreement, however, that our intuitions about partiality are best explained by holism.<sup>9</sup>

It is widely assumed that some forms of partiality—roughly, treating certain people or projects as more/less valuable than they really are—can be permissible and/or required. It seems that partiality or something associated with it, such as relationships, can be normatively relevant. A natural assumption is that partiality or relationships can be a distinct ground or source of reasons (e.g., Stroud 2010). Yet a holist might insist that sometimes the permissible/required partiality is explained, not by a difference of grounds, but by a difference in conditions or modifiers.

I can save my wife, I can save two strangers, but I can't do both. Intuitively, I am permitted, if not required, to save my wife. If we restrict our attention to "impartial reasons" (so that I have the same reason to promote my wife's wellbeing as I have to promote any other random person's wellbeing), it seems that I have double the reason to save the *two* strangers than I have to save my *one* wife. So how can we accommodate the intuition that it is permissible to save my wife? Let's restrict our attention to two hypotheses:

**Ground Only Hypothesis:** my relationship to my wife is a very weighty *reason* to save her (one at least as weighty as the life of another person) and the relationship does not modify my reason to promote my wife's wellbeing, and

**Modifier Only Hypothesis:** my relationship to my wife is a strong *amplifier* of my reason to save her (one that makes the reason provided by my wife's

---

<sup>8</sup> For example, holists Cullity (2002: 173-4) and Bader (2016: 29-30) criticize an early case of Dancy's (1993: 61).

<sup>9</sup> See, e.g., Keller (2013, especially 133-136, 152), Lord (2016), Bader (2016: 42-44), and Dorsey (2016: ch 6).

wellbeing at least as weighty as the lives of two people), and the relationship does not ground a reason to save my wife.

Both of these hypotheses claim that the relationship makes a big enough difference to explain how saving my wife is permissible, but they posit different explanations of how the difference is made. The Modifier Only Hypothesis provides the better explanation.

The Standard Holist Framework specifies the relation of each tool (functional role) to each of the other functional roles in the toolkit (enabler, disabler, ground, etc.). Yet we think there is more to many of these roles than the mere functional relations they bear to the other tools. And the holist does *not* take it to be arbitrary whether a consideration plays a modifier or grounding role. Some considerations, she thinks, are better suited to play the grounding role than others.

The ground of a reason plays two roles beyond its functional relations to the other holist tools.

First, the ground determines what is to be valued and for the sake of which we are to act...Second, the ground determines how we are to value and respond to it. (Bader 2016: 50)

For example, my wife's wellbeing is presumably a ground. It is to be valued for its own sake, and it calls for its promotion. Modifiers (qua modifiers) play neither of the ground roles. They play a secondary role: they affect "to what extent [the ground] is to be valued and how strong the reason is to which it gives rise" (Bader 49).

Suppose that I save my wife instead of the two strangers, and that my motivations perfectly track my normative reasons. The two hypotheses agree that my marital relationship makes saving my wife permissible. Yet they make rival predictions about my motivating reasons and phenomenology. The Modifier Only Hypothesis predicts that I have only *one* aim in acting (save her). I am acting only for my wife's great value. The relationship merely affects the extent to which I value saving her. I have one object of concern which, due to our relationship, is at least as valuable and weighty to me as both of the stranger's lives. I might deliberate: "I love my wife so much that her death is worse than the deaths of two strangers."

The Ground Only Hypothesis predicts that I have *two* aims in acting (save her, save marriage) *and* the relationship *doesn't* affect how I value her wellbeing. In addition to acting for the sake of my wife, I am also acting for the sake of our relationship's value, a relationship which will end if my wife dies. I have two distinct objects of concern—my wife and the relationship—that each are as valuable and weighty to me as the life of a stranger. I might deliberate: "If all I had to lose was my wife, I should save the two strangers. But I won't just lose my wife. I'll lose my marriage too!"

The latter deliberation is comically awkward. The former better captures what my (appropriate) motivations and phenomenology would be. I save my wife because, thanks to my relationship with her, I have at least as much concern for her wellbeing as I do for the combined wellbeing of the two strangers. If we have to choose between the Modifier Only and Ground Only Hypotheses, then we should choose the former. A more complete defense of this argument would consider other rival hypotheses, but we should stress two points. First, holism has more flexibility than atomism. Suppose you hold that the relationship is, strictly speaking, both a ground and a modifier and



you reject the Modifier Only Hypothesis for that reason. You still endorse holism. The holist can explain normative differences by appealing to conditions and modifiers and the atomist can't (§2).

Second, the atomist must find a way to explain why I appropriately have as much concern for my wife as the strangers combined. The holist explains this appropriate greater concern by appealing to amplification. If the traditional ground-based approaches to permissible partiality remain atomistic, they cannot explain it (Keller 2013). It is not clear how else the atomist could explain it. So the Wife or Two Strangers case seems to be an effective Contrast Case for holism. (For further defense of the idea that partiality involves modifiers and/or enablers—and so supports holism—see the references in note 9. For sustained argument that there exist disablers, see Cullity 2013).

#### 4. Holist Single Scale

If holism is both true and incompatible with scale-based models of weighing reasons, then we have a reason to reject those models. In this section, however, I show that Single Scale is compatible with holism *by* explaining how to make Prior Fixed Weights, Additivity, and Weight Comparativism compatible with holism.

##### 4.1. Prior Embedded Fixed Weights

Prior Fixed Weights makes the above version of Single Scale atomistic and incompatible with holism. It is correctly motivated by the idea that weight has counterfactual stability. For the metaphor of a reason's weight to be apt, a reason's weight in one context must systematically correlate with its weight in other counterfactual contexts. The atomist construal of Single Scale assumes that this stability is absolute: there is no difference in a ground's weight no matter how different the context.

The atomistic interpretation of stability is, in fact, a *departure* from physical weight. If the physical weight of my phone is 10 ounces in one context, then its weight will be 10 ounces in a wide variety of counterfactual contexts. Yet the physical weight of my phone is nonetheless sensitive to context. Cullity explains:

[Physical weight] is circumstance-relative, on both of its main definitions. On the gravitational definition (as the force exerted on a body by gravity), objects weigh less out in space than they do on the surface of the Earth. On the operational definition (as the force exerted by an object on its support), you weigh slightly less when the Moon is overhead. (2018: 426)

An object's physical weight, then, is sensitive to a limited number of contextual factors, such as the mass, distance, and relative location of other objects. Physical weight nonetheless has counterfactual stability: once you fix those specific contextual features, an object's weight remains invariant across counterfactual contexts. Physical weight is both context-sensitive and counterfactually stable.

The weight of a reason is likewise both context-sensitive and counterfactually stable. The weight of a ground is sensitive to a limited range of contextual factors, namely conditions (enablers, disablers) and modifiers (amplifiers, attenuators). A ground's weight nonetheless has counterfactual stability: once you fix those specific

contextual features, a ground's weight remains invariant across counterfactual contexts.

Counterfactual stability requires that weight be, at most, *finitely* context-sensitive (Gert 2007: 552-561). This rules out *particularism about weight*, the claim that the context-sensitivity of weight is *not* finitely specifiable. Hence, despite Dancy's protests to the contrary (e.g., 10, 78-85), we have good reason to deny that holism about weight supports particularism about weight. See Väyrynen (2006) for a different, more detailed argument that holism doesn't support particularism.

The original Prior Fixed Weights respects the counterfactual stability of a ground's weight, but not its context-sensitivity. It can respect both when refined as follows:

**Prior Fixed *Embedded* Weight:** the same *embedded* ground R (R + a given set of conditions and modifiers) has the same *embedded* weight in all contexts, and what particular *embedded* weight it has is prior to the deontic status of any particular action.

This refinement respects counterfactual stability, because a ground's weight can't vary when a limited number of contextual features (the features that count as conditions and modifiers) are held fixed. It respects context-sensitivity, because it allows a ground's weight to vary when those specific contextual features aren't held fixed.

#### 4.2. *An Improvement to Additivity*

The analogy with physical weight also motivates a replacement to Additivity. The combined weight of my three children will always be the sum of their weights in a given context, but the value of the sum may be different in different contexts (their combined weight is less in space than in my house). This change gives us:

The *embedded* weight of all relevant reasons for  $\phi$  is the summed *embedded* weight of each reason for  $\phi$ .

This change is enough to make additive aggregation compatible with holism. Yet we'll need to make further refinements in §5. These refinements concern the inputs to aggregation.

#### 4.3. *Embedded Weight Comparativism*

Now that we have the distinction between default and embedded reasons in mind, it should be easy to see how a holist-friendly Single Scale should modify Weight Comparativism. We put grounds on the scale to determine an act's deontic status only when those grounds have practical significance in the context. Only when they have embedded weight. For example, suppose that my promise to  $\phi$  has substantial default weight for  $\phi$ , but that it has no embedded weight in this context (perhaps because it was coerced and coercion is a disabler). In this context, the promise and its default weight are irrelevant concerning whether to  $\phi$ . We should replace Weight Comparativism with:

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

Atomistic Single Scale entailed the conjunction of Weight Comparativism, Prior Fixed Weights, and Additivity. In contrast, *Embedded Single Scale* entails the conjunction of Embedded Weight Comparativism, Embedded Prior Fixed Weights, and what we'll eventually call 'Embedded Additivity'. Embedded Single Scale retains

the scale metaphor because deontic statuses are determined solely by relative *embedded* weights. It is compatible with holism because embedded weights themselves are partly determined by context.

Conditions and modifiers are defined by their functional roles in fixing a reason's embedded weight. If there are no conditions or modifiers—if there is nothing that plays those functional roles—then a reason's default weight will always be equivalent to its embedded weight and atomism will be true (§2). Consequently, if there are no conditions or modifiers, the deontic verdicts of Embedded Single Scale will be extensionally equivalent to those of Atomistic Single Scale. While Embedded Single Scale is friendly to holism, it is just as friendly to atomism. **Holist Single Scale** is just the conjunction of holism and Embedded Single Scale.

## 5. Embedded Additivity

### 5.1. Individuating the Inputs to Aggregation

We have not yet finalized the principle that captures when Holist Single Scale is committed to additive aggregation. Any aggregation function—whether additive or not—needs an account of which reasons are inputs to aggregation (more precisely: which *embedded weights* are inputs to aggregation). An aggregation function needs to take into account the embedded weight, or practical significance, of all reasons in the context. Yet, to do so, it doesn't need to take every reason *as an input*.

Consider:

**Heat&Rain:** Suppose I am deliberating about an afternoon run, and that both heat and rain, taken individually, function as reason to not run; still, the combination of heat and rain together might function as a weaker reason to not run (say, because the heat is less onerous when there is rain). (Horty 2012: 61; cf. Prakken 2005: 86, Draai 65-7; Hawthorne and Magidor 133-5)

Even within a single context, two reasons to not run have more weight individually than they do when combined. Hence, Heat&Rain seems to be a counterexample to an embedded version of additive aggregation.

An increasingly common response to such cases is to restrict the inputs of aggregation to reasons that are *basic* or, in other terminology, *non-derived*, *fundamental*, or *load-bearing*. Heat and rain aren't basic or non-derived reasons. They are derived reasons: they are reasons only because they are causally relevant to something else that is a basic or non-derived reason, such as pleasure/pain or the respect of rights. It would be less miserable to run in dry, moderate weather than to run in both heat and rain; but running in both heat and rain would be less miserable than running in only heat or only rain. As the case was described, the total reason to not run is plausibly equivalent to the total misery of running in that context. If we restrict the inputs of aggregation to basic or fundamental reasons, Heat&Rain poses no threat to additive aggregation at all. By taking into account the embedded weight of all *non-derived* reasons, we thereby take into account the embedded weight of all *derived* reasons.

This restriction to basic/non-derived reasons is not enough to properly individuate the inputs to aggregation. Here are two genuine reasons not to shoot Vic: it would *cause pain* and it would *cause severe pain* (cf. Kearns' 2016: 186; Hawthorne &

Magidor 2018: 133-4; Maguire & Snedegar 2021: 368). Both reasons are basic or non-derived reasons. Its false that causing pain is a reason only because causing severe pain is a reason, and *vice versa*. An aggregation function *shouldn't* take the embedded weight of both reasons as inputs, at least not without subtracting something out. If an aggregation function takes the embedded weight of *causing severe pain* as an input, it thereby indirectly takes into account the embedded weight of *causing pain*. To take both as inputs is to double count the weight of causing pain. We need some restriction to ensure that the embedded weights of basic/non-derived reasons get taken into account *exactly once*. To see what form this restriction should take, let's consider when physical weight is additive.

Physical weight is not additive regardless of how you individuate the objects. I put my phone on the scale. I thereby put a second object on the scale, namely the phone's processor. The total physical weight on the scale is *not* the sum of my phone and the phone's processor. The weights of the phone and processor overlap. Physical weight is additive *if* you individuate objects so that their weights don't overlap.

What if, in some case, there is no way to individuate basic reasons so that their embedded weights don't overlap? Sher's sophisticated model of weight reasons raises this possibility and gives an answer to it (2019: 144-6; cf. Nair 2021: 48-49, 52). My answer to this question is more general. It appeals to the weight metaphor itself rather than the details of a specific model of weighing reasons. We get an answer to this question by thinking about the total physical weight of two objects when their weights overlap. Consider the weight of the top 60% of my phone and the weight of the bottom 60% of my phone. You don't just sum these weights to get the total weight of the phone. That procedure would double count the weight of the middle 20%. You sum these weights and then *subtract* the weight that you counted twice. In effect, what gets summed is the *non-redundant weight* of the two parts of the phone.

The probabilities of disjuncts are additive in exactly the same respect. If the probability of two disjuncts don't overlap, then their probabilities are additive. If the probability of two disjuncts *do* overlap, then you sum only their *non-redundant* probabilities (equivalently: you add their probabilities and then you subtract out the probability that you double-counted).

The Heat&Rain and pain/severe pain cases reveal, then, that the inputs of aggregation should be restricted to the *non-redundant* embedded weights of *basic* reasons (equivalently: the inputs of aggregation are restricted to the embedded weights of basic reasons and then you subtract out the embedded weight that you double counted). When we add these restrictions to the embedded aggregation principle from the end of §4.2, we get:

**Embedded Additivity:** the embedded weight of all relevant reasons for  $\varphi$  is the sum of the *non-redundant* embedded weight of each *basic* reason for  $\varphi$ . Embedded Additivity ensures that the embedded weights of all reasons are taken into account *exactly once*. Physical weight and the probability of disjuncts are additive in the same respect.

To be sure, Embedded Additivity does not require that there always be a *unique* "correct" input. Consider three reasons for  $\varphi$ : A, B, and the conjunction of A&B. Embedded Additivity demands that all embedded weight be counted *exactly once* during aggregation. Yet it might be neutral between one input (A&B), two inputs (A,

B), and even three inputs (A, B, A&B), as long as you subtract out the embedded weight that gets double counted with three inputs.

Now that we have formulated Embedded Additivity, you will likely want a theory concerning which reasons are the basic reasons whose weight is non-redundant. That's the job of first order normative theory, and different versions of Single Scale will give you different theories. That's hard work, and we won't have a complete defense of Embedded Additivity until we do it. But don't discount the progress we've made. We've seen that the stock examples that supposedly threaten Single Scale's aggregation principle attack additive aggregation when Single Scale isn't committed to it. If you want to show that Single Scale's commitment to additive aggregation is subject to counterexample, then you need to show that the counterexample applies to Embedded Additivity.

### 5.2. *Holism as a Shield*

Embedded Additivity ensures that there is one respect in which reasons combine in a regular way: non-redundant embedded weights of basic reasons always aggregate by addition. Recall our earlier example of reasons that "combine in peculiar and irregular ways" (Dancy 15):

**Bert&Ernie:** In contexts where Bert<sub>B</sub> (Ernie<sub>B</sub>) is the only reason for  $\phi$  (only he can be benefited), the total weight for  $\phi$  is 10 units. In contexts where they are both reasons for  $\phi$  (both can be benefited), the total weight for  $\phi$  is 18 units.

Cases of diminishing marginal weight, such as Bert&Ernie, threaten Embedded Additivity. They cannot be addressed by the restrictions discussed in §5.1. Both Bert<sub>B</sub> and Ernie<sub>B</sub> (the benefits to Bert and Ernie, respectively) will be inputs to aggregation, as long as the benefits are certain amounts of pleasure, wellbeing, or something else that is plausibly basic or non-derived. Bert's pleasure/wellbeing doesn't overlap with Ernie's pleasure/wellbeing, so the weights of these benefits don't overlap at all. Embedded Additivity entails, then, that their total embedded weight is the sum of their individual embedded weights. At first glance, however, Bert&Ernie seems to violate Embedded Additivity.  $10 + 10$  is 20, not 18.

The holist can reinterpret this apparent violation of Embedded Additivity as just another instance of the context-sensitivity of a reason's weight. Just let Bert<sub>B</sub> and Ernie<sub>B</sub> attenuate each other. Bert<sub>B</sub> (Ernie<sub>B</sub>) has an embedded weight of 10 units when Bert (Ernie<sub>B</sub>) is the only person who can be benefitted. When both Bert and Ernie can be benefitted, Bert<sub>B</sub> and Ernie<sub>B</sub> attenuate each other so that they each have an embedded weight of 9 units. This preserves Embedded Additivity.  $9 + 9 = 18$ .

Let us say that some version of holism is **inclusive** exactly when it includes the possibility that a ground can also be a condition or modifier. **Exclusive** holism excludes that possibility. I am not the first to use inclusive holism as a strategy to defend some sort of additive aggregation. Ironically, Dancy himself argues that apparent counterexamples to *restricted*<sup>10</sup> additive aggregation principles about value are better understood as further examples of value's context-sensitivity (e.g. 176-84, 189, 199-202). And, in one of these cases, he explicitly allows that grounds of value might also count as disablers. Dancy takes values and reasons to be closely linked

---

<sup>10</sup> Dancy (181) rejects the perfectly general claim that value is always additive.

(177). So Dancy himself should have no objection to my general strategy (reinterpret a putative counterexample to additive aggregation as a further instance of context-sensitivity) or my assumption that grounds can also be conditions or modifiers of each other (inclusive holism).<sup>11</sup>

To deal with the pain/severe pain case, it is common for philosophers to restrict the inputs of aggregation to reasons that are ‘independent’ of each other (e.g., Kearns 2016: 186, Nair 2016: 63). Inclusive holism requires us to tread carefully here. Bert<sub>B</sub>’s embedded weight *depends* on the ground, Ernie<sub>B</sub>, and vice versa. Each ground attenuates the other’s embedded weight. This failure of independence is irrelevant to aggregation. The particular way in which Bert<sub>B</sub>’s embedded weight depends on Ernie<sub>B</sub> does not make their embedded weights redundant. When you take into account Bert<sub>B</sub>’s embedded weight, you don’t thereby take into account Ernie<sub>B</sub>’s embedded weight. Hence, the correct aggregation function must take into account both of their embedded weights.

Although Bert<sub>B</sub>’s *embedded weight* depends on Ernie<sub>B</sub>, the grounds themselves are independent: they are not even partially constituted by each other. Perhaps that’s important for aggregation. But if it is, it is only because overlapping grounds have redundant embedded weights. What really matters is that Embedded Additivity sum the non-redundant embedded weight of all basic reasons.

You may be worried about other potential counterexamples to additive aggregation. Just keep in mind that there are two general strategies for defending Single Scale’s commitment to additive aggregation. The first general strategy is to reinterpret putative counterexamples as failures to rely on the correct account of aggregation inputs. If you take the pain/severe pain or Heat&Rain case to be a counterexample to Single Scale’s additive aggregation component, then you are relying on a faulty account of which reasons are inputs to aggregation. This first strategy is now well represented in the literature.<sup>12</sup> It led us to restrict additive aggregation to the *non-redundant* embedded weights of *basic* reasons.

This paper’s discussion of Bert&Ernie reveals a second general strategy that turns the Holist Challenge on its head: reinterpret putative counterexamples to Embedded Additivity as further instances of the context-sensitivity of a reason’s weight.

A referee worries that neither of these strategies can handle the following sort of case (adapted from Nair 2016: 66):

I’m considering a night out on the town. To get to town, I must pay a \$25 toll to cross the bridge. If I could both meet my friends at a restaurant and watch

---

<sup>11</sup> On the other hand, Bader (manuscript; cf. 2016: 52-4) argues that inclusive holism must be rejected to ensure that the debate between atomism and holism remains substantive. A factorization carves, or individuates, normatively relevant considerations into reasons, conditions, modifiers, etc. Bader argues that inclusive holism allows some normative theories to be compatible with both atomist and holist factorizations (manuscript: §3). He (§§1, 3) assumes that such compatibility would trivialize the debate between atomism and holism. This assumption is implausible. Even if a theory is compatible with both atomist and holist factorizations, one factorization can be normatively more adequate than another. In my §3, we considered an atomist factorization of permissible partiality (my marriage is a ground) and a holist factorization (my marriage is a modifier). The debate between atomists and holists is about which factorization, when combined with permissible partiality, provides the overall best account. Such a debate is clearly substantive.

<sup>12</sup> See, e.g., Gert (2004: 77-79); Bader (2016: 29); Kearns (2016: 186); Nair (2016: 63-70, 94-5); Cullity (2018: 424-5); and Maguire and Snedegar (2021).

a movie, the combined weight of the reasons to cross the bridge would outweigh the toll. But, as it happens, the restaurant reservation and the movie are at the same time. Hence, the total weight of the reason to cross the bridge is not any stronger than the individual reasons to do so (enjoying a meal with friends, enjoying a movie).

The challenge is to explain why, in this case, these individual reasons do not aggregate. The first strategy meets the challenge. The only reasons to cross the bridge are *derived* from your reasons to pursue some further ends. As the case was described, there is no way to enjoy both the meal and the movie. The individual reasons to take the means (cross the bridge) do not aggregate when those individual reasons are derived from reasons to take *incompatible* ends. To take both reasons as inputs to aggregation would be tantamount to aggregating the reason for  $\phi$  (enjoy the meal) and the reason for  $\sim\phi$  (enjoy the movie). Reasons for  $\phi$  and  $\sim\phi$  don't aggregate; they compete or oppose. Perhaps there is some alleged counterexample that can't be addressed by the above two strategies. And perhaps it also can't be addressed by a third strategy suggested by Wedgwood (2013: 52-3), namely reinterpret apparent counterexamples as cases in which complex interactions between reasons introduce new reasons. I still have shown that holism itself is no threat to Embedded Additivity.

### **6. Reverse Engineering and Holist Single Scale**

Drai (2018) raises a methodological challenge to the combination of holism and Single Scale. Consider how a proponent of Single Scale (whether atomist or holist) might determine what the (embedded) weights of reasons are. Single Scale holds that the relative weights of reasons explain deontic status. Yet the epistemic order often reverses the explanatory order. Our intuitions about whether an act is permissible are usually clearer and more plentiful than our intuitions about which reasons have more (embedded) weight than others. When our judgments about deontic status are more certain, we can *reverse engineer* the (embedded) weights of reasons: we know which acts are permissible and which aren't, we identify the weights of reasons which would explain that pattern of (im)permissibility, and then infer that those weights obtain. In a wide range of cases, we are permitted to sacrifice a lot to save five people from death. We have good reason to infer, then, that saving five people from death has lots of embedded weight in these cases.

Reverse engineering also allows us to assign relative values to different reasons. If you had to choose between an enjoyable cookie and saving the lives of 5 people, you are required to save the lives of five people. So, we should assign more embedded weight to the lives of 5 people than the enjoyment of a cookie.

The holist can extend this methodology to determine the *default* weights of reasons, as well as which factors are conditions and modifiers. Since we are permitted to sacrifice a lot to save five people from death in a wide range of cases, we have reason to infer that saving five people from death has lots of *default* weight too. Yet this inference is bound to be more tentative. Default weights, conditions, and modifiers have more *indirect* roles in fixing deontic status than the embedded weights of reasons. If we miss the presence of an enabler or amplifier in the wide range of cases that we consider, our reverse engineering can overestimate the default weight of saving five people from death.

Reverse engineering also can be applied to aggregation, which helps explain how holism can be a shield for Embedded Additivity. If diminishing marginal weight is genuine, then there is a systematic relation between the weight of each benefit when it occurs in isolation (10 units each) and how they contribute to the total embedded weight when they occur together (they contribute less than 10 units to the total). Attenuation is tailor-made to explain this systematic relation. Thus, if diminishing marginal weight is genuine, reverse engineering gives us some reason to infer inclusive holism and, more specifically, that the benefits to Bert and Ernie are both grounds and attenuators of each other.

Drai (67-8) worries that reverse engineering would *trivially* allow the proponent of Holist Single Scale to avoid any putative counterexample. No matter how plausible the counterexample, the holist could avoid it in an *ad hoc* way by inferring just the right combination of default weights, conditions, and modifiers. In reply, reverse engineering can be done well or badly. §3 provides a model for how to reverse engineer well. We saw how *non-trivial* it was to infer that partiality is a modifier rather than a ground.

Furthermore, you don't reverse engineer on a case-by-case basis. According to the Standard Holist Framework, default weights, conditions, and modifiers play *systematic* roles in fixing deontic status (§2). If partiality modifies in one case, then it modifies in a wide range of cases. If some holist reverse engineering were problematically *ad hoc*, this will reveal itself in how implausible its commitments are (or in arbitrary differences in how the particular holist appeals to modification). In short, it is the *principled* and *non-trivial* use of reverse engineering that allows us to infer the relevant default weights, conditions, and modifiers.

Drai will counter that even the principled and non-trivial use of reverse engineering is inadequate. To reverse engineer default weights is for our understanding of the "value of the part" (default weights) to depend on our understanding of the "value of the whole" (deontic status of actions). It follows, on pain of circularity, that the latter doesn't depend on the former. Consequently, our epistemology of default values "is empty in that it cannot explain how one's recognition of the value of the whole is determined by one's recognition of the values of the parts" (69-70).

Thankfully, it is no problem for a moral epistemology if it is empty in the relevant sense. It is hardly a problem for scientists that they can just see that there is a table, without first recognizing the protons, neutrons, and electrons that compose it. Nor is it any problem for the proponent of Holist Single Scale if she can just see that it is permissible to tell the truth in some circumstances, without first recognizing the default weights, conditions, etc. that would explain why telling the truth would be permissible.

Drai still worries that, if all we have is reverse engineering, then we "lack understanding of" default weight, or what she calls "prima facie weight" (72). If we don't have a direct access to default weights, can we really understand what a default weight is and know which default weights apply? Yes. Consider how scientists understand what electrons are and how they discriminate differences in electrons. Individual electrons can't be observed directly. They are theoretical posits. The concept *electron* is defined by its functional role in scientific theory. Scientists understand what electrons are by understanding that functional role, by understanding



electrons' role in explaining what we can observe more directly. This functional role allows us to *reverse engineer* the behavior of electrons. We infer what the electrons are doing based on our direct observations (e.g., there is lightning) and the role that links our observations with electron behavior.

Likewise, the Standard Holist Framework (and the further role specifications mentioned in §3) help us understand *what default weights are* and how to *discriminate differences in default weights*. (The parallel point holds for conditions and modifiers. I focus on default weights because Drai singles them out as especially problematic.) To understand what a default weight is you need only understand its functional role in the broader framework, including its role in explaining things we know more directly, such as whether an act is permissible. This functional role allows us to reverse engineer the default weights of grounds. We infer what a ground's default weight is by (i) considering the deontic status of actions in a wide range of cases that involve the ground and (ii) the role that default weight has in explaining those deontic statuses. This paper's use of reverse engineering is modeled on standard scientific methodology (at least as Lewis 1970 understands it). If scientific methodology is okay, then so is reverse engineering default weights.

Drai does have a separate objection that doesn't attack reverse engineering. Take two grounds, A and B. Which ground has more default weight? Drai holds that "The only way to answer this question is by looking at the counterfactual context where we put A next to [B] and compare between them" (73). As she rightly notes, however, this comparison only gives us A and B's *embedded* values in that specific situation. She's mistaken, however, when she claims that the only way to compare the weights of two grounds is by putting them *in the same context*. Just consider my wife's wellbeing in the actual context in which she is my wife and the counterfactual context in which we never met and remain strangers for our entire lives. Insofar as I'm the agent, it seems plausible that her wellbeing has much more (embedded) weight in the actual context than in the counterfactual one. Or compare the default weight of getting a little pleasure over the next five minutes with the default weight of getting a lot of pleasure over the next five minutes. Intuitively, getting a lot of pleasure would have more default weight than getting a little.

### Conclusion

Dancy and I agree that holism is true. We agree that a reason's weight is sensitive to context. Dancy further claims that holism is incompatible with Single Scale, roughly, the idea that an act's deontic status is determined by the relative weights of all the reasons for and against it. He is mistaken. The Standard Holist Framework is an account of how a ground gets its embedded weight: it is a function of varying context (conditions, modifiers) and a ground's invariant default weight. Once those embedded weights are established, Single Scale says that deontic status is determined by relative *embedded* weights. There is no conflict between these two ideas.

Holism is widely regarded as a threat to Single Scale's additive aggregation component. Many objections to additive aggregation dissipate when we realize that Single Scale is committed to additive aggregation only for the non-redundant embedded weights of basic reasons. Moreover, I showed that (inclusive) holism is not a threat, but a shield. Putative counterexamples to additive aggregation can be

reinterpreted as further instances of weight's context-sensitivity. This sort of reinterpretation relies on reverse engineering, the same sort of methodology scientists use to determine what's happening with particles that they can't observe directly.

Although I focused on the familiar Single Scale, my basic strategy is available to anyone who makes deontic status a function of weight, force, pressure, or strength. For example, Gert argues that deontic status is a function of justifying and requiring strength and that, consequently, Single Scale is false (2004, 2007; cf. Tucker 2022, forthcoming). Yet he could insist (i) that *embedded* justifying and requiring strengths are a function of varying context and invariant default justifying and requiring strengths, while also insisting (ii) that deontic status is a function of *embedded* justifying and requiring strength. This simple move would make his (currently atomistic) alternative to Single Scale compatible with holism.<sup>13</sup>

### References

- Bader, Ralf. (2016) 'Conditions, Modifiers, and Holism'. In Errol Lord and Barry Maguire (eds.), pp. 27-55.
- \_\_\_\_\_. (Manuscript) 'Admissible Factorisations and Hypertensional Independence'.
- Berker, Selim. (2007) 'Particular Reasons'. *Ethics*, 118, 109-139.
- Broome, John. (2013) *Rationality Through Reasoning*. Oxford: Oxford University Press.
- Chang, Ruth. (2016) 'Comparativism: The Grounds of Rational Choice'. In Errol Lord and Barry Maguire (eds.), pp. 213-240.
- Cullity, Garrett (2018) 'Weighing Reasons'. In Daniel Star (ed.), *The Oxford Handbook of Reasons and Normativity*. (Oxford: Oxford University Press), pp. 423-442.
- \_\_\_\_\_. (2013) 'The Context Undermining of Practical Reasons'. *Ethics*, 124, 8-34.
- \_\_\_\_\_. (2002) 'Particularism and Presumptive Reasons'. *Proceedings of the Aristotelian Society, Supplementary Volumes*, 76, 169-209.
- Dancy, Jonathan. (2004) *Ethics Without Principles*. Oxford University Press.
- Dorsey, Dale. (2016) *The Limits of Moral Authority*. Oxford University Press.
- Drai, Dalia. (2018) 'Reasons Have no Weight'. *Philosophical Quarterly*, 68, 60-76.
- Gert, Josh. (2007) 'Normative Strength and the Balance of Reasons', *Philosophical Review*, 116, 533-562.
- \_\_\_\_\_. (2004) *Brute Rationality: Normativity and Human Action*. New York: Cambridge University Press.
- Hawthorne, John & Ofra Magidor. (2018) 'Reflections on the Ideology of Reasons'. In Daniel Star (ed.), *The Oxford Handbook of Reasons and Normativity*. (Oxford: Oxford University Press), pp. 113-140.
- Horty, John. (2012) *Reasons as Defaults*. Oxford: Oxford University Press.
- \_\_\_\_\_. 'Reasons as Defaults'. *Philosophers' Imprint*, 7, 1-28.

---

<sup>13</sup> I wrote the initial draft of this paper using a generous summer research grant and sabbatical from William & Mary. Daniel Muñoz, Wes Siscoe, and anonymous referees provided very helpful comments on earlier drafts. Conversations with Philip Swenson and Walter Sinnott-Armstrong were also helpful, especially Walter's suggestion that diminishing marginal utility would usefully illustrate some of the issues in the paper. I am immensely grateful for this help from my institution and fellow philosophers.

- Kearns, Stephen. (2016) 'Bearing the Weight of Reasons'. In Errol Lord and Barry Maguire (eds.), pp. 173-90.
- Keller, Simon. (2013) *Partiality*. Princeton: Princeton University Press.
- Lance, Mark and Margaret Little. (2007) 'Particularism and Antitheory'. In David Copp (ed.), *The Oxford Handbook of Ethical Theory* (Oxford: Oxford University Press), pp. 567-94.
- Lechler, Andrea. (2012) 'Do Particularists Have a Coherent Notion of a Reason for Action?'. *Ethics*, 122, 763-772.
- Lewis, David. (1970) 'How to Define Theoretical Terms'. *Journal of Philosophy*, 67, 427-446.
- Little, Margaret. (2000) 'Moral Generalities Revisited'. In Brad Hooker and Margaret Little (eds.), *Moral Particularism*, (Oxford: Oxford University Press), pp. 276-304.
- Lord, Errol. (2016) 'Justifying Partiality'. *Ethical Theory and Moral Practice*, 19, 569-90.
- Lord, Errol and Barry Maguire (eds.). (2016) *Weighing Reasons*. New York: Oxford University Press.
- Maguire, Barry and Justin Snedegar. (2021) 'Normative Metaphysics for Accountants'. *Philosophical Studies*, 178, 363-84.
- Nair, Shyam. (2021) "'Adding Up" Reasons: Lessons for Reductive and Nonreductive Approaches'. *Ethics*, 132, 38-88.
- \_\_\_\_\_. (2016) 'How Reasons Accrue'. In Brad Hooker and Margaret Little (eds.), pp. 56-73.
- Portmore, Douglas. (2019) *Opting for the Best: Ought and Options*. New York: Oxford University Press.
- Prakken, Henry. (2005) 'A Study of Accrual of Arguments, with Applications to Evidential Reasoning'. *Proceedings of the 10<sup>th</sup> International Conference on Artificial Intelligence and Law, Bologna 2005* (New York: ACM Press), 85-94.
- Sher, Itai. (2019) 'Comparative Value and the Weight of Reasons'. *Economics and Philosophy*, 35, 103-58.
- Snedegar, Justin. (2018) 'Reasons for and Reasons Against'. *Philosophical Studies*, 175, 725-743.
- Titelbaum, Michael. (2019) 'Reason without Reasons For'. *Oxford Studies in Metaethics*, 14, 189-215.
- Tucker, Chris. (Manuscript) *The Weight of Reasons*.
- \_\_\_\_\_. (Forthcoming) 'The Dual Scale Model of Weighing Reasons'. *Noûs*. doi: <https://doi.org/10.1111/nous.12361>.
- \_\_\_\_\_. (2022) 'Weighing Reasons Against'. *Oxford Studies in Metaethics* 17, 77-107.
- Väyrynen, Pekka. (2006) 'Moral Generalism: Enjoy in Moderation'. *Ethics*, 116, 707-741.
- Wedgwood, Ralph. (Forthcoming) 'The Reasons Aggregation Theorem'. *Oxford Studies in Normative Ethics*.
- \_\_\_\_\_. (2013) 'The Weight of Moral Reasons'. *Oxford Studies in Normative Ethics*, 3, 35-58.