

***The Weight of Reasons***  
***A Framework for Ethics***

Chris Tucker  
Francis S. Haserot Professor of Philosophy  
William & Mary  
[cstucker@wm.edu](mailto:cstucker@wm.edu)

# Table of Contents

Acknowledgements .....	3
0. Introduction .....	7
0.1. Prologue: A Lesson from Little League .....	7
0.2. The Issues of Weighing and Weighting Reasons .....	8
0.3. The Aims of the Book .....	9
0.4. Some Comments on Complexity .....	13
Part I: Weighty Matters .....	15
Chapter 1: Minding the Metaphors of Weight and Weighing .....	16
1.1. Weight and Weighing.....	16
1.2. Weights and Weighing as Functional Structure .....	20
1.3. Two Functional Roles of Reasons.....	23
1.4. A “Forceful” Model of how Reasons Interact.....	26
1.5. The Dual Scale Model of Weighing Reasons .....	30
1.6. Metaphor Mongering or Minding the Metaphor? .....	34
Chapter 2: A Holist Balance Scale .....	39
2.0. Introduction .....	39
2.1. Atomistic Single Scale .....	40
2.2. A Standard Holist Framework.....	43
2.3. A Motivation for Holism.....	45
2.4. Holist Single Scale .....	48
2.5. Embedded Additivity .....	50
2.6. Reverse Engineering and Holist Single Scale .....	54
2.7. Holism and Scales: The Big Picture.....	57
Part II: From Single to Dual Scale .....	59
Chapter 3: The Representational Limits of Single Scale .....	60
3.0 Introduction .....	60
3.1. Justifying and Requiring Weight.....	61
3.2. The Limits of What Single Scale can Represent.....	68
3.3. What Single Scale Can’t Represent.....	71

3.4. Why Single Scale Can't Be the Fundamental Model.....	76
Chapter 4: The Representational Power of Dual Scale .....	81
4.1. Dual Scale .....	81
4.2. The Advantages of Dual Scale .....	86
4.3. Counteracting Single Scale's Distorting Influence .....	94
4.4. Alternatives to Dual Scale? .....	96
4.5. Plan B .....	103
4.6. How to Tell Which Model is Fundamental .....	105
Chapter 5: Weighing Reasons Against.....	106
5.1. Reasons For vs Reasons Against.....	107
5.2. Dual Scale and Reasons For/Against .....	110
5.3. Equivalence Established.....	111
5.4. Interlude: What are my Options? .....	113
5.5. The Which Alternatives Question .....	117
5.6. Wine-ing About Weight .....	122
5.7. From Weight to Dual Scale .....	123
5.8. Wrapping Up Part II.....	126
Part III: Ethical Puzzles Resolved .....	128
6. Against Weight Monism .....	129
6.1. Against Simple Weight Monism .....	130
6.2. Against Parity Monism.....	134
6.3. Generalizing the Attack on Weight Monism.....	138
7. Moral Options and the Normative Significance of Small Improvements .....	141
7.1. A Quick Review .....	141
7.2. Pluralist Permissibility and the Justification Transitivity Paradox .....	142
7.3. How to Explain Stable Moral Options .....	145
7.4. Altruistic Reasons: Pro Tanto Maximizing but Permissive .....	150
7.5. The Normative Significance of Small Improvements.....	153
7.6. Generalizing Pluralist Permissibility .....	156
7.7. Scoring Recap .....	159
8. The All or Nothing Ranking Reversal and the Unity of Morality.....	162
8.0. Introduction .....	162
8.1. Morally Better .....	164
8.2. Moral Commending Weight.....	167

8.3. The All or Nothing Problem for Pairwise Precedence .....	169
8.4. The All or Nothing Ranking Reversal Explained .....	171
8.5. The Unity of Morality .....	175
8.6. Conflicting Intuitions?.....	176
8.7. The Unity Explained .....	179
Part IV. Wrapping Up .....	181
9: Modeling Requirement for Any Number of Options .....	182
9.1. Generalizing the Permission and Commitment Assignments .....	183
9.2. The Fundamental Model of Weighing Reasons .....	189
9.3. Dual Scale’s Reliability.....	194
9.4. Better Tool: Single Scale or Dual Scale? .....	199
10. Conclusion.....	202
10.1. The Book’s Five Aims Revisited .....	202
10.2. Too Complicated to take Seriously? .....	204
References .....	206

## 0. Introduction

### 0.1. Prologue: A Lesson from Little League

It was a beautiful day for little league baseball on St. Simons Island, Georgia. Hot, but just overcast and breezy enough to take the edge off. Two men—well, 6<sup>th</sup> grade boys—were on base. The pitch was right where I wanted it. I swung, and everyone everywhere heard the most glorious *pling!*, the sound that only occurs when an aluminum bat perfectly strikes a baseball. The line drive barely went 15 feet high, but it shot over the left-center fence and, as I like to remember it, across the street where it rested among the majestic, mossy oaks that always line up to watch the games at Mallery Park. I was ecstatic. My teammates surrounded home plate, waiting to congratulate me. As I entered the throng, I picked up the bat, stepped over home plate, and continued to celebrate in the dugout. Until...until I heard those heartbreaking words: *Batter out!*

It turns out that I'm an accomplished athlete. I accomplished something that virtually no competent athlete in the history of the world has ever accomplished. I hit the ball over the fence in fair play, scored two runs, and yet somehow failed to get the home run. When I said I stepped over home plate, I meant it. I stepped *over* it. Not on it. (This was an early warning sign that I was destined for a career as an absent-minded professor.) And, if you don't step on home plate, you don't get the home run. You get called out and cry in the dugout like a middle-school boy who just had his dream ripped from his grasp.

I had always watched home runs with a spectator's awe, the sort of detached admiration in which you assume that such feats will never apply to you and so never bother to think them through. I knew the rules of baseball reasonably well. Had I ever bothered to think through what the rules of baseball imply about such feats, I would have concluded that one must step on every base *and* home plate to score the home run. But I never bothered to think through the rules. In my spectator's awe, the running around the bases always seemed like a victory lap. A celebration. And maybe that's true as far as it goes. But running around the bases is also a necessary part of the home run itself.

Sometimes we don't put 2 and 2 together. Sometimes things can be obvious upon reflection but they aren't obvious because we don't reflect. This book, especially the first couple chapters, is sometimes little more than the reflection needed to make obvious conclusions obvious. We ethicists frequently talk about weighing reasons, the force of reasons, one reason's being stronger than another, and so on. But there has been little systematic reflection on what this metaphorical talk implies or what the relation is between the metaphor of a reason's weight and the metaphor of a reason's force. As we'll see in the first two chapters, many criticisms of the weight metaphor result from not thinking the metaphor through.

If you want to make sure that you aren't missing anything about the rules for baseball, there is a single place you can go: the rulebook. Wouldn't it be nice if there was a single place that philosophers could go to understand the "rules" for weighing and weighting reasons? A single place that has thought through the metaphors well enough that it provides you with everything you need to know to

use *or criticize* them effectively. That's one thing this book aims to be, especially through Part I.

## 0.2. The Issues of Weighing and Weighting Reasons

Implicit in this first aim is both a subject matter and a methodology. The first aim is to be a relatively comprehensive treatment of certain general issues, weighing and weighting reasons, from a methodology that takes the metaphorical talk to be illuminating. Chapter 1 clarifies and defends this 'minding the metaphor' methodology. Here I want to clarify the general issues of weighing and weighting reasons. These general issues can be characterized independently of the metaphors, and a variety of methodological perspectives can be brought to bear on them, including perspectives that take the metaphorical talk to be misleading or unhelpful.

It appears that a lot of things matter for what to do, such as promises, rights, wellbeing, and so on. It also appears that a lot of things matter for what to believe, such as what perceptual experiences you have, what the prior probabilities are, what the experts say, and so on. When a lot of things matter for what to do or believe, the result doesn't have to be chaos. These things might interact in a systematic way to provide you with some all-things-considered, or all-in, guidance about what to do or believe.

An account of *weighing* reasons is, roughly, an account of how what matters interacts in a systematic way to determine what to do or believe. Such accounts abstract away from what matters and how much it matters. By itself, an account of weighing reasons is like an equation where none of the variables (the things that matter and how much they matter) have been specified. This equation might reveal general structural features of, e.g., morality or rationality, but the equation/structure by itself cannot tell us what to do or believe.

When the book is focused on weighing rather than *weighting* reasons, it is generally focused on clarifying how opposing considerations (e.g., the reasons for  $\phi$  and the reasons for some alternative) systematically interact to determine a normative status, such as permissible, impermissible, or supererogatory. Yet chapter 2 concerns two other two kinds of systematic interaction: how an individual consideration's weight might itself systematically result from context, and how the weights of individual considerations for an option systematically aggregate to determine the total, combined weight for that option.

An account of a consideration's *weight* is, roughly, an account of how much some consideration (e.g., rights) matters compared to other considerations (e.g., wellbeing). By itself, an account of *weighing* reasons gives us the values of the variables without giving us the equation to plug them into. It tells us how much each consideration matters individually, but it does not tell us how the weights of these individual considerations interact to give an overall verdict about what to do or believe. We don't get all-in guidance on what to do or believe until we combine an account of *how considerations are weighed* with an account of *how weighty* the various considerations are.

These distinct issues are nonetheless closely linked. For example, a primary way to test an account of weighing reasons is to see whether it can be combined with some

plausible account of weights to deliver plausible all-in guidance. After all, we care about weighing reasons partly because we want to understand what, e.g., morality tells us to do.

As I've characterized the two issues, a book on weighing and weighting reasons will inevitably concern both normative and meta-ethics. *Normative ethics* is, roughly, systematic theorizing about what morality, rationality, prudence, etc. tell us to do or aspire to. *Metaethics* is, roughly, systematic theorizing about the assumptions of normative ethical theories. Insofar as we are investigating the weights of reasons, we are doing *normative ethics* and, more specifically, we are trying to determine which things matter for what to do. Imogen has a cool iPhone, and you really want it. You shouldn't steal her phone, in part because her right to her own property matters more—it outweighs—the fact that you want it. Such normative ethical concerns are a focus of Part III.

Most of the book focuses, not on what the weights of reasons are, but on how they are weighed. This issue is more metaethical. A model of weighing reasons, by itself, takes no stand on what morality tells us to do. Nor does it take a stand on which things matter or how much they matter to morality's guidance. It is concerned with how morality must be structured to be in the business of giving this sort of guidance at all, or at least how morality must be structured when multiple things matter to its guidance (see §1.2 for why this caveat matters). Structural issues are arguably metaphysical ones, but they are distinct from the more traditional metaphysical issues about morality, such as whether it exists, whether it is objective, and whether it is non-natural. Weighing reasons falls under a subbranch of metaphysical issues we can call *structural metaethics*, the subbranch that concerns the structure of normativity and normative perspectives.

Since the book's focus is on weighing reasons, it is not directly concerned with moral epistemology. For example, it is *not* primarily concerned with how we know what's right/wrong, what our moral reasons are, or the methodology to be used when doing ethics. Nonetheless, the book needs to say something about these issues if it aspires to be a comprehensive rule book for weighing and weighting reasons. In this vein, I provide an account of how we often know what the weights of reasons are (§2.6) and a piecemeal methodology for evaluating models of weighing reasons (e.g., §2.7, §4.6).

### **0.3. The Aims of the Book**

In total the book has five aims. The first, as I mentioned already, is to provide a relatively comprehensive treatment of the issues of weighting and especially weighing reasons, and to do so *by* taking the metaphors of weighing reasons seriously. I think the metaphors of weight and weighing provide pretty good guidance as to how reasons function and interact. But, even if you disagree with me on this point, the approach can still be illuminating. Carefully examining the metaphors can help us understand the ways they mislead us. Chapter 1 defends this methodology and uses it to foreshadow some main conclusions of Part II.

The aim at comprehensiveness does mean that some sections delve into nitty-gritty details that can be skipped without losing the main narrative. I try to indicate those sections when they arise.

The book's second aim is largely completed in Part II and chapter 9. 'Single Scale' is my name for the model that comes most easily to mind when folks think about weighing reasons. The second aim is to show that Single Scale should be replaced with a model of weighing reasons that uses more than a single scale. I argue that the fundamental truths about weighing reasons are determined by a 'Dynamic Scale' model that has a scale for each alternative. The more alternatives there are, the more tedious it can be to apply this model. Hence, I tend to focus on a Dual Scale model that always uses two scales. Part II explains why Dual Scale is superior to Single Scale. Chapter 9 explores the difference between Dynamic and Dual Scale, and it reinforces the problems with relying on Single Scale.

Of course, no one will care about Dual Scale or the nuances of weighing reasons if these things aren't useful. The book's third aim, and the goal of Part III, is to apply the insights in the first two parts of the book to resolve a variety of ethical puzzles concerning the full range of permissible options and supererogation (roughly, going beyond the call of duty in a morally good way)

The final two aims of the book are themes throughout, and they concern a certain distinction. When we think through the metaphors of weight and weighing, we'll see that reasons have at least two weight values. A reason's *justifying weight* is how good the reason is at making acts permissible/okay. A reason's *requiring weight* is how good the reason is at making permissible acts required. This distinction is important, because some reasons are better at making acts permissible than making them required.

Suppose that, on a given occasion, you could do more good in the world by donating some of your money to an effective charity rather than using it to buy yourself a nice dinner out. You are not required to do the most good in this case. It would be ok to eat out and it would be ok to send the money to an effective charity. Neither action is required. Why? The self-interested benefits of eating out (it's enjoyable!) are better at making the act permissible than making it required. Self-interested benefits have more justifying than requiring weight.

By saying that self-interested reasons have more justifying than requiring weight, I thereby endorse **Weight Pluralism**, i.e., that reasons have at least *two* weight values and these two values aren't always equivalent. Reasons are multi-dimensional. The alternative, **Weight Monism**, holds that reasons have just a single weight value. Reasons are one dimensional. There is just *the* weight of reasons. Before I wrote this book, I thought Weight Pluralism was the runaway winner in this debate. So, I was puzzled by how much of the ethics literature seems to assume Weight Monism even when the authors were clearly aware of Weight Pluralism. If Pluralism is the runaway winner, then why did so many authors stick with Monism?

I found two partial answers to this question. The first is that I realized (sometimes thanks to referee 2) that the distinction between justifying and requiring weight needed further development and defense. It wasn't obvious, for example, how to apply the distinction to cases in which you have more than two options (see §4.1.3 and chapter 9). Nor was it immediately obvious how the justifying/requiring weight distinction mapped onto the reason for/against



distinction (see [ch 5](#)). Worse still, I discovered that there were some unanswered objections. Most importantly, a phenomenon associated with incommensurable goods—what I call the normative significance of small improvements—seems to give the distinction fits ([chs 6 and 7](#)).

The book's fourth aim swings for the fence (if I connect, I'll try to remember to step on home plate this time): develop and defend Weight Pluralism so that, when you add this book to the existing literature, the debate between Weight Monism and Pluralism will be completely one-sided. The existing literature—including the work of Josh Gert, Patricia Greenspan, Margaret Little, Coleen Macnamara, Daniel Muñoz, and Doug Portmore, to name just a few—already contains promising arguments for Weight Pluralism.

The existing arguments usually amount to some case or normative phenomena that Weight Pluralism explains better than Weight Monism. I discuss many of these cases, in part to have a comprehensive case for Weight Pluralism. To minimize redundancy, I usually discuss such examples only when I have some further aim in mind. Here are a couple examples. It is widely acknowledged that supererogation poses a problem for Weight Monism, assuming that it exists. I show that it poses a problem for Single Scale models of weighing reasons even if it *doesn't* exist (chs. 3-4). Muñoz uses the All or Nothing Problem as an argument for Weight Pluralism. I use it in this vein as well; however, I also use it to explore whether there is a third weight value (commending weight), give an account of 'ought', and defend Right > Wrong, the idea that permissible acts are better than their impermissible alternatives.

To the extent that I add new arguments to support Weight Pluralism, it is largely in connection to explaining why Weight Pluralism is better than a version of Weight Monism that appeals to incommensurability or Chang's notion of parity (chs 6, 7). My other contributions to the case for Weight Pluralism largely emerge out of grounding it in a systematic treatment of weighing reasons. This systematic treatment addresses (I hope) all previously unanswered objections that Monism is better than Pluralism.

There is one caveat to the goal of making the case for Weight Pluralism completely one-sided. If you think morality (or rationality) demands that you do the best that you can, then you can get by without appealing to Weight Pluralism. The *do the best* views have an attractive simplicity and impressive explanatory power; however, they don't fit well with our intuitions about the full range of permissible actions. It is permissible to buy yourself dinner out even though you could do even more good by sending the money to an effective charity. If you take that idea seriously, then it's false that you are always (morally) required to do the best that you can. The fourth aim, more precisely, is that adding this book to the existing literature makes the debate between Monism and Pluralism completely one-sided *given that* our intuitions about which actions are permissible are on the right track. Once you allow that we often have a wide range of permissible options, it is inevitable that we must appeal to reasons that are better at making acts permissible than making them required (i.e., reasons with more justifying than requiring weight).

In the end, though, I don't think that the apparent problems for the justifying/requiring weight distinction are the main reason that ethicists tend to assume Weight Monism. The main reason is that ethicists who are aware of the distinction often don't know how to *think with it*. It is one thing to be able to rattle off the difference between justifying and requiring weight. It is another to be able to apply it. In conversation, ethicists will emphatically tell me that they reject the justifying and requiring weight distinction and yet, when they present their own explanation of the full range of permissible actions, they commit themselves to the distinction in different terminology. They are aware of the distinction but don't understand it well enough to apply it to their own views, much less apply it in novel ways. The fifth and final aim is to help ethicists think with the justifying and requiring weight distinction well enough that they can apply it to resolve ethical puzzles themselves.

If you are reading this book (thanks, you are amazing!), you have a choice to make. You probably planned to start at the beginning and read through until you finish, become too bored to continue, or the zombie apocalypse occurs. That could be a good plan, depending on how much you prefer *normative ethics* (systematic theorizing about what morality and rationality tell us to do and aspire to) to *metaethics* (systematic theorizing about the assumptions of normative ethical theories). Starting at the beginning and reading through is a good plan if you like metaethics or you don't mind metaethics when you know that there is a normative payoff in the offing.

For those who want to do as little metaethics as possible, then the most minimal plan would be to read §1.3 and then chapters **6, 7, and 8**. A more moderate plan would involve:

*Chapter 1*: just read §§1.3-5, which clarify the justifying/requiring weight distinction and introduce Single and Dual Scale.

*Chapter 2*: this chapter is heavy on metaethics and can be saved for last; however, I recommend reading this chapter eventually if you aren't already familiar with how the context sensitivity of weight can explain normative phenomena.

*Chapter 4*: this is a deeper dive into Dual Scale, which largely develops the justifying/requiring distinction into a complete model of weighing reasons. This chapter includes discussions of satisficing (choosing an option because it is good enough), supererogation (going beyond the call of duty), and prohibition dilemmas (all options are impermissible). §4.2.3 and the final two sections, 4.5-6, can be skipped.

*Chapters 6 and 7*: The simplest Weight Monist views are incompatible with the full range of permissible actions. We must pay the price of complication. Weight Monists try to complicate the single weight relation by allowing weights to be vague or on a par, where parity is a fourth comparative beyond the traditional three (<, >, =). These particular complications can't get the job done. The Weight Pluralist's complication can explain the normative phenomena they explain and more besides. If you don't want to complicate normative theory beyond what is necessary, you should prefer Weight Pluralism.

*Chapter 8:* this chapter applies Dual Scale to resolve the All or Nothing Problem, it explains how one permissible action (the supererogatory one) can be morally better than a permissible alternative, and thus explains why what we ought to do need not always be what we are required to do. It also explains why Right > Wrong, i.e., why permissible actions are always better than their impermissible alternatives.

*Chapter 9:* Normative ethicists will have interest in Kamm's Café or Kid Case (§9.1.3), so reading through §9.1 *may* be worthwhile. Unless you like nerding out on the intricacies of modeling requirement when there are more than just a few options, then you should skip the rest of the chapter 9.

*Chapter 10/Conclusion:* I would read the whole thing (it's very short), especially §10.2.

This reading plan will give you the central elements of my model of weighing reasons and how that model can be applied to resolve ethical puzzles. It will spare you the extended discussion of what the metaphors of weight and weighing amount to (§§1.1-2); why we should carefully unpack these metaphors (§1.6); my extended attack on Single Scale (ch 3, 9.4); my account of the relation between reasons for and against (ch 5); and some of the finer details of and arguments for Dual Scale (e.g., §4.5-6, §5.6, and §§9.2-3).

#### **0.4. Some Comments on Complexity**

The idea at the heart of this book sounds simple: reasons are weighed to determine what to do. But it's not as simple as it sounds. At some point in the book, everyone is going to wonder whether the complexity is worth it. I address this worry at the very end of the book (§10.2), but it is worth addressing it here too.

Most of the book's complexity is conceptual, and almost every ethicist is committed to this complexity. The *do the best views* are, for example, committed to the conceptual distinction between justifying and requiring weight whether they realize it or not. The conceptual complexity provides a neutral playing field that allows us to properly assess the merits of competing normative theories. It also provides a model of weighing reasons that won't risk distorting the very phenomena it is used to model. Even proponents of the simplest views should appreciate such benefits.

Still, my normative views have more moving parts than the *do the best views* which generally make permissibility a function of some single factor. Such views hold that 'justifying weight' and 'requiring weight' are two different ways to refer to the same moving part (e.g., net pleasure). I hold that justifying and requiring weight are different moving parts, and so my view is more complex than theirs. The greater complexity is, in my view, an overall virtue. It allows my normative views to make much more reliable predictions about which actions are permissible.

A college wide receiver's speed, strength, and catching ability are conceptually different attributes, and they are all relevant to how the receiver would perform in the professional ranks. No one who understands how these attributes affect a

receiver's performance would claim that they can be reduced to a single factor. For example, no professional scout would claim that receivers who are equally fast are also equally strong and equally good at catching the ball. Such a claim runs contrary to what we observe when we watch receivers play.

Likewise, no one who understands how a reason's justifying and requiring weight affect deontic status will claim that a reason's justifying weight is always equal to its requiring weight. At least, no such person will make such a claim without *radically* revising our intuitive sense of which actions are permissible. I'm not after the simplest normative theory I can find. I'm after the simplest normative theory that doesn't require radical revision in our intuitive judgements about which actions are permissible. If the arguments of §2.3, Part III, and §9.1 are on the right track, then the simplest such theory will have lots of moving parts. But lots of moving parts doesn't mean an absence of systematicity. Part of the point of having a detailed model of weighing reasons is to structure the moving parts so that we can better understand their systematic interaction.

The best scout is not the one with the simplest theory about which receivers perform best in the professional ranks. The best scout is the one who best understands how the many relevant factors systematically interact to affect a receiver's performance. The best normative theory—whether it is mine or some other—is the one which best represents the many relevant factors and how they systematically interact to make acts permissible (required, supererogatory). Even if my normative theories are wrong-headed, this book still provides a framework that you can take with you. I hope it helps you better understand the systematic interaction of the (morally and/or rationally) relevant factors, whatever you take them to be.

## **Part I: Weighty Matters**

## Chapter 1: Minding the Metaphors of Weight and Weighing

### 1.1. Weight and Weighing

#### 1.1.1. The Basic Metaphors

In case you haven't noticed, 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 the act is permissible, impermissible, or required). These metaphors model the relation of reasons to deontic status on the physical *weight* of objects and *weighing* such objects on a (double pan) balance scale. Start with weight.

Physical weight is gradable and explanatory. It hurt when I dropped my laptop on my toe, but not when I dropped my pen on my toe. Why? The laptop *weighs more than* the pen by a wide margin. If reasons have weight, then they are both gradable and explanatory. Javon promises his friend to drive him to the airport. It would be permissible to cancel at the last minute if he has to rush his child to the hospital. It would be *impermissible* to do so if the gas would cost two cents higher than expected. Why? Saving his child's life *is weightier as reason* than an additional cost of two cents. More specifically, saving the life is *weighty enough* to justify breaking the promise and an additional cost of two cents is not.

The balancing metaphor further specifies exactly what a reason's weight explains, and how it explains it. In the abstract, there is no competition or opposition between the weights of different objects. The laptop weighs more than the pen, but, in the absence of some special context, those weights don't oppose or compete and the laptop doesn't win anything. The balance scale provides a special context in which weights *do* oppose or compete. It provides a special context in which weights behave like vectors, like they have both magnitude and direction. If you put the laptop on the left pan and the pen on the right, their weights are in direct competition. They push the pans in opposing directions. The laptop's weight pushes the left pan down and the right pan up. The pen pushes the *right* pan down and the *left* pan up.

The relative weight of physical objects determines which, if any, of the two pans is lower than the other. The left pan is lower than the right, because the laptop is weightier (i.e., the laptop pushes the left pan down harder than the pen pushes the left pan up). The relative weights of the reasons for an action and its alternative(s) determine whether it is permissible. If I surreptitiously swipe Sam's snack, I will feel a little happier but would violate Sam's right to his own snack. It would be impermissible to swipe the snack, because Sam's right to his snack is *weightier as a reason than* the small increase to my happiness. More specifically, it is impermissible to swipe the snack because my reason to do so is *outweighed* by the reason to do something else.

The language of one reason *outweighing* another and of "the balance of reasons" generally invoke the metaphors of weight and weighing. Such language is ubiquitous in ethical theory. This book is about these metaphors, what we can learn from them,

and how we can apply them to resolve ethical puzzles. It arose out of my own confusion.

My previous work on satisficing (roughly, rejecting the better for the good enough) made certain assumptions concerning how reasons are to be weighed. When I tried to articulate these assumptions to myself, I realized how little I understood the metaphors of weight and weighing. No worries, I thought. Given how ubiquitous these metaphors are, surely someone else would be able to explain them to me. Disappointment ensued. Neither proponents nor critics of the metaphors seemed to understand them any more than I did. Disappointment turned into opportunity. The opportunity for this book. Whether we use or reject the metaphors of weight and weighing, we should understand them first. We may find, as I argue in this book, that we can construct a promising model of how reasons interact to determine deontic status simply by cashing them out carefully. Let's continue this careful cashing out.

Actual balance scales have physical limitations and proper conditions for use. The laptop might not push its pan lower than the pen's pan because the scale is rusty or cleverly manipulated with magnets. If you put an elephant on the left pan and a tiger on the right pan, things will end badly for the scale...and for you. The metaphors of weight and weighing ignore such complications by focusing on an ideal scale in ideal conditions. An ideal scale can weigh any object(s) with physical weight against any other(s) and is sensitive to any difference of weight. The ideal context screens off any explanation of the relative heights of the pans besides (i) the starting positions of the pans, and (ii) the relative weights of the objects on the pans. It is usually assumed that the two pans of the balance scale begin at the same height, which indicates that the scale is neither biased toward the left pan nor the right. We will revisit this assumption in §3.2, but I work with it until then. Given an ideal context (and a non-biased starting position), there is no difference in the relative heights of the pans without a difference in the relative weight of the objects on the pans.

When we use balance scales to model how reasons determine deontic status when there are only two options, one pan represents an act,  $\phi$ , and the other pan represents the alternative,  $\sim\phi$ . Deontic status is assigned on the basis of the pans' relative heights, and so also on the basis of the relative weights of the reasons for  $\phi$  (the reasons on the  $\phi$  pan) and the reasons for  $\sim\phi$  (the reasons on the  $\sim\phi$  pan). A common view is that  $\phi$  is permissible just when the reasons for  $\phi$  are at least as weighty as the reasons for  $\sim\phi$ .

Given the balancing metaphor, then, there is no difference in deontic status without a difference in the relative weights of the reasons for and against  $\phi$ .<sup>1</sup> If today swiping the snack is impermissible (Swipe) and tomorrow kissing the cook is permissible (Kiss), then there is some difference in the relative weights of the reasons in those two cases. We previously noted that Swipe is impermissible because the reasons for Swipe (small increase in my happiness) are *less weighty* than the reasons against Swipe (violating Sam's right to his own snack). Perhaps Kiss is permissible because the

---

<sup>1</sup> It should go without saying that not just any difference between relative weights is relevant to deontic status. It would be *boring* if my reasons to eat spinach today were a little weightier than my reasons to eat lettuce. It would be *amazing* if my reasons to eat pie on pi day were  $\pi$  as weighty as my reasons to not do so. But such differences are irrelevant to deontic status, except insofar as they may or may not count as reasons for  $\phi$  or  $\sim\phi$ . It is only quantitative differences between the reasons for  $\phi$  and  $\sim\phi$  that determine deontic status, such as that the reasons for  $\phi$  are equally weighty, weightier, or less weighty than the reasons against  $\phi$ .

reasons for Kiss (the cook is my wife and she just asked for a kiss) are *weightier than* the reasons against Kiss (the dog will get jealous that I'm kissing her rather than him—true story, by the way).

### 1.1.2. *The Metaphors and Contemporary Ethical Theory*

The aptness of the weight and weighing metaphors is not trivial. For one thing, a normative perspective need not assign deontic status in virtue of how some relevant reasons interact. The US legal system is a normative perspective insofar as it issues verdicts of permissible (legal) and impermissible (illegal, prohibited). It is legally permissible to drink alcohol if you are 21 but not if you are 20 (sorry sophomores!). Yet the legal deontic status of an act is not, at least not normally, a function of reasons.<sup>2</sup> Lawmakers presumably weigh reasons when deciding whether to set the drinking age at 21 rather than 18. Yet once they make 21 the law, reasons play no role in making it illegal for my sophomores to join me for a dram of whiskey (I didn't really want to share anyway). All that matters is whether a law, or a rule, was broken or not. The metaphors of weight and weighing aren't apt when we talk about the deontic verdicts of the US legal system, and we don't talk like they are either.

Yet many of us do talk as though weight and weighing are apt metaphors when we consider the normative perspectives of morality and (practical) **rationality**, where rationality is the unique normative perspective that has final authority over what to do and aspire to.<sup>3</sup> (There is debate about whether morality and rationality are the same thing, i.e., whether morality has final authority over what to do and aspire to. For the purposes of this book, I am neutral on this issue.)

A *natural thought* is that what I ought to do morally or what I ought to do (period) is determined by the relative weights of the reasons for and against the acts that are available to me. But this natural thought is not trivial either. One way to deny its triviality is by insisting on certain fancy views about what reasons are. Given a fancy view of reasons, it might turn out that the things that get weighed to determine what is im/permissible don't count as reasons. For the purposes of this book, however, I don't care whether the things that get weighed are always best thought of as reasons. The idea of *weighing considerations* or *weighing pros and cons* works just as well for my purposes. For now, just think of reasons as those things that get "weighed" to determine what is im/permissible. I'll say more about how I'm using the term 'reason' in §§1.2-1.3.

I deny the triviality of the natural thought in a different way. It isn't trivial insofar as the metaphors of weight and weighing suppose that morality and/or practical rationality have a certain structure. Standard rule utilitarianism makes morality like the US legal system. The interaction of reasons doesn't determine whether your act is morally (im)permissible. All that matters is whether, e.g., the rule or moral code that

---

<sup>2</sup> Perhaps an exception is that weighing reasons determines which law trumps the other when two laws disagree about whether  $\phi$ -ing is permissible. Where are the jurisprudence scholars when you need them?

<sup>3</sup> What I say in this book is neutral on whether morality and practical rationality are objective or subjective. Roughly, if what we know can affect what our reasons are from these perspectives, then they are subjective. They are objective otherwise. The structural features that I'm interested in are separate from whether and how our knowledge affects which things are reasons.



brings about the best consequences, were it universally followed, permits or prohibits the act in question.<sup>4</sup>

Act utilitarianism is, however, a perfect example of how morality is structured if the weight and weighing metaphors are apt. Consider (*hedonistic maximizing*) *act utilitarianism*, the view that says  $\phi$  is permissible just when  $\phi$  causes at least as much net pleasure as  $\sim\phi$ , where net pleasure is just pleasure – pain. This claim is not stated in terms of the relative weights of reasons (or, if you prefer, the relative weight of considerations), but it entails such a claim. Pleasures and pains are “weighed” to determine which acts are im/permissible. The pleasure caused by  $\phi$  functions as a reason for  $\phi$  and the pain caused by  $\phi$  functions as a reason for  $\sim\phi$ . (It is reversed for  $\sim\phi$ : the pleasure caused by  $\sim\phi$  functions as a reason for  $\sim\phi$  and the pain caused by  $\sim\phi$  functions as a reason for  $\phi$ .) The more significant the pleasure/pain, the weightier the reason. So understood, if  $\phi$  causes at least as much net pleasure as  $\sim\phi$ , then the reasons in its favor are at least as weighty as the reasons for  $\sim\phi$ . Given act utilitarianism, then, relative weights determine deontic status even if proponents of act utilitarianism don’t use the terms “weights” and “weighing”. (See Gert 2004a: 73-7 and Berker 2007: 117 for similar takes on act utilitarianism.)

Decision theory (aka: expected utility theory) is normally couched in terms of *expected utility* rather than the weights of reasons. Titelbaum (2019: 200, 210) infers that the truth of decision theory would threaten the aptness of the weight and weighing metaphors. This inference is mistaken. Decision theory is a perfect example of how rationality would be structured if the metaphors of weight and weighing are apt. It is compatible with different views about which things are reasons or which things get weighed, but the simplest suggestion is that the expected utilities are themselves the reasons (cf. Sher 2019). Expected utilities represent a potential outcome’s *probability-adjusted* value (the value of some potential outcome multiplied by the probability of its occurring). Decision theorists think that a bird in the hand can be better than two in the bush. More precisely, they think that a guaranteed small good can *outweigh* a more significant but highly unlikely good.

(Maximizing hedonistic) act utilitarianism assumes that reasons are just pleasures and pains. The more significant the pleasure (pain), the weightier the reason. A (hedonistic) decision theory might claim that reasons are the *expected utilities* of pleasures and pains. The greater the expected utility (which takes into account both the significance *and* probability of the pleasure/pain), the weightier the reason. Instead of putting the pleasures and pains themselves on the pans, you put their expected utilities on the pans.  $\phi$  is permissible just when the reasons in favor of  $\phi$  (the total expected utility of  $\phi$ ) are at least as weighty as the reasons on the  $\sim\phi$  pan (the total expected utility of  $\sim\phi$ ).<sup>5</sup>

---

<sup>4</sup> Perhaps standard rule utilitarianism and/or the US legal system just entail a very *uninteresting* balance scale model. If  $\phi$  is in accordance with the relevant rule, then that fact goes on the  $\phi$  pan, making  $\phi$  permissible. If  $\phi$  breaks the relevant rule, then that fact goes on the  $\sim\phi$  pan, making  $\phi$  impermissible. The more that the weight and weighing metaphors are apt for standard rule utilitarianism and the US legal system, the more plausible it is that the weight and weighing metaphors are apt for morality and rationality.

<sup>5</sup> Alternatively (and probably preferably), if we appeal to the sort of holism explained in chapter 2, the decision theorist can say that potential pleasures and pains are still the reasons that go on the pans, but their weights are attenuated when it is less than certain that they will occur. The weight of the pleasures

I assume that weight and weighing are apt metaphors for morality and practical rationality. (In addition to what I say in this chapter, I provide an additional reason to think the metaphors are apt in §7.3.3.) *If* the metaphors of weight and weighing are apt for morality and practical rationality, then reasons interact to determine deontic status (contra, standard rule utilitarianism) *and* the best way of cashing out these metaphors is a contender for the best model of how reasons determine deontic status.

It is also common to assume that the weight and weighing metaphors apply to *epistemic* rationality—e.g., that what you ought to believe is determined by the relative weights of reasons. I set epistemic rationality aside for the purposes of the book. As the above suggests, I think most (but not all) prominent positions in ethics are compatible with the aptness of the metaphors even when they don't explicitly invoke them. In contrast, reliabilism, epistemic virtue theory, and proper functionalism are all widely endorsed in epistemology, and I'm not sure that any of them fit well with the aptness of the metaphors. I want to focus on ethics where the metaphors have the best chance of being apt, and I save discussion of weighing reasons in epistemology for another occasion.<sup>6</sup>

## 1.2. Weights and Weighing as Functional Structure

This book is about the structure of morality and rationality, given that the metaphors of weight and weighing are apt. It is *not* about deliberation or how we ought to decide what to do. I do *not* assume that proper deliberation requires that one engage in some activity of weighing reasons. Insofar as deliberation is an act, it is subject to moral and rational norms just like any other act. Still, we might wonder: if morality has a balance scale structure—if “morality weighs reasons”—should our deliberative activity reflect that structure, i.e., should *we* weigh reasons in our deliberations?

When you face a difficult and life-changing decision, it makes sense to spend a few days writing out all the relevant reasons and then weighing them in your deliberations. But it would be silly to go through this process every time you consider having a second cup of coffee. I do not need to deliberate to tell that it would be a bad idea to take a pellet gun and shoot the annoying neighborhood children who have, yet again, begun loudly playing in my backyard without permission. I think we rarely need to deliberate by engaging in some activity of weighing reasons. In any event, deliberation is separate from what I'm talking about in the book. (See Cullity 2018 for discussion of how what I'm talking about might be relevant to deliberation.) Whenever I talk about the weight and weighing metaphors, they are to be understood as claims about the structure of a normative perspective, such as morality, rationality, or the US legal system.

*Normative perspectives* take certain inputs (e.g., reasons, whether a law was broken) and output verdicts of permissible or impermissible. (Strictly speaking, a normative perspective might only issue verdicts of better/worse to actions, but I assume that morality and rationality are deontic insofar as they assign im/permissible

---

and pains would then mirror their expected utility. This holist decision theory would be a second way to show that decision theories have the structure that makes the metaphors apt.

<sup>6</sup> For some discussion of how the ideas of this book relate to epistemology, see my 2024a and 2024b.

to actions.) They are functions in the sense of having inputs and associated outputs.<sup>7</sup> Different normative perspectives can have different functional structures. Consider an analogy. Put certain things in either of two vending machines, and it will give you a coke. Nonetheless, they may be structured to take different kinds of inputs. One might take only coins and the other only credit cards. Or they may take the same kinds of inputs but be structured to give you the coke via very different causal processes—imagine that Rube Goldberg invented one of the two vending machines (i.e., imagine the causal process is unnecessarily and comically complex). Functional structures determine what kinds of inputs matter and how those inputs give you the output.

To say that the weight and weighing metaphors are apt for a certain normative perspective is to make a claim about that normative system's functional structure. It is to say that the normative system takes certain things as inputs (the weights of reasons) and goes from inputs to deontic outputs in a specific way (in virtue of the relative weights of those reasons). The US legal system seems to lack this structure. It doesn't (normally) take reasons as inputs, and it doesn't weigh reasons to determine what's permissible. I assume that the metaphors of weight and weighing apply to both morality and rationality, and so I assume that they have much functional structure in common.

If morality and rationality share the same functional structure, they still might be distinct normative perspectives—i.e., morality may fail to be the normative perspective that has final authority over what to do—if they assign weights differently. For example, morality might assign less weight to self-interested reasons than rationality, and so those perspectives may disagree about whether it is permissible for me to spend my summer salary on a new car rather than an effective charity. Such a disagreement between morality and rationality would be a substantive normative disagreement rather than a structural one. They determine deontic status in the same way, by weighing the reasons for and against; they just disagree about which things are reasons or how weighty certain reasons are. (When I use examples of some normative phenomena or take a stand on how to weight reasons, I generally have in mind how morality weights reasons. The main exception is that my potential examples of prohibition dilemmas are generally more promising as rational rather than moral dilemmas.)

Where there is functional structure, there are things with functions, or roles to play, in a larger system. The legislative branch has a function in the larger US government: it makes the laws. The executive branch has a distinct function: it enforces those laws. Reasons—or whatever we want to call the things that get weighed—have their own role to play within morality and rationality. They push an act toward (im)permissibility, i.e., they make an act (im)permissible in the absence of sufficiently weighty countervailing considerations.<sup>8</sup> If there is such a thing as a disabler, then it

---

<sup>7</sup> See Dorsey (2016, ch 1) for development of this basic idea.

<sup>8</sup> A couple clarifications. First a caveat: derived reasons are reasons in virtue of their relation to basic or non-derived reasons; they explain or push an act toward a deontic status only in a derived way. I usually ignore them because they make no difference to an act's deontic status, as we'll see in §2.5.1. Second, my definition amounts to an *explanatory role* approach to defining reasons. Standard explanatory approaches assume that there must be a single kind of thing that (non-derived) *reasons for* always explain. Fogal and Risberg (2023) provide the most sophisticated defense of such a view. I can tweak my view to fit Fogal and Risberg's framework; however, my working view is that *reasons for*

too will have a characteristic functional role, roughly, preventing something that would ordinarily be a reason from being a reason in the context at issue. The arguments of this book assume *bare* functionalism about normative perspectives, the claim that normative perspectives have functional structure and that normative concepts (e.g., *reasons* and *disablers*) have associated functional roles in that structure.

I'm not sure that the weight and weighing metaphors make sense outside of bare functionalism. If these metaphors aren't claims about the functional structure of a normative perspective, then I don't what they are (at least not when they are understood as claims about a normative perspective rather than as claims about deliberation). I hold, then, that bare functionalism is built into the weight and weighing metaphors and so doesn't need separate defense.

When writing this book I found it convenient to go beyond bare functionalism and assume *conceptual* functionalism, which adds that many normative concepts, such as reasons and disablers, are *defined* by their functional roles. One might think, for example, that the executive branch is, by definition, that branch which enforces the US laws. Likewise, I work with the idea that a reason is, by definition, to be something that pushes an act toward (im)permissibility. This idea can be a source of versatility and frustration. It is versatile insofar as it takes no stand on the metaphysics of reasons. It is compatible with the idea that reasons are normatively fundamental entities (Schroeder 2021), that they are always values (Maguire 2016), that they boil down to fittingness (McHugh & Way 2016), and so on. You tell me which things you think play the relevant systematic role, and that turns out to be your account of which things are reasons, as I've defined 'reasons'. That's why, in §1.1.2, I said that claims about, say, expected utility end up entailing claims about reasons.

This versatility can be frustrating if you don't think 'reasons' is best defined by certain roles in making things permissible and impermissible. Just remember that, for the purposes of this book, I don't care about the label 'reasons'. I care about the structure and the relevant functional roles. What is most helpful about the metaphors is that they can be a springboard for systematic theorizing about such functional structure. Ideally, the minding the metaphor methodology will help you think directly about this structure and understand it well enough that you will eventually no longer need the metaphors.

Functional role concepts come cheaply. You can just define a 'schmiggle' as the unique thing which wiggles everyday and twice on Tuesday, and *voila!* you have a new concept. But not all functional role concepts refer or are interesting. I take it that 'schmiggle' fails to refer because there is no single thing that wiggles in the relevant way, and even if there is, nobody gives a flying flip about it. The minding the metaphor methodology is in the service of developing a functional structure. There is a risk—especially if views like rule utilitarianism or error theory are true—that this structure fails to refer to any genuine normative reality. But I take it that this structure is at least interesting because, as I argued in §1.1.2, something like this structure is arguably implicit in most ethical theories.

---

need not always explain some specific kind thing; they need only play a specific explanatory role in fixing deontic status. This latter explanatory role approach may be behind Broome's talk of 'weighing explanations' (Broome 2013: 52), but there are some differences between our accounts of reasons (§3.1).

I work with *conceptual* functionalism because of how tractable it makes the relations between normative concepts. It makes it very easy to see what follows by definition and what doesn't. Titelbaum (2019: 194, including nt 5) is willing to assume bare but not conceptual functionalism. He is willing to assume that normative concepts have associated functional roles but not that these roles *define* those concepts. This might lead to some disagreement about which truths are true by definition, but that's it. For example, I argue that *it is true by definition that every reason for  $\phi$  is a reason against some alternative* (§5.1). Titelbaum's reluctance to assume conceptual functionalism may lead him to wonder whether the claim is true by definition; however, his reluctance will give him no reason to deny that the claim is true, i.e., that every reason for  $\phi$  is a reason against some alternative. I find it convenient to work with conceptual functionalism, but bare functionalism is all I need.

Functionalism in ethics is often associated with naturalism, and more specifically, the Canberra Plan. The US legal system specifies the definitional role of the US President. To find the US President, we identify the thing in the natural world that plays that role. The Canberra Plan says the same thing holds for figuring out which things are reasons. First, you first define the functional role of reasons. Then just find the things in the natural world that (best) play the reason role. (See, e.g., the essays in Braddon-Mitchell and Nola 2008, especially their introductory chapter.)

The functionalism in this book simply constrains the way that normative concepts relate to one another, and it takes no stand on how normative concepts relate to naturalistic or non-normative ones. I do not assume that reasons are fundamentally natural things but nor do I assume that they are fundamentally non-natural. Nor is the functionalism in this book intended to constrain our moral epistemology. I do not assume, for example, that to know that R is a reason, you must know what the functional role of a reason is. My kids couldn't tell you what the functional role of a reason is. Yet they seek out and bring to my attention every conceivable reason for the claim that they should be given even more tech time each day. (If they had to choose between another moment on their devices and my life, I'm not sure I would make it.)

### 1.3. Two Functional Roles of Reasons

A number of theorists hold that oughts and requirements are distinct. Perhaps you ought to complete the referee report this week even though you are only required to submit it next week. Perhaps you ought to stay an extra 10 minutes to continue helping a student even though you already did what you were required, namely help them during office hours. I discuss weighing reasons for oughts and supererogation in chapter 8. Until then, I'm focused on how to weigh reasons for permissibility and requirement. This work will help us build up to accounts of ought and supererogation, as they have a deontic component that concerns permissibility and requirement (an option's being morally permissible but not required) and a relative ranking component (e.g., an action's being morally better than some permissible alternative or, in the case of ought, better than all alternatives).

I said that to be a reason is, by definition, to be something that pushes an act toward (im)permissibility. Notice that there are two different functional roles here, pushing an act toward permissibility and pushing an act toward impermissibility. Following

my dear colleague Josh Gert (2003, 2004a, 2007, 2016), I use the terms ‘justifying weight’ and ‘requiring weight’ to track these two roles, respectively. This distinction is central to the book. Part II largely just provides a formal characterization of the distinction that is independent of the metaphors, and then it unpacks the functional structure that emerges from it. Part III applies the distinction to explain certain puzzling normative phenomena. For the moment, let’s just try to get an intuitive grip on what it amounts to and how justifying and requiring weight interact. (If you already accept the justifying/requiring weight distinction and find yourself getting annoyed by my way of characterizing it, you might consider reading §3.1.2 before continuing. The differences between characterizations usually don’t matter, but they sometimes do.)

A reason’s *justifying weight* is how good the reason is at making an act permissible. The better it is at making an act permissible, the harder it pushes the act toward permissibility. A reason’s *requiring weight* is how good the reason is at making a permissible act required. Permissible acts are required when the alternative is impermissible, when it is ok to do the act and not ok to do anything else. Hence, the better a reason is at making an act required, the harder it pushes *the alternative* toward impermissibility. (As we’ll see in the next section, we can use the terminology of justifying and requiring *force* if you prefer. The metaphors of weight and force get at the same thing.)

These characterizations imply that justifying and requiring weight oppose to determine whether an act is permissible. To be sure, there is no interesting opposition between the justifying and requiring weight for the same option,  $\phi$ . A reason’s justifying weight for  $\phi$  is how hard it pushes  $\phi$  toward permissibility. A reason’s requiring weight for  $\phi$  is how hard it pushes the alternative,  $\sim\phi$ , toward impermissibility. Hence, justifying and requiring weight for  $\phi$  push different things ( $\phi/\sim\phi$ ) in different directions (permissible/impermissible).

In contrast, justifying weight for  $\phi$  and requiring weight *for*  $\sim\phi$  push the same thing,  $\phi$ , in different deontic directions. A reason’s justifying weight for  $\phi$  is how hard the reason pushes  $\phi$  toward permissibility. A reason’s requiring weight *for*  $\sim\phi$  is how hard the reason pushes  $\phi$  (i.e., the alternative to  $\sim\phi$ ) toward impermissibility. It is this competition that determines whether  $\phi$  is permissible:  $\phi$  is permissible just when the justifying weight for  $\phi$  is at least as great as the requiring weight for  $\sim\phi$ .

This criterion for permissibility won’t be useful unless we can tell how much justifying and requiring weight reasons have. How do we do that? Consider an analogy. A player’s performance in some sport might depend on both his upper and lower body strength. We measure a player’s lower body strength by asking *how much weight can be on the bar and the player still squat it?* We measure a player’s upper body strength by asking *how much weight can be on the bar and the player still bench press it?* These questions provide us with an easy way to measure how much stronger a player’s lower body is than their upper body. We determine how much weight he can squat and then we keep taking weight off the bar until he can bench press it too. Perhaps a first player can squat 500 pounds but can press only 250 pounds. His lower body is about twice as strong as his upper body. (This same strategy leads us to conclude that different players can have different proportions of lower and upper body strength. A second player’s upper body might be almost as strong as his lower body, squatting 500 and pressing 400 pounds.)

Just as some players have more lower than upper body strength, some reasons have more justifying than requiring weight, i.e., some reasons are better at making acts permissible than making them required. The proportion of a reason's justifying and requiring weight allows us to track just how much better. Let's revisit an example from §0.3. The self-interested reason *that you would enjoy it* makes it permissible to eat out, even though you could do even more good by donating the money to an effective charity. Suppose that the charity feeds hungry children from developing nations. Now ask two questions (cf. Gert 2016: 158):

How many children could you otherwise feed and it still be permissible to eat out?

How many children could you otherwise feed and it still be *required* to eat out?

The answer to the first question tells us how good self-interested reasons are at making acts permissible, and so it tells us how much justifying weight they have. The answer to the second question tells us how good self-interested reasons are at making (permissible) acts required, and so it tells us how much requiring weight self-interested reasons have.

We can do a *lot* of good by giving to effective charities, so self-interested reasons have to be *really* good at making acts permissible (cf. §6.1.2). Presumably, it is permissible to eat out even if you could otherwise feed 5 hungry children for a single meal. Yet you are hardly *required* to eat out in such circumstances. You do nothing wrong if you donate to feed those hungry children. You are not even required to eat out if you could otherwise feed a single hungry child. This simple example suggests that self-interested reasons have *at least* 5 times more justifying weight than requiring weight. They are good at making acts permissible but bad at making them required.

How bad are self-interested reasons at making acts required? Some ethicists think that self-interested reasons have no capacity to make actions (morally) required at all. If you had to choose between eating out on this occasion and burning the money in your fireplace (which you wouldn't enjoy as much), they think it would be morally permissible to burn the money. I often work with this position because it is simple and popular. But, in §7.6.2, we'll see that self-interested reasons must have some ability to make acts required (have at least a little requiring weight) to account for certain normative phenomena.

Those who explicitly allow reasons to have more justifying weight than requiring weight include Archer (2016); Gert (2003, 2004a, 2007, 2016); Little & Macnamara (2017, 2020); Massoud (2016); Muñoz (2021); Muñoz & Pummer (2022); Murphy (2017: ch 3); Portmore (2008, 2011); Tucker (2017, 2020, 2022a-c, 2023a); and Whiting (2022: ch 3). Others make the same distinction in different terminology. 'Prerogatives' and 'mere permissions' are common ways to refer to things with justifying but no requiring weight.<sup>9</sup> Those who, *in some terminology or another*, allow justifying weight to outstrip requiring weight include Greenspan (2005); Hurka and Shubert (2012); and Scheffler (1982: ch 3).

---

<sup>9</sup> Agent-neutral value may serve as another way to talk about things that can justify but not require (see §6.2.2n9).

The distinction between justifying and requiring weight can sound like a dramatic departure from the most straightforward ways of unpacking the weight metaphor. There is just one kind of physical weight, so it may seem that reasons can't have two weight values. These impressions are understandable but mistaken. We'll see that the justifying/requiring weight distinction is built into the metaphors of weight and weighing, as well as other common metaphors.

#### 1.4. A “Forceful” Model of how Reasons Interact

##### 1.4.1. Other Metaphors for how Reasons Interact

Weight and weighing are not the only metaphors used to understand the interaction of reasons. Other common metaphors include a reason's force, pressure, support, strength, or power. Some philosophers tentatively suggest that force provides a better metaphor than weight (e.g., Hawthorne & Magidor 2018: 134; Fogal & Risberg 2023: 70-1). Does it matter which of these metaphors we mind? Not really. In this section, I argue that my main conclusions don't depend on some idiosyncratic features of the weight and weighing metaphors.

My view is that these other common metaphors pick out essentially the same functional structure as weight. In a way, a reason's *force* is the master metaphor. You put an object on the  $\phi$  pan. Its *weight* just is its downward force or pressure on that pan. The *strength* or *power* I exert on an object is the force I bring to bear on it. You should feel free to focus on whichever of these metaphors you find most useful. I'll just explain why I focus on the metaphors that I do.

I rely most directly on the metaphors of weight, force, and pressure (or pushing). I generally ignore the metaphors of support, strength, and power. Consider support. Reasons oppose, and it is nice to have a metaphor, such as a balance scale, that can illuminate this opposition. The metaphor of a reason's *providing support* (e.g., Fogal and Risberg 65)—is less illuminating in this respect. Two objects, A and B, can both be supported. But it is not obvious how to imagine the physical support for A as opposing the physical support for B.<sup>10</sup>

I ignore the metaphors of strength and power, because they have an additional connotation that makes them better for some applications than others. There is no question as to whether the object's weight exerts force on the pan. Yet, if I have the strength/power to push a pan down, there is still a question as to whether I exert that force. If we are trying to model claims about reason-possession, the metaphors of

---

<sup>10</sup> Similar points explain why it *isn't* natural to use the metaphor of reason's mass, even though the physical weight of an object is tightly tied to its mass (and the mass of surrounding objects). Mass does not provide a natural illustration of how reasons oppose or compete. The technical description of how masses behave goes something like this: things with mass attract each other, becoming friends, until so many friends congregate together that they collapse in on each other, igniting into a humongous, flaming ball of fury that incinerates anything within a million miles. There is a lesson to be learned here—don't have too many friends lest you become a flaming ball of fury—but it doesn't have anything to do with how reasons determine deontic status. Yes, there are contexts in which masses oppose, e.g., when placed on opposite pans of the balance scale. Yet they are also contexts in which it is more salient and accurate to say that *weights* or *forces* oppose. For example, if you take a balance scale into outer space, the items on the two pans will both be weightless even if they have different masses. The pans of the balance scale are at equal height (thus measuring the equal weight of the two items) rather than one pan being lower than the other (as it would if the scale were measuring unequal masses).




strength and power are handy. Perhaps *possessed* reasons for  $\phi$  are reasons that have the strength or power to push  $\phi$  toward permissibility *and* actually exert that power/strength, they actually push  $\phi$  toward permissibility. In contrast, perhaps *unpossessed* reasons for  $\phi$  have the strength/power to push  $\phi$  toward permissibility but that strength/power is *not* exerted on  $\phi$ . I ignore unpossessed reasons, because they make no difference to an act's deontic status. A reason has to exert force or pressure to (partly) explain why an act has the deontic status it does.

#### 1.4.2. From Single Vector Sum to Dual Vector Sum

The next section introduces the model of weighing reasons that structures the rest of the book. In the rest of this section, I explain the book's model of weighing reasons in terms of the (master) metaphor of force. I include this discussion for two reasons. First, some philosophers have found the ideas of the book more intuitive when I explain them in terms of force and vector sums. Second, contra Hawthorne & Magidor (2018: 134) and Fogal & Risberg (2023: 70-1), I show that the metaphors of force and weight take us to the same basic view about how reasons (considerations, pros and cons) interact to determine deontic status. You may prefer to skip the rest of this section if you are already happy with metaphors of weight and weighing and, more specifically, the justifying/requiring weight distinction.

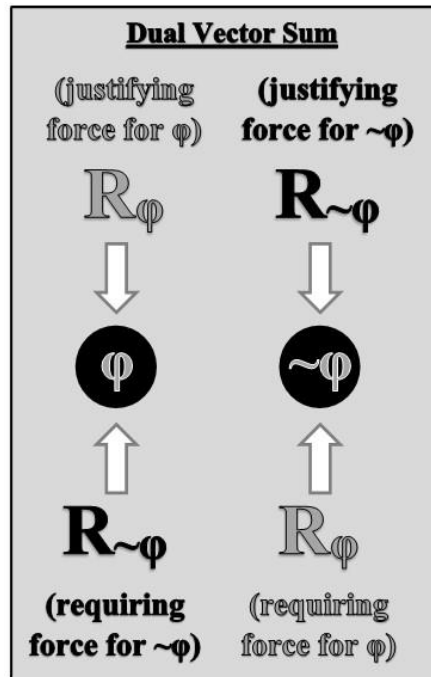
Forces push particles around, and forces oppose when they push the same particle in opposite directions. Suppose that a particle is stationary and then one force pushes it down and the other force pushes it up. After taking into account these forces, which direction is the particle moving? Up, down, or still stationary? Well, that depends on the relative forcefulness of the upward and downward forces. In other words, it depends on the value of the *net* downward force, or the *downward vector sum*: the downward force – the upward force. If the downward force is more forceful than the upward one, the particle goes down; if the two forces are equally forceful, then the particle is stationary; and if the downward force is less forceful, then the particle goes up. In short, the interaction of opposing forces determines the overall direction of a particle's movement.

The metaphor of a reason's force holds that, likewise, the interaction of opposing reasons determines an action's overall deontic status. Let a single particle represent an option  $\phi$ . It begins in a stationary state. In the physical world, forces can be exerted on an object in any direction: up, down, left, right, backwards, forwards, and everything in between. When we are concerned with how a reason determines whether an act is permissible, only two directions matter. The reasons for  $\phi$  ( $R_\phi$ ) push  $\phi$  down toward permissibility and the reasons for  $\sim\phi$  ( $R_{\sim\phi}$ ) push  $\phi$  up toward impermissibility. The vector sum of the downward force determines the deontic status of an act. We'll assume that  $\phi$  is permissible just when it is stationary or going down, and that it is impermissible just when it is going up. This gives us the Single Vector Sum model (cf. Ross 2002: 29 and Berker 2007: 122-3). If you rely on the metaphor of force, this is probably the model that you work with in your head.



Some permissible acts are also required, where  $\phi$  is **required** exactly when both  $\phi$  is permissible and the alternatives are impermissible. If morality requires you to keep your promise, it is permissible (ok) to keep the promise and impermissible (not ok) to break it. How can we model a requirement to  $\phi$  using forces and vector sums? One way is to introduce a second particle,  $\sim\phi$ , that gets pushed around by the same reasons, the reasons for  $\phi$  and  $\sim\phi$ . The reasons for  $\sim\phi$  push  $\sim\phi$  down toward permissibility. The reasons for  $\phi$  push  $\sim\phi$  up toward impermissibility.  $\phi$  is required just when  $\phi$  is stationary or going down *and*  $\sim\phi$  is going up. Dual Vector Sum is the idea that *two* vector sums, one for  $\phi$  and one for  $\sim\phi$ , represent the full deontic status of  $\phi$  (e.g., whether  $\phi$  is not just permissible but also required). If you rely on the metaphor of force, I recommend that you work with this model.

Notice that, on Dual Vector Sum, the upward and downward forces play two distinct functional roles in fixing  $\phi$ 's deontic status. A reason's *justifying force* for  $\phi$  is its downward force on  $\phi$ , making  $\phi$  permissible (as long as the opposing reasons don't push  $\phi$  up more forcefully). A reason's *requiring force* for  $\phi$  is its *upward* force on  $\sim\phi$ , making  $\sim\phi$  impermissible (as long as the opposing reasons don't push  $\sim\phi$  down at least as forcefully). The distinction between justifying and requiring force is not a distinction between fundamentally different kinds of forces. It is just a difference of direction. Justifying force always pushes something down toward permissibility and requiring force always pushes something up toward impermissibility. (In §5.1, we'll see that the reason for/against distinction specifies *which option* is pushed in the relevant direction,  $\phi$  or  $\sim\phi$ .<sup>11</sup>)



You might wonder whether the move from Single Vector Sum to Dual Vector Sum was really necessary. Single Vector Sum can hold that  $\phi$  is permissible but not required just when it remains stationary, and it is required just when it is moving down. So Single Vector Sum can distinguish between acts that are merely permissible and those that are also required. So why switch to Dual Vector Sum?

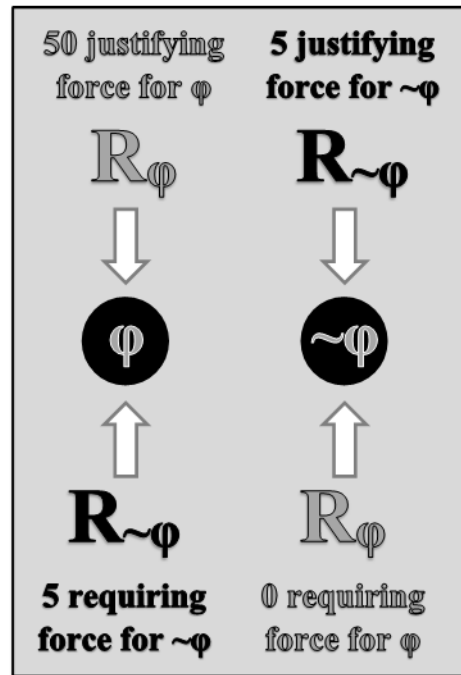
In the abstract, forces on one particle tell us nothing about forces on another particle. Yet when we are using force and vector sums to model deontic status, there must be some connection between the forces on  $\phi$  and  $\sim\phi$ . If  $\phi$  is required, then  $\sim\phi$  is impermissible and so  $\sim\phi$  must be going up. Hence, even if Single Vector Sum is, by itself, adequate to represent requirement to  $\phi$ , it still entails a second vector sum concerning  $\sim\phi$ . Dual Vector Sum just makes this second vector sum explicit. It is just Single Vector Sum applied to both  $\phi$  and  $\sim\phi$  at the same time. Single Vector Sum

<sup>11</sup> Here's the summary in terms of weight (rather than force). Justifying weight for  $\phi$  pushes  $\phi$  toward permissibility, and justifying weight *against*  $\phi$  pushes  $\sim\phi$  toward permissibility. Requiring weight for  $\phi$  pushes  $\sim\phi$  toward impermissibility and requiring weight *against*  $\phi$  pushes  $\phi$  toward impermissibility.

entails Dual Vector Sum, and so *Dual Vector Sum can't be less plausible than Single Vector Sum*.

In contrast, any model that represents requirement on a single vector sum is *less plausible* than Dual Vector Sum. I push up on Baby Benji to carry him and down on his balloon that I don't want to fly away; however, I push Baby Benji up a *lot* harder than I push the balloon down (what has he been eating!?). If I'm exerting force on two objects, there's nothing about force *per se* that requires that I push down on the first only as hard as I push up on the second. The forces on two objects are, in general, independent variables: you can't calculate the value of one from the value of the other.

If we allow justifying and requiring weight to be independent variables, then both  $\phi$  and  $\sim\phi$  can go down. Suppose that the reasons for  $\phi$  have 50 units of justifying force for  $\phi$  but 0 requiring force for  $\phi$  whereas the reasons for  $\sim\phi$  have 5 units of both justifying and requiring force for  $\sim\phi$ . This example is illustrated by the diagram to the right.  $\phi$ 's direction is determined by the competition between the justifying force for  $\phi$  and the requiring force for  $\sim\phi$ . 50 units of justifying force for  $\phi$  is greater than 5 units of requiring force for  $\sim\phi$ , so  $\phi$  will be permissible (go down).  $\sim\phi$ 's direction is determined by the competition between the justifying force for  $\sim\phi$  and the requiring force for  $\phi$ . 5 units of justifying force for  $\sim\phi$  is greater than 0 requiring force for  $\phi$ , so  $\sim\phi$  will also be permissible (it also will go down).



If you model requirement on a single vector sum, you must deny that both  $\phi$  and  $\sim\phi$  can go down at the same time. On such a model, the directions of  $\phi$  and  $\sim\phi$  are inversely related. If  $\phi$  is going down, then it is required. If  $\phi$  is required, then  $\sim\phi$  is impermissible and, therefore, must be going up. The underlying problem is that, if you model requirement on a single vector sum, you are committed to holding that justifying and requiring weight are *dependent variables*: you can calculate the value of one from the other. This is a controversial normative judgment about the forces of reasons.

The simplest way for justifying and requiring weight to be dependent variables is for them to always be equal (a reason always has the same amount of justifying and requiring weight); however, they will be dependent variables as long as justifying and requiring weight always come in the same proportion (e.g., reasons always have twice as much justifying as requiring weight). In other words, if you try to model requirement on a single vector sum, you are committed to *Single Proportion (first pass, force version)*: a reason's downward justifying force on  $\phi$  and its upward requiring force on  $\sim\phi$  always come in the same proportion.

If you use a single vector sum to model the full deontic status of  $\phi$  (e.g., whether it is required), you have some explaining to do. You are committed to Single

Proportion, which cannot be derived from the metaphors of force or vector sums. If anything, these metaphors suggest that Single Proportion is false. You need an argument for Single Proportion. In contrast, Dual Vector Sum does not depend on a controversial normative judgment like Single Proportion. It can represent requirement whether or not Single Proportion is true. If you want to use vector sums without potentially distorting the normative phenomena that you use them to model, then you must use Dual Vector Sum.

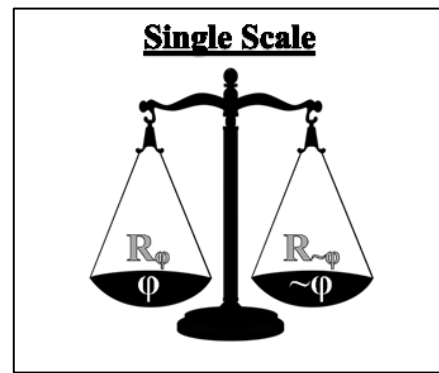
## 1.5. The Dual Scale Model of Weighing Reasons

### 1.5.1. From Single Scale to Dual Scale

The weight and weighing metaphors *just are* the force and vector sum metaphors applied to the balance scale context. I tend to use the weight/weighing metaphors, in part, because this specific context makes it easier to visualize two features that emerged in our discussion of force and vector sums. First, balance scales provide a context in which it is natural to focus on forces in only two opposing directions, up and down. Second, balance scales illustrate how forces on  $\phi$  can thereby be forces on  $\sim\phi$ : a reason for  $\phi$ 's weight on the  $\phi$  pan is just the downward force on the  $\phi$  pan and the upward force on the  $\sim\phi$  pan.

If truth be told, however, the main reason I focus on the weight/weighing metaphors, rather than the force/vector sum metaphors, is stylistic and rhetorical. The phrase 'weighing reasons' is familiar and easily brings to mind the book's subject matter. This is especially useful when my family and friends work up the courage to ask me what my book is about. I tend to say that the book develops and defends the idea that what we should do is determined by weighing the reasons for and against an action. Then an amazing thing happens: my family and friends actually understand what I'm talking about. They can even see why such a book might matter. Were I to talk about vector summing reasons or some such, they would become even more convinced that philosophers are totally useless people.

I argued in the previous section that, if we use vector sums to model how reasons determine deontic status, then we should use *two* vector sums. In this section, I establish the parallel claim which will structure the rest of the book: if we use balance scales to model how reasons determine deontic status, we should use *two* balance scales. On the standard balance scale model, the reasons for  $\phi$  ( $R_\phi$ ) go on one pan and the reasons for  $\sim\phi$  ( $R_{\sim\phi}$ ) 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. 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. For the moment, assume that  $\phi$  is permissible exactly when the reasons for  $\phi$  are at least as weighty as the reasons for  $\sim\phi$  and that  $\phi$  is required exactly when the reasons for  $\phi$  are weightier than the reasons for  $\sim\phi$ .



Recall that, if you use Single Vector Sum to model both whether  $\phi$  is permissible and whether it is required, then you assume Single Proportion, roughly, that all reasons have the same proportion of justifying and requiring weight/force. Single Scale proudly illustrates this assumption, for the balance scale creates a context in which the reasons for  $\phi$  push the  $\phi$  pan down toward permissibility (have justifying weight) only in proportion to how they push  $\sim\phi$  up toward impermissibility (have requiring weight).

To reject Single Proportion is to deny that reasons always come in the same proportion of justifying and requiring weight. For example, the standard account of supererogation (the standard account of going beyond the call of moral duty) holds that altruistic reasons have both justifying and requiring weight and self-interested reasons have only justifying weight. When values don't always come in the same proportion, they are *independent variables*. If justifying and requiring weight are independent variables, then Single Proportion is false and Single Scale can't model the full range of normative phenomena.

The underlying problem isn't the image of the scale, but the image of a *single* scale. To represent justifying and requiring weight as independent variables, we need (at least) *two* scales. But before I present this model, I need to explain one more term, 'commitment'. If your only goal in life is to eat every rock that you find, I might hold up a rock and remark that your aim *commits* you to eating this rock. I am not necessarily implying that eating the rock is the *right* thing to do, but only that your stupid aim makes not eating the rock a *wrong* thing to do. In this sense of commitment,  $\phi$  is a **commitment** just when every alternative is impermissible. *If* you can be committed to eating this rock without it being permissible to do so, then you are in what is often called a prohibition dilemma (both  $\phi$  and  $\sim\phi$  are *impermissible*).


For now ignore prohibition dilemmas and note that commitment to  $\phi$  is what you add to a permission to  $\phi$  to make  $\phi$  required. It is permissible to use some money to go out to eat tonight, but you aren't required to do so. For it is also permissible to donate that money to an effective charity. Permissible acts are required exactly when all other alternatives are *impermissible*, exactly when  $\phi$  is a commitment. A permission to  $\phi$  + a commitment to  $\phi$  (i.e.,  $\sim\phi$  is impermissible) = a requirement to  $\phi$ .

Dual Scale uses one scale to determine whether  $\phi$  is permissible and another to determine whether  $\phi$  is a commitment (whether  $\sim\phi$  is *impermissible*). The two scales work together to tell us whether  $\phi$  is required. Permission Scale holds, roughly, that

$\phi$  is permissible just when the justifying weight for  $\phi$  ( $JW_\phi$ )  $\geq$  requiring weight for  $\sim\phi$  ( $RW_{\sim\phi}$ ). Commitment Scale holds, roughly, that  $\phi$  is a commitment ( $\sim\phi$  is impermissible) just when the requiring weight for  $\phi$  ( $RW_\phi$ )  $>$  the justifying weight for  $\sim\phi$  ( $JW_{\sim\phi}$ ).

Dual Vector Sum modeled requirement to  $\phi$

by applying a single vector sum to both  $\phi$  and  $\sim\phi$  at the same time. The same principle



applies to Dual Scale at least in two option cases. Commitment Scale is just Permission Scale as it applies to  $\sim\varphi$ . You can see this if you look again at the illustration of Dual Scale. Commitment Scale is the mirror image of Permission Scale after swapping  $\varphi$  and  $\sim\varphi$ : after the swap, the *left* side of Permission Scale is the *right* side of Commitment Scale. In two option cases, then, Dual Scale models requirement by applying Permission Scale to both  $\varphi$  and  $\sim\varphi$  at the same time. (For reasons that will become clear in chapter 9, this does not hold in cases with more than two options. In such cases, the Permission and Commitment Scales will be distinct.)

Most of this book focuses on how to model permissibility and/or requirement in pairwise choices, choices between exactly two options. Almost any real-life choice will have more than two options. But no matter how many options there are, we'll see that permissibility and requirement still boil down to what happens in pairwise competitions (§4.1.3, §7.6.1, chapter 9). Once we understand all the complexities of weighing reasons for a single pairwise competition, it will be much easier to understand weighing reasons for any number of options.


### 1.5.2. Dual Scale Cannot be Less Plausible than Single Scale

While Dual Scale can represent justifying and requiring weight as independent variables, it can also represent them as dependent variables. If you think Single Proportion is true, then you just assign the same number to a reason's justifying and requiring weight. This will make Dual Scale's deontic verdicts extensionally equivalent to those of Single Scale (§4.2.1). The only difference between Single and Dual Scale is that Single Scale makes a controversial, substantive normative judgment about the relation between justifying and requiring weight (Single Proportion) and Dual Scale doesn't.

You might think that Single Scale avoids the distinction between justifying and requiring weight, and so still comes out more plausible. This is incorrect. The weight metaphor entails the conceptual distinction between justifying and requiring weight.

Justifying and requiring weight are not different *kinds* of weight. Think about the single kind of physical weight. You put my laptop on the left pan and its weight pushes the left pan down and the right pan up. One kind of physical weight can push different pans in different directions. Likewise, a single kind of normative weight can push different *options* in different *deontic* directions. Justifying weight for  $\varphi$  pushes  $\varphi$  toward permissibility. Requiring weight for  $\varphi$  pushes  $\sim\varphi$  toward impermissibility. Your promise to read a paper—or my book!—pushes reading it toward permissibility (has justifying weight for reading) and not reading it toward impermissibility (has requiring weight for reading). In the absence of sufficiently weighty countervailing reasons, you are required to keep your promise.

We can easily map justifying and requiring weight onto the diagram of Single Scale. The left side of Single Scale represents the permissibility of  $\phi$  as a competition between justifying weight for  $\phi$  and requiring weight for  $\sim\phi$ . The right side represents the permissibility of  $\sim\phi$  as a competition between the justifying weight for  $\sim\phi$  and the requiring weight for  $\phi$ . The bar between the left and right sides is what forces the left pan to go down only insofar as the right goes up, and so what forces justifying and requiring weight to come in the same proportion.



Now compare the map of justifying and requiring weight onto Single Scale with the diagram of Dual Scale. Permission Scale represents the same competition as the left side of Single Scale ( $JW_\phi$  vs  $RW_{\sim\phi}$ ). Commitment Scale represents the same competition as the right side of Single Scale ( $JW_{\sim\phi}$  vs  $RW_\phi$ ). The only difference is that there is no bar between the Permission and Commitment Scales that forces justifying and requiring weight to be dependent variables. The only difference between Single and Dual Scale is that only Single Scale takes a substantive, controversial stand on the relation between justifying and requiring weight. Therefore, *Single Scale entails Dual Scale and Dual Scale cannot be less plausible than Single Scale*.

Minding the metaphor leads us to Dual Scale in two option cases, because Dual Scale uses one permission scale for each alternative,  $\phi$  and  $\sim\phi$ . To apply this same sort of reasoning in cases with more than two options, we'll need one permission scale per alternative, giving us a "Dynamic Scale" model. If there are three options, we'll need three permission scales. If there are 100 options, we'll need 100 permission scales, one for each option. Each permission scale tells us whether its option is permissible, and an option is required just when it is the only permissible option. When there are more than two options, Dynamic Scale is what follows from minding the metaphors. Dual Scale is a simplification that can reduce computational complexity: it allows us to get by with only two scales no matter how many alternatives there are (ch. 9). For now, just ignore the difference between Dual and Dynamic Scale.<sup>12</sup>

Whether you want to use the force metaphor or the weight metaphor, we get the same moral for how reasons determine deontic status, at least for two option cases. If Single Vector Sum is true, then so is Dual Vector Sum. If *Single Scale* is true, then *Dual Scale* is true. And if you try to model both permissibility and requirement on a single vector sum or a single scale, you have some explaining to do. You are assuming that Single Proportion is true, even though there is nothing about force, and so nothing about weight, which would lead us to expect it to be true. It has to be defended independently of these metaphors. (Chapter 3 confirms the point in this section without assuming the aptness of the metaphors.)

<sup>12</sup> When we get to cases that involve more than two options, the individuation of options will be salient. See §2.1 and especially §5.4 for my account of option individuation and how it relates to weighing reason.

Why assume that justifying and requiring weight always come in the same proportion *come what may*? Why not start off with Dual Scale (Dual Vector Sum), go look at the full range of normative phenomena, and see whether there are any normative phenomena that are incompatible with Single Proportion. If there are, stick with Dual Scale. If there aren't, then you have a good reason to endorse Single Proportion. A good reason for Single Proportion won't disprove Dual Scale (or Dynamic Scale), but it will justify your reliance on Single Scale. For the conjunction of Dual Scale and Single Proportion gives us Single Scale. But here is the thing: I bet that once you start thinking with Dual Scale, you will find Single Proportion (and so Single Scale) increasingly dubious.

### 1.6. Metaphor Mongering or Minding the Metaphor?

The general issue of weighing reasons is just the question of how normatively relevant considerations systematically interact to determine an all-in normative verdict (§0.2). We can approach this issue from a variety of methodological perspectives, such as measurement theory, decision theory (e.g., Sher 2019), or nonmonotonic logic (Horty 2012: ix). Every methodology has its limitations, and it may take multiple methodologies working together to fully unpack what it is to weigh and weight reasons.

A key feature of the book's methodology amounts to *minding the metaphors* of weight and weighing, i.e., unpacking these metaphors and then explaining what account emerges of the relation between reasons and deontic status. There's no guarantee that what emerges from this methodology will be worth the paper it's printed on. Nonetheless, I argue that the account that emerges, Dual Scale, is plausible and worth serious consideration. It is admittedly unusual for a philosophy book to be so devoted to minding a metaphor. Some of you are bound to worry that this methodology is some sort of problematic metaphor mongering in disguise. I have five replies to this understandable concern.

*First*, the most common ways of understanding how reasons interact to determine deontic status are by appealing to the weight, force, pressure, support, strength, or power of reasons. These metaphors are ubiquitous, and they converge on essentially the same account of how reasons determine deontic status (§1.4.1 above). When our common ways of talking converge on the same or similar idea, it is typically worthwhile to flesh the idea out and see whether it is worth exploring further.

*Second*, the metaphors of weight and weighing are not, in general, well understood. Criticisms of these metaphors (almost?) always result from not sufficiently understanding them.<sup>13</sup> For example, Hawthorne & Magidor give the following counterexample to the additive aggregation principle they take to be built into the weighing metaphor:

if you like red items, then the fact that an item is bright red might be a reason (of a certain strength) for you to buy it, and the fact that it's red might also be a reason (of a certain strength) for you to buy it. But the strengths of the two

---

<sup>13</sup> This includes the criticisms from Dancy (2004), Draï (2018), Snedegar (2018), Hawthorne & Magidor (2018: 133-5), and Titelbaum (2019), all of which I discuss at various places in the book.



reasons do not add up (you don't have twice as much reason to buy it because it's both red and bright red). (133-4)

Notice that this counterexample attacks a type of additive aggregation that is *disanalogous* with physical weight. Physical weights are not additive *no matter how they are individuated*. The total weight of my phone and my phone's processor is not the sum of their individual weights. For you would double count the weight of my phone's processor. If we take the weight metaphor seriously, we should hardly be surprised that the total weight of *it's red* and *it's bright red* is not the sum of their individual weights. I have already addressed some other misunderstandings from critics of the metaphors, e.g., that decision theory is incompatible with the metaphors (§1.1.2) and that the force metaphor has different implications than the weight metaphor (§1.4). I address a range of other such misunderstandings in chapters 2 and 5.

Those who endorse the metaphors aren't much better. Reliance on the metaphors of weight and weighing is far more ubiquitous than reliance on the distinction between justifying and requiring weight. Yet we've already seen that the metaphor of weight entails the conceptual distinction between justifying and requiring weight (§§1.4-5). My impression from talking to ethicists is that they do not, in general, recognize this entailment.

Other ethicists may endorse the conceptual distinction but assume that Single Proportion is too plausible to deny. This would explain why so little of the ethics literature relies on the distinction between justifying and requiring weight. If they aren't independent variables, then the distinction won't be particularly useful for explaining normative phenomena. But why does Single Proportion seem too plausible to deny? There are certainly reasons to reject it: this book will show that we can resolve a range of otherwise puzzling phenomena by denying it. And arguments for it are *very* hard to come by.<sup>14</sup> In §4.6, I submit that its widespread acceptance is due to an implicit—but mistaken (§1.5)—assumption that the metaphor of weighing reasons entails Single Scale.

Arguably, then, ethicists often don't have a deep understanding of these metaphors. Whether we endorse or reject them, we should first understand the best way of cashing them out. The only way to get such understanding is to mind the metaphors.

*Third*, minding the metaphors allows us to avoid obscurantism. If you mind them, you can convert them into a testable philosophical hypothesis about how reasons interact to determine deontic status. You identify structural features of physical weight and then reformulate them for the weight of reasons. By explicitly formulating these features, you bring them into the open to be tested using the normal methods of philosophy. Admittedly, you can give an account of weighing reasons without using

---

<sup>14</sup> You can argue for Single Proportion by arguing first for, say, maximizing act utilitarianism. Then you can point out that such a view is true only if Single Proportion is true. A limitation of such arguments is that they don't explain why philosophers continue to find Single Proportion plausible even after abandoning maximizing act utilitarianism. Tenenbaum (2007: 162–174) and Cullity (2018: 431, nt 9) give the best direct arguments for something like Single Proportion that I'm aware of. I address Tenenbaum's argument in my 2017: 1375, nt 16. I address Cullity's objection in chapter 7.

the metaphors at all (§0.2). But how can we tell that ignoring the metaphors is better than minding them, if we've never consider what theory results from minding them?

Perhaps the clearest case of metaphor mongering would be using the metaphors without ever minding them. If the metaphors are inapt, then their continued usage may obscure the normative phenomena. Indeed, I argue that an implicit commitment to Single Scale—to a certain mistaken way of unpacking the metaphors—has made it hard for philosophers to see that a range of normative views even make sense, including supererogation, the existence of merely justifying reasons, and dilemmas (§§3.3, 4.3). By the end of this book, it will be so obvious to you that these ideas *make sense* that it would be hard to understand why folks found them so mysterious in the first place, were it not for the mystery being caused by an implicit commitment to Single Scale.

*Fourth*, by minding the metaphors, the resulting account of weighing reasons is motivated by analogical reasoning. The weight of physical objects and reasons are both gradable and explanatory. The weight of two objects can combine to have an overall weight, and likewise the weight of two reasons can combine to have an overall weight. The weight of my laptop and the weight of my pen oppose when you put them on different pans of a balance scale. Likewise, reasons for  $\phi$  oppose reasons for  $\sim\phi$  (1.1.1). These structural similarities make it plausible that the behavior of physical and normative weight will share other similarities too. Furthermore, we know that objects have weight and can be weighed. To the extent that weighing reasons shares significant structural features with weighing objects, we know that those structural features make sense or are possible.

Of course, analogical reasoning can take us only so far. This analogical reasoning suggests the model, Dual Scale. Part II arrives at the same model by appealing to a formal characterization of justifying and requiring weight and standard philosophical methodology used to explore conceptual connections. Part III provides further confirmation by considering how well it handles a range of otherwise puzzling phenomena, such as supererogation (§4.2, chapter 8), the Justification Intransitivity Paradox (ch 7, §9.1.3), the normative significance of small improvements (ch 7), and the All or Nothing Problem (ch 8).

In principle, we could disconfirm the model by showing that there is some behavior of reasons (alternatively: considerations, pros and cons) or their interaction with deontic status that fits poorly with the metaphors. Many argue that there is some such disconfirming behavior. For example, the balance scale metaphor is thought to be incompatible with holism (the context sensitivity of weight) and certain kinds of reasons against. I show that these thoughts are mistaken, and arguably arise out of *misunderstanding* the metaphors of weight and weighing.

Analogical reasoning is always slippery. Analogies always have disanalogies. If there weren't disanalogies, it wouldn't be an analogy but an identity. The hard part is figuring out which disanalogies are salient or relevant or would count against the aptness of the weight and weighing metaphors. Proponents and critics of the metaphors are likely to disagree about whether and how much a disanalogy counts against the aptness of the metaphors. But we should respect one general rule: a

disanalogy is salient to the extent that it is connected to whatever is driving the alleged analogy.

What's driving the analogies with weight and weighing, *as I develop them*, is just the way that weights function on balance scales in everyday contexts. For the purposes of this book, then, the salience of a disanalogy is determined, not by what physicists have to say about the nature of weight, but the way that weights function in these balance scale contexts. A *disanalogy* between physical and normative weight provides less disconfirmation when it concerns abstract physics than when it concerns everyday reflection on balance scales. This cuts both ways. A *similarity* also provides less analogical support when it emerges out of abstract physics than balance scales.

You don't need a PhD in physics to understand the idea that reasons are weighed to determine what to do. That's a good thing. It allows people with humble backgrounds in physics—people like me, your friends and family, many world leaders—to gain a basic understanding of how reasons interact to determine what to do.

Consider two potential disanalogies and how they might bear on the aptness of the metaphors. The first is that an object's physical weight is a function of its mass, but a reason's weight is not a function of anything analogous to a reason's mass. If this potential disanalogy is genuine,<sup>15</sup> I don't think it counts much, if at all, against the aptness of the metaphor. The relation between mass and weight is more a product of physics than everyday reflection on balance scales.

The second potential disanalogy is that the weight of reasons is not additive in the way that physical weight is additive. This potential disanalogy would be a much bigger problem for the aptness of the metaphors than the first. For the second disanalogy shows that a salient feature of our everyday understanding of weight is disanalogous to the weight of reasons.

This second potential disanalogy is, indeed, the most challenging alleged disanalogy between physical and normative weight. Physical weights aggregate by addition and there are a number of alleged counterexamples to the claim that the weights of reasons aggregate in a parallel way. Many of these counterexamples, like the red/bright red case above, fail because they attack a kind of additive aggregation that doesn't even apply to physical weight. When we pay attention to the way in which physical weight is additive, we'll discover some general strategies for addressing alleged counterexamples (§§2.4-5).

These strategies address every alleged counterexample that I'm aware of, but we all know the ingenuity of philosophers with an axe to grind. New cases will be devised, and I wouldn't be surprised if a new case proves fatal. So I need a Plan B that strips Dual Scale down to what's essential, and additive aggregation is not essential.<sup>16</sup> But Plan B is plan b for a reason: the more significant disanalogies there

---

<sup>15</sup> I doubt that the disanalogy is genuine, as the notion of 'presumptive weight' from §2.2 seems to be an analogue of mass. Just as the physical weight of an object is determined by its mass and context (e.g., whether it is in your living room or in outer space), the weight of a reason is determined by its presumptive weight and context (conditions and modifiers).

<sup>16</sup> Plan B also drops talk of "reasons for  $\varphi$ " (count noun) and sticks with talk of "reason for  $\varphi$ " (mass noun). If nothing else, this will discourage Titelbaum (2019), Fogal (2016), and Fogal & Risberg (forthcoming) from sending me hate mail.

are for physical and normative weight, the less my model is motivated by analogical reasoning.

Maybe it's still hard for you to see how any self-respecting analytic philosopher would rely on a metaphor to support their view; that's the sort of mischief Continental philosophers get themselves into. We needn't battle this out (or malign our Continental colleagues). If you take the metaphors as *mere illustration* of my views, then we lose some analogical support and thereby have less reason to believe that Dual Scale is true. No worries. Once you see everything that Dual Scale can do and understand the conceptually necessary ways in which justifying and requiring weight interact, you still will have plenty of reason to take it seriously.

*Finally*, recall that the metaphors are most useful as a springboard for systematic theorizing about the functional structure of normative perspectives (§1.2). Minding the metaphors is, in my hands anyway, largely an attempt to uncover the structural features that are shared by many normative theories. If these metaphors aren't apt, then, say, morality doesn't have the structure that many of us take it to have.

Perhaps maximizing act utilitarianism and standard decision theory are normatively inadequate for some reason or another. Perhaps they are too demanding or don't allow enough flexibility for reasonable differences in risk aversion or whatever. They nonetheless seem to be paradigms of a well-structured normative theory. They are structurally sound. In fact, they share exactly the same (functional) structure, the structure we have referred to metaphorically as weighing reasons or considerations on a balance scale (§1.1.2). I know of no better way to talk about this shared structure than one of our familiar metaphors (weight, force, pressure, support, strength, power).

The more important point, though, isn't about our talk. If the metaphors of weight and weighing aren't apt, then morality does *not* have the functional structure that maximizing act utilitarianism and expected utility theory assume that it has. It would turn out that these theories have structural problems after all. That would come as a shock to both proponents and critics of these theories.

Maximizing act utilitarianism and expected utility are perfectly intelligible normative theories even when they don't explicitly invoke the metaphors of weight and weighing (or force, support, etc.). We don't always need these metaphors to articulate a specific normative theory. But they are useful when we articulate the general structure we expect from normative theories of morality and practical rationality. If the metaphors *aren't* apt, our widely shared expectations of which normative theories are structurally sound need heavy revision. Rule utilitarians might survive unscathed. Most of the rest of us would need to rethink how we structure our normative theories.

The next chapter continues a theme that emerged in this one: by carefully cashing out the metaphors of weight and weighing, we'll see that critics tend to misunderstand them. I will address two general questions. How, if at all, do we make scale-based models compatible with Dancy-style holism, i.e., the idea that a reason's weight is context sensitive? Second, what sort of additive aggregation principle is built into the metaphors and is that principle as dumb as its critics take it to be?

## References

- Archer, Alfred. 2016. "Moral Obligation, Self-Interest, and the Transitivity Problem." *Utilitas* 28: 441-464. doi: [doi.org/10.1017/S0953820816000091](https://doi.org/10.1017/S0953820816000091)
- Bader, Ralf. 2016. "Conditions, Modifiers, and Holism." In Lord, Errol and Barry Maguire (eds.), 27-55.
- Benn, Claire. 2018. "Supererogation, Optionality, and Cost." *Philosophical Studies* 175: 2399-2417. doi: <https://doi.org/10.1007/s11098-017-0965-7>
- Berker, Selim. 2018. "A Combinatorial Argument Against Practical Reasons for Belief." *Analytic Philosophy* 59: 427-70.
- \_\_\_\_\_. 2007. "Particular Reasons." *Ethics* 118: 109-139.
- Braddon-Mitchell, David and Robert Nola. 2008. *Conceptual Analysis and Philosophical Naturalism*. MIT Press.
- Broome, John. 2013. *Rationality Through Reasoning*. Oxford University Press.
- \_\_\_\_\_. 2004. *Weighing Lives*. Oxford University Press.
- Buchak, Lara. 2017. *Risk and Rationality*. Oxford University Press.
- Chang, Ruth. "Hard Choices." *Journal of the American Philosophical Association* 3: 1-21. doi: <https://doi.org/10.1017/apa.2017.7>
- \_\_\_\_\_. 2016. "Comparativism: The Grounds of Rational Choice." In Lord, Errol and Barry Maguire (eds.), 213-240.
- \_\_\_\_\_. 2002. "The Possibility of Parity." *Ethics* 112: 659-688.
- \_\_\_\_\_. 2001. "Against Constitutive Incommensurability or Buying and Selling Friends." *Nous* 35: 33-60.
- 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.
- \_\_\_\_\_. 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.
- Curtis, Barry. 1981. "The Supererogatory, the Foolish, and the Morally Required." *Journal of Value Inquiry* 15: 311-318.
- Dancy, Jonathan. 2014. "Review of Thomas Hurka's *British Ethical Theorists from Sidgwick to Ewing*." *Notre Dame Philosophical Reviews*. Stable URL: <https://ndpr.nd.edu/news/british-ethical-theorists-from-sidgwick-to-ewing/>.
- \_\_\_\_\_. 2004. *Ethics Without Principles*. Oxford University Press.
- \_\_\_\_\_. 1993. *Moral Reasons*. Blackwell Publishers.
- Dorr, Cian, Jacob M. Nebel and Jake Zuehl. Forthcoming. "The Case for Comparability." *Nous*. doi: [doi.org/10.1111/nous.12407](https://doi.org/10.1111/nous.12407)
- Dorsey, Dale. 2016. *The Limits of Moral Authority*. Oxford University Press.

- Drai, Dalia. 2018. "Reasons Have no Weight." *Philosophical Quarterly* 68: 60-76.
- Ferguson, Benjamin. 2016. "The Paradox of Exploitation." *Erkenntnis* 81: 951-72.
- Ferguson, Benjamin and Sebastian Köhler. 2020. "Betterness of Permissibility." *Philosophical Studies* 177: 2451-2469.
- Fogal, Daniel. 2020. "Rational Requirements and the Primacy of Pressure." *Mind* 129, 1033-70. doi: [doi.org/10.1093/mind/fzz038](https://doi.org/10.1093/mind/fzz038)
- \_\_\_\_\_. "Reasons, Reason, and Context." In Lord, Errol and Barry Maguire (eds.) 2016: 74-103.
- Fogal, Daniel and Olle Risberg. 2023. "Explaining Normative Reasons". *Nous* 57: 51-80. doi: <https://doi.org/10.1111/nous.12393>
- Frowe, Helen. 2021. "The Limited Use View of the Duty to Save." *Oxford Studies in Political Philosophy* 7.
- Gert, Josh. 2016. "The Distinction between Justifying and Requiring: Nothing to Fear." In Lord, Errol and Barry Maguire (eds.), 27-55.
- \_\_\_\_\_. 2015. "Parity, Preference, and Puzzlement." *Theoria* 81: 249-71. doi: [doi.org/10.1111/theo.12069](https://doi.org/10.1111/theo.12069)
- \_\_\_\_\_. 2007. "Normative Strength and the Balance of Reasons." *Philosophical Review* 116: 533-562.
- \_\_\_\_\_. 2004a. *Brute Rationality: Normativity and Human Action*. Oxford University Press.
- \_\_\_\_\_. 2004b. "Value and Parity." *Ethics* 114: 492-510.
- \_\_\_\_\_. 2003. "Requiring and Justifying: Two Dimensions of Normative Strength." *Erkenntnis* 59: 5-36.
- Givewell. 2020. GiveWell cost-effectiveness analysis — version 1. [https://docs.google.com/spreadsheets/d/1BmFwVYeGMkpys6hG0tnfHyq\\_ZFzf-bmXYLSHODGpLY/edit#gid=1034883018](https://docs.google.com/spreadsheets/d/1BmFwVYeGMkpys6hG0tnfHyq_ZFzf-bmXYLSHODGpLY/edit#gid=1034883018)
- Greenspan, Patricia. 2005. "Asymmetrical Practical Reasons." *Experience and analysis: Proceedings of the 27th International Wittgenstein Symposium*. Maria Reicher, & Johann Marek (eds.). (pp. 387–394). Öbv&hpt.
- Hawthorne, John & Ofra Magidor. 2018. "Reflections on the Ideology of Reasons." *The Oxford Handbook of Reasons and Normativity*. Daniel Star (ed.). 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. doi: <https://doi.org/10.1017/S026505250999015X>
- Horton, Joe. 2017. "The All or Nothing Problem." *Journal of Philosophy* 114: 94-104.
- Horty, John. 2012. *Reasons as Defaults*. Oxford University Press.
- \_\_\_\_\_. "Reasons as Defaults." *Philosophers' Imprint* 7: 1-28.
- 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.
- Kahneman, Daniel. 2011. *Thinking Fast and Slow*. Farrar, Straus, and Giroux.
- Kamm, Frances. 1996. *Morality, Mortality, Volume II: Rights, Duties, and Status*. Oxford University Press.
- \_\_\_\_\_. 1985. "Supererogation and Obligation." *Journal of Philosophy* 82: 112-38.

- Kawall, Jason. 2003. "Self-Regarding Supererogatory Actions." *Journal of Social Philosophy* 34: 487-98.
- Kearns, Stephen. 2016. "Bearing the Weight of Reasons." In Lord, Errol and Barry Maguire (eds.), 173-90.
- Kearns, Stephen & Star, Daniel. 2009. "Reasons as Evidence." In Russ Shafer-Landau (Ed.), *Oxford Studies in Metaethics* (Vol. 4 pp. 215-242). Oxford: Oxford University Press.
- \_\_\_\_\_. 2008. "Reasons: explanations or evidence?" *Ethics*, 119, 31-56. <https://doi.org/10.1086/592587>
- Keller, Simon. 2013. *Partiality*. Princeton University Press.
- Kiesewetter, Benjamin. 2018. "Contrary-to-Duty Scenarios, Deontic Dilemmas, and Transmission Principles." *Ethics* 129: 98-115.
- \_\_\_\_\_. 2015. "Instrumental Normativity: In Defense of the Transmission Principle." *Ethics* 125: 921-946.
- Lance, Mark and Margaret Little. 2007. "Particularism and Antitheory." *The Oxford Handbook of Ethical Theory*. David Copp (ed.). Oxford University Press, 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(13): 427-446.
- Little, Margaret. 2000. "Moral Generalities Revisited." *Moral Particularism*. Brad Hooker and Margaret Little (eds.). Oxford University Press, 276-304.
- Little, Margaret Olivia and Coleen Macnamara. 2020. "Non-Requiring Reasons." *The Routledge Handbook of Practical Reason*. Ruth Chang and Kurt Sylvan (eds.). Routledge, 393-404.
- \_\_\_\_\_. 2017. "For Better or Worse: Commendatory Reasons and Latitude." *Oxford Studies in Normative Ethics* 7: 138-60.
- Lord, Errol. 2016. "Justifying Partiality." *Ethical Theory and Moral Practice* 19: 569-90.
- Lord, Errol and Barry Maguire (eds.). 2016. *Weighing Reasons*. Oxford University Press.
- Maguire, Barry. 2016. "The Value-Based Theory of Reasons." *Ergo* 3: 233-62. <https://doi.org/10.3998/ergo.12405314.0003.009>
- Maguire, Barry and Justin Snedegar. 2021. "Normative Metaphysics for Accountants." *Philosophical Studies* 178: 363-84.
- Massoud, Amy. 2016. "Moral Worth and Supererogation." *Ethics* 126: 690-710.
- McHugh, Conor & Jonathan Way. 2016. "Fittingness First." *Ethics* 126: 575-606. <https://doi.org/10.1086/684712>
- McMahan, Jeff. 2018. "Doing Good and Doing the Best." *The Ethics of Giving: Philosophers' Perspectives on Philanthropy*. Paul Woodruff (ed.). Oxford, 78-102.
- McNamara, Paul. 2011. "Supererogation, Inside and Out: Toward an Adequate Scheme for Common Sense Morality." *Oxford Studies in Normative Ethics*, Volume I: 202-35.

- Muñoz, Daniel. 2021. "Three Paradoxes of Supererogation." *Nous* 55: 699–716. <https://doi.org/10.1111/nous.12326>
- Muñoz, Daniel & Theron Pummer. 2022. "Supererogation and Conditional Obligation." *Philosophical Studies* 179: 1429-43. doi: [doi.org/10.1007/s11098-021-01724-y](https://doi.org/10.1007/s11098-021-01724-y)
- Murphy, Mark. 2017. *God's Own Ethics: Norms of Divine Agency and the Argument from Evil*. Oxford University Press.
- Nair, Shyam. 2021. "Adding Up" Reasons: Lessons for Reductive and Nonreductive Approaches'. *Ethics* 132: 38-88.
- \_\_\_\_\_. 2016. "How Reasons Accrue." In Lord, Errol and Barry Maguire (eds.) 2016: 56-73.
- Nebel, Jacob M. "The Good, the Bad, and the Transitivity of *Better Than*". *Nous* 52: 874-899. doi: [doi.org/10.1111/nous.12198](https://doi.org/10.1111/nous.12198)
- Parfit, Derek. 1982. "Future Generations: Further Problems." *Philosophy & Public Affairs* 11: 113-72.
- Portmore, Douglas. 2019. *Opting for the Best: Ought and Options*. Oxford University Press.
- \_\_\_\_\_. 2017. "Transitivity, Moral Latitude, and Supererogation." *Utilitas* 29: 286-298.
- \_\_\_\_\_. 2013. "Perform Your Best Option." *Journal of Philosophy* 110: 436-59.
- \_\_\_\_\_. 2011. *Commonsense Consequentialism: Wherein Morality Meets Rationality*. Oxford University Press.
- \_\_\_\_\_. 2008. "Are Moral Reasons Morally Overriding?" *Ethical Theory and Moral Practice* 11: 369-388.
- 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.
- Pummer, Theron. 2021. "Impermissible Yet Praiseworthy." *Ethics* 131: 697-726.
- \_\_\_\_\_. 2019. "All or Nothing, but if Not All, Next Best or Nothing". *The Journal of Philosophy* 116: 278-291.
- Rabinowicz, W. 2012. Value Relations Revisited. *Economics and Philosophy*, 28, 133-64. doi: [doi.org/10.1017/S0266267112000144](https://doi.org/10.1017/S0266267112000144)
- \_\_\_\_\_. Value Relations. *Theoria*, 74, 18-49. doi: [doi.org/10.1111/j.1755-2567.2008.00008.x](https://doi.org/10.1111/j.1755-2567.2008.00008.x)
- Raz, Joseph. 1990. *Practical Reason and Norms*. 2<sup>nd</sup> ed. Oxford University Press.
- Rinard, Susanna. 2018. "Reasoning One's Way Out of Skepticism." Brill, 240-64.
- Ross, W. D. 2002. *The Right and the Good*. Philip Stratton-Lake (ed.). *Oxford University Press*.
- Rubio, Daniel. 2018. "God meets Satan's Apple: the Paradox of Creation." *Philosophical Studies* 175: 2987-3004.
- Rulli, Tina. "Conditional Obligations." *Social Theory and Practice* 46: 365-90. DOI: <https://doi.org/10.5840/soctheorpract20204189>
- Rulli, Tina and Alex Worsnip. 2016. "IIA, Rationality, and the Individuation of Options." *Philosophical Studies* 173: 205-221.
- Scheffler, Samuel. 1982. *The rejection of consequentialism: A philosophical investigation of the considerations underlying rival moral conceptions*. Oxford: Oxford University Press.



- Schmidt, Thomas. 2023. "How Reasons Determine Moral Requirements." *Oxford Studies in Metaethics* 18: XX-XX.
- Schroeder, Mark. 2021. "The Fundamental Reason for Reasons Fundamentalism." *Philosophical Studies* 178: 3107-27. <https://doi.org/10.1007/s11098-020-01572-2>
- Sher, Itai. 2019. "Comparative Value and the Weight of Reasons." *Economics and Philosophy* 35: 103-58.
- Silver, Nate. 2015. *The Signal and the Noise: Why so Many Predictions Fail—but Some Don't*. Penguin Books.
- Slote, Michael. 1989. *Beyond Optimizing*. Harvard University Press.
- Snedegar, Justin. 2021. "Reasons, Competition, and Latitude." *Oxford Studies in Metaethics* 16: 134-56. <https://doi.org/10.1093/oso/9780192897466.003.0006>
- \_\_\_\_\_. 2018. "Reasons for and Reasons Against." *Philosophical Studies* 175: 725-743.
- \_\_\_\_\_. 2017. *Contrastive Reasons*. Oxford University Press.
- \_\_\_\_\_. 2016. "Reasons, Oughts, and Requirements." *Oxford Studies in Metaethics, Volume 11*. Oxford University Press, 155-181.
- Swanton, Christine. 2005. *Virtue Ethics: A Pluralist View*. Oxford University Press.
- Temkin, Larry. 2012. *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning*. Oxford University Press.
- Tenenbaum, S. 2007. "Brute requirements". *Canadian Journal of Philosophy* 37: 153–171.
- Tetlock, Philip E. 2006. *Expert Political Judgment: How Good Is It? How Can We Know?* Princeton University Press.
- Titelbaum, Michael. 2019. "Reason Without Reasons For." *Oxford Studies in Metaethics, Volume 14*.
- Tucker, Chris. 2024a. "Withhold By Default: A Difference Between Epistemic and Practical Rationality." *Philosophical Studies*.
- \_\_\_\_\_. 2024b. "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>.
- \_\_\_\_\_. 2023a. "Parity, Moral Options, and the Weights of Reasons". *Nous* 57: 454-480. doi: [doi.org/10.1111/nous.12410](https://doi.org/10.1111/nous.12410)
- \_\_\_\_\_. 2023b. "A Holist Balance Scale." *Journal of the American Philosophical Association* 9: 533-553. <https://doi.org/10.1017/apa.2022.13>
- \_\_\_\_\_. 2022a. "The Dual Scale Model of Weighing Reasons." *Nous* 56: 366-92. doi: <https://doi.org/10.1111/nous.12361>.
- \_\_\_\_\_. 2022b. "Weighing Reasons Against". *Oxford Studies in Metaethics* 17: 77-107.
- \_\_\_\_\_. 2022c. "Too Far Beyond the Call of Duty: Moral Rationalism and Weighing Reasons." *Philosophical Studies* 179: 2029-52. doi: <https://doi.org/10.1007/s11098-021-01747-5>
- \_\_\_\_\_. 2020. "Divine Satisficing and the Ethics of the Problem of Evil." *Faith and Philosophy* 37: 32-56.
- \_\_\_\_\_. 2017. "How to Think About Satisficing." *Philosophical Studies* 174: 1365-1384. doi: [doi.org/10.1007/s11098-016-0763-7](https://doi.org/10.1007/s11098-016-0763-7)

- Väyrynen, Pekka. 2006. "Moral Generalism: Enjoy in Moderation." *Ethics* 116: 707-741.
- Way, Jonathan. 2017. Reasons as premises of good reasoning. *Pacific Philosophical Quarterly*, 98, 251-270. <https://doi.org/10.1111/papq.12135>
- 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.
- White, Stephen. 2017. "Transmission Failures." *Ethics* 127: 719-732.
- Whiting, Daniel. 2022. *The Range of Reasons in Ethics and Epistemology*. Oxford University Press.