

1

An Opinionated Guide to the Weight of Reasons

Errol Lord and Barry Maguire

1 On the need for Weighted Notions in Normative Theory

It is more important that our theory fit the facts than that it be simple, and the account we have given . . . corresponds (it seems to me) better than either of the simpler theories [from Kant and Moore] with what we really think, viz. that normally promise-keeping, for example, should come before benevolence, but that when and only when the good to be produced by the benevolent act is very great and the promise comparatively trivial, the act of benevolence becomes our duty.

(W.D. Ross, 1930, p. 19).

Any normative theory comes with its own normative toolkit: some distinctive normative ideology such as value, rights, obligations, the all-things-considered ought. The power and flexibility of the theory and the structure of its explanations are determined by the nature of this ideology. For instance, the simplest version of act consequentialism restricts itself to neutral value—a monadic property of states of affairs—and the all-things-considered ought—a single directive based on a full evaluation of consequential states of affairs.¹ We will open this introduction and this book by arguing, with Ross, that any decent theory needs to have at least one *weighted* notion in its normative toolkit. But first, in order to locate our topic we'll start with some ideological taxonomy.

We find it most helpful to classify normative ideology by two central cross-cutting distinctions. First we have a distinction between strict and non-strict notions. Within the class of strict normative notions we have directives, such as

¹Such a theory needs further ideology, such as the notions of properties and of states of affairs, but we are focusing on its evaluative and normative ideology.

4 Introduction

requirements, oughts, and prohibitions; and non-directives, such as permissions and supererogation. The strict notions are generally marked by their relationship to the appropriateness or inappropriateness of criticism. When you do something that is permitted, you are immune from a certain type of criticism. When you fail to do something that is required or when you do something that is prohibited, you are open to criticism (or at least there is some reason to criticize you).² The non-strict notions are generally marked by their role in the explanation of strict facts.

Second, we have a distinction between weighted and non-weighted notions. Within the class of weighted notions we have normative reasons—considerations that count in favor of doing or feeling or thinking this or that. We have modifiers: considerations that increase or decrease the strength of your reasons to some degree. We also have thresholds of various kinds: perhaps it is permissible to violate some right only if some threshold of value is met or disvalue avoided; or perhaps you would be criticizable for performing some action only if your evidence about your act instantiating the wrong-making feature met some threshold.³

Typically, non-weighted notions are strict—e.g., the all-things-considered ought—and the weighted notions are non-strict—e.g., normative reasons—but there are exceptions to both generalizations. Obligations are strict but plausibly weighted. If you break your promise you are apt for criticism, even if in the circumstances you had much more reason to do so than not. Perhaps even rights are like this; some rights are more important than others. Enabling conditions and disabling conditions are plausibly non-strict and non-weighted. My promise to read your paper gives me a reason to do so only if you take me up on it.

It is common to think that the non-strict notions will explain the strict. It is also increasingly common to think that the weighted notions will explain the non-weighted, and that within the weighted notions, reasons will explain the rest. It is clear at least that reasons are the paradigmatic weighted notion.

In the quoted passage, Ross is pointing out a shortcoming with theories that restrict themselves to strict notions like requirements and permissions. He is pointing out that in addition we need some notion of a consideration that has some weight, which can in principle be outweighed. As he says, the fact that some action would fulfill a promise or bring about some pleasure speaks in favor of performing that action. However, such considerations by themselves do not decide the strict normative status of the action; for instance, whether the action is permitted or required. The promise or the consequent good might be comparatively trivial; then some other consideration might well decide the matter.

Ross assumes that the following is a commitment of both Kantianism and Act Consequentialism: If some act-token instantiates some property (such as being the breaking of a promise, or leading to great harm), and if in virtue of that act-token

²It is natural to think that any epistemic indexing applicable to oughts will be matched by some appropriately indexed fact about criticism-worthiness.

³On this topic, see the contribution by Michael Smith and Frank Jackson.

instantiating that property, that act-token has some strict normative property, then any other act-token of that act-type would instantiate the same strict normative property. He additionally assumes that these theories restrict normative relevance to a particular, situation-invariant, set of act-types. For Kantianism, promise-keeping and the like; for Utilitarianism, pleasure-maximizing. Ross persuasively argues that whether one ought to promise-keep or pleasure-maximize in some situation may depend on further facts. In particular, it might depend on the presence of other normatively relevant act-tokens, and on the *relative importance* of the promise-keeping or pleasure-maximizing in that situation.⁴ Talk about the ‘importance’ of considerations is usually talk about their weight.

Various kinds of ethical theories attempt to avoid weighted notions. Some do avoid weighted notions altogether; we’ll argue such theories are substantively implausible. Others purport to avoid weighted notions, but (perhaps unwittingly) sneak them into the theory in disguise. These theories might make the right predictions, but their explicit ideologies are misleading.

One strategy for avoiding weighted notions is to accept deep normative dilemmas. Perhaps in some situation one ought to keep one’s promise *and* one ought to maximize pleasure, even though one cannot do both. There are two problems with this position. One is that some such “deep normative dilemmas” do not look like normative dilemmas at all. This will be the case when, intuitively, one normatively relevant consideration is relatively insignificant and the other is extremely important. Since this is an instance of the simplest and most powerful argument for the conclusion that a certain normative notion is weighted, we’ll give it a name: the *Extreme Comparison Argument*. Suppose you arranged to meet me in the pub at lunchtime but instead you are called upon to perform some life-saving procedure. One can insist that you ought to meet me in the pub and you ought to perform the procedure, though you cannot do both, and that neither has any ethical priority. One would thereby avoid weighted notions, but only at the cost of implausibility.⁵

⁴Here’s a slightly different argument for weighted notions: we need weighted notions to account for the difference between two epistemically possible cases, in both of which we should not “level down” to improve the equality of the distribution of some good, but in one of which the availability of this equal distribution provides an outweighed reason to level down, and in the other of which it does not.

⁵What if one analyzed these strict notions in terms of the appropriateness of criticism, allowing this latter property to vary with the severity of the violation, so that in the example in the main text it would be appropriate to blame one less for breaking the pub promise than for failing to perform the procedure? This reply clearly involves a *gradable* property, viz. the degree of blame that would be appropriate. But it doesn’t obviously involve a gradable *normative* property, since we are assuming that the property of appropriateness is not gradable. (The most natural explanation of the amount of criticism that would be appropriate is the weight of the reason to do the thing you ought to have done but didn’t do. But some other non-normative explanans might be available.) And more importantly it doesn’t involve any *weighted* notions, since these facts about the different amounts of criticism that would be appropriate do not generate any further normative facts about what you ought all things considered to do. We think this is an odd view, but we concede that it is conceptually available, even if as yet theoretically undermotivated.

6 Introduction

The second problem concerns proliferating dilemmas. Few are the situations in which there is only one normatively relevant consideration. Far more common are situations in which there are not only different types of normatively relevant consideration, but also many different tokens of each of these types. Without weighted notions, we would face trenchant incommensurability in a great many situations. But this is just a problem with the theory. In the real world we usually have a good idea about what we ought to do, given the many different normatively relevant considerations in our situations. So again the strict-only theory is implausible.

One way for a strict-only theorist to attempt to avoid these objections is to embrace a more particularistic explanation of facts about requirements, permissions and the like. She can reject Ross's claim that if some act-token has some strict property, then any other act-token of the same type has the same strict property. She can insist that the strict notions are situation specific, fixing themselves to different act-tokens in different situations, or extremely specific. This would enable her to avoid these objections, but at the cost of relinquishing the explanatory power of her theory. She has no further resources with which to explain why these all-things-considered notions affix themselves to different act-types in different situations. This is one of the main explanatory tasks fulfilled by weighted notions. They explain the differential normative significance of different considerations in different situations. So this reply simply underlines the need for weighted notions.

Once we see the need for weighted notions, we need a better account of their nature and diversity. We can start with reasons, and with some clear examples. The fact that you like steamed goat burritos is a reason to go to the food cart on 2nd and 2nd for lunch. The fact that the Boston Red Sox won the World Series is a reason for the President of the United States to call their manager this evening. If there is a bomb near the food cart, or an international crisis, clearly these reasons would be outweighed.

However, examples such as these may not be sufficient to precisely pick out one single weighted property. There might be various essential differences between kinds of reasons. We will introduce one: concerning different kinds of relationships between reasons and strict notions. It is natural to think that it is in the *nature* of reasons to play some role in the explanation of strict facts. Certainly they often do this. In the case above, it is natural to think that the President *ought* to attend to the international crisis.

But it has been alleged that some reasons do not generate oughts, even if they are not outweighed. Suppose we're between floors in an elevator, and someone would be mildly amused by your pulling a funny face. That provides you with a reason to do so, but it is odd to think that you ought to do this, even if you have nothing better to do. The pulling of the funny face is *optional*.⁶ A related view maintains that there is a distinction not between kinds of reasons but between

⁶Jonathan Dancy (2004, p. 21) calls these kinds of reasons "enticing reasons."

kinds of weight that a given reason can have, as between that reason's *justifying* weight, which can make an action rational but not required; and that reason's *requiring* weight, which can make non-performance criticism-worthy. In his contribution, Joshua Gert showcases his arguments for this position.

In response to the former position, John Broome (2004, p. 39) has argued that the relevant ought is more like "should" than the moral "must." It isn't too implausible that if one has slightly more reason to A than B, and no reason to do anything else, then even though the reasons supporting A and B are really weak, there is *some* sense in which one ought to A, and some sense in which one would be mildly worth of criticism for B-ing. Plausibly here our intuitions are weak because not much is at stake. For those who insist that there is no ought here, another strategy is to deny that there is any reason either; the relevant intuition can be accommodated by appealing to the *value* of performing the act in question.

Another issue that arises in characterizing the relationship between reasons and strict notions is that some considerations are weighted *and* strict. One class of normative notions with this feature are obligations or *commitments*. For example, promissory commitments seem to have to be weighted and strict. They require but they can also be outweighed, as in the earlier example of the promise to meet in the pub. One project in the theory of weight is to give an account of these *prima facie* different weighted notions; perhaps to explain them all in terms of a single general weighted property. Alida Liberman and Mark Schroeder, in their contribution, spell out some features of a theory of weighing commitments and suggest that this has implications for a general theory of weight.

It will be helpful for what follows to highlight some further important distinctions. First and second, we have our two distinctions between strict and non-strict and weighted and non-weighted. Third, we have a distinction between *local* and *all-things-considered* normative notions. Some normative fact is *local* just in case it obtains in virtue of some particular fact or facts. Facts about promissory commitments, facts about reasons to realize particular episodes of pleasures, and so forth are local normative facts. A normative fact is *all-things-considered* just in case it obtains in virtue of *all* the normatively relevant facts in some situation, usually together with the fact—the *normative totality fact*—that these are all the normatively relevant facts. The consequentialist's ought is all-things-considered; so are facts about what you have most reason to do.⁷ Some local facts are strict, such as facts about commitments and obligations. It is natural to think that no all-things-considered facts are non-strict, though facts about the overall value of a world may be. Fourth, we distinguish *defeasibility* from *indefeasibility*. This is a property

⁷Given this definition, all-things-considered facts are also local. This is a consequence of the fact that it is difficult to provide some precise account of locality, but we have a sufficiently good imprecise grasp on the notion for it to do theoretical work for us. One might alternatively say that a fact is *global* just in case it obtains partly in virtue of some normative totality fact (on which see, for instance, Richard Holton [2002]).

8 Introduction

of attributions of normative notions to act-types. Lying is *defeasibly* wrong only if some possible token lies are wrong and other possible token lies are not wrong. Lying is *indefeasibly* wrong only if every possible token lie is wrong. The gravest sin in normative taxonomy is to confuse defeasibility with weightedness. Again, all these distinctions are clearly distinct. As we have seen, promissory commitments are plausibly weighted, strict, and local, and it is an open question whether or not they are defeasible. For utilitarians, certain facts about pleasure-maximization are strict, all-things-considered, indefeasible normative facts.

Once we isolate any class of weighted notions, all sorts of more particular questions arise. Are all weighted facts gradable (thresholds may not be)? Are they all comparable (perhaps right and wrong kinds of reasons for something are not)? How exactly do they interact? Do some sets systematically *override* others? For instance, do moral reasons systematically override non-moral reasons? The Extreme Comparison Argument will be helpful here again. Plausibly weak moral reasons are outweighed by weighty non-moral reasons in some situations. Are all weighted facts normative facts—e.g., facts about being good *as* something are perhaps weighted but non-normative facts, e.g. being good as a forest fire? Do all weighted facts ground normative facts? How? What kinds of normative facts are grounded by weighted facts? Restricting ourselves to gradable weighted facts, what explains their weight? What is the relationship between the grounds of some weighted fact—e.g., some reason) and the grounds for the fact about its weight? We'll now proceed to sketch some different answers to some of these questions. More detailed answers are to be found in the contributions that follow.

2 Machinery in the Theory of Weight

The goal of the last section was to make a preliminary case for thinking that we need to supplement our ideology consisting of some all-things-considered notions with a weighted notion, namely reasons. Assuming that we successfully carried this off, attention should now be turned to *how* reasons explain the all-things-considered. In this section we will introduce some tools that philosophers have used to explain the all-things-considered in terms of reasons.

The most familiar such explanation relates facts about the all-things-considered ought with facts about what one has most reason to do. The facts about most reason are explained in terms of the various facts about what reasons there are to do different things.⁸ Thus we have:

MOST REASON OUGHT: You ought to *x* if and only if your *x*-ing is supported by the greatest net weight of reasons.

⁸Throughout we'll assume that the count noun "reasons" is conceptually and explanatorily prior to the mass noun "reason." This is common and plausible but contentious; for some discussion, see Daniel Fogal's contribution.

We'll leave to the side questions of analysis, i.e., questions about whether ought is more fundamental than the weight of reasons or (as John Broome would have it) vice versa. The interesting question for our purposes is whether this thesis is substantively plausible, and more specifically, with which classes of ethical theories it might be inconsistent. We have two such classes of theories in mind. One class of theories has gradable reasons alongside "side-constraints" such as rights or simply further requirements or prohibitions, where these notions are strict and not gradable.⁹ So long as these side-constraints do not have weights, it may turn out that you ought to do something other than what you have most reason to do. Another class of theories *weakens* this condition on ought. It is easiest to see this by considering cases of supererogation, in which some action might be supererogatory, hence not required, and hence (one might think) not such that you ought to do it. This might be true (one might think) even if the reasons supporting it are weightier than the reasons supporting some alternative.

Suppose we accept this latter view. Then we will be driven to give some account of when some reason or reasons make it permissible for you to perform some action, and when some reason or reasons make some action required. The most common way of understanding all-things-considered notions like permission and requirement is in terms of *sufficient reasons* and *decisive reasons*.

SUFFICIENT REASONS: A is permitted to x just in case there is sufficient reason to x .

DECISIVE REASONS: A is required to x just in case there is decisive reason to x .

Among those with a proclivity toward such things, there are some generally agreed-upon claims about how to understand sufficiency and decisiveness in terms of reasons and their weights. These claims are simple. Let's start with sufficiency. At least as Weighty seems to be on the right track when it comes to analyzing sufficiency:

AT LEAST AS WEIGHTY: The reasons to x are sufficient just in case the reasons to x are at least as weighty as the reasons to do anything else.

This thesis is quite plausible, and has some prominent support.¹⁰ After all, there many things that one is permitted but not required to do. Right now you could continue to read this chapter or reply to some emails or take a break. Those things are all permitted, and plausibly none of them is required. Moreover, it seems as if you have sufficient reasons to do each of these things. Why is this? It's because your reasons for doing each of the things are important enough to justify your doing any of these things *and* because the weight of your reasons to do any of the particular things is on a par with the weight of your reasons to do the other things.

⁹On side-constraints, see Robert Nozick (1968).

¹⁰For instance, see the contribution by Ruth Chang (as well as Schroeder [2007], Schroeder [forthcoming]).

10 Introduction

We have a couple of worries about this proposal. One problem concerns the property of “being at least as weighty as.” It is an open question whether weights are *sharp*. If so, then there will not be many cases in which sets of reasons are equally weighty. And in that case, this property of *sufficient* reason would have been made practically redundant. One way to respond is to understand “at least as weighty as” in terms of “parity” rather than equality. But another problem with this proposal is that intuitively you sometimes have sufficient reason to do something even though you have *more* reason to do something else—for instance, when this alternative is supererogatory. This second problem suggests that we need some additional ideology that provides the *threshold* beyond which some set of reasons favoring some option counts as sufficient. (As we’ll consider shortly, another option is to explain things the other way around: to explain the weights of particular reasons in terms of sufficient reasons.)

When I have *decisive* reasons to do something, things are different. When I have decisive reasons to do something, the reasons I have to do those things *win out* against the reasons there are to do anything else. This supports Weightier Than:

WEIGHTIER THAN: The reasons to x are decisive just in case they are weightier than the reasons to do anything else.

If Weightier Than and Decisive Reasons are true, then you are required to x just in case you have most reason to x . This seems like a platitude of contemporary normative theory.

But similar worries resurface. If you have a small handful of optional reasons, one of which is weightier than the others, plausibly you don’t have decisive reason to do anything. Even if you have a small handful of sets of weighty reasons favoring different options, one of which sets is very slightly weightier than the others, plausibly you don’t have decisive reason to do that thing. Intuitively again, a set of reasons favors an option *decisively* when its net weight is greater than the net weight of other sets of reasons by some threshold. This threshold cannot be explained in terms of reasons of the same kind.

2.1 FILLING THE STORY IN, PART 1: CONDITIONS AND MODIFIERS

With these basics in place, the next step is to explore some tools one might use to understand what it takes for some reasons to be just as weighty or weightier than the reasons in favor of the other options.

We have seen that reasons are considerations that count in favor. But the same consideration in different situations can count more or less in favor of some action. For instance, the fact that Harold will be pleased seems less weighty, other things equal, than the fact that Harold will be utterly delighted. Leave aside substantive questions about which considerations are reasons with which weight. Our interest for now is in characterizing the natures of these normative notions.

The first tools are *conditions* and *modifiers*. A condition or modifier of a reason to x affects the weight of the reason to x even though it is not itself a reason to x or a reason to not x . For instance the fact that you are the only person around who can save the child is not a reason to save the child, but it plausibly increases the weight of your reason. Conditions affect the weight of a reason indirectly. This is because they operate on the existence of reasons. Conditions come in two forms: enablers and disablers. An enabler is a fact that must obtain in order for some reason to exist. So, for example, in order for the fact that I'd win \$1 million by dunking on an NBA-sized basketball hoop to be a reason for me to dunk, I have to have the physiological ability to dunk. Having the physiological ability is an *enabler* of that fact constituting a reason for me to dunk.

Enablers, rather obviously, can affect the balance of reasons in two ways. First, they can go from not obtaining to obtaining. In that case, a new reason will come into existence. This will strengthen the case for reacting in the way that reason recommends. Second, they can go from obtaining to not obtaining. In that case, a reason will go out of existence, which will weaken the case for reacting in the way that fact previously recommended.

Disabling conditions, on the other hand, are facts that, when they obtain, prevent some fact from constituting a reason. For example, the fact that some particular promise was coerced prevents the fact that I promised to x from being a reason to x . Disablers can also affect the balance of reasons in two ways. First, they can go from not obtaining to obtaining. In this case, a reason will go out of existence, thus weakening the case for the reaction that reason recommended. Second, they can go from obtaining to not obtaining. In this case, a reason will come into existence, thus strengthening the case for the reaction it recommends.

Unlike conditions, *modifiers* directly affect the weight of particular reasons even though they themselves are not reasons (at least not reasons for or against the options they modify the case for). There are also two kinds of modifier. The first kind are intensifiers. Intensifiers are facts that make the weight of some reason greater without themselves being reasons. For example, the fact that today is the exhibit's last day is not itself a reason to go to the museum. After all, if the exhibit is terrible, the case for going is not weightier given that it's the last day. However, the fact that it's the last day can intensify the weight of the other reasons to go, when there are other reasons to go.

The second kind of modifier is an attenuator. Attenuators are facts that make the weight of some reason weaker without themselves being reasons. To take an epistemic example, the fact that your color perceptions are only correct 50% of the time attenuates the reason provided by the fact that the wall appears blue to you. It isn't itself a reason to believe the wall is not blue. But it does make the reason provided by your perception weaker than it would be if you were more reliable.

One worry here is that the distinctions between reasons and conditions and modifiers are shallow or context-sensitive. The thought is that we can only separate these notions within a particular context, and the means by which we do this

12 Introduction

will be purely pragmatic. If this is right, then it seems dubious to think that there are these entities—reasons—the weights of which determine what it is permitted or forbidden to do. Instead, we should think that there is a group of facts—the members of which we context-sensitively distinguish in some pragmatic way as reasons, conditions, and modifiers—that collectively determine what is permitted and forbidden. Those facts will together provide sufficient reason, but it's false that any of them individually will be sufficient *reasons*. We think that this is a serious challenge. This challenge is developed in Daniel Fogal's contribution. Ralf Bader's contribution is largely dedicated to meeting this challenge. Bader does this by explicating a metaphysical framework that makes room for and motivates the distinctions between enabling and disabling on the one hand and attenuators and intensifiers on the other.

2.2 FILLING THE STORY IN, PART 2: INTERACTIONS BETWEEN REASONS

Once we have some story about how reasons, conditions, and modifiers combine to generate the weights of reasons—or at least their preliminary weights—we need to look at how reasons interact.

There are various ways in which the weight of one reason can seem to depend upon the weights of others. Some reasons depend on others in straightforward ways: for instance, reasons to apologize plausibly depend upon having reasons not to do the thing you are apologizing for; reasons to blame plausibly depend upon someone having reasons not to do what they are being blamed for doing. Sometimes distinct reasons *overlap*. That is, they don't all contribute distinct weight. To give an example, the fact that Ronnie likes to dance and the fact that there will be dancing at the party both provide distinct reasons but it's implausible that each reason contributes a distinct amount of weight. Sometimes reasons seem to attenuate other reasons. For instance, your reason to try your best to beat your opponent at tennis attenuates your reason not to ruin your opponent's day.¹¹ And sometimes situations are just awkwardly set up. Jonathan Dancy asks us to:

[C]onsider a case in which I am thinking of doing something for a friend. My action, were I to do it, would be good, and partly good because it is an expression of our friendship. But now, if I were to be doing the action and not doing it for a friend, I would presumably be doing it for someone who is not a friend, and it might be that doing it for someone who is not a friend is even better than doing it for a friend. . . . [O]ur friendship seems to be a reason to do the action even though if we were not friends I would have even more reason to do it.

All these kinds of cases challenge *separability*, the thesis that the grounds for the weight of any given reason never include facts about the weights of any other

¹¹Scanlon (2004). See also Joseph Raz's discussion of "exclusionary reasons" [1990 and elsewhere].

reasons. Separability is important since it is *prima facie* required in order to defend *additivity*, the thesis that the total net weight of the reasons favoring any action is the sum of the reasons in favor minus the sum of the reasons against. Additivity, or something like it, is attractive because it is intuitive, workable, and theoretically straightforward. Work is required to be done either way to give an account of the interaction between reasons, whether to defend additivity from these kinds of cases, or to defend the most systematic alternative that seems to be available. Different substantive or metaphysical accounts of weight may also deal with these issues in different ways.¹² Both Ralf Bader and Shyam Nair's contributions examine how we should understand how and when the weight of separate reasons add up.

Some philosophers are interested in systematizing the ways in which the weights of some reasons are dependent or derivative on the weight of other reasons (Shyam Nair's contribution makes progress on how to understand such systematizations). They do this by offering *weight transmission* principles. These principles tell us when (at least some of) the weight of one reason is transmitted to the weight of another reason. We will mention two of the most talked about and relied upon transmission principles.

Both principles are about the relationship between the weight of the reasons to take some end and the weight of the reasons to take the means to that end. The first principle is the Necessary Means Transmission Principle (NMTP):

NECESSARY MEANS TRANSMISSION PRINCIPLE: If the weight of *A*'s reason to *x* is *n* and *y*-ing is a necessary means to *x*-ing, then there is a reason for *A* to *y* that has at least weight *n*.

The NMTP is intuitively compelling. If one must *y* in order to *x*, then it's hard to see how one could have less weighty reasons to *y* than to *x*. Any considerations that weaken the case for *y*-ing seem to weaken the case for *x*-ing, as well. And this is precisely *because* *y*-ing is necessary to *x*-ing.

The other principle is the Sufficient Means Transmission Principle (SMTP):

SUFFICIENT MEANS TRANSMISSION PRINCIPLE: If the weight of *A*'s reason to *x* is *n* and *y*-ing is a sufficient means to *x*-ing, then there is a reason to *y* with weight *m*, where *m* is not greater than *n*.

The SMTP differs from the NMTP in two important ways. First, it concerns sufficient means, not necessary means. There can be multiple sufficient but unnecessary means and there can be necessary but insufficient means. Second, the SMTP only guarantees that a reason to take the sufficient means exists. It doesn't guarantee that that reason is as weighty as the reason to pursue the end. This is

¹²For instance, Maguire (2013) argues that a value-based theory of reasons does a better job than reasons fundamentalism in accounting for separability and overlap.

14 Introduction

for good reason. Cases abound where one has less reason to take some sufficient means than they do to take the end—e.g., cases where *y*-ing is a sufficient means to *x*-ing but it involves significant costs that other sufficient means do not involve.

Despite the fact that the SMTP only guarantees a reason to take the sufficient means, it can be seen as a weight transfer principle. This is because the reason to take the sufficient means guaranteed by the fact that there is a reason for the end will bolster the case for performing the action that constitutes the sufficient means. Thus, the SMTP has consequences for the weight of the reasons to perform the act that constitutes the sufficient means. There is clearly plenty more that needs to be said about these issues.

3 Analytical Issues in the Theory of Weight

In the last section we introduced some of the paradigm features associated with the weight of reasons. Now our question becomes: What type of general theories can one adopt in order to explain these features?

We will organize the discussion around an important choice point for anyone in the business of theorizing about weight. So far we have played fast and loose with *weight* on the one hand and *weightier than* on the other. Our organizing choice point is whether or not one takes *weight* to be more fundamental than *weightier than* or vice versa.

Those that take *weight* to be more fundamental think that we can analyze what it takes for particular reasons to have strength of a certain magnitude in isolation from the strengths of the other reasons. Once we determine what weight each individual reason has, we can then come up with an ordering of these weights. This ordering will determine the relational *weightier than* facts. We'll call these theorists weight fundamentalists.

Those that take *weightier than* to be fundamental think that the ordering of relative weights comes first. They think that we first analyze what it is for some reason (or set of reasons) to be weightier than another reason (or set of reasons). Once we do that, we can determine which weights individual reasons have. We'll call theorists of this type weightier-than fundamentalists.

This section will have three parts. First, we'll discuss some weight-first views. Then we'll discuss some weightier-than-first views. Finally, we'll discuss two views that are neither weight-first views nor weightier-than-first views. Our goal here is not to settle which view is best. Instead, we'd like to get the views on the table and point out some virtues and some bugs with each view.

3.1 WEIGHT IS PRIOR TO WEIGHTIER-THAN

This subsection is dedicated to *weight* fundamentalist views. According to weight fundamentalism, we first determine what each reason's weight is. Once we have

etermined these facts about weight, there will be some function that takes us to an ordering of the various reasons. The most natural weight fundamentalist views hold that reasons aren't fundamental. These views seek to understand reasons in terms of some other notion. We'll focus on the two most prominent views.

3.1.1 The Desire-Based View

According to desire-based theories of reasons, in order for some fact f to be a reason for A to x , f must (at least partly) explain why x -ing would promote at least one of A 's desires. Since desires themselves have strengths, desire-based theories have a natural place to turn when it comes to the weight of reasons. A very natural view for the desire-based theorist to hold is that the weight of some reason r for A to is determined by the strength of the desire promoted by A 's x -ing and the extent to which x -ing promotes that desire. Following Schroeder (2007), we'll call this view *Proportionalism*.

The main virtue of Proportionalism is its simplicity and concreteness. Since it reduces weight to something (1) we have a pre-theoretical grasp on and (2) we need to posit anyway, it takes much of the mystery out of what weight is. It also is in line with traditional naturalistic motivations for desire-based theories of reasons.

That said, there are several problems with Proportionalism. First, it makes counterintuitive predictions. This is because one can have strong desires to bring about bad states of affairs. It is counterintuitive that one *thereby* has a weighty reason to bring about that state of affairs. As Schroeder (2007) points out, the weight some reason has seems to be a normative matter that is not settled simply by how strongly one happens to desire something. These types of considerations lead Schroeder to reject Proportionalism even though he accepts a desire-based view of reasons.

Proportionalism also has trouble accounting for all of the machinery laid out in section 2. The most pressing problem has to do with modifiers. It is unclear how the view can account for modifiers, as they are not themselves reasons even though they can affect the weight of reasons. It doesn't seem like we can get modifiers from the strength of desires or degrees of promotion in any extensionally intuitive way. Thus, if Proportionalism is going to account for them, it will need to add some ideology. It is mysterious to us what could be added that would maintain the simplicity, concreteness, and naturalistic friendly virtues that attracts people to the account to begin with. We mention this problem here not to settle whether it can be done, but to make it clear that defenders of Proportionalism need to innovate in order to account for some of the phenomena (or explain away the appearance of the phenomena).

In her contribution, Kate Manne presents a novel desire-based theory of reasons, and goes into some detail to show how it has the resources to account for modifiers and conditions.

16 Introduction

3.1.2 The Value-Based View

The value-based theory of reasons starts with the idea that reasons for action obtain in virtue of facts about the relations between actions available to you and states of affairs that are finally neutrally valuable—that is, valuable objectively and for their own sake. According to one version of the view, for some fact to be a *basic reason* to x is for it to be a fact of the form $\langle x$ -ing would realize $S \rangle$, where x is an action available to you and S is some finally valuable state of affairs. *Non-basic reasons* are facts that partly or fully explain basic reasons. For instance, the fact that you will please Alfred immensely by spending the afternoon playing chess with him is a basic reason to do so, since Alfred's being pleased is valuable. Moreover, the fact that Alfred really likes chess is a non-basic reason to spend the afternoon playing chess with him, since the fact that he really likes chess partly explains the fact that your playing a game of chess with him would immensely please him.¹³ We distinguish both basic and non-basic reasons from facts explained or entailed by these facts, e.g., the fact that there is at least one reason to x , and from evidence that some basic reason or non-basic reason obtains, where that evidence is not itself a basic reason or a non-basic reason.

This makes way for a rather natural account of weight, namely one according to which the weight of a reason is some function of the value of the state of affairs S that features in the basic reason. It is simplest to assume that this will be an increasing function.¹⁴ Then, the more valuable Alfred's pleasure, the greater the weight of your reasons to play chess with him, other things being equal. Similarly, the greater the value of your studying the theory of weight, the weightier the relevant reason to do so, other things being equal; the greater the disvalue of eating that cheeseburger, the weightier your reason against doing so. The distinction between basic and non-basic reasons provides one with an account of when reasons *overlap*, which in turn makes the business of weighing reasons simpler. Two reasons overlap just in case they are explained by the same finally valuable state of affairs. For instance, the fact that playing chess will please Alfred immensely overlaps with the fact that Alfred really likes playing chess. Consequently, only one counts toward what you ought to do.

The value-based theory of weight faces two main objections: the objection from partiality, and the deontic objection.

¹³Another version of the view holds that in order for certain descriptive facts to be reasons, some background facts connecting the act favored and value must obtain. Importantly, the facts connecting the act favored and value are merely background conditions that must obtain in order for the descriptive facts to provide reasons. On this view, then, the fact that playing chess with Alfred would instantiate some final value is not itself a reason. Instead, that fact obtaining is merely a background condition that must be met in order for some other fact—say, the fact that Alfred likes to play chess—to provide a reason. (For a view with a similar structure to this view, see Schroeder 2007). See also Errol Lord (forthcoming).

¹⁴In his contribution, Joseph Raz presents some important challenges to this assumption. We'll leave these worries aside for now, to discuss the simpler version of the value-first account.

The objection from partiality maintains that the weights of our reasons vary not just with the values of states of affairs realizable by our actions, but also with facts about the nature of our relationship to this state of affairs. Sometimes the person whose life is at stake is our lover, sometimes our mortal enemy. Their deaths might be objectively just as tragic, but intuitively you have a weightier reason to save your lover than your enemy.

A familiar attempt to solve this problem rejects the notion of agent-neutral value for an agent-relative conception of value, one according to which notion is valuable simpliciter, but things are only ever valuable-relative-to-an-agent. It seems to us that this response throws out the baby with the bathwater, since it abandons the central idea motivating value-first theories, namely that Alfred's pleasure, your ethical understanding, the child's not drowning in the pond, and so on are valuable objectively.

A more promising attempt to solve this problem starts with the idea that the function from value to weight might contain other factors. For instance, one might want to add a place in the function for the (subjective or objective) likelihood that your action will be successful in realizing the relevant state of affairs. Or one might want to add a place for temporal discounting, or perhaps for giving absolute priority to those who are badly off. To deal with partiality, we add a place to the function for the "moral distance" between the agent and the state of affairs to be realized. This notion of moral distance is to be accounted for in natural or axiological terms. Ralf Bader works out some of the details of this way of accounting for partiality in his contribution.¹⁵

The second objection stems from a broader objection to the value-based theory of reasons, namely that some reasons seem not to have their source in value, e.g., reasons arising from promises and from people's rights. This objection seems even more problematic when applied to the value-based theory of the weight of reasons, since promises and rights not only seem to be normatively significant, but also seem to generate reasons with a kind of stringency that seems not to be gradable in the same way as value, nor to be easily compared with value.

The idea that deontic reasons are not gradable in the same way as value can be exaggerated. Starting with an "extreme comparison" point, it is clear that in cases of mild rights and extreme values that we can have more reason to violate the right, and similarly with promises. More generally, a plausible line of response to such worries is to maintain that promises and rights are normative only in virtue of their relations to values. T. M. Scanlon's work (2004) is one prominent attempt to elucidate the values that explain the normativity of promises, e.g., the values and disvalues associated with expectations, trust, and cooperative planning.¹⁶

¹⁵A similar strategy is developed in Maguire (2013).

¹⁶See also Walter Sinnott-Armstrong (2009).

Correspondingly, the weight of one's reason to keep a promise will be a function of the amount of whatever value it is that explains the normativity of the promise in question. In the case of rights, this *will* have implications for whether you have more reason to save five people than one, for instance. It will turn out, as against John Taurek (1977), that you do.

Finally, it is important to point out that there is room to develop a value-based theory of weight, according to which weightier-than is explanatorily prior to weight, by appealing to the comparative notion of better-than.

3.2 WEIGHTIER-THAN IS PRIOR TO WEIGHT

3.2.1 Higher-Order Reasons Views

The most worked-out *weightier-than* fundamentalist view is the *higher-order reasons view*. According to HOR views, the (relative) weight of the set of reasons S to is determined by the weight of one's reasons to do something else.¹⁷ In the version of the view due to Schroeder (2007), the higher-order reasons that are relevant are reasons to place certain weight on S in deliberation. This gives us HOR-First Pass:

HOR-FIRST PASS: The set of reasons S for A to x is weightier than the set of reasons S' to perform some alternative act y just in case A has more reason to place more weight on S in deliberation than S' .

A very nice feature of HOR views is that they provide an elegant treatment of modifiers. This is because HOR theorists can (and do) hold that attenuators and intensifiers are reasons to either put less weight on a set of reasons in deliberation (attenuators) or more weight on a set of reasons in deliberation (intensifiers). Thus, on HOR views, modifiers are themselves reasons—reasons to place less or more weight on other sets of reasons. This is attractive mainly because it allows the view to be very flexible in its treatment of modifiers, which seems demanded by the phenomena.

That said, HOR-First Pass cannot be right. This is because it gives rise to both a regress and a charge of vicious circularity. Let's start with the regress. The first pass holds that one's reasons to x are weightier than one's reasons to y just in case one's reasons to place weight on one's reasons to x are weightier than one's reasons to place weight on one's reasons to y . But, of course, we need to understand what it takes for one's reasons to place weight on one's reasons to x to be weightier than one's reasons to place weight on one's reasons to y . This, according to HOR, is true just in case one's reasons to place weight on the reasons to place weight on S are weightier than one's reasons to place weight on the reasons to place weight on S' . But then we need to know what it takes for *this* to be true. The worry is that this

¹⁷One needn't think that the things with weights are sets of reasons to hold a higher-order reasons view. We will set things up this way just because this is the way that Schroeder sets them up.

can go on *ad infinitum*. If so, then it's not clear that HOR views have any explanatory power at all.

The circularity worry arises because HOR-First Pass analyzes *weightier than* in terms of *weightier than*. This looks bad since we expect the analysans of *weightier than* to be things other than *weightier than*! To put it more metaphysically, it seems like HOR-First Pass holds that all *weightier than* facts are dependent on other *weightier than* facts. This is obviously problematic if the HOR-First Pass is supposed to be an analysis of *weightier than*. If it is an analysis, the *weightier than* facts better turn out to only be dependent on some more fundamental facts.

Schroeder's (2007) response to these worries is two-fold. First, he makes his view *recursive*. That is, he adds a base case clause. So instead of it always being the case that when a set of reasons *S* is *weightier than* a set of reasons *S'*, the set of reasons to place weight on *S* is *weightier than* the set of reasons to place weight on *S'*, *S* can *also* be *weightier than* *S'* if there are *no* reasons to place weight on *S'* and some reason to place weight on *S*. So the analysis becomes a standard disjunctive recursive analysis. It becomes HOR-Recursive:

HOR-RECURSIVE: A set of reasons *S* is *weightier than* a set of reasons *S'* just in case either (1) *S* is non-empty and *S'* is empty or (2) the set of reasons *S''* to place weight on *S* is *weightier than* the set of reasons *S'''* to place weight on *S'*.

The second part of his response consists in the claim that it's not plausible for the chain of reasons to place weight on other reasons to go on *ad infinitum*. At some point in all weighings, he thinks, we'll get to a place where there are some reasons to place weight on *S* and no reasons to place weight on *S'*. These will be the base cases. The first disjunct of the analysis will kick in and then the second disjunct will handle all of the other links in the chain.

This move clearly solves the regress problem. It's not as clear that it solves the circularity problem. However, given plausible assumptions about the nature of recursion, it also solves the circularity worry. It's generally thought that recursive definitions are analyses (although, as far as we can tell, this is seriously undertheorized). This is because, in this case, all of the chains of weighings will bottom out in the base case, which doesn't invoke *weightier than*.

3.3 NEITHER WEIGHT NOR WEIGHTIER-THAN IS FUNDAMENTAL

This section is about two views represented in the literature that take neither *weight* nor *weightier than* to be fundamental.

3.3.1 Scanlon's Strict First View

Recently, T. M. Scanlon (2014) has advanced a novel theory of weight. According to Scanlon's view, the strict notion of sufficient reason is more fundamental than weight. On Scanlon's view, there are at least two fundamental reason relations: *pro tanto* reasons and sufficient reasons. Sufficient reasons, roughly, make the act they

recommend permitted. Scanlon understands *outweighing* (which, we assume, is intimately related to being weightier than) in terms of sufficient reasons. Scanlon analyzes outweighing in the following way:

OUTWEIGH: A reason r outweighs a reason r' to just in case r is a sufficient reason to x and r' is not a sufficient reason to x even though r' would be a sufficient reason to x if the conditions that give rise to r were not met (108)

Outweigh captures a core part of a theory of weight, but not all. Scanlon recognizes this and attempts to analyze other parts of the theory in terms of the notion of outweighing. He starts by analyzing what it takes for a reason r to be weightier than some other reason, r' . According to this analysis, r is weightier than r' if r would outweigh a reason r'' to perform some alternative act and r' wouldn't outweigh r'' were r'' to obtain.

This position avoids both weight-first and weightier-than-first fundamentalism. But it is far from clear why this is advantageous for its own sake, and it is not clear what other advantages Scanlon's view enjoys. Moreover there are several problems with this view. First, by making sufficient reason more fundamental than weight, Scanlon is forced to give up on the idea that all-things-considered notions like sufficiency can be understood in terms of reasons and their weight. As we pointed out in section 1, there is considerable theoretical pressure to understand the all-things-considered notions in terms of some weighted notion. It seems as if Scanlon has to give up on understanding sufficient reasons in terms of reasons and their weight.

A second problem for Scanlon's view has to do with its lack of generality as stated. A reason can only outweigh its competitors when it is a sufficient reason.¹⁸ But it's natural to think that nearly all reasons outweigh *some* reasons, even the ones that fail to be sufficient. Scanlon recognizes this point to some extent. That's why he gives a counterfactual account of when a reason to outweighs another reason to x . This is far from accounting for all the phenomena, though. We'll mention one other phenomenon and suggest an analysis in a Scanlonian spirit.

Let's focus on reasons that are insufficient for me to x right now. Some of those reasons are weightier than others. So, for example, I have some reason to take a nap right now. This reason is insufficient because I really need to get this writing done. I also have some reason to go to the airport and fly to Rome to see Bernini sculptures. This reason is also outweighed by my reason to write. It is also less weighty than my reason to take a nap—it would disrupt my life too much. Given what Scanlon has actually said, he has no resources for explaining why my reason to take a nap is stronger than my reason to fly to Rome.

A counterfactual account is not far away, though. The natural move for Scanlon to make is to hold that my reason to take a nap is stronger than my reason to

¹⁸This point is made independently in Schroeder (2015).

fly to Rome because were those the only reasons in play, my reason to take a nap would be sufficient, while my reason to fly to Rome would not.

The final problem we will mention is a very general problem for counterfactual accounts of almost anything: The Conditional Fallacy. The conditional fallacy arises because there is no way to guarantee in advance that nearby worlds will behave in the way one needs them to behave in order to make the right predictions. So, for example, imagine that in the nearest worlds where only the nap reason is in play and the Rome reason is in play, my semester is over, my writing obligations fulfilled, and my vacation impending. I'm much less tired in this world and much freer to go on vacation. In this world, it's not so clear that my reason to nap is a sufficient reason. But if it isn't, then it follows that *in the actual* world my reason to nap is not weightier than my reason to go to Rome.

Perhaps you don't like this case. That doesn't matter much. This is because we should still be confident there will be counterexamples of this kind. There just is no way to fix the nearby worlds in the right way in every case. Once you start dropping reasons that exist in the actual world, things might change in ways that affect which reasons are sufficient in the nearby worlds. One shouldn't hold facts about the weights of the *actual* reasons hostage to such counterfactual facts.

3.3.2 Horty's Austere View

The final view we will mention is John Horty's. Strictly speaking, Horty provides no account of *weight* or *weightier than*.¹⁹ These notions play no role in his view at all. Instead, he gets by just with reasons. Despite the sparseness of his view, it resembles the higher-order reasons view in many key ways.

Horty models reasons with default rules. Default rules tell you when you can draw conclusions, by default, when certain facts obtain. So, for example, you might have a default rule that tells you to conclude that Lindsey is happy when she is informed that she was granted tenure. If we let B stand for "Lindsey is happy" and A stand for "Lindsey is informed that she was granted tenure" and " \rightarrow " stand for the default rule operator, we can represent this by saying $A \rightarrow B$.

As we all know, reasons can conflict. Since Horty is modeling reasons with default rules, it follows that often the default rules that apply to one conflict. But, of course, some of those are better than others in the sense that they are more important. In order to model this, Horty introduces an ordering relation (" $<$ ") between default rules. Thus, when r and r' are two default rules, $r < r'$ means that r' takes priority over r . When the default rules that are applicable conflict, we turn to the ordering relation to see which ones win out. In cases where there are no conflicts between the default rules that are applicable, the option licensed by those rule(s) win(s) out.

¹⁹Horty downplays this feature of his view quite a bit—especially in his contribution to this volume. We can see that there are rhetorical reasons for him to do this. We think, however, that this part of his view brings with it some serious advantages, especially for those who think that reasons are the fundamental constituents of the normative (more on this below).

A natural question arises at this point: How do we determine the ordering? Horty models this by introducing default rules about giving priority to other default rules. So, to use an example from the last section, suppose I need to get this writing done (A), I'm tired (C), and Rome is a fabulous place to vacation (E). Given A, I can conclude by default that I write (B)— $A \rightarrow B$. Given C, I can conclude by default that I take a nap (D)— $C \rightarrow D$. And given E, I can conclude by default that I fly to Rome (F)— $E \rightarrow F$. Let's call the first default r , the second r' , and the third r'' . These defaults conflict. On Horty's view, there are defaults to give priority to some defaults over others. Given the stipulations of the case, it seems like given A and B and C (let's call this conjunction G), I can conclude by default $r'' < r' < r - G \rightarrow r'' < r' < r$.

In many ways Horty's view and Schroeder's HOR view are alike. Horty's default rules to give more priority are just like Schroeder's reasons to place weight. Since Horty models reasons with default rules, the default rules to give priority are *reasons* to give priority. Since Horty's view has a very similar structure to Schroeder's, it also has an easy time accounting for modifiers. Modifiers will be modeled by default rules to give priority to certain other default rules.

But Horty's view enjoys other advantages to Schroeder's. First, it is at least a little bit virtuous—or at least *cool*—that Horty's view has so much predictive power given the sparseness of its ideology. Second, Horty's view completely avoids the regress and circularity worries that arise for Schroeder. It avoids the circularity worry because Horty isn't trying to analyze anything with the higher-order defaults. They are just defaults like any other. So there is no threat of having a circular analysis.

One might worry, though, that the regress worry still looms. After all, we haven't settled whether the reason to place priority on r is itself undefeated. There could be other reasons to place priority in play. In that case, we'll need a yet higher-order default about which second-order defaults to prioritize. This might go on for a long time. The worry is that it will go on forever. If it does, then there won't be a fact of the matter about which first-order reasons take priority.

We agree with this conditional but think that reflection on Horty's view shows that it's quite implausible that this could go on forever. There are only so many relevant facts. At some point, we think, the conflict must end. Once it does, Horty's ordering gets off the ground as long as he assumes the plausible principle that a reason unconflicted is a reason undefeated. Since he does accept this principle, we don't see the regress worry as very worrisome.

Again, Horty's solution to the regress problem is very similar to Schroeder's. It relies on the claim that the conflict must end. It's worth pointing out, though, that the regress poses much less of a theoretical problem for Horty than it does for Schroeder. Recall that Schroeder had to make his account disjunctive in order to save his view from this problem. Horty doesn't. Horty just has to assume that a reason unconflicted is a reason undefeated.

The real upshot of the comparison between Horty's view and Schroeder's, we think, is that there is surprising theoretical advantage to restricting one's ideology

in the way Horty does. It's plausible that Horty's view retains the basic structure and thus the advantages of Schroeder's view, but it costs much less to maintain. Depending on your philosophical predilections, this might provide you with strong reason to dispense with the notions of *weight* and *weightier-than* altogether.²⁰

4 Conclusion

We have argued that ethical theory needs to account for weighted notions, and that there is plenty of theoretical work that needs to be done to sort out different candidate weighted notions, and working out how they relate to each other and to other normative facts.

It is clear that attending to whether or not one includes weighted notions will have serious substantive implications. Among other things you will, for instance, save your brother's life by lying to the murderer at the door. But it should also be clear that attending carefully to the theory of these weighted notions will also have serious substantive implications. Just within these pages you can find applications to the rational choice theory (Ruth Chang's contribution), to the dualism of practical reason (Karl Schafer's and Stephen Darwall's contributions), and to reasoning with precedent in the common law (John Horty).²¹

Works Cited

- Broome, J. 2004. "Reasons". In *Reason and Value: Themes from the Moral Philosophy of Joseph Raz*, edited by Wallace, Smith, Scheffler, and Pettit. Oxford: OUP.
- Dancy, J. 2004. *Ethics without Principles*. Oxford: OUP.
- Holton, R. 2002. "Particularism and Moral Theory." *Aristotelian Society Supplementary Volume 76* (1): 191–209.
- Lord, E. (forthcoming), 'Justifying Partiality.' In *Ethical Theory and Moral Practice*.
- Maguire, B. 2013. *Values, Reasons, & Ought*. PhD dissertation, Princeton University.

²⁰We think this might be particularly helpful for theorists who think that reasons are the most fundamental constituents of the normative. A simple reason for the fundamentalist to opt for Horty's view is that it eliminates one normative property one is required to analyze in terms of reasons. Moreover, as we've seen from Schroeder, it is not trivial analyzing weight (or weightier-than) in terms of reasons. Schroeder ultimately must embrace a disjunctive recursive analysis that, it seems to us, isn't obviously the type of metaphysical analysis that the reasons fundamentalist is after. Even if recursive analyses can be seen as reductive metaphysical analyses, it is a cost to have to opt for a disjunctive analysis (as Schroeder recognizes; see ch. 7 of Schroeder [2007]).

²¹For helpful comments, many thanks to Derek Baker, Brad Cokelet, Jeff Horty, Tristram McPherson, Travis Rieder, Kurt Sylvan, Jonathan Way, and Jack Woods.

24 Introduction

- Nozick, R. 1968. "Moral Complications and Moral Structures." *Natural Law Forum*. Paper 137.
- Raz, J. 1990. *Practical Reason and Norms*. Oxford: OUP.
- Ross, W. D. 1930. *The Right and the Good*. Oxford: OUP.
- Scanlon, T. M. 1998. *What We Owe to Each Other*. Cambridge: Harvard University Press.
- Scanlon, T.M. 2014. *Being Realistic about Reasons*. Oxford: OUP.
- Schroeder, M. 2007. *Slaves of the Passions*. Oxford: OUP.
- Schroeder, M. 2015. "Being Realistic about Reasons, by T. M. Scanlon." *Australasian Journal of Philosophy* 93(1).
- Schroeder, M. Forthcoming. "What Makes Reasons Sufficient?" *American Philosophical Quarterly*.
- Sinnott-Armstrong, W. 2009. "How Strong is this Obligation? An Argument for Consequentialism from Concomitant Variation." *Analysis* 69(3): 438–442.
- Taurek, J. M. 1977. "Should the numbers count?" *Philosophy and Public Affairs* 6 (4): 293–316.