

Commitment: Worth the Weight

To appear in *Weighing Reasons*, edited by Errol Lord and Barry Maguire, forthcoming from OUP

The last few decades of moral philosophy and epistemology have become increasingly preoccupied by the turn to reasons. Reasons are taken to be so central to normative inquiry that it is sometimes claimed that what makes normative inquiry normative is just that it is about reasons. The vernacular of reasons is used to elucidate theses as diverse as evidentialism in epistemology and egalitarianism in political philosophy.¹ But once we make the turn to reasons, we have to worry about their weight. For what makes reasons important is how they balance in competition with one another in order to determine what we ought to do. And what we ought to do is determined not by what we have the most reasons to do, but by what we have *most reason* to do – meaning the reasons that are *weightiest*. Until we know more about when reasons are weighty, what makes them weighty, and how the weight of reasons determines what we ought to do, therefore, we are still in only the very earliest stages of normative inquiry.

This paper takes a very indirect approach to the question of how reasons and their weights come together to determine what we ought to do. The approach is indirect, because we will be concerned not with the weight of reasons directly, but with two other important and, as we will argue, *weighty* normative concepts: obligation and commitment. Obligations, we will argue, are a distinct normative concept that also admits of weighing, just as reasons do. We call attention to W.D. Ross's [1930] natural, simple way – which we call *Simple Weighing* – of construing the manner in which both reasons and obligations are weighed.²

¹ For example, according to Jonathan Adler [2002, 9], evidentialism's "claim is that to (fully) believe that p one needs adequate reasons". Similarly, according to Derek Parfit [1997, 209], one important form of egalitarianism is the thesis that "we always have a reason to prevent or reduce inequality, if we can".

² Ross, of course, conducts his discussion in terms of duties, rather than reasons or obligations, so we are generalizing, here. If we are wrong to think of Ross's views as views about obligations, in our sense, then *Simple Weighing* is merely Ross-inspired, rather than strictly Ross's own view. Either way, we find the comparison to Ross instructive.

We then introduce the notion of commitment, as a third normative concept that admits of weighing. But we argue that *Simple Weighing* is inadequate for commitments, and thus that we need something more sophisticated. This is an important result because, as we go on to argue, commitments are actually a special case of obligations, so obligations in general need a more sophisticated weighing process, too. If all of this is right, then we'll have shown that weighing obligations is harder than it looks. The payoff of all of this for our understanding of the weight of reasons is a simple challenge: if Ross was wrong about how obligations weigh, could Simple Weighing also be wrong about how reasons weigh?

1.1 Reasons as Competitors

We take it as given that one of the most important features of reasons, and one of the chief reasons why they have come to play such a prominent role in both moral philosophy and epistemology, is that they compete with one another. There are often reasons in favor of each of multiple conflicting options, or reasons for and against the very same option. But the conflict among reasons does not leave us deeply puzzled about what to do; on the contrary, we think that often some reasons win out over the reasons with which they compete, and that when they do, we ought to do the action which they support. Though there may be reasons both for going to USC and reasons for going to UCLA, the reasons on one side are better, and that is what we ought to do (to go to USC, that is).

By the 'weight' of reasons, we simply mean whatever dimension along which reasons are ordered, that is determinative of which reasons win in their competitions among one another. We sometimes say that reasons are 'weighty', and at other times that they are 'strong'; we see no essential difference between these ways of talking, though each is more natural in some circumstances than in others. Both employ, we take it, extended or loosely metaphorical uses, as reasons are not 'weighty' or 'strong' in their primary literal senses. At other times we speak of some reasons being 'better' than others; we take it that in normal contexts this is true just in case the former are weightier than the latter. However, we caution against identifying the meaning of 'good' with 'weighty' or of 'better' with 'weightier' in this context for the same reasons that we

would caution against identifying the meaning of ‘good’ with ‘fast’ when applied to racecars.³ Reasons are good in virtue of being weighty, just as racecars are good in virtue of being fast. In what follows we will appeal indifferently to the language of strength or of weight, whichever appears to be most natural in the situation.

Some philosophers distinguish between ‘objective’ and ‘subjective’ senses of ‘reason’, but we will ignore this distinction in much of what follows. If our arguments in this paper are on the right track, however, it will be natural to conclude that objective and subjective reasons may be weighed in different ways.

1.2 Other Competitors: Reasons vs. Obligations

The concept of a reason is not the only important normative concept that allows for such competition. Another important normative concept that admits of conflict and competition in the determination of what an agent ought to do is that of *obligation*. In asserting that obligations admit of conflict, we do not mean to be weighing in on the important debate about whether there are conflicts in *all-things-considered* obligation – cases in which an agent is all-things-considered obligated to do each of two things that are jointly incompatible.⁴ We mean only to be pointing out that in our ordinary, common-sense use of ‘obligation’ as a count noun, an agent may have many different obligations, and sometimes these obligations may come into conflicts which need to be resolved.

For example, it is highly plausible that when an agent makes a fully-informed and uncoerced promise to do something that is otherwise morally indifferent or praiseworthy, she incurs an obligation to do that thing. If you have made many promises, then you will be under many obligations. But it is easy for promises to lead to conflicts. If you promise each of Bill and Ben to attend their birthday parties, but unbeknownst to you they are scheduled for the same time, then you have conflicting obligations. A naïve way to try to evade this conclusion might be to suppose that you have an obligation to whomever you promised first, and argue that once you’ve made that promise, it’s wrong to go to the second party (because of the conflict), and hence

³ Contrast Gregory [2013].

⁴ See, for example, Van Fraassen [1973], Marcus [1980], Brink [1994], and Horty [2003].

the promise to do so does not result in any obligation. But this naïve response doesn't work for simultaneous promises, or for promises that are consistent when made, but turn out, given later developments, to come into conflict. In fact, later developments can often generate obligations that conflict with a promise. If you promise Alex to meet her for lunch, but on your way there you encounter someone who has just been in a bad cycling accident and needs your immediate help to get to the hospital, your obligation to Alex doesn't go away, even though it is probably outweighed. So obligations can come into conflict with one another.⁵

But even though obligations come into conflict with one another, some of them clearly win out over others. Your promise to Alex is outweighed by the pressing needs of the injured cyclist, and your promise to Bill may outweigh your promise to Ben – for example, if you are married to Bill but Ben is just your brother-in-law. So not only do obligations conflict, but they *compete*, and weigh against one another in what looks much like the way that reasons do. So reasons, it seems, are not alone in having weights.

Of course, if obligations are just reasons – or conversely – then this conclusion would be too hasty. For then obligations would not be *another* normative concept that admits of weight, but just an interesting special case (or alternatively, the more general case) of reasons. However, we think this is unlikely. Though it may well be that either reasons or obligations can be analyzed in terms of the other, obligations cannot simply be a special case of reasons, and reasons cannot simply be a special case of obligations. For there are at least three important differences between obligations and reasons.

⁵ There is a variety of evidence that the obligations generated by promises are not all-things-considered but do persist in the face of contrary obligations. If promissory obligations were all-things-considered, then it would be too hard to obligate oneself by making a promise, because many circumstances can make it permissible, all-things-considered, not to keep a promise. And if they did not persist in the face of contrary obligations, we would encounter several puzzles. For example, suppose that you promise both Ben and Bill to attend their parties, but your promise to Ben is more important. If this made your lesser promise to Bill go away, then it would be puzzling why, when Ben cancels his party, you should go to Bill's. Or to take another case, suppose that your promises to Ben and Bill are equally important. If neither generates an obligation in this case, then it is puzzling why you ought to go to at least one of the parties (Brink [1994], Horty [2003], Nair [unpublished]). And finally, if your obligation to Alex goes away when you encounter the injured cyclist, it is puzzling why explaining to Alex that you had no obligation to her because the cyclist's injury was more important does not constitute an acceptable form of apology.

Our first point is suggestive: obligations are *strict* in a way that reasons are not.⁶ They make *demands* on us, and not just suggestions. When you have an obligation to go to Bill's party and an obligation to go to Ben's, and you go to Bill's party rather than Ben's, you are still subject to an important kind of criticism, even though you did the best that you could do.⁷ In contrast, you have reasons to eat lunch at Lemonade and reasons to eat lunch at Moreton Fig. But you can only eat at one place, so you choose to eat at Moreton Fig, since those reasons are better. You've made the right choice, and so you aren't criticizable for failing to eat lunch at Lemonade. This is true even if the reasons to eat at Lemonade are excellent, forceful reasons. So reasons seem to lack this 'strict' character of obligations.

More importantly, there is a striking difference in the ontology of reasons and obligations. If I tell you that you have a reason to eat lunch at Lemonade, it is natural for you to ask me what that reason is, and for me to respond by telling you that it is their excellent salads, or perhaps that this week they have a special on soups. In contrast, if I tell you that you have an obligation to attend Bill's party, it does not make great sense for you to ask me what your obligation is, because my only natural answer is that it is an obligation to attend Bill's party. Obligations, it seems, are just nominalizations of the *is obligated to* relation between agents and actions. There is no real intelligible question as to what they are. But reasons, in contrast, are genuine relata of the *is a reason for... to* relation. This fact by itself makes it very puzzling how obligations could just be reasons, or vice-versa.⁸

Finally, obligations are at least sometimes *directed*, but reasons are not. Your obligation to attend Bill's party is an obligation *to Bill*, and my obligation to attend Bill's party, since I promised you to spend the evening with you there, is an obligation *to you*. Plausibly not all obligations are obligations to anyone in particular, but at least many important obligations are, and promissory

⁶ Though contrast Dancy [2004], Greenspan [2005], and Gert [2005], who each argue that some reasons *are* strict in this way.

⁷ For a defense of and elaboration on this gloss on strictness in terms of criticizability, see Shpall [2013].

⁸ Note that though we cannot ask what your obligation is, we *can* ask what *obligates* you. Inversely, though we can ask what your reason is, we cannot ask what reasons you. This might be taken to suggest that the fundamental obligation relation is not the *is obligated to* relation, but the *obligates to* relation. If so, the thesis that obligations are reasons could be recast, less misleadingly, as the thesis that whenever A obligates B to do C, that is a matter of A being a reason for B to C, plus some further condition being met. Because 'obligation' is just the nominalization of the *is obligated* relation, it is misleading to cast this by saying 'obligations are reasons', but it is nevertheless a view on which obligations reduce to reasons plus some further condition. We still take our other objections to tell against this view.

obligations among them. In contrast, you never have a reason *to me* to eat lunch at Lemonade – that doesn't even make sense. So if obligations are at least sometimes directed and reasons are never directed, that is further evidence that obligations aren't just a special case of reasons.

None of these points show that obligations cannot be analyzed in some way in terms of reasons, or conversely. But we believe they do show that reasons are not the only important normative concept that admits of weight. They are joined by obligations. The question that drives us in the remainder of this paper is whether there is yet a *third* important normative concept that also admits of weight. To this we will ultimately answer a qualified 'yes': it is *commitment*.

2.1 Simple Weighing

Given that reasons are important because they compete with one another, it is natural to assume that if the reasons for X to do A are weightier than the reasons against X doing A, then X ought to do A. An apparently trivial corollary of this principle is that if there are reasons for X to do A, and no reasons against X doing A, then X ought to do A. This attractive principle, which we can call *Simple Weighing*, makes sense of the idea that your reasons support what you ought, other things being equal, to do.

It is also natural to assume that Simple Weighing extends in a straightforward way to explain how competing obligations are weighed, as well.⁹ If X's obligation to do A is stronger than any obligations X has not to do A, then X ought to do A. The same corollary follows, which is that if X has an obligation to do A and no directly competing obligation not to do A, then X ought to do A. A version of Simple Weighing for obligations can be found in the work of W.D. Ross, who argues that we are subject to competing *prima facie* moral duties, such as duties to keep promises and act beneficently, which must be weighed against each other to determine what we ought to do all-things-considered.¹⁰ Failing to satisfy a *prima facie* duty is permissible only when that duty

⁹ Note that our goal is ultimately to call into question whether obligations obey Simple Weighing, and thereby to raise the question as to whether reasons do after all.

¹⁰ 'Another instance of the same distinction [between 'parti-resultant' and 'toti-resultant' attributes such as *prima facie* duties and actual or absolute duty] may be found in the operation of natural laws. *Qua* subject to the force of gravitation towards some other body, each body tends to move in a particular direction with a particular velocity; but its actual movement depends on *all* the forces to which it is subject. It is only by recognizing this distinction that we can preserve the absoluteness of laws of nature, and only by recognizing a corresponding distinction that we can preserve the absoluteness of the general principles of morality' Ross [1930, 28-29].

conflicts with some other, more important duty or duties. If keeping your promise directly conflicts with a more important duty like acting beneficently towards the injured cyclist, you are permitted to break your promise. Otherwise, you ought to keep it. On the assumption that talk of 'duty' is not importantly different from talk of 'obligation', this is just the thesis of Simple Weighing for obligations.

A natural worry to have at this point is that given our arguments that reasons are not a special case of obligations and obligations are not a special case of reasons, Simple Weighing cannot be true both for reasons and for obligations. For given the distinctness of reasons and obligations, there could easily be cases in which one's reasons for A outweigh one's reasons against A, but one's obligations against A outweigh one's obligations in favor. In such a case, Simple Weighing for Reasons and Simple Weighing for Obligations cannot both provide true answers about what one ought to do.

There are several possible ways of preserving the spirit of Simple Weighing in light of such cases, which we will not decide among. For example, it could be that obligations and reasons go into a single weighing calculation. Or it could be that obligations first narrow down the options, and then reasons decide among those that remain. Or it could be that obligations rule some options out by default, which can then be justified by sufficient reasons. What is important to us is that on any of these revised pictures, our corollaries remain: if there is some reason or obligation to do A, and neither any reason nor any obligation against A, then you ought to do A. It is this corollary that we will be calling into question.

The way in which we weigh reasons or obligations can have important philosophical upshots. For example, we believe that John Brunero implicitly appeals to Simple Weighing in an argument against a narrow reading of the logical scope of conditional requirements of rationality. In general, we take agents who have no intention to do what they believe themselves to have conclusive reason to do, or agents who believe what they believe themselves to have insufficient reason to believe, to be irrational. And these thoughts are intuitively encapsulated by the following claims:¹¹

¹¹ Brunero [2010]. We have reformulated these requirements for the sake of a more streamlined discussion; the reformulated requirements are logically equivalent to Brunero's originals, which are due to Kolodny [2005].

(I+): If you believe you have conclusive reason to Φ , then you are rationally required to intend to Φ .

(B-): If you believe there is insufficient evidence that p, then you are rationally required not to believe that p.

According to the narrow-scope theories that are Brunero's target, these claims are true as stated, with 'required' taking scope within the consequent of each conditional. In contrast, wide-scope theorists like Brunero hold that these claims are true only if 'required' is interpreted to take scope over the whole conditional.

We take it to be at least a familiar idea that narrow-scope theories must hold that 'required' – or whichever normative concept figures in the consequents of true narrow-scope conditionals – expresses a weighty or *pro tanto*, rather than an all-things-considered, normative concept.¹² A narrow-scope theory which claims instead that the narrow-scope normative concept is all-things-considered will have to endorse ubiquitous all-things-considered conflicts – cases in which an agent is all-things-considered required to do each of incompatible things. For by (I+), an agent who believes that she has conclusive reasons both to Φ and not to Φ is rationally required both to intend to Φ and to intend not to Φ . But intending inconsistent things is also irrational, and so means violating some other requirement of rationality.¹³ In order to avoid further irrationality, the agent in this case must weigh the conflicting rational requirements to which she is subject. So it makes sense for Brunero to appeal to a principle governing weighing in an argument against narrow-scope theories, for weighing is an important component of such theories.¹⁴

Brunero offers a counterexample to narrow-scope readings of (I+) and (B-). He asks us to suppose that Annie believes that she lacks sufficient evidence to believe that she has conclusive reason to attend a lecture, but in violation of (B-) she believes this anyway. Yet though she believes that she has conclusive reason to attend the lecture, in violation of (I+) she has no intention to go. Brunero argues that under a narrow-scope view, the only rational way for Annie to satisfy both

¹² Compare Broome [2007], Shpall [2013].

¹³ This other requirement could be wide-scope, or it could be narrow-scope.

¹⁴ Note that Lord [2011] offers a different, but, we think, compatible response to Brunero's argument. Lord does not take Brunero's argument to rely on a principle about weighing, but rather to turn on conflating violation of a requirement with escape from that requirement.

(I+) and (B-) is for her to intend to go to the lecture, in accordance with (I+), and to give up her belief that she has conclusive reason to go, in accordance with (B-). And this sounds weird.

Brunero's counterexample relies on Simple Weighing about requirements. The corollary of Simple Weighing states that if X is required to do A and there is no competing requirement for X not to do A, then X ought to do A. (I+) yields a requirement to intend to go to the lecture. (B-) yields a requirement to give up the belief that there is conclusive reason to go. And these requirements do not conflict with each other – it is perfectly possible to fulfill both. So unless there are conflicting requirements that we have not mentioned in describing the case, Simple Weighing entails that Annie ought to satisfy both of them.¹⁵

But surely it would not be best for Annie to *both* intend to go to the lecture *and* to give up her belief that she has conclusive reason to do so. For the latter belief is the rational basis of the former intention. And it would be problematic if one of the narrow-scope requirements to which Annie is subject demanded that she have a certain intention while another demanded that she give up the rational basis of that very intention. Brunero argues that a narrow-scope reading of conditional requirements commits us to this problematic view, and should therefore be rejected.

2.2 Commitments and Sophisticated Weighing to the Rescue

If the requirements mentioned in the narrow-scooper's reading of (I+) and (B-) are understood as obligations that are subject to Simple Weighing, then Brunero's argument looks persuasive. But they need not be understood as obligations (or at least as ordinary obligations), for there is another concept to which the narrow-scooper could appeal. Sam Shpall [2013], [forthcoming] and Mark Schroeder [2013] have each advocated understanding rules of rationality such as (I+) and (B-) in terms of the normative concept of *commitment*.

Shpall [forthcoming] notes that we speak of commitment frequently and in a variety of ways. We often talk of one person being committed to another, or of someone being committed

¹⁵ Note that since only one case is required in order to generate Brunero's objection, the solution cannot turn on the existence of a contrary requirement not to intend to go to the lecture, unless there is some true principle which guarantees the existence of such a requirement, given the stipulations of the case. Though we cannot rule this possibility out, relying on the postulation of such a principle does not seem to us like the most promising strategy for the narrow-scooper to respond to Brunero's challenge. Thanks to Abelard Podgorski for pressing us on this point.

to a cause or ideal. This is commitment in what Shpall calls the *volitional* sense, and is not what concerns us here. But we also talk about commitment in another sense, as when we say that Plato is committed to a wacky ontology, or that Adam is committed to believing that the world was created in six days because he believes that everything in the Bible is true and that the Bible says that the world was created in six days. These are examples of *rational* commitments, cases in which holding one belief or intention commits the agent to holding another belief or intention, on pain of irrationality.

Crucially, rational commitments are always *agent-dependent*, meaning they are the result of some state or activity of the agent. Adam is committed to believing that the world was created in six days *because* he believes that everything the Bible says is true. An important consequence of this is that commitments are also *escapable*. If Adam gives up the ground of his commitment by giving up the belief that everything the Bible says is true, he will no longer be committed to believing that the world was created in six days. Many obligations are not agent-dependent and escapable in this way. For example, a Rossian prima facie duty of beneficence applies categorically to all agents and is not something from which they can escape. And nothing an agent does makes it the case that she is subject to wide-scope rational obligations, or could allow her to escape from these obligations.

We can shed light on Annie's case by looking at it through the lens of commitment. We can understand (I+) as describing a narrow-scope commitment, according to which Annie is committed to intending to go to the lecture on the basis of her belief that she has conclusive reason to do so. Like obligations, commitments can compete and allow of conflicts. If Adam believes that Genesis 1 states that man and woman were created simultaneously and also believes that Genesis 2 states that man was created first, then he is committed to believing both that man and woman were created simultaneously and that man was created first. Since these commitments are incompatible, they compete with one another.

Simple Weighing for commitments states that if X's commitment to believing A is stronger than any commitments X has to not believing A, then X ought to believe A. The corollary holds that if X is committed to believing A and is not committed to not believing A, then X ought to believe A. So if Simple Weighing is true of commitment, Brunero's objection succeeds: since

Annie's commitments generated by (I+) and (B-) do not directly conflict with each other, Annie ought to satisfy each.

But what if commitments are weighed in a different way? If we adopted a more nuanced method of weighing commitments, perhaps we could avoid the problematic result. And in fact, it does seem overly strict to claim that *every* unconflicted commitment yields a conclusion about what an agent ought to do, all-things-considered. After all, what agents are committed to is a matter, at least in part, of their beliefs and intentions. And we would not take any of the agent's beliefs or intentions as a fixed point in determining what she ought to do. So why should we take her commitments as such a fixed point? They should be negotiable along with the beliefs and intentions that are their basis.

Let us call any view that takes this insight on board a version of *Sophisticated Weighing*. When weighing a commitment, views that satisfy Sophisticated Weighing take into consideration whether there are other commitments, obligations, reasons, or other factors that undermine the initial commitment in some way. If there are, that commitment does not go away, but it does not factor in the ordinary way into determining what the agent ought to do. There may be different ways of making the fundamental ideas of Sophisticated Weighing precise, but any such view will share two important features. First, it must give answers about each of the things that an agent ought to do, believe, or intend, in a way that is determined, *inter alia*, by some among the agent's commitments. But second, commitments that the view predicts the agent should give up must not be among the ones that do this determining. So, for example, if a commitment is grounded in an irrational belief, then any version of Sophisticated Weighing should ignore that commitment, because that is a belief that the agent should give up.¹⁶

¹⁶ As we've described it, Sophisticated Weighing is very abstract. Here are two examples of how it might be made more concrete. According to a first picture, which we'll call the *Priority Picture*, commitments are ordered into higher and lower priorities. Simple Weighing is then applied to the highest-priority commitments, and whenever this tells us that the agent ought not to have any of her lower-order priorities, we cross them off the list. Then we apply Simple Weighing to the next-highest-priority commitments, and so on. So long as commitments about what to believe have higher priority than commitments about what to intend, this picture can deal correctly with Annie's case.

According to a second picture, which we'll call the *Equilibrium Picture*, there is no ordering of commitments into higher and lower priorities. Instead, we look at every possible subset of the agent's commitments. Some of these will be *unstable*, in the sense that if we applied Simple Weighing to the commitments in that set, it would tell us that the agent ought not to have all of those commitments. On this picture, what the agent ought to do is a result of applying Simple Weighing to the most inclusive subset of her commitments that is not unstable. Compare, for example, the technique of *scenarios* in Horty [2012].

The way that this works in Annie's case is simple: given that Annie has a commitment due to (B-) not to believe that she has conclusive reason to go to the lecture, and has no conflicting commitments, other things being equal we expect our weighing to tell us that Annie ought not to have this belief. But since this belief is the sole basis of her commitment due to (I+) to intend to go to the lecture, no view that satisfies Sophisticated Weighing can accept this conclusion and also tell us that Annie ought to intend to go to the lecture.

Sophisticated Weighing is a plausible and powerful approach to weighing commitments in general. Does Adam's commitment to believing that the world was created in six days entail that this is what he ought to believe all-things-considered? The answer depends on whether the grounds of Adam's commitment are rational. We must ask: what is Adam's evidence for his belief that everything the Bible says is true? Does this belief conflict with other beliefs he has that are well grounded in his evidence? Merely assessing whether Adam is also committed to believing the directly conflicting claim that the world was not created in six days is overly simple, and may not yield the right result.

If Sophisticated Weighing is correct for commitments, then Annie's commitments are not subject to Brunero's counterexample, which means that there is a coherent way of understanding (I+) and (B-) as narrow-scope requirements after all. This suggests to us that narrow-scopers, at least, should be very friendly to the existence of commitment as a normative concept that admits of weight, and that does not obey the Simple Weighing principle that reasons and obligations seem to obey.

3.1 Commitments in General

Commitments can help the narrow-scopers avoid Brunero's objection, but commitments are not just for narrow-scopers. The concept of a commitment is natural and highly intuitive. If Adam believes that the Bible says that the world was created in six days and that everything the Bible says is true, there indeed seems to be some kind of strict normative pressure for him to believe that the world was created in six days. But it seems problematic to claim that this pressure tells

Both of these pictures would require some refinement in order to ultimately be plausible, but our hope is that they provide some sense for the space of possible views that satisfy Sophisticated Weighing.

us what he rationally ought to do; surely in light of all of the evidence, Adam ought to abandon one of his grounding beliefs instead. Nor do we want to claim that Adam has any strong reason to believe that the world was created in six days, for Adam can have this commitment even in cases in which his belief is irrational. If the mere fact that he held those beliefs led to his having a strong reason to believe what follows from them, we would have normative bootstrapping of a very troubling sort, but there is no problem with recognizing that agents sometimes are committed to things they have no corresponding strong reason to do. This shows that commitments cannot just be reasons – at least, commitments can be forceful even when reasons are weak.

In fact, as philosophers we are most often interested in each other's commitments precisely because we recognize that someone may be committed to something that we believe she has no strong reason to believe. Indeed, we often try to exploit this fact by arguing that someone should give up her beliefs by showing that they commit her to things that we all agree that she ought not to believe. When we think as philosophers about our commitments and about what conclusions we should draw, we would make a mistake if we always held our commitments fixed, as is suggested by Simple Weighing; instead, we should leave open that we may give up some of our commitments, and on the condition that we do, we should no longer take them into account.

3.2 Are Commitments Epiphenomena?

We have argued that commitment is an important and general normative concept that admits of conflicts, that there may be forceful commitments even in cases in which there are no or at least no strong reasons, and that Simple Weighing is implausibly simple as an account of how our commitments contribute to what we ought to do or believe. And we have suggested that theorists with a particular set of commitments – narrow-scopers – ought to hold that some version of Sophisticated Weighing is true for commitments. But this does not yet show that Sophisticated Weighing *is* true for commitments, or even that commitments are a genuinely weighty normative concept, for we have not argued that narrow-scoping is true. One important objection that we must confront claims that commitments need not weigh against one another because they do not

make any contribution to what an agent ought to do. Instead, this is determined only by reasons and/or obligations. Call this the objection that commitments are *normative epiphenomena*.

Here is how the proponent of the normative epiphenomena objection is thinking about things in Annie's case: stripping away all talk of commitment, we could say that Annie is subject to the wide-scope obligation to either believe that she has sufficient evidence to believe that p , or to cease believing that p . However, her evidence is insufficient to support p , which means that Annie has reason to satisfy this obligation in the latter way. Annie is also subject to the wide-scope obligation to either cease believing that she has conclusive reason to Φ or to intend to Φ . And her evidence entails that she has reason to satisfy this obligation in the former way. We have explained what Annie should do, but have not needed to appeal to commitments. If this is right, then perhaps commitments do not play a unique or ineliminable role in the normative landscape after all.

Similarly, the proponent of the normative epiphenomena objection holds that these remarks generalize. Just as it is possible to describe what Annie ought to do without appealing to commitments, it is possible to describe how you ought to revise your philosophical views without appealing to your commitments. Whatever help we get from thinking about your commitments is really just a heuristic, on this view, and all that matters, at bottom, is the fact that you are required to satisfy some sort of closure norm.

It is not surprising that it should be possible to describe what agents ought to do and to believe without mentioning 'commitments', because philosophers have offered so many kinds of theories about what agents ought to do and to believe, and very few of them have ever discussed commitment as a normative concept worthy of study in its own right. Yet we believe that the fact that commitments *can* be ignored does not show that they *should* be.

To take a comparison that we find instructive, the United States Senate operates by a set of procedural rules for its own conduct. These rules are set by the Senate itself, and so they can be changed by the Senate at any time. They are, essentially, the Senate's commitments about how to proceed. It would be possible to argue that these commitments are really an epiphenomenon, because what is really doing all of the work in determining how the Senate ought to proceed is simply the requirement that they either follow their procedural rules or change

them. But we think that this approach would be misguided; it is obvious that the Senate's procedural rules do play an important explanatory role in determining how the Senate ought to proceed. Similarly, we think, for the rational commitments of ordinary agents.

If anything, the idea that there is a requirement on the Senate to either follow their procedural rules or change them looks to us like a normative epiphenomenon. The only sense in which this is clearly true is that if the Senate neither follows their rules nor changes them, then they are breaking the rules, which is clearly a bad thing. But if the claim that the Senate is required to either follow their rules or change them means anything other than this, it is actually quite mysterious where this requirement comes from or why the Senate has to follow it.

That commitments do vary in weight can be illustrated by a consideration of some examples. David Lewis, for example, was committed to many things – to the thesis of Humean Supervenience, to the existence of a plurality of possible worlds, to counterpart theory, to an analysis of counterfactuals, to a particular analysis of causation, to a particular account of what contributes to the proximity relation among possible worlds, and to much else. But some of these commitments are weightier than others, and it is no surprise to see Lewis or those interested in Lewisian metaphysics navigating carefully around those stronger commitments when evaluating which of the lesser commitments to give up.

3.2 Sophistication Without Weighing

Now we have argued that rational commitments do have weights, and that it is a mistake to think that commitments are a normative epiphenomenon. Yet it is still possible to worry that though commitments make a difference for what we ought to do, and though they do in some natural sense have 'weights', they do not make a difference *by* weighing against one another.

This objection can be given force by considering the case of Adam. Adam, recall, is committed to believing that man and woman were created simultaneously, and also committed to believing that man was created first. He is committed to the first because it is entailed by his belief that everything the Bible says is true and his belief about Genesis 1, and he is committed to the second because it is entailed by his belief that everything the Bible says is true and his belief about Genesis 2. We might think that these two commitments could weigh against one another,

because it can't be the case that man was created both before and simultaneously with woman, but for the same reason that Adam shouldn't have both of these beliefs, he also should not have all of the grounding beliefs which commit him to them. And so the question of which, if either, grounding belief to hold seems to line up directly with the question of which commitment to give up.

It is natural to conjecture based on this case that in general, once we remove commitments that an agent ought not to have, there will no longer be conflicts among her remaining commitments, and hence there will be no substantive weighing left to do.¹⁷ If this conjecture is correct, then though commitments require sophistication, they do not require weighing after all, in order to determine what you ought to do.

We believe that this conjecture is plausible for the case of rational commitments. We doubt that it is correct for the case of commitments more generally.¹⁸ But either way, it is enough for what we want to show that commitments at least require sophistication, even if they do not require weighing all by themselves. For as we will now argue, commitments are just a special case of obligations. So if commitments require sophistication, and obligations more generally require weighing, then obligations require Sophisticated Weighing.

4.1 Are Commitments Just a Special Case of Obligations?

So far, we've argued that commitments are an important normative concept, that they really do have genuine normative weights, and that those weights contribute in a different way to determining what we ought to do than reasons do. Commitments, it seems, have weights but are not reasons. But this doesn't yet answer our question as to whether they are really a third distinct important normative concept that admits of weight. To complete the answer to this question we must decide whether commitments are just a special case of obligations. We will now argue that they are.

¹⁷ Thanks to Abelard Podgorski for pressing a version of this point.

¹⁸ Compare Shpall [forthcoming], who helpfully contrasts rational commitments with what he calls *moral* commitments. A paradigm case of moral commitments are those arising from promises. But it is not always the case that you can get out of one of two conflicting promises, or even, we think, that you always should.

Recall that we distinguished obligations from reasons on the grounds that obligations but not reasons are strict, that obligations and reasons have different ontology, and that obligations may be directed but reasons are not. Commitments pattern with obligations in being strict – you are subject to criticism when you fail to do what you are committed to doing, even when what you actually do is all-things-considered better.¹⁹ They also pattern with obligations on ontology. We may ask about why you are committed to believing that the world was created in six days, or about what the *basis* is for this commitment, but it doesn't make sense to ask about what the commitment is – it's just the commitment to believe that the world was created in six days.

The question as to whether commitments may be directed is slightly more complicated, because none of our chief examples of commitments appears to be directed. 'Commitment' talk also appears, curiously, to be divided over whether commitments can be directed, depending on whether we use 'committed' as an adjective', or talk about 'having a commitment', in its nominalized form. Though talk about being committed to Lucy to meeting her for lunch is questionably grammatical, it makes perfect sense to say that you have a commitment to Lucy, to meet her for lunch. We tentatively conclude that this is some evidence that commitments, like obligations, may be directed.

You might worry about this conclusion that the rational commitment in which we've been interested in this paper is really quite different from the use of 'commitment' in talk about having a commitment to Lucy, to meet her for lunch. Perhaps, we might think, these two uses of 'commitment' have little in common. Fortunately, Shpall has taken up precisely this question in his insightful paper 'Moral and Rational Commitment', which argues that in fact the very same notion of commitment is at stake in each of these two cases. So delegating this question to Shpall, we conclude that *if* his arguments are correct, then commitments pattern with obligations on all three dimensions that we used to distinguish them from reasons.

If it is really true that commitments pattern with obligations on all three scores that we used to distinguish obligations from reasons, then that is evidence, we believe, that commitments are just a special case of obligations.²⁰ And intuitively, this is highly plausible – a commitment is

¹⁹ See Shpall [2013], [forthcoming] for more detailed discussion and arguments for the strictness of commitments.

²⁰ Obligations, of course, cannot all be a special case of commitments, because not all obligations are agent-dependent.

just an obligation that you impose on yourself. That is why promissory obligations are such a natural example of commitments. If this is right, then it turns out that we haven't located a third distinct weighty normative concept after all, but only spent our time discussing a particularly interesting special case of obligations.

4.2 Back to Reasons

If this is right, then it turns out that we can learn something important about obligations through our study of their special case, commitments. Recall that we began our investigation of commitment with the observation that commitments do *not* combine in the Simple Weighing way that we *assumed* that reasons and obligations do, in order to determine what you ought to do. The reason that they cannot is that we cannot hold your commitments as fixed points in determining what you ought to do; instead, given that your commitments are agent-dependent and to at least some extent up to you, we need to allow that your commitments themselves are part of what is negotiable in considering the holistic question of what you ought to do. The Simple Weighing claim that if you have an obligation to do something and no conflicting obligation, then you ought to do it therefore cannot be extended to the case of commitments. So if Simple Weighing is correct about obligations, then commitments cannot be a special case of obligations, after all, despite their striking similarities and despite the intuitive appeal of the idea that commitments are just self-imposed obligations.

This suggests to us that Simple Weighing is wrong about obligations. This is not, however, to say that there is not something true about it. Recall that the reason it fails to hold for commitments is that it builds in the holding fixed of what your commitments are. Because your commitments are escapable and are therefore up for negotiation, this is a mistake. But in general, obligations that are not commitments are not up for negotiation when you are deciding what to do. They are not, after all, under your control. If some obligations *are* under your control, we would *expect* them to be up for negotiation. Consequently, we can see Simple Weighing as a plausible but false generalization from a limited diet of examples. Commitments *are* just self-imposed obligations, we infer, but we were wrong all along about how obligations combine to determine what we ought to do – at least in the general case.

This conclusion leaves us with one final thought that we won't be able to pursue, here: if we were wrong all along about how obligations are weighed and combine, could the same conclusions go for reasons? In this paper, we've glossed over distinctions that are sometimes made between *objective* and *subjective* reasons. Objective reasons, according to conventional ways of understanding this distinction, contribute to what it is *correct* to do, or it is sometimes said, what you ought in the objective sense to do. Whereas subjective reasons, it is said, contribute to what it is *rational* for you to do, sometimes equated with what you ought in the *subjective* sense to do.

An agent's objective reasons are generally taken to be given by the facts, and her subjective reasons to be given by her beliefs, or by some subset of her beliefs. The facts sound like they may be the right kind of thing not to be up for negotiation in the determination of what is correct for an agent. But beliefs are precisely the sort of thing that is up for negotiation in the determination of what it is rational for an agent. Given what we've learned about commitments so far, this makes us suspect that Sophistication may be called for.²¹

²¹ Special thanks to Josh Crabill, Errol Lord, Matt Lutz, Barry Maguire, Abelard Podgorski, Sam Shpall, Ryan Walsh, Aness Webster, and an anonymous referee for OUP.

References

- Adler, Jonathan [2002]. *Belief's Own Ethics*. Cambridge, MA: MIT Press.
- Brink, David [1994]. 'Moral Conflict and its Structure.' *The Philosophical Review* 103: 215-247.
- Broome, John [1999]. 'Normative Requirements.' *Ratio* 12(4) 398-419.
- _____ [2007]. 'Requirements.' In *Hommage à Wlodek. Philosophical Papers Dedicated to Wlodek Rabinowicz*. Ed. T. Rønnow-Rasmussen, B. Petersson, J. Josefsson & D. Egonsson. www.fil.lu.se/hommageawlodek.
- Brunero, John [2010]. 'The Scope of Rational Requirements.' *Philosophical Quarterly* 60(238): 28-49.
- Conee, Earl and Feldman, Richard [2004]. *Evidentialism: Essays in Epistemology*. Oxford University Press: New York.
- Dancy, Jonathan [2004]. 'Enticing Reasons.' *Reasons and Value: Themes from the Moral Philosophy of Joseph Raz*, ed. Jay Wallace, Philip Pettit, Samuel Scheffler, and Michael Smith. Oxford University Press: New York: 91-118.
- Gert, Joshua [2005]. 'A Functional Role Analysis of Reasons.' *Philosophical Studies* 124(3): 353-378.
- Greenspan, Patricia [2005]. 'Asymmetrical Practical Reasons.' *Experience and Analysis: Proceedings of the 27th International Wittgenstein Symposium*, ed. M. E. Reicher and J. C. Marek. öbv & Hpt: Vienna: 387-94.
- Gregory, Alex [forthcoming]. 'A Very Good Reason to Reject the Buck-Passing Account of Value.' Forthcoming in *Australasian Journal of Philosophy*.
- Horty, John F. [2003]. 'Reasoning with Moral Conflicts.' *Noûs* 37(4): 557-605.
- _____ [2012]. *Reasons as Defaults*. Oxford: Oxford University Press.
- Kolodny, Niko [2005]. 'Why Be Rational?' *Mind* 114(455): 509-563.
- Lord, Errol [2011]. 'Violating Requirements, Exiting from Requirements, and the Scope of Rationality.' *Philosophical Quarterly* 61(3): 392-399.
- Marcus, Ruth Barcan [1980]. 'Moral Dilemmas and Consistency.' *Journal of Philosophy* 77: 121-136.
- Nair, Shyam [unpublished]. 'Conflicting Reasons, Unconflicting Obligations.' Unpublished manuscript.

Parfit, Derek [1997]. 'Equality and Priority.' *Ratio* 10(3): 202-221.

Ross, W.D. [1930]. *The Right and the Good*. Oxford: Oxford University Press.

Schroeder, Mark [2013]. 'Scope for Rational Autonomy.' *Philosophical Issues* 23 (Epistemic Agency): 297-310.

Shpall, Sam [2013]. 'Wide and Narrow Scope.' *Philosophical Studies* 163(3): 717-736.

_____ [forthcoming]. 'Moral and Rational Commitments.' Forthcoming in *Philosophy and Phenomenological Research*.

van Fraassen, Bas [1973]. 'Values and the Heart's Command.' *The Journal of Philosophy* 70: 5-19.