



Normative metaphysics for accountants

Barry Maguire¹ · Justin Snedegar²

© The Author(s) 2020

Abstract We use normative reasons in a bewildering variety of different ways. And yet, as many recent theorists have shown, one can discern systematic distinctions underlying this complexity. This paper is a contribution to this project of constructive normative metaphysics. We aim to bring a black sheep back into the flock: the balancing model of weighing reasons. This model is threatened by a variety of cases in which distinct reasons overlap, in the sense that they do not contribute separate weight for or against an option. Our response is to distinguish between derivative reasons and load-bearing reasons, only the latter of which contribute non-overlapping weight to an option. This distinction is close at hand for analyses of reasons in terms of the promotion of significant outcomes. But we also develop an account of this distinction for fundamentalist theories of normative reasons.

Keywords Normative reasons · Weighing reasons · Reasons fundamentalism · Value

Can't escape the scale if I tried.
-Pusha T

The order is alphabetical; the authors contributed equally.

✉ Justin Snedegar
jlsnedegar@gmail.com
Barry Maguire
barry.maguire@stanford.edu

¹ Philosophy Department, Stanford University, Stanford, CA 94305, USA

² Department of Philosophy, University of St Andrews, St Andrews, Fife KY16 9AR, UK

1 Introduction

Since at least 1930, theorists working in ethical theory have explicitly developed constructive accounts of the competing normative considerations that bear on what we ought to do.¹ For each option available to an agent in a choice situation,² there are usually various different *reasons* in its favour and reasons against. These reasons interact to determine some fact about what the agent has most reason to do, all things considered. On many views, reasons for options interact by contributing normative weight for or against the relevant options.³ This interaction is often introduced with the metaphor of a balance or scale. Picture one scale for each option, with reasons as marbles, the reasons for going in one pan, and the reasons against going in the other. The net weight of reason for each option is represented by whether the scale tips towards the reasons for, or the reasons against, and how far. You have most reason to take the option whose scale tips the most towards the reasons for. This is the *balancing model* (for discussion, see Baier 1958; Broome 2004; Berker 2007; Lord and Maguire 2016).

The balancing model has several attractive features that we want from a model of how reasons interact. The reasons for different options contribute weight for each option.⁴ They compete and combine with others bearing on the same option. Reasons for the same option can combine to support the option more strongly than either does individually. Reasons bearing one option also interact, slightly less directly, with those bearing on other options. And various deontic properties can be represented by the various comparisons made available; for instance, it is common to maintain that one ought to take whichever option is supported by most reason.⁵ This model fits with a representation of significant chunks of ordinary practical thought in choice situations. We are responsive to lists of pros and cons for different options or plans, often implicitly. The balancing model also underlies influential ways of thinking in normative ethics. Kagan (1988) has argued, for instance, that arguments based on normative differences between similar cases generally presuppose some version of this model. Kelman (2016) has argued that this model also underlies standard cost benefit analysis, and, relatedly, an important strand of argument in public policy and developmental economics.

But the balancing model also faces some familiar challenges. Firstly, since Jonathan Dancy's pioneering work, it has been clear that reasons are subject to both

¹ This is to trace this line of thought back to Ross's (1930) *The Right and the Good*. In this essay, we focus on reasons for options, paradigmatically actions or associated items such as choices, plans, or activities. We set aside questions about the support for beliefs or other attitudes.

² Our topic is restricted to the bearing of contributory notions in the determination of deontic properties of options in pertinent choice situations. We set aside broader ethical questions about whether this ideology is the most ethically perspicuous, and what explains the pertinence of a given choice situation.

³ For an overview of this familiar way of thinking about the way in which reasons support conclusions about what we ought to do, see Lord and Maguire (2016). For some notable exceptions, see Greenspan (2005), Gert (2007) and Horty (2012).

⁴ There are disputes about the precise details of the relationship between being a reason and contributing weight (cf. Fogal 2016; Broome 2013; Schroeder 2007).

⁵ Cf. Lord and Maguire (2016). The weighing behavior of epistemic reasons may be different. We stick with reasons for actions in this essay.

conditionality and modification.⁶ That is, considerations may only get to count as reasons given that certain other facts—not necessarily themselves reasons—obtain, and the weight of these reasons can be modified by the presence or absence of certain other facts—again, not necessarily themselves reasons. These structures may enable the theorist to represent the normative role of such phenomena as rights, commitments, and obligations. However, these complexities needn't trouble the balancing model, so long as the existence and weights of marbles represent the existence and weights of duly conditioned and modified reasons. The balancing model need not model all of the facts about how reasons and contributory considerations more generally interact; adequately modeling the competition between reasons, once they've been conditioned and modified, is still significant.

Another challenge maintains that, unlike the weights of marbles, the weights of reasons may not combine in a straightforwardly *additive* way.⁷ To take Prakken (2005)'s and Horty's (2012: 61) example, that it's hot outside and that it's raining outside are each reasons not to go for a run, but together they support not going for a run *less* strongly than either does individually, since the rain lessens the unpleasantness of the heat, and vice versa. Or the old joke about the Catskill mountain resort, 'the food is terrible and the portions are too small.'⁸

We are interested in a challenge that is related to but in a way deeper than the additivity worries. The challenge is that distinct reasons can *overlap*, in the sense that they do not contribute separate weight to an option. This problem appears in passing in many different discussions of normative reasons, but it has not yet received any sustained discussion.⁹ Here is an example to give you a feel for the problem (Hawthorne and Magidor 2018):

...if I like red items, then the fact that an item is bright red might be a reason (of a certain strength) to buy it, and the fact that it's red might also be a reason (of a certain strength) to buy it. But the strengths of the two reasons do not add up (I don't have twice as much reason to buy it because it's both red and bright red) ... [the] weighing analogy is simply the wrong way to think about pro tanto reasons.

Such cases put pressure on the balancing model. For when we're weighing up marbles, each marble makes a single, distinct contribution to the total weight in the pan. The corresponding assumption about reasons is the *separability* thesis, that each reason contributes weight for or against an option independently of other reasons.¹⁰ But in cases of overlapping reasons, this assumption would appear to be

⁶ See in particular Dancy (2004, Chapter 3). Bader (2016) establishes the consistency of conditionality and modification with separability and additivity (which we'll discuss momentarily).

⁷ For a sophisticated discussion of additivity, see Berker (2007). He introduces a weaker notion of quasi-additivity, though not in much detail. We do not presuppose more than this weaker notion in this paper. However, we'll stick with talk of additivity for convenience.

⁸ This joke is cited in Dancy (2004). As far as we can tell, it is originally from *Annie Hall*.

⁹ In addition to Hawthorne and Magidor (2018), see Dreier (2007), Schroeder (2009), Fogal (2016), Nair (2016b), Sher (2019), Maguire (2016), King (2019), Chang (2009).

¹⁰ We set aside complications concerning conditionality and modification henceforth.

false. These cases seem to involve reasons that do not make single, distinct contributions of weight. Instead sets of distinct reasons all make *overlapping* contributions, or as we'll say more simply, they *overlap*.

The standard presumption is that the balancing model must be rejected. But we conjecture that more progress can be made by holding fixed the balancing model and instead using these cases as an opportunity to refine our understanding of how various different reasons might serve to make the same normative contribution, and specifically, of what that contribution is.

We propose that the phenomena of overlap draw our attention to an important distinction between *load-bearing* reasons and *derivative* reasons. The load-bearing reasons are the contributors of normative weight. The derivative reasons are only reasons in virtue of their relationships to the load-bearing ones.¹¹ Since these reasons are derivative on the load-bearing ones, and the load-bearing ones are the primary contributors of normative weight, the derivative reasons overlap with the load-bearing ones, and with other derivative reasons that are derivative on the same load-bearing reasons. We both motivate this picture in general, and explore how two important, competing metaphysical accounts of reasons should understand the distinction between load-bearing and derivative reasons in order to successfully explain overlap. These two kinds of accounts are *Analyses* of reasons in terms of the promotion of significant objectives (e.g. values or desires), and *Reasons Fundamentalism*, according to which there are facts about reasons that do not obtain in virtue of any further normative facts. Our aim is to show that each kind of theorist should be optimistic about explaining overlap.

At this point, a more general word about our methodology might be helpful. Many of our uses of normative reasons are connected in significant ways with our uses of non-normative reasons, for instance explaining why the universal constants are just so, or why there are no green zebras. The details of the relationship between normative reasons and explanatory reasons is contentious. But it would be unsurprising if some of the features of explanatory reasons practices—involving the use of metonyms and highly specific conversational pragmatics¹²—are also features of our normative reasons practices.

We conjecture that this messy façade conceals systematic structure. In particular, we maintain that ordinary judgements about reasons talk are rather sophisticated. Take your average person jotting down pros and cons. They will clearly avoid writing down both 'bright red' and 'red' as separate reasons in favour of buying the toy letterbox, unless perhaps they also liked bright things (in which case they would probably just write 'bright' and 'red' in their pro column). Neither would they token both in their reasoning, or offer both as advice to another person who likes red things. These are problems for theorists, not people.

Our discussion is based on optimism both about the availability of underlying structure and the probity of our considered judgements about it. But we also wish to

¹¹ For related distinctions in earlier work, see Nair (2016a, b) and Maguire (*op. cit.*).

¹² On which, see, especially, Fogal (2016, 86–87) and Broome (2013).

draw interesting distinctions where we see them. For this reason, we restrict our attention to practical reasons. We set aside questions about the normative support for beliefs and affective attitudes which might operate differently.¹³ We will also ignore the impact of commitments and constraints and priority relations of various kinds that might play some explanatory role prior to the interaction of reasons in our practical thinking.

It may help to point out an issue where the distinction between derivative and load-bearing reasons is important. It is plausible that, in an important range of cases, an agent's motivating reason is creditworthy only if it corresponds to a normative reason (cf. Arpaly 2002; Way 2017a, b). But one and the same reason may correspond to very different putative contributors. Moreover, this very indeterminacy may be exploited by those hoping to conceal less attractive motives. Consider Kant's shopkeeper motivated to give back \$3 by the fact that it is the correct change. A sharper account of load-bearing reasons will yield a clear theoretical standard against which to assess creditworthiness.¹⁴

2 The set-up

We follow the convention of assuming that reasons are facts; and that facts are true propositions.¹⁵ This fits well with the idea that practical reasoning often involves reasons, since most views of reasoning understand it to involve attitudes towards propositions. We help ourselves to a fine-grained account of facts, according to which [Pearl is happy] and [Pearl is very happy] are different facts.¹⁶

We distinguish the fact that is the reason, which will often be non-normative, from the fact that that fact is a reason, which is paradigmatically normative,¹⁷ from the fact that that fact is a reason for whatever it is a reason for, which is also paradigmatically normative. We'll call the first one the 'reason', the third one the 'reasons fact', and assume that the second one is really just a short-hand way of talking about the third one.

¹³ Of course, epistemic reasons and practical reasons themselves interact in interesting ways. A tells you to see the movie, and B also tells you to. That's may be more reason to see the movie. But it may be each piece of advice is based on the same review. We set such complications aside.

¹⁴ See Howard (ms) for a discussion of moral worth, or creditworthiness, in relation to issues similar to those discussed in this paper.

¹⁵ Though see Howard (ms), who argues on the basis of cases like the some of the one's we'll consider that this widespread assumption is false. We are convinced that analogous issues will arise for alternative ontologies.

¹⁶ Might one avoid some overlap worries by adopting a more coarse-grained individuation of facts/propositions? No. First, we suspect that it will be difficult to find a level of granularity for facts that cuts things at just the right level. Second, by adopting a coarse-grained individuation of facts, we introduce a different kind of overlap phenomenon to be explained, where one fact—so one reason—makes multiple distinct contributions. More on this later.

¹⁷ Even if not authoritatively normative; see Maguire and Woods (2020).

A (practical) reason is a reason for an option. We will assume that overlapping reasons support the very same option.¹⁸ In many cases, a reason is a fact about that very option, whether about its intrinsic features, or about what it promotes, or the capacities it would exercise, or something else. In other words, the syntactic form of such a reason to ϕ is that ϕ is F. This rather trivial observation makes it easier to see how we can set aside some mere loose talk about reasons. For clearly, we can successfully communicate, to ourselves or others, that ϕ is F, in many other ways that by expressing that ϕ is F. Suppose we both know that going to the party will be great fun only if Emily isn't going. (Simplify and assume that would be the only reason to go). Assume we also both know that you think Emily will be going. I know she isn't. So, I can offer you, as a reason to go, that Emily isn't going to be there. And you can reason as follows: Emily isn't going to be there, so I'll go. But we both know this isn't *really* a reason to go. We might well know it is really the absence of a disabler. And if pushed for our real reason, we'd say that the party was going to be great fun.

It will be tempting to respond to some of our cases by insisting that merely overlapping facts aren't *really* reasons. But there is a difference between saying, of the Emily fact, that it isn't really a reason, and saying of one reason, in the presence of another, that it isn't an *additional* reason. This latter impulse will often really express a sensitivity to our distinction. To support the idea that the merely overlapping reasons are really reasons, we rely on the usual 'earmarks' of reasonhood in what follows (cf. Schroeder 2007, Kearns 2016): the fact counts in favour of the option; plausibly plays the 'reasons-role' in explaining a strict or verdictive normative fact, e.g. about what the agent ought to do; plausibly plays the 'reasons role' in reasoning, advice, and justificatory contexts.¹⁹

We wish to be neutral concerning the specific relationships between the metaphysics of reasons, reason, and weight. One contentious issue concerns whether there can be reason to ϕ (with 'reason' used as a mass noun) without there being a reason (count noun) to ϕ . We restrict our attention to cases involving reasons (count), so this issue doesn't bear on our discussions. The balancing model does entail that the overall weight (mass) of reason for an option is a function of the individual weights of the reasons for and against the options. A separate issue concerns the relations between reasons and weight. We take it to be analytic that a reason has weight, even if not *separate* weight.²⁰ A final issue concerns the metaphor of 'weight'. Perhaps the metaphor of 'force' would better suit other behaviours of reasons, for instance their susceptibility to modification and

¹⁸ The 'giving the dog a treat tonight or giving the dog a treat this morning' case in Maguire (2015, 246) falls afoul of this restriction.

¹⁹ On reasons counting in favor, see Scanlon (1998), Dancy (2004), and Parfit (2011). On reasons as explanations of oughts, see Broome (2004). On reasons as evidence of oughts, see Kearns and Star (2009). On reasons as premises of good reasoning, see Setiya (2014) and Way (2017a, b). On reasons as appropriate things to offer as advice, see Manne (2014). The thought that reasons can be offered as justification is ubiquitous.

²⁰ Even if zero weight. On the difference between having zero weight and no weight, see Bader (2016).

conditionality; it would perhaps also generalize better to epistemic reasons.²¹ But we will stick with the language of ‘weight’ since it is (currently) more familiar.

3 Three kinds of overlap

We now turn to three different kinds of examples of overlap, and show how the distinction between load-bearing and derivative reasons can help explain why the reasons overlap. In general, we suggest that the explanation for why the reasons overlap will involve the fact that there’s really just one normatively significant thing at stake, and all the reasons are ways of drawing our attention to this one thing.²² This single normatively significant thing at stake is what is captured in the load-bearing reason.

3.1 Differences of grain

Consider the following facts about Pearl, a short-haired German pointer:

1. Giving the treat to Pearl will make Pearl very happy.
2. Giving the treat to Pearl will make Pearl happy.
3. Giving the treat to Pearl will make a sentient creature happy.
4. Giving the treat to Pearl will make a sentient creature happy on a Tuesday.

These are distinct facts. One concerns a degree of happiness, one mere happiness, others merely a sentient creature.²³ All of these facts are reasons to give the treat to Pearl. Any of them could felicitously be offered as advice in support of that option, or justification after the event; all are (defeasible) evidence that you ought to take the option; any might play a role in explaining²⁴ whether you ought to take the option, etc.

All of them support the same option, viz. giving the treat to Pearl. But they do not provide separate weight to this option. Suppose you had a slightly greater reason not to give Pearl the treat—perhaps that it is bad for her cholesterol. Adding another of these overlapping reasons would not increase the total weight of reason in favour of giving the treat to Pearl. Clearly, adding these reasons could not shift the balance of reasons in favour of giving Pearl the treat.²⁵

Our hypothesis is that in this case there is a single load-bearing reason, and that the others are derivative on this reason. Different normative theories will disagree about which facts are the primary sources of normative significance, but one

²¹ Many thanks here to Daniel Fogal.

²² Putting things this way shows why we agree with much of what Fogal (2016) says; we briefly discuss his view directly in Sect. 3.4.

²³ Cf. Sher (*op. cit.*).

²⁴ We default to the permissive notion of explanation employed in Broome (2004).

²⁵ Fogal (2016) makes a related point, that we could not typically felicitously offer more than one of these considerations as a reason (in giving advice, justifying ourselves, etc.) in a given context. The point here is rather the metaphysical one that these considerations do not contribute separate weight.

plausible candidate that it is that giving Pearl the treat will make Pearl happy. The other facts are then derivative reasons, derivative on this one: making Pearl *very* happy is a way of making Pearl happy, giving her the treat makes a sentient creature happy *in virtue of* making Pearl happy, and so on. Since these other facts are only reasons in virtue of their relationship to the load-bearing reason, which is the primary contributor of normative weight, they do not contribute independent weight, and so they overlap.

Which facts are the load-bearing reasons will be determined by the details of one's normative theory. We selected 2 for illustration, that giving the treat to Pearl would make her happy. A simple utilitarian theory that does not put much weight on the separateness of canines may choose 3, instead, that giving the treat to Pearl would make a sentient creature happy.

3.2 Alternative satisfiers

Suppose that Bill needs to be vaccinated against a disease. There are two equally effective vaccinations, V1 and V2. You have the opportunity to press a button which will dispense both V1 and V2 to Bill. We'll stipulate that there is no downside or upside to having both vaccines, rather than just one. Each vaccine is guaranteed to be effective and the button delivery mechanism always works flawlessly. Consider the following facts:

5. Pressing the button gives Bill V1.
6. Pressing the button gives Bill V2.
7. Pressing the button vaccinates Bill.

Each of these facts seems to be a reason to press the button, but they also overlap—learning 6, for example, does not make a stronger case for pressing the button than 5 alone, and similarly for 7. Given that these are distinct facts, and so distinct reasons to press the button, we need some explanation for why they do not make distinct contributions.

Once again, it is plausible that there is just one load-bearing reason here. Fact 7, that pressing the button vaccinates Bill, is a plausible candidate. It is only because giving Bill V1 vaccinates Bill that 5 is a reason to press the button, and only because giving Bill V2 vaccinates Bill that 6 is a reason to press the button. So these reasons are merely derivative, and so they overlap.

3.3 Alternative grounds

Our final kind of overlap case involve what we call *alternative grounds* of reasons. These cases illustrate an interestingly different kind of overlap. Consider:

8. Pressing button A will give the dog a treat.
9. Pressing button A will satisfy your promise to Jones to please a dog today.
10. Pressing button A will benefit the dog.

8 overlaps with 9, and 8 overlaps with 10.²⁶ Clearly, there will be many sets of overlapping reasons like these. Consider another set:

11. Going to the shop will let you buy doughnuts.
12. Going to the shop will let you have doughnuts to eat.
13. Going to the shop will let you give doughnuts to your mum.

Each of these facts is a reason to go to the shop. But reason 11 overlaps both with 12 and with 13.

Interestingly, cases like this raise a different kind of overlap issue. The cases we have been considering involve multiple distinct reasons making non-distinct contributions. These cases also display this kind of overlap: pairs of reasons, in this case, {8, 9}, {8, 10}, {11, 12}, {11, 13}, make a single contribution—each pair involves two reasons that make overlapping contributions. But the cases here also involve a *single* fact, and so a single reason (8 in the first case, and 11 in the second), making *multiple* distinct contributions. Each of these reasons makes *two* contributions, and should, in a way, be counted twice—once for each of these pairs in which it appears.²⁷

This kind of overlap becomes more pressing when we note that the same reason can contribute weight both for and against an option. Consider the fact that the exhibition is about [the artist you find most irresistible] in a context in which you will have at most only a short time to visit the museum. That fact is clearly a reason to go, since you love the topic, but also a reason not to, either because you should postpone until a more convenient time, or perhaps just because it will be agonizing to have to tear yourself away.

What needs explained in cases like this isn't (just) that multiple distinct reasons make a single contribution, but rather that a single reason makes multiple contributions. Ideally, our explanation of the previous kinds of cases—which involve the distinct reasons but single contributions—will straightforwardly generalize to explain these kinds of cases, in which single reasons make distinct contributions. A load-bearing reason should not only be something that makes a distinct contribution, but also something that makes a *single* contribution.

An initially tempting response to cases like this is to appeal just to the weight of reasons. The thought is that reasons like 8 and 11 don't make multiple distinct contributions, but instead that their contribution is just greater than it would otherwise be (or less than it would otherwise be, in cases in which one fact seems to be both a reason for and a reason against, as in the museum case above). So, for example, the single reason 11 makes a single contribution in favour of going to the shop, but it's just weightier than it would be in a case in which your mother didn't want a doughnut. If this is right, then we don't need to explain how one reason can

²⁶ Notice that we are not appealing to the putative reason 'that you have a reason to keep your promises' (cf. Sher forthcoming). We prefer to distinguish between reasons and reasons principles; on this more soon.

²⁷ Note that we aren't claiming that the only *real* reasons here are the pairs like {8, 9}. These are pairs of overlapping reasons, so both members of the pair are reasons in their own right.

make multiple distinct contributions. We just have the familiar problem of explaining how the weight of reasons is determined.

The central problem with this strategy is that it seems to obscure a normatively relevant distinction. The two contributions are importantly distinct. As a mark of this, note that different kinds of considerations could defeat them. If, for example, your mother calls and tells you that she doesn't want a doughnut after all, that defeats one contribution but not the other. On the view in question, this would have to be captured by saying merely that the weight of the reason to go to the shop becomes weaker. But this seems very different from a regular weight modification case, in which, for example, your desire for a doughnut gets weaker, which would plausibly weaken the reason.²⁸

Our hypothesis also offers a nice explanation of this kind of overlap phenomenon. In the previous cases, we appealed to the fact that distinct but overlapping reasons were derivative on a single load-bearing reason, which is the primary contributor of normative weight. In this case, we instead observe that the reason given in 11 is derivative on two different load-bearing reasons. That you would enjoy a doughnut is plausibly a load-bearing reason, as is the fact that your mother would enjoy a doughnut, or that you've promised your mother a doughnut. Since these load-bearing reasons are distinct sources of normative significance, we see why the single fact, 11, would be a reason "twice over".

3.4 Fogal on normative clusters

Before moving on to develop metaphysical accounts of the distinction between load-bearing and derivative reasons, we want to briefly consider a different reaction one may have to the problem of overlap. Fogal (2016) argues for a theory of reasons on the basis of phenomena similar to the overlap phenomena we discuss here. He holds that we should think of the things that we typically call 'reasons' as mere representatives of *normative clusters*, which are the genuine units of normative significance. This is analogous to a natural way to understand talk of *causes*: there is a whole cluster of facts relevant to some event occurring, and in different conversational contexts, we can cite any one (but usually *only* one) of them as the cause of the event. Similarly, Fogal thinks that there is a whole cluster of relevant facts that explain the normative force in favour of acting in some way, and whether or not a given fact belonging to the cluster can be felicitously cited as a reason will depend on (and vary with) conversational context.

Fogal's explanation of overlap, then, is that the overlapping reasons are all part of the same normative cluster. The reason why these considerations overlap is just that each represents exactly the same cluster, which is the genuinely normatively significant thing. Given Fogal's account of reason ascriptions, this also explains why it is typically infelicitous to cite more than one overlapping reason in support of some option: it is typically infelicitous to cite more than one member of the cluster

²⁸ This desire-based example of weight modification is merely illustrative—as are all of our examples. However, plausibly any decent theory will be able to explain why at least sometimes your reason to go to the shop for a doughnut gets weaker if your desire for a doughnut becomes less intense.

in a given conversational context, since the cited member stands in for, or represents, the whole cluster. And we can explain why it is nevertheless technically *true* that each is a reason for the option: ' p is a reason to ϕ ' is true as long as p is suitable representative of a cluster that supports ϕ .

We are sympathetic to much of Fogal's account. As noted above, we agree with him that these cases involve different ways of representing the single normatively significant thing. We also agree that in different conversational contexts, citing some overlapping, merely derivative reason may appropriate, and that it's typically infelicitous to cite more than one. But we think that just talking in terms of clusters and representatives of those clusters misses out on some important metaphysical structure.²⁹ One way to press this thought is that we'd like an explanation for which facts get in the cluster and which don't, at least in our full metanormative theory (which Fogal does not claim to offer). Those who accept Reasons Fundamentalism, as we'll see, hold that, roughly speaking, it is a brute fact; those who offer certain analyses of reasons, who we'll call *Analysts*, hold that it depends on the promotion of significant objectives. What we hope to show is that answering this question, about what determines whether a fact belongs in the cluster of normatively relevant facts bearing on ϕ , imposes some structure on the facts: some of them turn out to be load-bearing reasons and others turn out to be derivative. And our hypothesis is that this structure lets us give a promising account of overlap.

The load-bearing reasons are more basic or more fundamental than the derivative ones in an important sense: the derivative ones are just reasons *because* of their relationship to the load-bearing reasons. What we want now is an account of this distinction. This will involve two tasks: first, identifying the load-bearing reasons, and second, saying how the derivative reasons are related to these. Here we just focus on the first of these tasks; there is lots of interesting work on the second already, and we cannot engage with it properly here.³⁰

3.5 Metaphysical fundamentality versus normative fundamentality

We are looking for an account of which of these reasons are more fundamental and which are derivative on these. Since reasons are facts, and some facts are more fundamental than others, a natural idea is to appeal to an independent account of which facts are more fundamental than which others. For example, the fact that Tom is in the room is plausibly more fundamental than the fact that someone whose name begins with 'T' is in the room. But sometimes less fundamental facts are more fundamental reasons. Suppose you promise your sister that you will make someone smile. The fact that squirting water at Jack would make Jill smile is a reason to squirt the water, and so is the fact that squirting water at Jack would make someone

²⁹ Compare Bader (2016), who argues for a different kind of structure within the cluster of normatively relevant facts, namely distinctions between reasons, conditions, and modifiers. See section 5.3 of Fogal (2016) for discussion of the worry that his account objectionably 'flattens' the normative landscape. He does not address the distinction between load-bearing and derivative reasons, under that name or any other.

³⁰ For important recent discussion, see Kolodny (2018).

smile. The fact that squirting water at Jack would make Jill smile is a more fundamental fact, than the fact that squirting water at Jack would make someone smile, but a less fundamental reason.

The important point is that the kind of (relative) fundamentality we are looking for here is *normative* fundamentality, which may come apart from metaphysical fundamentality: there is a difference between a fact being a reason in virtue of its relationship to other reasons, and a fact holding in virtue of its relationship to another fact. Now we turn to two kinds of theories—Reasons Fundamentalism and Analyses—that offer accounts of which reasons are the load-bearing ones, and evaluate their explanations of overlap.

4 The analyst's solution

The Analyst holds that reasons are or are explained by facts about the promotion of significant objectives.³¹ Two prominent instances are *desire-based* theories and *value-based* theories. Desire-based theorists take the significant objectives to be desired states of affairs, where these are usually the (actual or idealized) desires of the agent (e.g. Smith 1994; Schroeder 2007; Evers 2014), but may also be the desires of anyone (e.g. Manne 2016). Value-based theorists take the significant objectives to be objectively valuable states of affairs (e.g. Wedgwood 2009; Finlay 2006, 2014; Maguire 2016).³²

Since the promotion of these significant objectives is the basic source of normative significance for the Analyst, a natural account of the load-bearing reasons is that they are facts about the promotion of these objectives. A load-bearing reason to ϕ , on this view, is a fact of the form [ϕ promotes O], where O is a significant objective, to be determined by the details of the Analysis. For example, on a familiar agent-centered desire-based theory, the relevant O's will be outcomes that the agent desires. On a value-based theory, these will be objectively valuable outcomes, e.g. that someone is happy or that a promise is kept.³³ Derivative, and so overlapping, reasons will be ones that are only reasons because of their relationship to these load-bearing reasons.

The cases of overlap, in which multiple distinct reasons make overlapping contributions, will be ones in which there is just a single load-bearing reason, and so

³¹ When we talk about 'explanation', we have in mind a metaphysical, rather than epistemic or pragmatic, sense (compare Schroeder 2007, p. 29, fn. 12 and Broome 2004, pp. 32–33). We want to remain as neutral as possible about the nature of this kind of explanation, but it is arguably the sort of relationship that's at issue in discussions of fundamentality, grounding, 'in virtue of', etc.

³² On the Analyst in general, compare also Snedegar (2014, 2017). Note here that we're focusing just on promotion-based reasons, as opposed to reasons explained in terms of honoring or respecting significant objectives; cf. Anderson (1993), Scanlon (1998), Pettit (1991). We do, however, adopt a maximally permissive construal of the promotion relation. For alternative analyses of reasons, see Broome (2004), Kearns and Star (2009). Notice that evidence-based theories of reasons will face a slightly different overlap challenge, due, in part, to the different weighing behavior of epistemic reasons.

³³ Note that the claim is not that facts *about* desires or values (de dicto) are the load-bearing reasons. The fact that the outcome is desired or valuable is what makes the promotion of that outcome a reason. Compare Schroeder (2007), Chapter 2, on the 'background conditions'.

a single significant objective to be promoted, in virtue of which all the overlapping reasons are reasons. This is a very plausible thing to say about these cases. There's just one dog to be made happy by being given a treat, just one guy who needs vaccinated against one disease, and so on. In the cases in which a single fact seems to be a reason "twice over", the explanation extends very naturally: there are multiple significant objectives to be promoted (e.g. you enjoying a doughnut and you keeping the promise to your mother to bring her a doughnut), and so multiple load-bearing reasons. The fact that's a reason twice over is derivative on both of them.

To focus on the case involving giving Pearl a treat, the significant objective is plausibly *that Pearl is happy*. So, the load-bearing reason to give Pearl a treat is that [giving Pearl a treat would make her happy]. That giving her the treat would make her *very* happy, that it would make her happy *on a Tuesday*, and so on are derivative, and so overlapping, reasons. Alternatively, if the load-bearing reason is that [giving Pearl the treat would make a sentient creature happy], then that giving her the treat would make her happy is a derivative reason. Again, the details will depend on the details of our theory, in particular which objectives are the ones the promotion of which explains reasons.

An important class of derivative reasons are facts that *explain why*, or are part of the explanation for why, the fact of the form [ϕ promotes O] holds. For example, that they have doughnuts at the shop explains why going to the shop promotes the objective of you enjoying some doughnuts. Such a fact will be a derivative reason to ϕ , and so its contribution will overlap with that of the load-bearing reason, that going to the shop promotes the objective of you enjoying some doughnuts, and with the contributions of any other reasons that are derivative on the same load-bearing reason. Some philosophers who defend these kinds of Analyses of reasons hold that what it is to be a reason to ϕ is to explain why ϕ -ing would promote some significant objective.³⁴ Importantly, these theories do not include a distinction between load-bearing reasons and derivative reasons: if some fact meets the condition, then it is a reason. We can agree with this as an account of what it is to be a reason, but we insist that explaining overlap requires drawing the distinction within the class of reasons.

The Analyst looks to have an explanation of overlap ready to hand. Since the sources of normative significance are certain kinds of objectives to be promoted, these theories come with a natural account of the load-bearing reasons: facts about which actions would promote these objectives. In cases of overlap, we either have multiple facts that are reasons in virtue of their relationship to a single load-bearing reason, or a single fact that is a reason twice over—in virtue of its relation to *two* load-bearing reasons.

But things are not quite so straightforward for the Analyst. We will discuss two problems. The first problem is that the significant objectives that ground reasons on the Analyst's view—desires, values, or something else—can themselves exhibit overlap.

³⁴ We have in mind Finlay (2006, 2014) and Schroeder (2007), in particular.

Value-based theorists of reasons who accept Value Fundamentalism (the view that the property of *being valuable* does not admit of full definition in other terms) face a version of this problem. Consider the value of the following facts, which involve differences in grain of the kind we saw in Sect. 3.1:

14. Pearl is a very happy dog.
15. Pearl is a happy dog.
16. Someone is happy.
17. Someone is happy on a Tuesday.

Adding 15–17 would not add to the value already given by fact number 14. The value of 15–17 overlap with that of 14. So, if values are what ground reasons, we have just pushed the problem of overlap back a step.

This problem for value is more well-known than the problem of overlapping reasons.³⁵ In response, axiologists have attempted to provide an account of ‘basic’ value, in terms of which such relations of ‘overlap’ between values are to be accounted for. Value-based theorists (e.g. Maguire (2016)) have argued that an account of overlapping reasons can be based on such an account of overlapping values. The Value Fundamentalist may argue that such an account of basic value has employment beyond this role in accounting for overlapping reasons, and consequently, that the resources required for a value-based theorist to explain overlapping reasons are motivated independently of anything about reasons.³⁶

Desire-based theorists also face a version of this overlap challenge. Suppose that you desire for someone to be happy. Then fact 14 satisfies your desire, as do facts 15–17. But you will not be any *more* satisfied by the addition of facts 15–17—unless, of course, you have some other desires that would thereby be satisfied.

But interestingly, desire-based theorists—including various idealized desire-based accounts—may have various advantages over value-based Analysts here. Firstly, there are surely facts about the grain of one’s desires, however vague or complex they might be. What’s more, these facts are clearly independent of one’s theory of reasons. We need an account of the grain of desires for purely non-normative purposes, e.g. in predicting others’ behavior. We can therefore assume that some independently motivated account of this grain is available free of charge, so to speak.

The second problem for the analyst is that their strategy relies on a so-far unargued-for assumption that the load-bearing reasons—the things that make single, non-overlapping contributions of normative weight—will line up in just the right ways with the significant objectives. For the value-based theorist, this is the assumption that the load-bearing reasons will be facts about the promotion of basic values. But this is contestable: load-bearing reasons to give money to charity presumably do not advert to basic values; rather they concern more general expected values. For the desire-based theorist, this is the assumption that the actual

³⁵ E.g. Chappell (2015, p. 327): “Whatever substantive disputes we may have about what is of value, we should all acknowledge the formal difference between (i) a pair of options serving distinct but equally weighty final values, and (ii) a pair of options serving literally one and the same final value”.

³⁶ See, e.g., Feldman (2000).

grain of the desire that grounds the reasons will coincide with the load-bearing reason, rather than with one of the overlapping, merely derivative reasons. One concern here is that people might have *overlapping desires*, but it wouldn't seem to follow they have twice as much reason to pursue their satisfaction.

So, the Analyst does have a promising solution to the overlap challenge, since a very natural explanation of overlap appeals to facts about the significant objectives to be promoted. More work remains to be done. But we do think that the questions that have arisen—about overlapping values and desires, and about how basic value connects up with load-bearing reasons—are of independent interest. The overlap challenge for reasons illustrates the importance of these questions for Analysts, and also plausibly provides a rich source of data for theorizing about them. We turn now to a competing metaphysical account of reasons, Reasons Fundamentalism.

5 The Reasons Fundamentalist's solution

The Fundamentalist is committed to maintaining that at least some facts about reasons are normatively fundamental, in sense that they are not fully grounded in other normative facts (that are not themselves fully grounded in facts about reasons). Sometimes Reasons Fundamentalism is called *Reasons Primitivism*; we prefer to stick with the former to emphasize that our topic is metaphysical rather than conceptual. The view (under either name) should be distinguished from what we'll call *Reasons First-ism*, which holds that all other normative properties can be explained in terms of reasons. Many Reasons Fundamentalists also accept Reasons First (see, e.g., Parfit 2011), though not everyone who accepts Reasons First accepts Reasons Fundamentalism or Reasons Primitivism (see, e.g., Schroeder 2007).

The simplest version of Reasons Fundamentalism conjoins the thesis that the reasons relation is fundamental with the thesis that nothing informative can be said about what it is to be a reason, other than that it counts in favour of the option for which it is a reason.³⁷ This is consistent with the phenomenon of overlap. For each reason contributes weight to an option. It is just that not each reason contributes *separate* weight to an option. Some distinct reasons make an identical contribution of weight to the option. The challenge for this simple view is to explain the distinction between merely derivative reasons and load-bearing reasons.

We have seen so far that merely pointing to the grounds of reasons posited by the Analyst is not sufficient to fully explain why reasons overlap. This may amount to simply pushing the problem back, given that the grounds themselves can exhibit overlap. Nevertheless, we think that we should be optimistic that the Analyst will be able to make progress here, since, as noted, it is very natural to appeal to facts about the promotion of significant objectives, e.g. that there is just one significant objective to be promoted, in explaining overlap.

This may seem to present a challenge for Reasons Fundamentalism. It is because the Analyst has an answer to the question of what it is to be a reason, given in terms

³⁷ Cf. Scanlon (1998), Parfit (2011).

of these significant objectives, that this (start of an) explanation of overlap is available. But though the Fundamentalist denies that we can explain what it is to be a reason in other terms, they need not deny that we can at least sometimes explain why some fact is a reason. In particular, it is open to the Fundamentalist to adopt our hypothesis that some reasons are load-bearing while others are merely derivative. We explain the derivative reasons in terms of their relationships to the load-bearing reasons. As long as we do not explain what it is to be a reason in other terms, as the Analyst does, this will be consistent with Reasons Fundamentalism.

The easiest way to see the idea behind this strategy is to see how it applies in a given case. The first task is to identify some load-bearing reason—perhaps the fact that giving the treat to Pearl will make a sentient creature happy—that explains why all the other facts are reasons. For example, that giving her the treat would make a sentient creature happy *on a Tuesday*, that it would make *Pearl* happy, and that it would make Pearl *very* happy are all plausibly derivative reasons, derivative on the fact that giving her the treat would make a sentient creature happy, which is something you have load-bearing reason to do. These load-bearing reasons are the primary sources of normativity, and so reasons that are derivative upon the same load-bearing reason will make overlapping contributions, just as on the Analyst’s solution.

The first thing to note about giving an account of load-bearing reasons is that we’d like to do so in a non-ad hoc way. In particular, we’d like some general answer about what the load-bearing reasons are that delivers the right results in cases of overlap, rather than taking them case-by-case and picking out just the right load-bearing reasons to give the intuitively correct results regarding overlap.

We think that the best strategy for the Reasons Fundamentalist is to follow Scanlon (2014) and Rosen (2017) in first identifying some *reasons principles*.³⁸ These are necessary principles that explain which facts are reasons. Presumably, some substantive account of reasons principles will need to be defended, at least in part, on the basis of substantive normative theorizing. There are various constraints on a plausible account of reasons principles. In addition to various plausible consistency constraints, there may be redundancy requirements, e.g. ruling out two distinct principles, one of which yields reasons to promote pleasure, the other of which yields reasons to promote pleasure on Tuesdays. Such a redundancy requirement would help with some cases of overlap that we’ve considered, the alternative grain cases in particular. Coming up with a suitable supply of reasons principles will undoubtedly be a complex task. We will focus just on one sort of problem involving overlap.

³⁸ Scanlon calls these ‘pure reasons.’ We prefer ‘reasons principles’ since they aren’t really reasons for anything. See also Crisp’s (2006) notion of an ‘ultimate reason.’ A less promising alternative to the pure reasons strategy is the *particular reasons* strategy, according to which, roughly, more general reasons explain less general reasons – without the need to appeal to any abstract principles like Scanlonian pure reasons (compare the discussion in Rosen 2017). We are not convinced the difference between generality and abstractness sustains a difference of much significance for our discussion.

The following is a candidate for a reasons principle:

RP: For all agents A, options O, and circumstances C, if A's O-ing in C would prevent the suffering of a sentient being, then that fact is a reason for A to O in C.

The load-bearing reasons will then be true instances of the antecedent of some reasons principle—in this case, that A's O-ing in C would prevent the suffering of a sentient being. To see how this works, consider the vaccination case from Sect. 4. There we argued that all of the following facts are overlapping reasons for you to press the button:

9. Pressing the button gives Bill V1.
10. Pressing the button gives Bill V2.
11. Pressing the button vaccinates Bill.

We could have added another apparently overlapping reason:

12. Pressing the button will prevent the suffering of a sentient being (since contracting the disease would cause Bill to suffer).

On the view we are suggesting here, according to the reasons principle we described above, this fact, that your pressing the button would prevent the suffering of a sentient being, is a load-bearing reason to press it. The other facts, 18–20, are derivative reasons, explained by their relationship to this load-bearing reason.³⁹ Thus, they overlap both with the load-bearing reason and with each other, since the load-bearing reason is the single primary contributor of normative weight.

This strategy gives an attractive explanation of cases like this vaccination case that has the kind of structure we proposed: we identify some basic source of normative significance as the load-bearing reason, and point out that the overlapping reasons are reasons because of their relationship to this basic source. Whereas the Analyst steps outside the domain of reasons to find this basic source of normative significance, the Reasons Fundamentalist privileges a certain class of reasons as the basic sources.

However, it turns out to be difficult to specify reasons principles that get the right results; we'll illustrate with the principle RP. The principle faces an *overgeneration* problem: it predicts that there will be overlap in cases in which there is no overlap. Suppose that instead of giving Bill both vaccines V1 and V2, pressing the button will give Bill V1 and Bob V2; again, both vaccines are equally effective and neither have any side-effects. Consider the following reasons:

22. Pressing the button will prevent the suffering of a sentient being.
23. Pressing the button will vaccinate Bill.
24. Pressing the button will vaccinate Bob.

³⁹ As noted above, specifying exactly what transmission relationship these derivative reasons must bear to a load-bearing reason is an important task for this strategy, but we set it aside here. But plausibly, at least often, when some fact F explains why A-ing is a way to comply with a load-bearing reason R, F is a derivative reason, derivative on R.

The relations of overlap are intuitively obvious: 22 overlaps with both 23 and 24. The strategy we're developing now explains each of these overlap relationships in the way we've just outlined: 22 is the load-bearing reason, explained by the reasons principle RP, and 23 is derivative (roughly, since it explains why pressing the button is a way of preventing the suffering of a sentient being), so 23 overlaps with 22. An analogous explanation explains why 22 overlaps with 24. But note that to explain these overlap relationships, we have to hold that 23 and 24 are derivative reasons explained by a single load-bearing reason, namely 22. In order to explain overlap in cases like the original vaccination case (and the other cases we've considered), we said that when two reasons are derivative on the same load-bearing reason, they overlap both with the load-bearing reason and with one another. So, this strategy predicts that, in this case, 23 overlaps with 24. But that's incorrect: these are two distinct reasons to press the button that make distinct, non-overlapping contributions in favour of doing so.

This problem relies on the claim that the load-bearing reason is given by a reasons principle like RP, which does not mention specific individuals. Since the reasons principles are meant to be necessary normative truths, it is reasonable to assume that they address more general properties, such as preventing the suffering of a sentient being.⁴⁰ But perhaps there is a better way to understand the reasons principle involved that can give us two distinct load-bearing reasons, and so two contributors of normative weight. For example, a revised version of the reasons principle RP is the following:

RP*: For all individuals I, agents A, options O, and circumstances C, if A's O-ing in C would prevent the suffering of I, then that fact is a load-bearing reason for A to O in C.

The load-bearing reasons explained by this principle will be facts about preventing the suffering of particular individuals, even though the principle itself doesn't mention any specific individuals.⁴¹ We still get an explanation for why 22 overlaps with both 23 and 24, but we deny that 22 is itself a load-bearing reason. There will be *two* load-bearing reasons: that pressing the button will prevent Bill's suffering, and that pressing the button will prevent Bob's suffering. The former overlaps with both 22 and 23, while the latter overlaps with both 22 and 24, but 23 and 24 do not overlap, just as we would expect.⁴²

But we haven't yet fully solved the overgeneration problem. Suppose that the button will give Bill both vaccine V1 and also vaccine V3, which prevents a different disease. Then we have the following two reasons:

⁴⁰ We don't mean to deny that proper names might feature in some reasons principles, for instance in cases of reasons of partiality, or perhaps reasons to respect God's commands. We simply deny that the reasons principles governing reasons of benevolence mention unremarkable individuals like Bill.

⁴¹ Note that we can motivate RP* over RP by appeal to the separateness of persons, as well.

⁴² This means that overlap is not a transitive relation, but given the existence of what we called "alternative grounds" cases, we should not have expected it to be transitive. In the case at hand, the reason in 22 is a reason that displays "alternative grounds" overlap.

25. Pressing the button gives Bill V1.
26. Pressing the button gives Bill V3.

These two reasons do *not* overlap, since each vaccine prevents a different disease. Note that each *does* overlap with reason 27:

27. Pressing the button prevents the suffering of Bill.

We can explain why 25 and 26 overlap with 27, given RP*, since 27 is a load-bearing reason and 25 and 26 are both derivative on 27. But then, since 25 and 26 are derivative on a single load-bearing reason, we get the incorrect result that 25 and 26 overlap with one another. To avoid this result, we would need to revise RP* so that we can distinguish not only between different individuals, but also between different ways in which those individuals may suffer. This is possible, of course. But we predict that there will likely be ways to construct overlap cases even for a principle like this. The general problem facing this strategy is this: we need to make our reasons principles sufficiently fine-grained to avoid the overgeneration problem, but in doing so, we must be careful to avoid an *undergeneration* problem—missing out explaining intuitive cases of overlap. We also should avoid the ad hoc method of simply picking and choosing our load-bearing reasons to generate the correct results about which reasons overlap, especially if we want these load-bearing reasons to follow from a supply of necessary reasons principles.

A pessimistic reaction is that this kind of ad hoc method will be the best we can hope for, undermining the appeal of Reasons Fundamentalism, or at least of the initially promising strategy of identifying load-bearing reasons to explain overlap. But we prefer a more optimistic reaction. It is not ad hoc to allow intuitions about which reasons overlap and which ones don't to inform our first-order theorizing about which reasons principles are true. Thus, the phenomenon (or phenomena) of overlapping reasons provides a rich source of data for further developing both Reasons Fundamentalism and first-order normative theory.

To take this one step further, the Reasons Fundamentalist may also develop the load-bearing reasons strategy by drawing various modal distinctions intended to preserve the thesis that the reasons relation is fundamental. In particular, Scanlon (2014) has argued that the desire-based theory of reasons is best understood not as an analysis of reasons, but as a substantive account of what reasons there are. (He thinks it is a false substantive account, but that's beside the point.) The same might be said of the value-based theory. If so, then the Reasons Fundamentalist can mimic the Analyst's explanation of overlap by holding that the load-bearing reasons are facts of the form [ϕ -ing would promote value V]/[ϕ -ing would promote the satisfaction of desire D], and holding that reasons that involve the promotion of the same value/desire overlap. As we argued above, there are important questions about the Analyst's explanation of overlap. But the point here is that if we can give a satisfactory explanation of overlap by appealing to analyses of reasons in terms of values or desires, we may also be able to give a satisfying explanation by appealing to a substantive view on which the load-bearing reasons are about values or desires.

More generally, it seems available to a reasons fundamentalist to appeal to a set of substantive principles concerning which facts are reasons. One job for such a set

of principles will be to pick out reasons with just the right grain. Just as a substantive axiology will determine whether (e.g.) happiness, or fulfilment, or pleasurable experiences, or perhaps pleasure on a Tuesday, is basically valuable, so a substantive theory of reasons will determine whether the promotion of happiness, or the fulfilment, or what have you is a load-bearing, or merely derivative, normative reason.

6 Conclusion

Developing the most natural explanation of overlap on behalf of Reasons Fundamentalism is not straightforward. The Analyst does have an easier explanation for all of these kinds of cases. In the original vaccination case, there is just one valuable or desired outcome—that Bill is vaccinated from the disease, so the reasons overlap. In the case involving both Bill and Bob, there are *two* valuable or desired outcomes, one each for Bill and Bob, and so the reasons do not overlap. In the case just involving Bill but involving two separate diseases, we have two valuable or desired outcomes, one for each disease, so the reasons do not overlap. And again, these explanations follow very naturally from the Analyst’s account of what it is to be a reason.

As we explained, however, there are also important open questions for the Analyst’s explanation of overlap. In general, these concern independent issues from either axiology (for the value-based theorist) or the philosophy of mind (for the desire-based theorist). This isn’t all good news for the Analyst, since there remains the possibility that theorizing in axiology or the philosophy of mind ends up not supporting the style of explanation of overlap we have proposed here—for example, if the most plausible account of basic value does not give us a plausible account of load-bearing reasons.

There remain various subtle differences between Fundamentalists and Analysts in the manner of defending the distinction between load-bearing and derivative reasons. But we hope to have done enough to establish some hope that, on either approach, this distinction can be upheld in a way that yields a satisfying explanation of a range of plausible cases of overlap. There is good reason to pursue this project, since the distinction yields a defence of separability, which, in turn, constitutes a reply to an important objection to the balancing model. It doesn’t follow that the balancing model is true, of course. That will turn, at least in part, on questions about which first-order ethical theory is true. Different ethical theories employ different ideological and logical infrastructures. Still, we hope to have done enough to show that the attractive features of the balancing model remain available, and that there remains hope for constructive metaphysics even in this messy part of practical philosophy.

Acknowledgements Many thanks to Daniel Fogal, Nathan Howard, Benjamin Kiesewetter, Shyam Nair, Itai Sher, Jonathan Way, Jack of the Woods, and an extremely helpful anonymous referee. Thanks also to audiences at the Harvard Normative Ethics and Welfare Economics Conference, the Southampton-Humboldt Normativity Conference, Duke University, Arizona State University, and the University of Zurich.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Anderson, E. (1993). *Value in ethics and economics*. Cambridge: Harvard University Press.
- Arpaly, N. (2002). *Unprincipled virtue*. Oxford: Oxford University Press.
- Bader, R. (2016). Conditions, modifiers, and holism. In E. Lord & B. Maguire (Eds.), *Weighing reasons* (pp. 27–55). Oxford: Oxford University Press.
- Baier, K. (1958). *The moral point of view*. New York: Random House.
- Berker, S. (2007). Particular reasons. *Ethics*, 118(1), 109–139.
- Broome, J. (2004). Reasons. In R. J. Wallace, M. Smith, S. Scheffler, & P. Pettit (Eds.), *Reason and value: Themes from the moral philosophy of Joseph Raz* (pp. 28–55). Oxford: Oxford University Press.
- Broome, J. (2013). *Rationality through reasoning*. London: Wiley Blackwell.
- Chang, R. (2009). Voluntarist reasons and the sources of normativity. In D. Sobel & S. Wall (Eds.), *Reasons for action*. New York: Cambridge University Press.
- Chappell, R. (2015). Value receptacles. *Noûs*, 49(2), 322–332.
- Crisp, R. (2006). *Reasons and the good*. Oxford: Oxford University Press.
- Dancy, J. (2004). *Ethics without principles*. Oxford: Oxford University Press.
- Dreier, J. (2007). Comment on Peasoup blog. Available at: http://peasoup.us/2007/01/the_buckpassing/. Accessed 17 Feb 2020.
- Evers, D. (2014). In defence of proportionalism. *European Journal of Philosophy*, 22(2), 313–320.
- Feldman, F. (2000). Basic intrinsic value. *Philosophical Studies*, 99(3), 319–346.
- Finlay, S. (2006). The reasons that matter. *Australasian Journal of Philosophy*, 84(1), 1–20.
- Finlay, S. (2014). *Confusion of tongues*. Oxford: Oxford University Press.
- Fogal, D. (2016). Reasons, reason, and context. In E. Lord & B. Maguire (Eds.), *Weighing reasons* (pp. 74–103). Oxford: Oxford University Press.
- Gert, J. (2007). Normative strength and the balance of reasons. *Philosophical Review*, 116(4), 533–562.
- Greenspan, P. (2005). Asymmetrical practical reasons. In M. E. Reicher, & J. C. Marek (Eds.), *Experience and analysis*. Öbv&Hpt, (pp. 387–394).
- Hawthorne, J., & Magidor, O. (2018). Reflections on the ideology of reasons. In D. Star (Ed.), *Oxford handbook of reasons and normativity*. Oxford: Oxford University Press.
- Horty, J. (2012). *Reasons as defaults*. Oxford: Oxford University Press.
- Howard, N. (ms). The Goals of Moral Worth. Manuscript, Texas A&M University.
- Johnson King, Z. (2019). We can have our buck and pass it, too. *Oxford Studies in Metaethics*, 14, 167–188.
- Kagan, S. (1988). The additive fallacy. *Ethics*, 99(1), 5–31.
- Kearns, S. (2016). Bearing the weight of reasons. In E. Lord & B. Maguire (Eds.), *Weighing reasons* (pp. 173–190). Oxford: Oxford University Press.
- Kearns, S., & Star, D. (2009). Reasons as evidence. In R. Shafer-Landau (Ed.), *Oxford Studies in Metaethics* (Vol. 4, pp. 215–242). Oxford: Oxford University Press.
- Kelman, S. (2016). An ethical critique of cost-benefit analysis. In J. Anomaly, et al. (Eds.), *Politics, philosophy, and economics, an anthology*. Oxford: Oxford University Press.
- Kolodny, N. (2018). Instrumental reasons. In D. Star (Ed.), *The Oxford handbook of reasons and rationality*. Oxford: Oxford University Press.

- Lord, E., & Maguire, B. (2016). An opinionated guide to the weight of reasons. In E. Lord & B. Maguire (Eds.), *Weighing reasons*. Oxford: Oxford University Press.
- Maguire, B. (2016). The value-based theory of reasons. *Ergo*, 3(9), 233–262.
- Maguire, B. and Woods, J. (2020). The game of belief. *The Philosophical Review*, vol. 130, issue 2.
- Manne, K. (2014). Internalism about reasons: Sad but true? *Philosophical Studies*, 167(1), 89–117.
- Manne, K. (2016). Democratizing Humeanism. In E. Lord & B. Maguire (Eds.), *Weighing reasons* (pp. 123–140). Oxford: Oxford University Press.
- Nair, S. (2016a). Conflicting reasons, unconflicting ‘ought’s. *Philosophical Studies*, 173(3), 629–663.
- Nair, S. (2016b). How do reasons accrue? In E. Lord & B. Maguire (Eds.), *Weighing reasons* (pp. 56–73). New York: Oxford University Press.
- Parfit, D. (2011). *On what matters*. Oxford: Oxford University Press.
- Pettit, P. (1991). Consequentialism. In P. Singer (Ed.), *A companion to ethics* (pp. 230–237). Oxford: Blackwell Publishers.
- Prakken, H. (2005). A study of accrual of arguments, with applications to evidential reasoning. In *Proceedings of the tenth international conference on artificial intelligence and law, Bologna 2005*. New York: ACM Press, pp. 85–94.
- Rosen, G. (2017). Scanlon’s modal metaphysics. *Canadian Journal of Philosophy*, 47(6), 856–876.
- Ross, W.D. (1930). *The right and the good*. Oxford: Oxford University Press
- Scanlon, T. M. (1998). *What we owe to each other*. Cambridge: Harvard University Press.
- Scanlon, T. M. (2014). *Being realistic about reasons*. Oxford: Oxford University Press.
- Schroeder, M. (2007). *Slaves of the passions*. Oxford: Oxford University Press.
- Schroeder, M. (2009). The Buck-passer’s negative thesis. *Philosophical Explorations*, 12(3), 341–347.
- Setiya, K. (2014). What is a reason to act? *Philosophical Studies*, 167(2), 221–235.
- Sher, I. (2019). Comparative value and the weight of reasons. *Economics and Philosophy*, 35(1), 103–158.
- Smith, M. (1994). *The moral problem*. London: Wiley Blackwell.
- Snedegar, J. (2014). Contrastive reasons and promotion. *Ethics*, 125(1), 39–63.
- Snedegar, J. (2017). *Contrastive reasons*. Oxford: Oxford University Press.
- Way, J. (2017a). Reasons as premises of good reasoning. *Pacific Philosophical Quarterly*, 98(2), 251–270.
- Way, J. (2017b). Creditworthiness and matching principles. In M. Timmons (Ed.), *Oxford studies in normative ethics*, Vol 7 (pp. 207–228).
- Wedgwood, R. (2009). Intrinsic values and reasons for action. *Philosophical Issues*, 19, 342–363.

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.