

# Overlapping Reasons

Barry Maguire (UNC) and Justin Snedegar (St Andrews)\*

*Draft for Harvard Conference on Normative Ethics and Welfare Economics*

*Note: This is co-authored work that is very much in-progress. We are still working out some kinks and tightening things up. We also need to add references and citations. So please do not quote or cite. I (JS) will also attempt to present a slightly more accessible and streamlined version at Harvard.*

Since at least 1930, theorists working in moral philosophy and normativity more generally have recognized the importance of accounting for *trade-offs* in normative theories of what we ought to do.<sup>1</sup> The most common way to do this is to appeal to contributory or *pro tanto* reasons. The thought is that the reasons for and against the options open to the agent interact and compete—in highly complex ways—to support a conclusion about what the agent ought to do, all things considered. On most views, the reasons compete by contributing normative weight for or against the relevant option.<sup>2</sup> This thought already raises many interesting questions, and this paper will interact directly with these issues. One goal of the paper, as we'll see, is to bring out an underappreciated complication for this picture. But our main goal concerns what is largely an orthogonal issue, about the nature of reasons themselves, rather than how they interact to support conclusions about what we ought to do.

We argue for what we call a 'bottom up' theory of normative reasons, according to which reasons are grounded in some underlying property or state. The most popular bottom up theories are *desire-based* theories (e.g., Schroeder (2007), Manne (2016)) and *value-based* theories (e.g., Wedgwood (2009), Maguire (2016)); we also include *duty-based* theories such as Ross (1930). These theories share a common structure (or at least certain structural features); since our argument is primarily about the structure of reasons, we remain neutral about which bottom up theory is correct. The main goal here is to argue in favor of bottom up theories, as compared to *Reasons Fundamentalism*, as defended by Scanlon (2014), Parfit (2011), and others. What's distinctive of Reasons Fundamentalism, for our purposes, is the claim that reasons are not grounded in anything at all, but are rather fundamental. (More carefully, reason relations are fundamental.)

The argument turns on a class of phenomena that we group under the heading *overlapping reasons*. As we'll see, we group several different kinds of cases under this heading. What they have in common is that they are cases in which the reasons—that is, the facts which are reasons—come apart in some way from what we call the *normative contributors*—the things that contribute normative weight for or against an option.

---

\* [maguirbj@email.unc.edu](mailto:maguirbj@email.unc.edu) and [js280@st-andrews.ac.uk](mailto:js280@st-andrews.ac.uk)

<sup>1</sup> This is to trace this line of thought back to W.D. Ross's (1930) *The Right and the Good*.

<sup>2</sup> For some notable exceptions to this way of understanding the way in which reasons support conclusions about what we ought to do, see Greenspan (2005), Gert (2007), and Horty (2012).

Bottom up theories easily explain the cases, but Reasons Fundamentalism struggles to do so. We'll present the phenomena in section 1, expanding on the distinction between reasons and contributors mentioned above, and then offer our preferred explanation in section 2. Our claim is that the phenomena should be explained by appealing to the grounds of reasons. In section 3, we develop the argument against Reasons Fundamentalism. This primarily involves arguing against an explanation that seeks to understand overlapping reasons in terms of derivative and non-derivative, or basic, reasons. There is much more to say about overlapping reasons, and about how the Reasons Fundamentalist may attempt to explain the phenomena, but given the time constraints, we take ourselves to have at least revealed an explanatory burden for Reasons Fundamentalism.

## **1 Overlapping Reasons**

### **1.1 One reason, one contribution**

The easiest way to introduce the phenomena of overlapping reasons is to begin with a familiar metaphor. As we said in the introduction, one of the main advantages of talking about reasons in practical philosophy, as well as epistemology, is that doing so allows us to easily accommodate the fact that thinking about what to do and what to think frequently involves making *trade offs*. For each option, there are usually some things to be said in favor and some things to be said against. Reasons for and against the options are exactly what we need to capture this. These reasons interact and compete with one another to determine which option is most strongly supported by the reasons, once all of them are taken into account. This competition is often represented by the metaphor of a balance or scale, with pans corresponding to each option and the reasons for those options corresponding to counters or marbles of various sizes.

This metaphor is popular, and for good reason. It has several features that we want from a model of how reasons compete, including representations of the fact that reasons for competing options compete with one another and the fact that reasons for the same option can combine to support the option more strongly than either does individually. But there are also well known limitations. For example, even though reasons for the same option seem to combine, they may not do so in any straightforward, additive way, unlike the weights of marbles. For another limitation, while the weights of marbles are always comparable, it is at least highly plausible that some reasons are not comparable. So the balancing metaphor builds in some questionable assumptions about the interaction of reasons. Philosophers are by and large aware of these limitations, though, and so they do not often lead to any obvious problems.

The phenomena of overlapping reasons put pressure on a different, and not widely recognized, assumption of this model. This is what we'll call the 'one reason, one contribution' assumption. When we weigh up some marbles, each marble makes a single, distinct contribution to the total weight in the pan. The corresponding

assumption about reasons, then, is that each reason for or against an option makes a single, distinct contribution to the determination of whether the agent ought to perform that option. What unites the cases of overlapping reasons is that they are cases in which this assumption appears to be false.

We want to emphasize that our claim is not that our opponents—in particular, Reasons Fundamentalists—are deeply committed to this assumption, or more generally to the balancing metaphor. Rather, framing things in this way will help to make clear the important features of the cases we are interested in. The challenge for our opponents is not to somehow fix up the balancing metaphor, but rather to show how to explain these cases on their preferred theory of reasons.

Cases of overlap fall into two different groups. On the one hand, there are cases in which distinct considerations, all of which are reasons, do not make distinct contributions. That is, the contributions of distinct reasons overlap. On the other hand, there are cases in which a single consideration, or a single reason, makes multiple distinct contributions. That is, the reasons providing distinct contributions overlap.

As should be clear, this way of describing things relies on drawing a distinction between reasons and what we call *normative contributors* (we'll often just say 'contributors'), where a normative contributor is just whatever it is that makes a single, distinct contribution of normative weight for or against an option. The 'one reason, one contribution' assumption built into the balancing metaphor is that there is no distinction between reasons and contributors. The challenge for a theory of normativity is to account for this distinction. We think this should be done by appealing to the grounds of reasons. Reasons Fundamentalists must take a different approach.

Put another way, the challenge raised by cases of overlap is to provide a way to individuate contributors. Individuating reasons is easy: since reasons (we're going to assume, along with our opponents) are facts, different facts are different reasons. But this is not, in general, true of contributors, as our cases below are meant to illustrate.

## 1.2 Examples of overlapping reasons

First, there are cases of what we will call *alternative satisfiers* of a reason. These are cases in which different facts seem to be reasons, but not do not make distinct normative contributions. One intuitive earmark for this is if one of the reasons were 'satisfied' — when the action it is a reason for is performed — the other reasons no longer apply. There are a variety of such cases. There are cases in which different facts are reasons for the same agent to do different things, but don't make distinct contributions. Suppose you can push button A or push button B. Either way, some particular child will be given an important vaccination. There are no other normatively relevant factors in the case. Do you have a reason to press button A? Yes: the fact that pressing button A would cause the child to be vaccinated is a reason to press it now. Do you have a reason to press button B? Again, yes: the fact that pressing button B would cause the child to be vaccinated is a reason to do so. But these reasons do not seem to make distinct contributions to answering the question of what you ought to do.

Change the example to one with the same action but two different relevantly similar outcomes. Suppose you can push a button now and one of two children, Bill and Ben, will be given a vaccination. Both children are strangers to you and neither needs the vaccination more than the other. The fact that pressing the button will increase the chances of Bill getting vaccinated is a reason to press. So is the fact that pressing the button will increase Ben's chances. But there is something odd about saying you have two reasons to press the button. After all, only one of these children will get the vaccination.

A related class of overlapping reasons is what we'll call cases of *alternative grain*. The fact that giving lemon cake with cherry frosting to Pearle the wonder dog on a Tuesday would make her a very happy dog is a reason to do so. So is the fact that giving cake to a dog on a Tuesday would make her very happy. So is the fact that giving a delicious treat to a dog on any day of the week would make the dog happy. But these reasons don't add weight to each other, surely. Suppose that first of these facts would be slightly outweighed by another reason: that you had promised to give the cake to Mildred. The further two considerations do not add weight to the first in a way that could outweigh this promissory reason.

We'll subsume under this class of cases another class of cases involving *alternative referrers*. Suppose you can bring your mother, Pam, some flowers. This will make her happy. So consider the following reason ascriptions: that Pam will be happy is a reason to bring your mother flowers; that your mother will be happy is a reason to bring your mother flowers; that someone will be happy is a reason to bring your mother flowers. All of these seem like perfectly good reasons to bring your mother flowers. But it is not as if you can combine these three reasons. If you've already cited the fact that your mother will be happy as a reason to bring some flowers, it doesn't help your case (if I'm wondering whether you should really be bringing flowers) to point out that not only will your mother be happy, but in addition *someone* will be happy, and also that *Pam* will be happy. So this seems to be a case in which these three reasons overlap: they don't make distinct contributions.

There are situations in which these considerations could make different contributions: perhaps you promised to make someone called "Pam" happy today. Then the fact that bringing flowers will make Pam happy does seem to contribute separately from the fact that bringing flowers will make your mother happy. But notice what we did to make this sound plausible: we appealed to an independent explanation of why these reasons make distinct contributions. As we'll see shortly, we think this observation is very important for explaining overlap. For now, the point is that it is clear that there are cases in which there are no further such considerations, in which these distinct considerations do not make distinct contributions to what you ought to do.

These sets of cases involve different ways in which different reasons make a single contribution. One response is to deny that we've really identified distinct reasons in these cases. Perhaps some of these considerations are not, in fact, reasons. But, to focus for concreteness on the case of giving Pearle the cake, it is hard to deny that the facts cited in this case are indeed reasons to give the cake to Pearle on standard theories

of reasons. These facts all figure in explanations of why giving the cake to Pearle would promote the object of one of your desires (if you, for example, desire to make dogs happy), or promote or respect some objective value. They also potentially figure in explanations of why you ought to give the cake to Pearle—or at least, the play the “forgiving her the cake” role in an explanation of why you ought, ought not, or neither ought or ought not give her the cake. They are evidence that you ought to give her the cake. And they certainly seem to count in favor of giving her the cake.<sup>3</sup>

The last class of cases involves situations in which a single reason makes multiple contributions. We call these cases of *alternative grounds* of reasons, though this name is admittedly biased in the direction of our preferred explanation. Suppose that you want a doughnut. Then the fact that they sell doughnuts at the bakery is a reason to go there. Suppose you also want to buy your mother a doughnut. Then, again, the fact that they sell doughnuts at the bakery is a reason to go there. This one fact, that they sell doughnuts at the bakery, seems to be a reason twice over. For another case, suppose you see someone in serious need of help. Since people should get help when they need it, the fact that they are in serious need of help is a reason to help them. Suppose that helping those in need improves your character. Then, again, the fact that this person is in serious need of help is a reason to help them. Once again, we have one fact, that the person is in serious need of help, that seems to be a reason twice over. It contributes normative weight in favor of helping twice, in two different ways.

One tempting response here is to just hold that in this kind of case, we have one reason and one contributor, but that the contribution is just greater. That the bakery sells doughnuts is an even weightier reason to go there, given that in addition to wanting one for yourself, you also want to buy one for your mother. One problem with this strategy is that it seems to smudge over a normatively relevant distinction. The two contributions are importantly distinct. A sharper problem is that there are cases of alternative grounds where the two contributions are not contributions in favor of the same action. Suppose you would really enjoy a doughnut, but also have promised your mother that you’ll stay away from doughnuts. In this kind of case, the fact that they sell doughnuts at the bakery is plausibly a reason to go to the bakery and also a reason not to go to the bakery, or a reason to stay away. We cannot simply combine the weights of the two contributions, since one of them favors going to the bakery and the other favors staying away, even though there seems to be just the one reason. In terms of the balancing metaphor, for example, we’d have to put a marble corresponding to this one fact—that they sell doughnuts at the bakery—on two competing pans.

### 1.3 Desideratum for a theory of normativity

The phenomena of overlapping reasons suggest a desideratum for a theory of normativity. Given the distinction between reasons and contributors, a theory of

---

<sup>3</sup> On the desire-based theory, see, e.g., Schroeder (2007). On the value-based theory, see, e.g., Maguire (2016). On the reasons-as-ought-explainers view, see Broome (2004). On the reasons-as-evidence view, see Kearns and Star (20xx). For the reasons-as-favorers view, see Scanlon (1998, 2014), Dancy (2004), and Parfit (2011).

normativity should not only tell us when some consideration is a reason, but also tell us when a consideration is a contributor—that is, when a consideration makes a distinct contribution of normative weight.

There are various reasons why this is important. The first is simply that this is an important structural feature of the contributory normative domain that calls out for explanation. Another is that for determining what an agent ought to do, what matters most are the contributors. This is because what the agent ought to do is a function of the weight that has been contributed for and against her options, not (or at least not most directly) of the considerations that are reasons. Cases of overlap show that specifying what the reasons are, and even what the weights of the individual reasons are, underdetermines what the agent ought to do, since sometimes multiple reasons make a single contribution and sometimes a single reason makes multiple contributions. When two reasons overlap, for example, it is not as if *really* there is just one reason, or one of the reasons has no weight. It is rather that they don't contribute weight separately. A third reason is that an account of contributors will help us systematize the explanatory relationship between the contributory considerations and the all things considered status of an option. A fourth, looking ahead, is that attending carefully to the distinction between reasons and contributors will provide evidence for a substantive account of reasons—and perhaps an analysis.

In the next section we'll make a preliminary case for the claim that in order to satisfy this desideratum, a theory of normativity needs to specify the grounds of reasons. The grounds of a reason provide the metaphysical explanation for why it is a reason.

## 2 Grounds of reasons

In the last section we argued that a theory of normativity needs to provide an account of when some consideration makes a distinct contribution of normative weight. If what we've said so far is correct, the answer cannot be as simple as “if R1 and R2 are different facts, then they are different contributors”, as the balancing metaphor may suggest. The specific account will depend on one's more general theory of normativity, so we will not provide a full-fledged view here. Rather, we will argue that theories that can appeal to grounds of reasons—specifically, bottom up theories of reasons—can give an account of overlap.

When some consideration is a reason for or against an option, we might ask *why*. We answer this question by providing grounds for the reason. For example, why is the fact that there are doughnuts at the bakery a reason for you to go? Plausibly, because you want a doughnut. This fact about your desires grounds the reason for you to go to the bakery, constituted by the fact that they have doughnuts there. Why is the fact that the child in the pond is drowning a reason for you to save him? One plausible answer is that the child's life has value. Why is the fact that I promised to pick you up at the airport a reason to pick you up? Because I have a duty to keep my promises. These particular answers are controversial, but the general point is that there are lots of familiar answers to the question *why* some fact is a reason to act in some way. That is,

there are familiar theories of the grounds of reasons (in this paragraph, we've seen desires, values, and duties).

These are what we call bottom-up theories of reasons. There is some other property—promoting a desire, promoting or respecting a value, fulfilling a duty—that underlies and explains when and why some consideration is a reason for or against a given action. As we'll argue in this section, this seems to be exactly what we need to sort out our problems with overlap. The central cases of overlap are when distinct considerations don't make distinct contributions, and when one consideration makes more than one contribution. In both cases we can explain why this is so by appealing to the grounds of the reasons.

Distinct considerations can both indicate the same one way to satisfy a single desire, or the same one way to promote a single value, or the same one way to fulfill a duty. Or distinct considerations can indicate different ways of satisfying the same desire once, promoting the same value once, or fulfilling the same duty once. In any of these cases, these different considerations would fail to contribute separate weight, and pointing out these facts seems sufficient to explain why. Alternatively, a single consideration can indicate the same one way to satisfy multiple desires, promote or respect multiple values, or fulfill multiple duties. In these cases, the same consideration would contribute weight (at least) twice. Again, pointing out these facts about desire satisfaction, value promotion or respect, or duty fulfillment would be sufficient to explain why.

First, consider *alternative satisfiers*. Pushing button A would vaccinate the child from the horrible disease. That's a reason to do so. Pushing button B would vaccinate the child from the horrible disease. That's a reason to do so, as well. But these reasons overlap because the underlying value in question is the same: helping this child to avoid the horrible disease. If pushing the button twice would help a little more than once, there would be two reasons. But it wouldn't, so there aren't. So we see that by specifying the grounds of the reasons, we can sort out when the reasons overlap, and when they don't.

In *alternative grain*, the fact that Joe would like a doughnut is a reason to go to the bakery. So is the fact that one of your friends would like a doughnut, and, so is the fact that someone would like a doughnut. Do these contribute separately to your going to the bakery? That depends on *which* of your friends you are referring to, and who is the someone. When they are all the same person, these reasons overlap.<sup>4</sup> Again, we see that the grounds of the reasons help to explain overlap.

As these cases show, sometimes we have to know a lot about the details of the case to figure out whether two reasons make overlapping contributions. The fact that the dog needs exercise is a reason to walk her in the forest now, and a reason to walk her in

---

<sup>4</sup> An important complication: The question here is not 'why is the fact that Joe would like a doughnut a reason to go to the bakery?' – to which the answer might be: you promised to get Joe something he would like. It is 'what makes it the case that someone would like a donut' – to which the answer is: 'the fact that Joe would like a donut'.

the forest an hour from now. Is this one reason or two? Well, that depends on how much exercise the dog needs, wants, or would benefit from.

Finally, consider the cases of *alternative grounds*, in which one reason makes multiple contributions, as in the doughnut case above. That they sell doughnuts at the bakery is a reason to go. But there seem to be two explanations or grounds for this reason: that you want a doughnut and that you want to buy your mother a doughnut. By appealing to these separate grounds, then, we can explain the sense in which this one fact—that they sell doughnuts at the bakery—is a reason to go to the bakery twice over. That is, appealing to the two grounds of the reason lets us explain why it makes multiple contributions.

So here is our hypothesis. Giving an account of overlapping reasons—of when a consideration makes a unique, distinct contribution of normative weight—requires specifying the grounds of the reasons. When one reason has multiple grounds, it makes multiple contributions, one for each ground. When multiple reasons for the same action have the same ground, their contributions overlap. When multiple reasons for distinct actions make overlapping contributions, we can often explain this by pointing out that the reasons have the same ground—indicate why acting in the relevant way would satisfy the same desire, promote or respect the same (instance of the same) value, or fulfill the same (instance of the same) duty.

The particular account of what a contributor, as opposed to a reason, is will depend on one's more general theory of normativity. Our hypothesis at this point is just that whatever account we give of contributors, or of when reasons overlap, we will need to appeal to the grounds of reasons.

### **3 The Challenge for Reasons Fundamentalism**

In this section, we'll develop our argument against Reasons Fundamentalism. After sketching the *prima facie* challenge, we'll consider what we take to be the most promising reply. This will involve appealing to a distinction between basic, or non-derivative, reasons and derivative reasons. We argue that this strategy faces serious problems.

#### **3.1 The challenge**

Here is an initial gloss on Reasons Fundamentalism, due to T.M Scanlon:

...any attempt to explain what it is to be a reason for something seems to me to lead back to the same idea: a consideration that counts in favor of it. "Counts in favor how?" one might ask. "By providing a reason for it" seems to be the only answer (Scanlon (1998), p. 1)

The challenge for this view is to account for the fact that not all reasons count in favor separately. In light of the phenomena of overlapping reasons, it is more natural to answer, "Counts in favor how?" with "By contributing weight". But not all reasons



contribute weight separately. Consequently, this answer does not provide the “same idea” as “By providing a reason” in reply to the question—or at least, more needs to be said. The challenge for Reasons Fundamentalism is to account for the distinction between reasons and contributors without undermining the consistency or plausibility of the claim that reasons are fundamental—that there’s no more to be said about what reasons are, other than that they count in favor (or, we would add, against).

Reasons Fundamentalism holds that at least some reasons facts (facts of the form [R is a reason for A]) do not obtain in virtue of any other facts. We will add, partly as a matter of stipulation, but also in keeping with the views of Reasons Fundamentalists, that Reasons Fundamentalism entails *Reasons First-ism*. This is the thesis that all normative and evaluative facts—facts about duties, values, and so on—are to be explained in terms of facts about reasons.<sup>5</sup>

The challenge for Reasons Fundamentalism, then, is to give an account of when two reasons make the same contribution, and when one reason makes more than one contribution, without appealing to any facts that explain why the reasons involved are reasons. In the previous section, we argued that one nice way to explain these cases of overlap is to appeal to the grounds of the reasons—for example, the desires, values, or duties underlying the reasons. But again, what is distinctive about Reasons Fundamentalism is that reasons do not have grounds like this—and in fact, facts about reasons are meant to ground at least values and duties. So the Reasons Fundamentalist cannot adopt the strategy we suggested in the previous section. In this section, we will consider another strategy.

### 3.2 The Basic Reasons Strategy

Given our characterization of Reasons Fundamentalism above, we can see that it is consistent with the view that some reasons *are* explained, or *do* have grounds. What’s important is that these grounds are *themselves* facts about reasons.<sup>6</sup> That is, reasons can be explained by other reasons, consistently with Reasons Fundamentalism. That observation is the starting point for the *Basic Reasons Strategy*, which we will develop on behalf of Reasons Fundamentalism. This strategy holds that when reasons R1 and R2 overlap, that is because there is a more basic reason that grounds both R1 and R2. The contributors are these basic reasons. Derivative reasons like R1 and R2 overlap because they only contribute, if at all, in virtue of their relationship to the basic reason. It makes sense to talk about the weight of these derivative reasons, but (i) they only have this weight in virtue of their relationship to the basic reason, and (ii) they do not contribute

---

<sup>5</sup> See Scanlon (1998, 2014), Parfit (2011), Dancy (2004), Skorupski (2010). Reasons First-ism does not entail Reasons Fundamentalism, since one may hold that all normative and evaluative facts are explained by facts about reasons, but hold that facts about reasons are themselves explained by other facts; see, for example, Schroeder (2007), who accepts Reasons First-ism, but holds that reasons are explained by desires.

<sup>6</sup> It’s also compatible with the view that the most direct grounds of some reasons are facts of other kinds—e.g., facts about values or duties—as long as these facts are themselves ultimately explained by facts about reasons. We will focus on the strategy of grounding some reasons in other reasons.

this weight, at least not separately from the basic reason, for or against the relevant options.<sup>7</sup>

For example, the fact that they sell doughnuts in the shop is a reason to go to the shop, and so is the fact that you are hungry, and so is the fact that you really want a doughnut. Plausibly, the fact that you really want a doughnut explains why the fact that they sell doughnuts in the shop is a reason to go, and why the fact that you are hungry is a reason to go. In that case, the Basic Reasons Strategy will explain why these reasons overlap as follows. The basic reason here is that you really want a doughnut—this is the contributor. That they sell doughnuts at the shop and that you are hungry only contribute weight, if at all, in a derivative way, due to their relationship to the basic reason. So they do not contribute separately from each other, or from the basic reason. More generally, when one reason explains another, or two reasons are explained by the same basic reason, they do not contribute weight separately.

This strategy can be extended to account for cases of what we called *alternative grounds*. Recall that these are cases in which we have one consideration—one reason—that seems to make multiple distinct contributions. The Basic Reasons Strategy will hold that in such cases, the single consideration constitutes two distinct derivative reasons, grounded in two distinct basic reasons. For example, the fact that you really want a doughnut is a basic reason to go to the bakery, as is the fact that you want to buy your mother a doughnut.<sup>8</sup> That they sell doughnuts at the bakery constitutes two derivative reasons to go, one grounded in the first fact and one grounded in the second. Since these basic reasons are the contributors, we explain the sense in which the derivative reason seems to contribute twice. But in fact it doesn't contribute—or at least doesn't make its own contribution—at all.

This strategy is certainly *prima facie* plausible. A natural reaction to many of the cases of overlapping reasons above is that the reasons cited are derivative or somehow non-basic. Moreover, this kind of distinction between derivative and basic reasons mirrors other important distinctions, like that between basic and derivative value. There are also plausibly further theoretical roles for the basic/derivative reasons distinction.<sup>9</sup> We will argue, however, that this strategy is not one that will ultimately help the Reasons Fundamentalists explain overlapping reasons.

Note first that implementing this strategy requires two things. First, we need to identify a class of basic reasons. Second, we need to specify some principles telling us when some consideration is a derivative reason, explained by a given basic reason. We argue that there are problems with both parts of the strategy.

### 3.2.1 Identifying the basic reasons

---

<sup>7</sup> Compare Nair (2016a), who hypothesizes that derivative reasons for the same action do not *accrue* to support the action more strongly than either does individually.

<sup>8</sup> It is unlikely that, in the end, these will count as basic reasons once we develop a substantive normative theory. But we'll stick with these for illustrative purposes.

<sup>9</sup> See, for example, Nair (2016b) and Portmore (manuscript).

The first task is to identify a class of basic reasons, which will serve as the contributors. The first challenge is one of rationale: we would like an explanation for why *these* are the basic reasons rather than *those*—some theory of what distinguishes the basic reasons from the non-basic ones.

Since reasons are just facts (or so we're assuming here), we might try appealing to an independent conception of fundamentality for facts. Some facts are more fundamental than others. So we might think that if the fact that constitutes reason R1 is more fundamental than the fact that constitutes reason R2, then R1 is a basic, or at least more basic, reason.

But this hypothesis is incorrect. Sometimes less fundamental facts are more fundamental, or basic, reasons. Suppose you promise to make someone smile. The fact that squirting water at Jack would make Jill smile is a reason to squirt the water, but this is a less fundamental reason than the fact that squirting water would make someone smile, even though the latter *fact* (the fact that is the reason) obtains in virtue of the former. So giving an account of basic reasons and of the explanatory relations between reasons won't be quite so simple.

One option is to rely on substantive normative theorizing to identify the class of basic reasons. For example, if as a first-order normative fact, hedonism is true, then the basic reasons will be facts about pleasure promotion. The Reasons Fundamentalist can in principle even accept, as a matter of substantive normative fact, that all the basic reasons are about promoting intrinsically valuable states of affairs, or about promoting things that are desired for their own sake. In this case, the Reasons Fundamentalist would agree with the value-based or desire-based theorist about what reasons there were, but would just disagree about whether reasons were to be *analyzed* or *explained* by values or desires.<sup>10</sup>

The problem with this strategy is not that it is incoherent. Rather, it threatens to undermine the plausibility or rationale for Reasons Fundamentalism. If, for example, the basic reasons turn out to be, again, as a matter of substantive normative fact, facts about desire-promotion, we'd like to know *why*. The desire-based theory of reasons, according to which the property of being a reason is analyzed in terms of desire-promotion, seems to have an easy explanation. Reasons Fundamentalism, on the other hand, does not. Generally, if we have some sort of true generalization about the basic reasons—e.g., that they are all about pleasure, desire, or value promotion—we look to be on our way to an *analysis* of reasons, which the Reasons Fundamentalist cannot accept.

This kind of argument against the Basic Reasons Strategy won't work if the substantive normative theory is some form of Rossian pluralism. In that case, the basic reasons would be, or be about, Rossian *prima facie* duties (duties of gratitude, fidelity, beneficence, and so on). But in this case, we wouldn't have made much progress on the question of why *these* are the basic reasons rather than *those*. Ross (1930) was happy, or at least willing, to accept this kind of brute pluralism, but many commentators have been with us in finding this unsatisfying.

---

<sup>10</sup> Compare Scanlon (2014), Chapter 2, on the consistency of a *normative* desire theory of reasons (like the position described in the text) with Reasons Fundamentalism.

So there is a challenge identifying the class of basic reasons. We want to press a related challenge that is more *structural* in nature. The challenge is to say not *which* reasons are the basic ones, but what those reasons are like, structurally. Two approaches naturally present themselves. First, we may identify the basic reasons as some class of *particular* reasons—particular facts that are reasons for particular actions. For example, the fact that A-ing would produce pleasure may be a basic reason to A, and explain other, derivative reasons—e.g., reasons to do things that would facilitate A-ing. This approach would hold that what the basic reasons are can vary from case to case, since the particular facts and the particular options can of course vary from case to case. Second, we may identify the basic reasons as what Rosen (2016), following Scanlon (2014), calls *pure reasons*, or *reasons principles*. These are principles like ‘That some action would relieve pain is always a reason to perform that action’. These pure reasons would then explain other, derivative reasons. For example, this pure reason would explain why the fact that John has a headache is a reason to give him an aspirin, since aspirin relieves headaches, which are painful.

Take our case of *alternative satisfiers*. You can push button A or push button B. Either way, some particular child will be given an important vaccination. There are no other normatively relevant factors in the case. Do you have a reason to press button A? Yes: the fact that pressing button A would cause the child to be vaccinated is a reason to press it. Do you have a reason to press button B? Again, yes: the fact that pressing button B would cause the child to be vaccinated is a reason to do so. But these reasons do not seem to make distinct contributions; they overlap. The Basic Reasons Strategist claims that the two facts mentioned above—that pressing button A would cause the child to be vaccinated and that pressing button B would cause the child to be vaccinated—are derivative reasons. So there must be some basic reason that is the real normative contributor.

Start with the reasons principles version. Suppose there is just the one relevant pure reason: that an action reduces risk of serious disease is a reason for that action. (More plausibly, this will be explained by some more fundamental pure reason to reduce harm or ill-being.) Take the two facts: the fact that pressing button A would vaccinate the child and the fact that pressing button B would vaccinate the child. Both are explained by the same pure reason. So, you might think, the real contributor here is the pure reason, and the other reasons overlap because they are explained by this one pure reason. But in fact this will not work. That explanation would overgeneralize. For suppose that pressing button A would vaccinate one child and pressing button B would vaccinate another child. In that case, you have two non-overlapping reasons to press the buttons. But both are explained by the same pure reason, namely ‘that an action reduces risk of serious disease is a reason for that action’.

It doesn’t help to turn to the first version of the Basic Reasons Strategy, explaining the derivative reasons by other particular, but basic, reasons. The fact that pressing button A would vaccinate the child is a reason to press it. And the fact that pressing button B would vaccinate the child is a reason to press it. These reasons overlap, so this version of the Basic Reasons Strategy must point to a particular, basic

reason that explains both. The problem is that a particular reason, whether it is basic or not, is a reason for a particular action. But in this case we have two actions: pressing button A and pressing button B. So though there may be more basic reasons that could explain each of these overlapping reasons, there will not be a single basic reason that explains both. So this does not seem to be a promising strategy for explaining why these two reasons overlap.

By contrast, bottom-up theories, like the value-based theory or the desire-based theory, can explain this data nicely. In the original, one-child case, there is just one valuable or desired outcome—that *this* child is vaccinated. Since there's just one valuable or desired outcome, even though two different actions would bring it about, there's just one contribution. In the two-child case, on the other hand, there are two valuable or desired outcomes, so there will be two contributions.<sup>11</sup>

### 3.2.2 Explanatory relations among reasons

Implementing the Basic Reasons Strategy for explaining overlapping reasons requires (i) identifying a class of basic reasons, and (ii) specifying principles telling us what the derivative reasons are, which are explained by a given basic reason. We've just argued that there are some challenges carrying out the first part of the strategy. Now we will argue that there are also challenges carrying out the second part. In particular, we'll argue that capturing the intuitive explanatory relations among reasons itself seems to require appealing to the grounds of reasons, which, again, is not something the Reasons Fundamentalist will want to accept.

Let's start with a simple thought. If you have a reason R1 to make Mildred happy, and giving Mildred flowers is a necessary and sufficient means to making Mildred happy, then you have a reason R2 to give Mildred flowers. There might well be plenty of reasons not to give Mildred flowers, or to do other things with your time and money. But there is at least one reason to give Mildred flowers, and it's explained by R1. Some evidence that this reason R2 is indeed explained, in some important sense, by R1 is that the weight of this particular reason R2 is not greater than the weight of the reason R1 to make Mildred happy. Again, there might be some further reason R3 to give Mildred flowers – perhaps you promised Alfred that you would. So the total weight of reason to give Mildred flowers might be greater than the weight of R1. But there is one particular reason to give Mildred flowers the weight of which is not greater than the weight of the reason R1.

A full version of the Basic Reasons Strategy will likely need to appeal to several different kinds of explanatory relations to capture the full range of derivative reasons.

---

<sup>11</sup> A possible solution here, in the case of alternative satisfiers, is to follow Skorupski (2010) and Snedegar (forthcoming) in relativizing reasons to *choice sets*. This is to make reasons sensitive to the ways in which the options are individuated. So in the vaccination case, what we have basic reason to do is to press a button. This is a *coarse-grained* action that does not distinguish between pressing A or pressing B. There is much work to be done spelling this out in a way that's consistent with Reasons Fundamentalism. But even if this works for cases of alternative satisfiers, it is not clear how it will help with, e.g., alternative grounds.

We will focus here on means-end transmission principles, where we have a derivative reason for a means to some end which is explained by a basic reason (or at least a *more* basic reason) for the end.

For concreteness, consider the simple, existential transmission principle that when there is a reason for an end there is a reason for any necessary and sufficient means to that end. Return to the case above: you have a reason to make Mildred happy. Giving Mildred flowers is necessary and sufficient for making her happy, so you have a reason for the means, giving Mildred flowers. Start with the question: what is the reason to take the means? The fact that you promised to give Mildred flowers is a reason for the means—giving Mildred flowers—but it doesn't satisfy the intuitive relationship between reasons for ends and reasons for means. The idea is that there is a reason for the means *because* there is a reason for the end. There is no explanatory relationship between the promissory reason and the "end-given" reason to make Mildred happy. There may be lots of reasons to take the means: you really enjoy giving Mildred flowers, you want to give her flowers, it would realise tremendous value to do so. All of these reasons satisfy the existential principle. But none of them satisfies the intuitive transmission *from* the reason to make Mildred happy through the fact that giving her flowers is a means to her happiness *to* a reason to give her flowers.

The same kind of problem afflicts a slightly more sophisticated approach: namely to maintain that the instrumental reason is the very conjunctive fact that some instrumental relationship holds between the means and the end and that there is a reason for the end. Using our example: the instrumental reason would be the fact that you have a reason to make Mildred happy and that giving her flowers is a sure-fire way to make her happy. But there will be cases of overlapping reasons that show that even this more sophisticated existential principle fails to account for the intuitive transmission from the end-given reason to the reason for the means. For suppose there are two reasons to make Mildred happy: that her happiness is valuable and that you promised John you would. In this case, you have (at least) two reasons to give Mildred flowers, one arising from each reason to promote the end in question, namely making Mildred happy. The more sophisticated existential principle only gets us one reason: the conjunctive fact [that *there is* a reason to make Mildred happy and giving her flowers would make her happy]. This is an example of what we called *alternative grounds* overlapping reasons above, where we have one consideration that seems to make multiple contributions.

Think about what is intuitively going on in this case in which we have one means to two ends. Intuitively the reasons for the means are these: the fact that giving Mildred flowers is a way to keep your promise, and that giving Mildred flowers would make her happy. Ignore overlap challenges bearing on the reason for the end. To respond to the overlap challenge specific to instrumental transmission we need to specify the non-overlapping reason to promote the end. Moreover it is the weight of each of these 'end' reasons that constrains the weight of each of these 'means' reasons. Importantly, it seems that these two 'means' reasons may have different weights, depending on the weight of the corresponding 'end' reasons. If the only 'means' reason we have is the one

conjunctive fact [that there is a reason to make Mildred happy and giving her flowers would make her happy], it is hard to see how to capture this difference in weight.

The problem is not merely that the principles we've considered so far are merely existential. Suppose we make the transmission principle much stronger, so that it tells us not just that *there is* a reason for the means, but that that reason is the very fact that is the reason for the end. So, suppose that the fact that you love Mildred is a reason to make her happy. Since giving her flowers is (we're assuming) necessary and sufficient for making her happy, this strong principle would tell us that this very same fact—that you love Mildred—is a reason to give her flowers. But even this does not fully capture the sense in which this fact is a reason for the means *because* it's a reason for the end. To see this, suppose you've promised to give someone you love flowers. Then the fact that you love Mildred is a reason to give her flowers. So we could have a case in which the fact that you love Mildred is a reason to make her happy, giving her flowers is necessary and sufficient for making her happy, and the fact that you love Mildred is a reason to give her flowers, but still not have captured the thought that the reason for the means is a reason *because* of the reason for the end.

Giving a satisfactory transmission principle requires capturing the sense in which the reason for the means is explained by the reason for the end. We've seen that existential principles, saying that *there is* a reason for the means when there's a reason for the end fail to do this. We've also seen that a stronger principle that says that the reason for the means is the very same fact as the reason for the end do not suffice. Now we will argue that doing better requires specifying the grounds of the transmitted reason. If this is right, then we have a further challenge for the Reasons Fundamentalist's strategy of explaining overlap. This Basic Reasons Strategy, recall, is to specify a class of basic or fundamental reasons along with a set of explanatory principles that determine the class of derivative reasons, given the basic reasons, and hold that overlapping reasons are just derivative reasons. Our argument here challenges this strategy because the Reasons Fundamentalist cannot provide grounds for reasons (at least not grounds which are not themselves reasons), and so, if our argument is successful, they cannot specify satisfactory explanatory principles.

We suggest that a transmission principle needs to take this form: *If R1 is a reason for E, grounded in G, and M is a necessary means for E, then there is a reason R2 for M, grounded (partially) in G.* We leave it an open question what other conditions must be written into the principle, whether R2 is identical to R1, and so on. The important points for our purposes now are that the principle references the grounds of the reasons, and that the grounds of the reason for the end is the same as the grounds of the reason for the means.

Apply this to cases like those above. That you love Mildred is a reason to make Mildred happy, and giving her flowers is necessary and sufficient for doing so. Thus, there should be a transmitted reason to give her flowers. We saw above that this mere existential claim is too easy to satisfy: suppose, for example, that you promised to give Mildred flowers. Then it's true that there is a reason to give Mildred flowers, but this isn't the transmitted reason we're looking for. A natural explanation for why is that the

*ground* of this reason differs from the ground of the reason for the end, of making Mildred happy, that you love her. The ground of the reason for the end, plausibly at least, is that making those you love happy properly respects the value of loving relationships. The ground of the reason for the means of giving Mildred flowers is that you've promised to do so.

On the other hand, a reason to give Mildred flowers that's explained (at least in part) by the value of loving relationships *would* seem to be the kind of transmitted reason we're looking for. For concreteness, consider the kind of strong transmission principle we considered above, which says that the reason for the means is the very same fact as the reason for the end. The fact that you love Mildred is a reason to make her happy, explained by the value of your loving relationship with Mildred. Giving Mildred flowers is necessary and sufficient for making her happy. So the fact that you love Mildred is a reason to give her flowers, also explained by the value of your loving relationship with her. Insisting that the ground of the reason for the means is the same as the ground of the reason for the end at least takes us much closer to capturing the sense in which the reason for the means is transmitted from the reason for the end.

We've just argued that giving satisfactory transmission principles requires specifying the ground of the reasons involved. Combined with our previous argument that specifying a class of basic reasons, which are meant to explain all other reasons, pushes us towards the view that reasons have grounds, we have a kind of unified challenge to both parts of the Basic Reasons Strategy, as employed by the Reasons Fundamentalist.

#### **4 Conclusion**

The phenomena of overlapping reasons are, we think, interesting and important features of the contributory normative domain that call out for explanation. We've suggested here that explaining these phenomena will best be accomplished by a bottom-up theory of reasons, according to which reasons have *grounds*. In particular, we've argued that a bottom-up theory does better than Reasons Fundamentalism, which is the view that (at least some) reasons do not have grounds.

#### **References**

- Broome, J. (2004). Reasons. In xx, xx, and xx, eds., *Reason and Value: Themes from the Moral Philosophy of Joseph Raz*. Oxford University Press, pp. xx-xx.
- Dancy, J. (2004). *Ethics Without Principles*. 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 J. C. Marek and M. E. Reicher, eds., *Experience and Analysis: Proceedings of the 27<sup>th</sup> Annual Wittgenstein Symposium*. ÖBV and HPT, pp. 387-394.

Horty, J. (2012). *Reasons as Defaults*. Oxford University Press.

Kearns, S. and Star, D. (2009). Reasons as evidence. *Oxford Studies in Metaethics*, Vol. 4. Oxford University Press, pp. 215-242.

Magure, B. (2016). The value-based theory of reasons. *Ergo* 3(9).

Manne, K. (2016). Democratizing Humeanism. In E. Lord and B. Maguire, eds., *Weighing Reasons*. Oxford University Press, pp. 123-140.

Nair, S. (2016a). How do reasons accrue? In E. Lord and B. Maguire, eds., *Weighing Reasons*. Oxford University Press, pp. 56-73.

Nair, S. (2016b). Conflicting reasons, unconflicting 'ought's. *Philosophical Studies* 173:629-663.

Parfit, D. (2011). *On What Matters*. Oxford University Press.

Portmore, D. (manuscript). Maximalism vs. omnism about reasons. Unpublished manuscript, Arizona State University.

Rosen, G. (2016). Scanlon's modal metaphysics. *Canadian Journal of Philosophy*.

Ross, W.D. (1930). *The Right and the Good*. Oxford University Press.

Scanlon, T. (2014). *Being Realistic about Reasons*. Oxford University Press.

Scanlon, T. (1998). *What We Owe to Each Other*. Harvard University Press.

Schroeder, M. (2007). *Slaves of the Passions*. Oxford University Press.

Skorupski, J. (2010). *The Domain of Reasons*. Oxford University Press.

Snedegar, J. (forthcoming). *Contrastive Reasons*. Oxford University Press.

Wedgwood, R. (2009). Intrinsic value and reasons for action. *Philosophical Issues* 19:321-342.