

OVERLAPPING REASONS

Justin Snedegar and Barry Maguire¹

Draft as of 3 August 2017.

Abstract

It is a commonplace that reasons for and against options interact to explain what you ought to do. However, it has not been widely recognized that some reasons – perhaps the majority – fail to make any difference at all in this explanation. This is because many reasons *overlap* with each other. In an overlapping set of reasons, only one reason, the contributor, plays an explanatory role in determining what you ought to do. We argue that ‘bottom-up’ theories of reasons – those which explain reasons in terms of some kind of facilitation relation to some normatively significant property of the action or its outcome – are well-placed to explain overlap and a range of closely related phenomena: the transmission of reasons from ends to means, the transmission of weight from ends to means, and the sometimes transmission of reasons from means to means. We argue that Reasons Fundamentalism – the view that some reasons facts do not admit of explanation – has serious difficulties explaining these phenomena. We don’t argue that Reasons Fundamentalism *cannot* account for these cases. Much of the constructive work of the paper is in the service of Reasons Fundamentalism. But, in the end, an abductive argument supports explanations of reasons from the bottom up.

1 Overlapping Reasons

1.1 One reason, one contribution

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

¹ Both authors equally to blame.

do.² The idea is that for each option open to an agent, there are some considerations, or *reasons*, in its favor, and some against. These reasons for and against the options interact to determine some fact about what the agent ought to do, all things considered.

On many views, the reasons interact by contributing normative weight for or against the relevant option.³ 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.

It is widely held that there are limitations to this model. Since Jonathan Dancy's pioneering work it has been clear that reasons are subject to both conditionality and modification.⁴ But, this feature of reasons is not really a limitation of the balancing model, since the relevant analogues to the marbles are facts about reasons *taking into account* all the relevant facts about modifiers and conditions. A would-be reason that is not 'enabled' never makes it into the pan; and the 'weight' of a reason is its weight once it has been duly attenuated or intensified. Another feature often thought to be a limitation of the simple weighing model is that even though reasons for the same option seem to compete and combine, they may not do so *additively* – unlike the weights of marbles.⁵

We are interested in a set of phenomena that put pressure on a different, and not widely recognized, assumption of the simple weighing model, which we call the 'one reason, one contribution'

² This is to trace this line of thought back to W.D. Ross's (1930) *The Right and the Good*.

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

⁴ The seminal account is Dancy (2004).

⁵ A third feature of reasons talk that isn't really a limitation of this model is that the facts we *call* reasons are often the salient parts, or *representatives*, of a larger whole (Broome (2013), Fogal (2015)). For instance, you might felicitously say that your reason for going home is that the clouds are getting very dark; the contributor you thereby pick out is that by going home now you will be able to cover your furniture before the storm.

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. The ‘one reason, one contribution’ assumption is just that, however we combine these contributions, each reason makes a single, distinct contribution. What unites the cases of overlapping reasons is that they are cases in which this assumption appears to be false. This is different from the assumption that the weights of reasons combine in an additive way, since this says nothing about how to combine the contributions made by the reasons. However, we conjecture that when contributors, not reasons, are recognized as the proper analogues of the marbles in the pan, many putative failures of additivity will be resolved.

We are not here to defend the simple weighing model. Rather, framing things in this way helps to make clear the important features of the cases we are interested in. The challenge for a theory of reasons is not to somehow fix up the balancing metaphor, but rather to show how to explain cases of overlap. Individuating reasons is easy. Reasons are facts – specifically, the facts that favor or disfavor the options in question. So different facts are different reasons, and a single fact is a single reason. But this is not, in general, true of contributors, as our cases will illustrate.

We will argue that reflection on cases of overlap provides some support for a class of views, which we’ll call ‘bottom up’ theories of normative reasons. According to bottom up theories, reasons facts – minimally, facts of the form ‘this fact is a reason for that option’ are grounded in facts about some kind of facilitation relation between the option and some normatively significant property of the action or an outcome. We will mostly have in mind *desire-based* theories of reasons (e.g., Schroeder (2007), Manne (2016)) and *value-based* theories of reasons (e.g., Wedgwood (2009), Finlay (2014), Maguire (2016)).⁶ We’ll show that these kinds of views can offer nice explanations of the phenomena in question. The main competitor to bottom up theories that we’ll consider is what we’ll call *Reasons Fundamentalism*, as defended by Scanlon (1998, 2014), Parfit (2011),

⁶ Our arguments may extend to some value *honouring* theories (cf. Pettit 1989), some *duty-based* theories of reasons such as (with a bit of interpretation) Ross (1930), and some *fittingness*-based accounts such as Howard (forthcoming). We remain agnostic about this.

and others. According to Reasons Fundamentalism, some facts about reasons admit of no further explanation. Facts about reasons are, in this sense, fundamental. Much of the constructive part of the paper will consist in showing how a Reasons Fundamentalist might explain the cases of overlapping reasons. In the end, we do not think it is impossible for them to do so. But we also think that bottom up theories offer more satisfying explanations.

The three goals for this paper, then, are the following: to introduce the overlapping phenomena as such and the associated distinction between reasons and contributors; to show how a Reasons Fundamentalist might explain these phenomena; and to suggest that bottom-up theories of reasons explain these phenomena more handsomely.

1.2 Overlapping Reasons: Some cases

First there are what we'll call cases of *alternative grain*. The fact that giving lemon cake with cherry frosting to Pearle the shorthaired German pointer would make her a very happy dog is a reason to do so. So is the fact that giving cake to a dog would make the dog very happy. So is the fact that giving a delicious treat to a dog would make that dog happy. These are distinct facts, and they are reasons, so they are distinct reasons. But these distinct reasons presumably don't add weight to each other. Suppose that the first of these reasons were 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.

A related class of cases involve *alternative referrers*. Suppose you can bring your mother, Pam, some flowers. This will make her happy. So consider the following reasons claims: 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 reasons 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 cases involve distinct reasons making an indistinct 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—either pre-theoretically, or on any of the 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 "for-giving 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.⁷

Next, there are cases of what we will call *alternative satisfiers*. Like the previous examples, these are cases in which we have multiple distinct reasons that do not seem to make distinct contributions to what the agent ought to do. An intuitive earmark for this is that when one of the reasons is 'satisfied'—when the action it is a reason

⁷ For the idea that reasons count in favor, see Scanlon (1998), Parfit (2011), Skorupski (2010). For the reasons as explanations of oughts view, see Broome (2004, 2013). For the reasons as evidence view, see Kearns and Star (2009). For desire-based views, see Schroeder (2007) and Manne (2016). And for value-based views, see Wedgwood (2009), Finlay (2014), and Maguire (2016).

for is performed—the other, overlapping reasons are no longer in force, but not because they've been undermined. We will focus on a case in which there are different reasons for an agent to perform different actions, but which do not make distinct contributions. Suppose that pressing button A would ensure that Bill receives an important vaccination. Pressing button B would also ensure that Bill receives the important vaccination. You can press either button, or both buttons. But there is only one vaccination available. The fact that pressing A ensures that Bill receives the vaccination is a reason to press A, and the fact that pressing B ensures that Bill receives the vaccination is a reason to press B. Similarly, the fact that pressing a button ensures that Bill receives the vaccination is a reason to press a button.

But all three of these reasons seem to overlap: they do not make distinct contributions to what you ought to do. Some evidence that they do not make distinct contributions is that the reasons do not compete with one another—you will not, for example, be torn between pressing A and pressing B, especially since you could just press both. And you certainly will not be torn between pressing A and pressing a button. But the reasons do not combine, either: your reason to press both buttons is not stronger than your reason to press A or your reason to press B. If the reasons neither compete nor combine, that is at least some evidence that they do not make distinct contributions, much like the other examples of overlapping reasons we described above.

There are some similarities between this case and certain ‘Buridan’ cases. In traditional Buridan cases, there are two options that are equally desirable. Buridan’s donkey is torn between these two options. This is clear when we compare the traditional Buridan case with another predicament described by Montaigne: “And who should place us betweene a bottle of wine and a gammon of bacon, with an equall appetite to eat and drinke, doubtlesse there were noe remedy, but to die of thirst and of hunger.”⁸ Equal support, as in these cases, is not the same as the overlap in our cases, where the support for the two options is indistinct.⁹ The current task for a theory of reasons, then, is to explain why the reasons in our cases do

⁸ Montaigne 1603 / 1924.

⁹ Our two buttons case is more like the case due to Bratman (1987, 22), in which he must decide on one of two equally desirable routes to San Francisco. The goal of getting to San Francisco is the analog of getting Bill the vaccination, and the two equally desirable routes are analogs of the two buttons.

not make distinct contributions. If we can do this, then we will be able to give a theoretical explanation of this intuitive difference between equal contributions and indistinct contributions. At this point, though, it is enough for our purposes to make the case that there is something here to be explained.

Another response to the two buttons case is to say that you really only have a reason to press *a* button, but no reason to press button A and no reason to press button B. If that is right, then we don't have a case of overlap, since we don't have distinct reasons failing to make distinct contributions. But we are very sceptical about this. On many ways of understanding reasons, the fact that pressing A will ensure that Bill receives an important vaccination is a reason to press A. As before, this fact counts in favor of pressing A, it can figure in an explanation of what you ought to do, it is evidence that you ought to press A, it explains why pressing A would promote a valuable or desired outcome, and so on. Importantly, this response—that in fact there is no reason to press A and no reason to press B, but only reason to press a button—strikes us as a very theory-driven thing to say.¹⁰ Part of what is driving it in this case, we hypothesize, is the feeling that the reason to press A and the reason to press B aren't really distinct from the reason to press a button (or from one another, perhaps). But if we can respect the very plausible thought that in fact there is a reason to press A and a reason to press B but explain this feeling of non-distinctness, then we need not make the theory-driven claim that really there's just reason to press a button. A more concessive version of this response to our case is to say that the reason to press A and the reason to press B are merely derivative reasons, whereas the reason to press a button is basic, or at least more basic. We will explore this strategy in detail in the following section.

The last class of overlapping reasons involves situations in which a single reason—one particular fact—makes multiple contributions. We call these cases of *alternative grounds* of reasons, though this name is admittedly skewed 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

¹⁰ Moreover, it is plausible that if you have a reason to do A and doing B is sufficient for doing A, then you at least often thereby have reason to do B. There are counterexamples to this kind of sufficient means transmission if it is taken to be a general principle, but it seems undeniable that often we do have reason to do things that are sufficient for other things we have reason to do. We will return to this point about reasons transmission later in the paper.

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.¹¹

One tempting response here is to distinguish the accounts of reasons and their weight, and say that the fact that the bakery sells doughnuts is reason to go, and its weight is greater in the case in which in addition to wanting one for yourself, you also promised to buy one for your mother.¹² The central problem with this strategy is that it seems to smudge over a normatively relevant distinction. The two contributions are importantly distinct. As a mark of this, different kinds of considerations could defeat them. If, for example, your mother calls and tells you 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 bakery 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.

2 Bottom-Up Explanations of Overlap

When some fact is a reason for or against an option, we can ask *why*. We can answer this question by providing grounds for the reason: the facts in virtue of which the fact that is the reason is a reason. For example, why is the fact that you can get doughnuts at the bakery a reason for you to go? Plausibly, because you want a doughnut. This fact about your desires grounds the fact that the fact that you can get doughnuts at the bakery is a reason for you to go to the bakery.¹³

¹¹ See Howard (manuscript) for similar cases.

¹² Thanks to <redacted> for this suggestion.

¹³ We distinguish between grounds and the underlying principles in virtue of which grounds ground what they ground. An underlying principle is of the form: a fact of the form [ϕ promotes S] is a reason if and only if S (or the fact that ϕ promotes S)

Why is the fact that the child in the pond is drowning a reason for you wade in to save him? One plausible answer is that saving the child is valuable. 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.

These are what we call bottom-up theories of reasons. We will illustrate this with a simple regimentation of reasons facts, of the form [some action ϕ available to an agent A stands in some relation R to some state of affairs S].¹⁴ The ground of such a reason is a fact about this fact—that S is valuable, or that S satisfies a desire of A's. 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 distinct considerations can indicate different ways of satisfying the same desire once or promoting the same valuable state of affairs once. In any of these cases, these different considerations would fail to contribute separate weight, and pointing out these facts seems at least necessary, if not sufficient, to explain why. Alternatively, a single consideration can indicate the same one way to satisfy multiple desires or to promote or respect multiple values. In these cases, the same consideration would contribute weight more than once. Again, pointing out these facts about desire satisfaction or value promotion seems at least necessary, if not 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

meets condition N (is valuable, is desired, is duty-fulfilling, etc.). We say the *ground* of the fact *that* [ϕ promotes S] is a reason to ϕ is the fact that S meets N (that S is valuable, desired, etc.). Compare Bader's (2015) distinction between sources and grounding principles. On this distinction more generally, see Dasgupta (2016).

¹⁴ Crucially, 'S' may be a state of affairs consisting in the action itself instantiating some intrinsic property, such as 'walking by the cliffs would be an enjoyable way to spend the afternoon' (cf. Wedgwood (2015)).

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. That's because there is only one valuable state of affairs at stake or only one desire-satisfaction (depending on your preferred account). 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 the forest an hour from now. Is this one reason or two? Well, that depends on how much exercise the dog would benefit from, or how much you want to give her.

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 or promote the same (instance of the same) value.

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

3.1 The challenge

Here is an initial gloss on Reasons Fundamentalism, due to T.M Scanlon (1998, p.1):

...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.

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.

3.2 The Basic Reasons Strategy

Reasons Fundamentalism holds that at least some reasons facts (facts of the form [F is a reason for A]) do not obtain in virtue of any other facts. (This is, of course, not to deny that the facts that are reasons, here the fact F, obtain in virtue of other facts.) Reasons Fundamentalism is usually paired with *Reasons First*, the thesis that all normative and axiological facts are explained in terms of facts

about reasons.¹⁵ Strictly speaking, Reasons Fundamentalism just entails the weaker thesis that no facts about reasons are explained by normative or axiological facts that are not themselves reasons facts or facts that are explained by reasons facts.¹⁶ The challenge for Reasons Fundamentalism is to account for the distinction between reasons and contributors without undermining the 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 against).

A natural thought, especially in light of the success of the bottom-up style explanation, is to appeal to this possibility that some reasons facts explain others. When one reasons fact explains—which is to say, metaphysically grounds—another, the first is *more basic* than the second; the reasons fact that explains others but is not explained by any is the basic reasons fact (or equivalently, the most basic fact). This yields the *basic reasons strategy*, which maintains that the 'basic' reasons fact is the contributor, and all reasons facts explained by that basic reasons fact overlap with it, and so with one another.¹⁷ There are two parts to this strategy, and so two tasks for the basic reasons strategist: (A) to identify the basic reasons, and (B) to account for the explanatory relations between basic reasons and derivative reasons.

3.2.1 The Basic Reasons

The first task for the basic reasons strategist is to identify the basic reasons, which serve as the genuine contributors and explain all of the other derivative reasons. This must be done in some principled way, that yields the intuitively correct verdicts about our cases.

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: for instance, the fact that John is happy obtains in virtue of the fact that John is happy to some specific degree. 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,

¹⁵ See Scanlon (1998, 2014), Parfit (2011), Dancy (2004), Skorupski (2010).

¹⁶ Reasons Fundamentalism is compatible with other primitive normative properties featuring in explanatorily isolated chains of normative facts. Also notice that Reasons First is compatible with a 'bottom up' theory of reasons, as long as the relevant grounds are non-normative (as in Schroeder 2007).

¹⁷ Here, on behalf of the Reasons Fundamentalist, we develop the strategy suggested by Nair (2016a, 2016b).

reason. But this hypothesis does not succeed. Sometimes less fundamental facts are more fundamental reasons. An example will help. Suppose you promise your sister that you will make someone smile today. 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 smile. The fact that squirting water at Jack would make Jill smile is a more fundamental fact than the other, but it is a less fundamental reason. Intuitively, the most fundamental reason is the fact that squirting water at Jack would make someone smile. It just so happens that making Jill smile is a way of making someone smile. So the fact that squirting water at Jack would make Jill smile is a reason to squirt water at Jack in virtue of the fact that squirting water at Jack would make someone smile.^{18,19}

We think that the best way for the Reasons Fundamentalist to do this is to follow Scanlon (2014) and Rosen (forthcoming) in first identifying some reasons principles, or ‘pure reasons’.²⁰ These are necessary principles about when considerations of certain kinds are reasons for actions of certain kinds. For example, the following is a candidate for a reasons principle: For all agents S, actions A, and circumstances C, if S’s A-ing in C would prevent the suffering of a sentient being, then that fact is a reason for S to A in C. We can leave it to first-order normative theory to identify the reasons principles. The basic reasons will then be true instances of the antecedent of the reasons principle. So the fact that your pressing a button would prevent the suffering of a sentient being is a basic reason to press it, in the circumstances described in the two buttons case above. This basic reason is the genuine contributor.²¹

¹⁸ We don’t think this example relies on anything quirky about promising as such. Think about your reasons to donate some percentage of your salary to effective charities. These reasons are grounded by the range of benefits to lots of people; they are not, in the first instance, grounded by particular benefits to particular people.

¹⁹ The inference in the final sentence of this paragraph relies on an almost invisible transmission principle. We’ll get to examining these presently.

²⁰ See also Roger Crisp’s (2006) notion of an ‘ultimate reason’.

²¹ 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, though at no point are there abstract principles like Scanlonian pure reasons (compare the discussion in Rosen (forthcoming)). We are not convinced the difference between generality and abstractness sustains much of a difference here. Moreover, defenders of this strategy are required to defend transmission principles from more to less general principles along the lines of those criticized later in the paper.

To explain overlap, the basic reasons strategist needs to show that the overlapping reasons stand in some pertinent relation to this single basic reason. Recall that the overlapping reasons are the following: (i) that pressing button A would vaccinate Bill is a reason to press A, and (ii) that pressing button B would vaccinate Bill is a reason to press B. Let's suppose the basic reason is (iii) that pressing a button would prevent the suffering of a sentient being is a reason to press it. In the following section we'll discuss the explanatory relations that hold between basic and derivative reasons. But let's set this aside for now. According to the pure reasons version of the basic reasons strategy, the reason (iii) is the instantiation of the pure reason, and so it is basic, and so it is the contributor; the reasons (i) and (ii) are derivative, and explained by reason (iii), and so they overlap with the reason (iii), and with one another.

However, this strategy faces an *over-generation* problem: it predicts that there will be overlap in cases in which there is no overlap. Consider a different case in which there are two vaccinations available, one attached to each button, and two patients who need the vaccinations. Pressing button A will give Bill one of the vaccinations, and pressing button B will give Bob the other. As before, you can press both buttons. In this case, there is a reason to press A and a reason to press B, but they do not overlap. But note that both of these reasons, (iv) that pressing A will vaccinate Bill is a reason to press A, and (v) that pressing B will vaccinate Bob is a reason to press B, are going to be explained by the single basic reason we identified above, that pressing a button will prevent the suffering of a sentient being is a reason to press it. So if this is what explains overlap in the previous case, it seems that we will incorrectly predict that there is overlap in this case.

To prevent this, we need to find two different basic reasons in this second case, one of which will explain the reason to press A (that it will vaccinate Bill) and the other of which will explain the reason to press B (that it will vaccinate Bob). It does not seem promising to look for a second necessary reasons principle, so we must find a way to draw a distinction one level up, at the level of basic reasons—which we get from applying the reasons principle to concrete situations. To do this, we need to make the basic reasons more fine-grained, so that we can distinguish between two basic reasons that are both subsumed under the more coarse-grained fact, that pressing a button will prevent the suffering of a sentient being. Getting more fine-grained about the action mentioned in the basic reason (e.g.,

instead of 'pressing a button', one basic reason will mention pressing button A, the other will mention pressing button B) will not work, since then we will not be able to explain why there is overlap in the original two buttons, one vaccination case—remember that we want to identify two distinct basic reasons to explain why there is no overlap in the two buttons, two vaccinations case. So that only leaves the option of being more fine-grained about the property of the action—that it will prevent the suffering of a sentient being. For example, we could distinguish two basic reasons: that pressing a button will prevent Bill's suffering is a reason to press it, and that pressing a button will relieve Bob's suffering is a reason to press it.

We think this is the best way for the basic reasons strategist to handle these cases, but we do not think that it is particularly promising. In each case, we will need to be careful to choose basic reasons that distinguish between these kinds of properties of the action at just the right level of detail. If we go too coarse-grained, and so have too few basic reasons, then we will overgenerate cases of overlap. If we go too fine-grained, and so have too many basic reasons, then we will fail to explain some cases of overlap. For example, if in the original two buttons, one vaccination case, we distinguished between the fact that pressing a button will prevent the suffering of a sentient being and the fact that pressing a button will prevent the suffering caused by illness of a sentient being, then we would have too many basic reasons, and so too many contributors. We would thus lack an explanation of some overlapping reasons. It would be illegitimate, in this context, to rely on intuitions about overlap to guide our selection of basic reasons. So it is hard to see how to give a principled, satisfying way of identifying the basic reasons.²²

This last point is an important one for the purposes of this paper. To identify the basic reasons in a way that lets them explain overlap, the

²² A possible solution here, in the case of alternative satisfiers, is to follow Skorupski (2010) and Snedegar (2014, 2017) in relativizing reasons to *choice sets*. This is to make reasons sensitive to the ways in which the options are individuated. For example, in giving the basic reasons in the original two buttons case, we could insist on a coarse-grained individuation that does not distinguish between pressing A and pressing B. In the two buttons, two vaccinations case, on the other hand, we could insist on an individuation of options that does make this distinction. There is much work to be done spelling this out in a way that's consistent with Reasons Fundamentalism. For example, Snedegar (2014, 2017) argues that we need to capture certain relationships between reasons relativized to different choice sets, and shows how to do this by appealing to the grounds of reasons. It is not at all clear how to do this on a Reasons Fundamentalist view. But even if this works for cases of alternative satisfiers, it is not clear how it will help with, e.g., alternative grounds.

Reasons Fundamentalist has had to appeal to properties of the actions or their outcomes. This is exactly what the various bottom up theories do: they explain why certain reasons overlap by pointing out that the actions in question promote just the same values or desires, and why other reasons do not overlap by pointing out that the actions promote different values or desires. 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. So the bottom up theories seem to provide a better account of these cases.

3.2.2 Transmission from Reasons to Reasons

Implementing the basic reasons strategy for explaining overlapping reasons requires (A) identifying a class of basic reasons, and (B) specifying principles telling us what the derivative reasons are, which are explained by a given basic reason; at least some of these will be what are often called *transmission principles*. We've just argued that there are some challenges to carrying out the first part of the strategy. Now we will argue that there are also challenges to carrying out the second part. In particular, we'll argue that the bottom up theorist can provide a more satisfying account of these transmission principles than the Reasons Fundamentalist.

Let's start with a simple thought. If you have a reason R1 to make Mildred happy, and giving Mildred flowers is a 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. 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 means to that end.²³ Return to the case above: you have a reason to make Mildred happy. Giving Mildred flowers is a means to 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* of the means-end relationship—this is the sense in which it is a *derivative* reason. There is no explanatory relationship between the promissory reason and the “end-given” reason to make Mildred happy; you promised to give her flowers, not to make her 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 making her happy *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

²³ Some people believe that reasons only transmit to necessary means, and not to sufficient or merely facilitating means. This seems overly restrictive to us—in many cases, reasons do indeed seem to transmit to sufficient and even merely facilitating means (see Schroeder (2009), Bedke (2009)). One major reason for this restriction to necessary means is the problem of subversion (Bedke (2009), Millsap (unpublished)), which we discuss in more detail below. Moreover, the main problem we raise for the Reasons Fundamentalist in this section will go through even if we restrict transmission to necessary means.

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 for now. 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], we cannot 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 a means to 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* of the means-end relationship. 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, which is explained by the promise. 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 a means to 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 means-end relationship.

At this point, we suspect the Reasons Fundamentalist will be required to build the explanatory dependency of the means reason on the end reason into the transmission principles. So a simple version of a means-end principle would say something like this: "If R is a reason for E, and M is a means for E, then *because of that*, R is a reason for M". And in fact, this is how these principles are sometimes formulated.²⁴ To distinguish the transmitted reasons for the means from the other reasons for the means, then, the Reasons Fundamentalist will just say that though the other considerations *are* reasons for the means, they are not reasons *because* the means is a means to the end in question.

Our objection to this is simply that it would be preferable to give an explanation for *why* certain reasons for the means are explained by the means-end relationship and others aren't. We believe that bottom up theories are well placed to do this, unlike Reasons Fundamentalism. We will now sketch a bottom up explanation of means-end transmission of reasons.

What we'd like is an explanation why the reason for the means is explained by the means-end relationship, instead of a stipulation. 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, and neither does a principle that says that the reason for the means is (or includes) the fact about the means-end relationship.²⁵ Now we will argue that we can do better by specifying the grounds of the transmitted reason. If this is correct, then this is some reason to prefer a bottom up theory to Reasons Fundamentalism, since (i) Reasons Fundamentalism requires specifying adequate transmission principles, but (ii) bottom up theories are better placed to do this.

²⁴ See, for example, Kolodny (forthcoming).

²⁵ This is to reject Kolodny's claim (forthcoming, footnote 8) that specifying that the reason for the means is the fact about the means-end relationship captures the sense in which the reason for the means is a reason *because* of the reason for the end.

We suggest that a transmission principle should take this form: *If R1 is a reason for E, grounded in G, and M is a 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. The important points for our purposes now are that the principle references the ground of the reasons, and that the ground of the reason for the end is the same as the ground 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 a means 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. The reason for the end is that you love Mildred. The ground of this reason, 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 mentioned above 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 a means to 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.

To sum up the discussion so far: suppose that action M is a means to action E, and that there are some reasons for E, and some reasons for M. A satisfactory transmission principle should allow us to distinguish between the reasons for M that are transmitted from the

reasons for E via the means-end relationship from those reasons for M that are not. We suggested that a bottom up theory could do this by appealing to the grounds of the reasons: the transmitted reasons for M will be grounded by the same values, desires, etc. as the corresponding reasons for E. The Reasons Fundamentalist, on the other hand, simply builds into the transmission principle the claim that the transmitted reasons for the means are reasons in virtue of the means-end relationship. So she will distinguish the two kinds of reasons for the means by saying that some are grounded in the reasons for the end, along with the transmission principle, and some are not. Again, this is not incoherent, but insofar as we want to give explanations of phenomena about reasons, we think the bottom up theory looks preferable on this score.

Moreover, the bottom up theory allows us to explain other phenomena related to the transmission of reasons. As Bedke (2009), Rippon (2011), Millsap (unpublished), and others have pointed out, it seems that a satisfactory principle of instrumental transmission must rule out some chaining of instrumental reasons, or iterative applications of the principle. That is, we should not straightforwardly apply the instrumental principle to the instrumental reasons it generates (or at least, the reasons the existence of which it predicts). For example, suppose I have a reason to play tennis. A means to doing this, let's suppose, is buying some tennis balls. So the instrumental principle will tell us that I have a reason to buy some tennis balls. So far, so good. But now try applying the instrumental principle to *this* reason. A means to buying tennis balls is selling my tennis racket, and using the money to buy the balls. So on this iterative application of the instrumental principle, I have a reason to sell my tennis racket, explained by my reason to buy tennis balls, which is itself explained by my reason to play tennis. But of course if I sell my tennis racket, I can't play tennis. So it isn't true that I have a reason to sell my tennis racket, or at least that I have one that is explained by my reason to buy tennis balls, since the only reason I have (let's assume) to buy tennis balls is explained by my reason to play tennis. Bedke (2009, footnote 12) calls this the *problem of subversion*.

One response to this problem is to restrict the instrumental principle so that it applies only to basic reasons, which rules out iterative applications of the principle altogether (Bedke 2009, Kolodny forthcoming; compare Millsap manuscript). This fits with the kind of Reasons Fundamentalist approach we're sketching here, since it

already draws a distinction between basic and derivative reasons, in order to account for overlap.

This response, however, faces a challenge. For there are many cases in which instrumental reasons *can* be chained together in the way described above. That is, there are cases of *non-subversive* chaining. Indeed, these are the most familiar kinds of cases. Restricting the application of the principle to basic reasons rules this out. For example, assume that the fact that the dog would enjoy it is a basic reason to press the button that dispenses the treat to the dog. A means to this is to walk over to the treat-dispensing machine, so I have a transmitted reason to walk over to the machine. A means to walking over to the treat-dispensing machine is getting out of the swimming pool. So I have a transmitted reason to get out of the swimming pool. And so on. Explaining all of these reasons seems to require us to allow iterative applications of the instrumental principle. But we saw above, allowing this in an unrestricted way leads to the problem of subversion.

To account for non-subversive chaining, the Reasons Fundamentalist, following the basic reasons strategy, would need to say that each action in a non-subversive chain is favored by the fact that the action in question is a means, by that very chain, to the end in question.²⁶ If action A1 promotes action A2 which promotes action A3 which promotes the state of affairs of the dog enjoying the treat, the reason for each of these actions is explained directly by its relationship to that state of affairs. The fact that getting out of the pool promotes the dog's enjoying the treat is a reason to get out of the pool. The fact that walking over to the treat dispenser promotes the dog's getting a treat is a reason to walk over to the treat dispenser. The Reasons Fundamentalist following the basic reasons strategy maintains that each of these reasons is explained by direct transmission from the most basic reason, which is that pressing the button on the treat dispenser would promote the dog's enjoying the treat.

Our main worry with this strategy is that the final action in the sequence is given undue importance in this explanation. If things had been a little different, and after pressing the button one needed to turn some lever, it would follow, on this account, that all the reasons in the chain would instead be explained by the basic reason to turn

²⁶ Compare Kolodny (forthcoming), p. 28: the "reason for a means to a means is transmitted directly from the "ultimate" end, not from the means to which it is a means" (see also Bedke 2009).

that lever. Indeed, all the transmitted reasons would *be* facts of the form [getting out of the pool / or walking to the machine / etc. would promote the turning of the lever]. But this is counterintuitive. The reasons to get out of the pool and to walk to the machine are not plausibly explained by the reason to press the button, or the reason to press the lever. They are explained by the fact that one is in a position thereby to enable the dog to enjoy the treat. Perhaps the Reasons Fundamentalist would say instead that the ‘last action in the chain’ isn’t turning the lever but giving the dog the treat. But this concedes our point, for this abstract description picks out an action just by the extrinsic property of facilitating the normatively significant outcome in question. At this abstraction, the details of the last action in the chain drop out of the explanation. What is left is the facilitation relation and the normatively significant outcome. This is precisely what is predicted by the bottom up account.²⁷

The bottom up theorist gives a straightforward account of non-subversive chaining. The explanation for the reason to press the button and the reason to get out of the pool are the same: both are ways of promoting the significant state of affairs, namely the dog’s enjoying the treat. If things had been a little different, and after pressing the button one needed to turn some lever, it would not follow that all the reasons in the chain would instead be explained by the basic reason to turn that lever. All would be explained by the very same thing, the fact that they facilitate the dog’s getting the treat.

Moreover, appealing to the state of affairs and the nature of the facilitation in question is what plausibly explains the weight of these reasons, both in the case of the ‘basic’ or ‘ultimate’ reason itself, and down the chain of transmitted reasons. The bottom up theorists have simple and intuitive explanations of the transmission of reasons and weight in all of these cases.

4 Objections and Replies

4.1 Not all reasons are promotion-based

²⁷ Bear in mind the state of affairs in question may be the performance of some deontologically significant action, or some intrinsically valuable action, so this structure doesn’t beg any questions against deontology.

The objection runs as follows. Not all reasons are facts about facilitating outcomes. Some of them are facts about having fitting responses to evaluatively significant objects or states of affairs.²⁸ The bottom-up accounts assume that reasons are all explained by facts about the *facilitation* of normatively significant properties or actions or their outcomes. But this model better suits promotion-based reasons than fittingness-based reasons. Moreover, overlap issues also afflict fitness-based reasons. The fact that you are in the Sistine Chapel is a reason to take your hat off. The fact that you are inside a church is a reason to take your hat off. And the fact that wearing a hat in a church is a sign of disrespect is a reason to take your hat off. These are, presumably, fittingness-based reasons for an action. And they do, presumably, overlap.

We believe that this objection raises a range of interesting, contentious, and underexplored issues. The literature on the distinction between promotion-based and fittingness-based reasons has focused more on reasons for attitudes than for actions. The broader ethical literature on reasons for action does not give much account of the interaction of these two kinds of reasons. We postpone our account of overlapping fitness-based reasons until more theory is available. For now, we explicitly restrict our thesis to broadly promotion-based reasons for action. The notion of ‘promotion’ is extremely deflationary. So this is a very large class of reasons for action—bearing in mind an action can promote something by instantiating it, as with the walk by the cliffs being an enjoyable way to spend the afternoon.²⁹

4.2 Overlap Worries Afflict Value- and Desire-based views

Here is an example of ‘overlap’ involving value. The fact that Jones is happy to degree five is valuable, and the fact that Jones is happy is valuable, and the fact that someone is happy is valuable—but (assuming the someone is Jones) the value of all three facts ‘overlaps’—presumably with the value of the first of them. Here is an example involving ‘overlap’ with desire satisfaction. The fact that Henry Fonda is placing his cool left hand on your fevered brow satisfies a desire of yours. The fact that Henry Fonda is placing his cool hand on your fevered brow would also satisfy a desire of yours.

²⁸ For the canonical account of the promotion/honouring distinction, see Pettit (1989). On the distinction between promotion-based and fittingness-based reasons, see Howard (forthcoming).

²⁹ Compare Wedgwood (2009).

But there is overlap here too: they are both the very same satisfaction of the same desire. Issues of overlap are not a distinctive problem for reasons as such. So—the objector infers—any problems faced by Reasons Fundamentalism will also be faced by Value-based or Desire-based theories of reasons.

This is a fair challenge. Questions of grain, at least, arise for all normative properties and relations. There is a history of inquiry into the distinction between basic and derivative value.³⁰ One promising hypothesis, roughly, is that the basically valuable states of affairs are those involving individuals picked out *de re* instantiating whichever properties are picked out by the true principles of substantive axiology. The attractiveness of this proposal appears to help the value-based theory of reasons to yield a satisfactory account of overlapping reasons.³¹

There are many subtleties bearing on the individuation of desires and their satisfaction conditions. We suggest that reflection on relations of overlap between desires provide some support for desire-based theories of reasons over both value-based theories and Reasons Fundamentalism. Firstly, basic reasons seem not always to coincide with basic value. For instance, in cases of reasons to give to the *Against Malaria Foundation*, the basic values are plausibly facts about the well-being of particular individuals. But the basic reasons to give seem rather to be of the form ‘Giving would prevent *someone’s* suffering’. It is plausible in such cases that the grain of the basic reasons coincides with the grain of some psychological attitude towards the outcome. Secondly, the hierarchical structure of desires seems to correspond well to the hierarchical structure of reasons. Reflect on how to distinguish the overlapping reasons in a case where you want to spend time with your father, relax for the evening, and play some chess. Playing chess with your father would satisfy these various desires; but they also overlap a little. Playing some chess is a way of spending time with your father, and relaxing for the evening; but it is also something you would have independent desire to do. Here we suggest that there might be cases of *partially* overlapping reasons—and that this gradability is well-captured by advertting to the grounds from the bottom up.

The difficulties faced by Reasons Fundamentalism appear to be more profound. For each relatum in the reasons relation—and there have

³⁰ See especially Feldman (2000); also Zimmerman (2001).

³¹ Maguire (2015) makes this case.

been hypothesized to be as many as five³²--there are questions of alternative grain and questions about alternative referrers. But more problematically, in the case of reasons there are also the problems with alternative satisfiers, and alternative grounds.

The distinctiveness of this problem is clearest when we see that there are analogues of the two buttons case for one person at different times, and for two people at the same time. Suppose you are planning your walk, and you know that you need to press the button either on your way up the hill, or on your way down, and it doesn't matter which. You have a reason to press it on the way up, and one to press it on the way down, but these two reasons overlap when comparing different plans (route A, which enables you to press just on the way up, route B which enables you to press just on the way down, and route C which enables you to press on both). Similarly, if only one of the two of us needs to press the button, and we are considering the matter together, then our reasons overlap. These further issues can be resolved by appealing to the facilitation of a normatively significant outcome in these cases, as in the earlier cases. The Reasons Fundamentalist has all sorts of difficulties explaining all this. These difficulties are different in kind, we want to say, from those afflicting value fundamentalists or philosophers or mind working on desire. At the very least, we hope to have introduced a research project for Reasons Fundamentalists analogous to the long-standing project in axiology.

4.3 What the hell kind of abductive argument are you running here?

Response: A weak one. The evidence is mostly situational. But it does all point in the same direction. There are lots of overlap cases. These have traditionally been taken to be problems for the simple theory of weighing. We have been suggesting, in effect, that they are problems for the simple theory of reasons. And they all point towards an account of contributors in terms of the facilitation of some normatively significant property of an action or its outcome.

The bottom up theorist can easily explain the various structural features we have discussed. She can explain which reasons are basic, if she needs to. But this issue is also less pressing, since the overlap of reasons is accounted for not by appealing to some special reason, but

³² Cf. Skorupski (2010).

to some feature of an action or outcome indicated by overlapping reasons. She can also easily explain transmission—indeed, transmission is almost a theorem of her account. As subversion cases show, facilitation is not transitive. But, as we saw, by appealing to some property of outcomes, bottom up theories are in a position to distinguish the good chaining cases from the bad in a simple way. For each of these features, the Reasons Fundamentalist must simply add more substance to the essence of being a reason—a methodology that stands in some dialectical tension with the initial insistence that reasons lack any substantive essence.

This brings us to a final explanatory advantage of bottom-up theories. If indeed it turns out to be substantively plausible that the basic reasons are all facts about value promotion or desire promotion this would not be inconsistent with Reasons Fundamentalism, as Scanlon points out.³³ But it asks for explanation; and this is something bottom up *analyses* of reasons are uniquely well-placed to provide.³⁴

³³ See Scanlon (2014), p. 5.

³⁴ Thanks to <many people>.

References

- Broome, J. (2004). Reasons. In Wallace, Smith, Scheffler, eds., *Reason and Value: Themes from the Moral Philosophy of Joseph Raz*. Oxford University Press, pp. 2004-28.
- Crisp, R. (2006). *Reasons and the Good*. Oxford University Press.
- Dancy, J. (2004). *Ethics Without Principles*. Oxford University Press.
- Dasgupta, S. (2015). Metaphysical rationalism. *Noûs* 50.2: 379-418.
- Feldman, F. (2000). Basic intrinsic value. *Philosophical Studies*, 99(3), 319–346.
- Finlay, S. (2014). *Confusion of Tongues*. 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 27th Annual Wittgenstein Symposium*. ÖBV and HPT, pp. 387-394.
- Horty, J. (2012). *Reasons as Defaults*. Oxford University Press.
- Howard, C. (forthcoming). The fundamentality of fit. In *Oxford Studies in Metaethics*, volume 14.
- Howard, N. (manuscript). Reasons First and the problem of mixed reasons'. Unpublished manuscript, University of Southern California.
- Kearns, S. and Star, D. (2009). Reasons as evidence. *Oxford Studies in Metaethics*, Vol. 4. Oxford University Press, pp. 215-242.
- Lord, E. & Maguire, B. (2016). An opinionated guide to the weight of reasons.' In E. Lord and B. Maguire, eds., *Weighing Reasons*. Oxford University Press, pp. 1 – 22.
- Maguire, 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.

Millsap, R. (unpublished). Transmission principles. Unpublished manuscript, University of Maryland.

Montaigne, M. 1604 / 1924. *Essays, book 2*. Public Domain, accessed here: <http://www.luminarium.org/renascence-editions/montaigne/2xiv.htm>

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.

Pettit, P. (1989). Consequentialism and respect for persons. *Ethics*, vol. 100, no. 1, pp.116-126.

Portmore, D. (forthcoming). Maximalism vs. omnism about reasons. Forthcoming in *Philosophical Studies*.

Rippon, S. (2011). In defense of the wide-scope instrumental principle. *Journal of Ethics and Social Philosophy*, 5(2): 1-21.

Rosen, G. (forthcoming). Scanlon's modal metaphysics. Forthcoming in *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. (2017). *Contrastive Reasons*. Oxford University Press.

Snedegar, J. (2014). Contrastive reasons and promotion. *Ethics* 125(1): 39-63.

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

Zimmerman, M. (2001). *The Nature of Intrinsic Value*. Rowman & Littlefield.