

The Value-Based Theory of Reasons

Barry Maguire

The University of North Carolina at Chapel Hill

Forthcoming in *Ergo*.

Abstract

This paper develops the Value-Based Theory of Reasons in some detail. The central part of the paper introduces a number of theoretically puzzling features of normative reasons. These include weight, transmission, overlap, and the promiscuity of reasons. It is argued that the Value-Based Theory of Reasons elegantly accounts for these features. This paper is programmatic. Its goal is to put the promising but surprisingly overlooked Value-Based Theory of Reasons on the table in discussions of normative reasons, and to draw attention to a number of areas for fruitful further research.

Section One: Introduction

“If I am walking past a shallow pond and see a child drowning in it, I ought to wade in and pull the child out. This will mean getting my clothes muddy, but this is insignificant, while the death of the child would presumably be a very bad thing.”¹

Here’s the first axiom of value-first ethical theory. If some action available to me would bring about some valuable state of affairs, that fact is a reason for me to do that thing; the weight of this reason is an increasing function of the amount of that value. This seems to be motivating Peter Singer’s response to his example, one of the most well known in ethical theory. The fact that by wading into the pond you can save the life of a child is a reason to do so. The fact that by doing so you’ll get your clothes muddy is a reason not to wade into the pond. However, compared to the value of the saved life, the disvalue of the muddy clothes pales in significance.

Too many sympathetic philosophers, alas, have failed to pause at this first axiom, and have jumped headlong into some variety of consequentialism. There are a

¹ Singer, 1971.

great many varieties of consequentialism available, and there are compelling counterexamples to many of them. One thing uniting existing consequentialisms is that they explain facts about the all-things-considered ought (or other strict facts such as rightness) immediately in terms of facts about value, without any intermediate explanatory level of other normative facts.

My working hypothesis is that value-first ethical theorists stand to benefit from attending separately to two different levels of explanation. At the first level we explain *contributory* facts – in particular facts about reasons and their weights – in terms of facts about value (along with other facts). Then at the second level, the facts about the weights of reasons for and against various options explain the net weight of reason in favour of each option; these facts in turn explain the *strict* facts – facts about ought, permissibility, blameworthiness, decisiveness, etc.² In this paper I focus on and work within the first stage of the explanation. I will characterise and provide some motivation for the Value-Based Theory of Reasons. Oddly enough, this work has not been done.³

² For thorough discussion of these two levels, and the relations between strict and contributory facts more generally, see Lord & Maguire 2016.

³ In recent times Joseph Raz is most often cited as a defender of a value-based theory of reasons. This attribution is disputable. In his 1999 he does say that “reasons are facts in virtue of which ...actions are good in some respect and to some degree” (23) – though Raz is more concerned with motivating reasons than normative reasons in this essay. In his 2011 he says “any explanation of [the notion of a normative reason] in which the word ‘reason’ does not occur will include another term or phrase whose meaning is close to that of ‘a reason’, so that those who puzzle over the nature of reasons will not be helped by the definition” (18). In his 2016 he explicitly rejects the thesis that the weight of a reason is proportional to the value of its object state. Ralph Wedgwood has an account of the relations between reasons for action and neutral values that has affinities with my account (especially his 2009 and 2014). However, Wedgwood denies the explanatory priority of values or reasons. Others have relinquished our notion of neutral value for some kind of *relative* value (on this distinction, see footnote 11). Most prominently, Stephen Finlay (2012, 2014) has developed a *semantics* of reasons according to which reasons are explanations of why acting would be good in a respect; good is analyzed as a relation to individual ends. A reason for S to do A is an explanation of why it would be good (relative to some end E) that S do A – where it’s good relative to E that S do A just if S’s doing A promotes E. There are two contrasts with the view I propose that are worth highlighting. (1) On my view, the relevant goodness is the value of the *outcome* that would be promoted by the act; on Finlay’s view, the relevant goodness is the goodness (relative to some end R) of the *act* itself, which is explained in terms of the promotion of E. (2) Since Finlay’s view places no constraints on ends, his view allows intuitively bad or worthless ends to generate reasons. Your turning on the doomsday machine would promote the end of destroying all sentient life, so your performing that act is good relative to that end; thus, there is a reason for you to perform that act—even if there’s nothing good about the destruction of all sentient life, and even if you have no desire to destroy all sentient life. On my view, there is no reason to turn on the machine. The disvalue of the destruction of all sentient life would entail is a weighty reason for you not to turn on the doomsday machine. Douglas Portmore’s ‘commonsense consequentialism’ takes as fundamental facts about reasons for desires rather than neutral value. There are other snippets and suggestions of a value-based theory of reasons elsewhere, mostly in the context of objecting to other views: in Pekka Väyrynen 2006; in G.A. Cohen 2008; in Andrew Reisner 2009; in Eric Wiland 2012; and in my (forthcoming), which focuses on a puzzle concerning extrinsic value. For some criticisms, see Jonathan Way 2013 and John Brunero 2013.

I will say nothing about the second level of explanation. It is worth emphasizing that an account of reasons is in no way committed to a *maximising* theory of ought, i.e. to the thesis that you ought to do whatever you have most reason to do. It may be what you have most reason to do is supererogatory, and that if some option is supererogatory it is not the case that you ought to perform it.

One often encounters impatience with reasons theory. Many philosophers not working directly on these issues fail to see the attractiveness of the Parfit/Scanlon-style 'Reasons Fundamentalism' research program.⁴ But it is important to see that one can motivate the importance of theorising about reasons independently of that program. One gets to theorising about reasons by way of two thoughts: first, that there are often conflicting considerations bearing on what one ought to do; and second, that this conflict is often resolved by appealing to facts about the relative weights of these competing considerations. Anyone who might once have looked for exceptionless principles relating options with oughts should look instead for an account of the character and content of considerations that count in favour of doing things. This is what I shall do here.

It is not the purport of this paper to argue that other theories of reasons are false – whether the desire-based theory of reasons or reasons fundamentalism or the evidence-based theory.⁵ I'll occasionally make reference to these other theories in order to situate my own view, and I make some suggestive critical remarks in the conclusion. Less ambitiously, the point of this paper is to put the Value-Based Theory of Reasons on the table and to show you some of its virtues. I will show that the Value-Based Theory of Reasons has the resources to explain a large number of puzzling features about reasons. These features have not been much discussed; drawing attention to them is a secondary goal of the paper.

I begin with some quick-and-dirty motivation for the Value-Based Theory of Reasons as a substantive account of which facts are reasons for which options. But the main goal of this paper is to show the *theoretical* rather than the substantive virtues of the account. I introduce a number of puzzling features of reasons, the most central of which are these. Reasons have *weights*. The fact that by wading into the pond you can save the child is a weightier reason to wade into the pond than the fact that by wading into the pond you will ruin your shoes is a reason not to. The weights of reasons are *modified* by a variety of facts, for instance by various probabilities, and they *transmit* in interesting ways, in particular in cases of instrumental reasons. Different reasons also *overlap*. The fact that by wading into the pond you will ruin your shoes, and the fact that by wading into the pond you will have cost yourself \$90 in footwear are clearly not considerations that contribute separate weight against wading into the pond. Finally, reasons are *promiscuous*. In regular deliberation, advice, and justificatory

⁴ Scanlon 1998, Parfit 2011.

⁵ See respectively Mark Schroeder 2007, T.M. Scanlon 1998, Kearns & Star 2008.

contexts it is felicitous to say of an astonishingly wide variety of different facts that they are “reasons” to do this or that. The Value-Based Theory of Reasons fully explains these features.

I will close this introductory section with a quick sketch of the theory. It comes in three stages. We start with facts of the form [ϕ would promote S], where ϕ is an option (e.g. an action) and S is a state of affairs. I call facts of this form *canonical* facts. Not all canonical facts are reasons. In the first stage, the Value-Based Theory of Reasons explains which canonical facts are reasons, and their weights. In the second stage, the value-based theory draws a further distinction within this class of canonical facts that are reasons, between the *basic* reasons and others. The basic reasons are the considerations that contribute non-overlapping weight to an option. In the third stage we turn to non-canonical facts that are reasons. These are considerations that play a pragmatic role in reasons-giving practices, usually by drawing attention to basic reasons or some other part of the explanatory structure rooted in basic reasons. This third stage explains the promiscuity of our reasons talk.

In the next section, Section Two, I will introduce canonical facts, explain which are reasons, and introduce the account of their weight. In Section Three I provide some brief substantive motivation for the value-based theory of reasons. We turn then to the theoretical arguments. In Section Four I explain various facts about modification and transmission of weight. In Section Five I present the account of basic reasons. In Section Six I introduce an account of which non-canonical facts are reasons.

Section Two: The Value-Based Theory of Reasons

Let a *canonical fact* be a fact of the form [ϕ would promote S], where ϕ is some agent's taking some option in some situation and S is a state of affairs.^{6,7} Not all canonical facts are reasons. The fact that your diving out of the window would promote the breaking of your bones is surely not a reason to do so. The Value-

⁶ There are two main regimentation options for the canonical reasons. There is the *property of option* regimentation: for some option ϕ and property F, a ‘canonical’ reason is a fact of the form [ϕ is F]. For example, the fact that your meeting me for lunch is a way of keeping your promise is a reason to do so. There is the *promotion* regimentation: for some option ϕ and state of affairs S, a ‘canonical’ reason is a fact of the form [ϕ would promote S]. For example, the fact that giving flowers would please Alfred is a reason to do so. The property of option regimentation of reasons has been popular (Cf. Raz 1999, p.23, Wedgwood 2009, Brunero 2013). The two are consistent: one way that an option can promote an outcome is by instantiating it. Hence any state of affairs consisting in an option instantiating a property can be rendered as an option promoting a state of affairs. I prefer the promotion regimentation. It is more suggestive: reasons are often facts about outcomes of actions not about properties of actions. It is also more useful: it draws our attention to some pertinent state of affairs, and to some class of relations between options and states of affairs.

⁷ For convenience I omit an *additional* placeholder for the agent – e.g. “A’s ϕ -ing.” We can always include reference to the agent in the description of the action.

Based Theory of Reasons maintains that a canonical fact is a reason if and only if and because the promoted state of affairs is valuable. Thus:

VALUE-BASED THEORY OF REASONS: Some fact of the form [ϕ would promote S] is a reason to ϕ if and only if and if so in virtue of the facts that ϕ would promote S and that S is valuable.

For example, the fact that giving \$100 to GiveWell would promote an improvement in child health in Sub-Saharan Africa is a reason to do so, since an improvement in children's health in Sub-Saharan Africa is valuable. A canonical fact is a reason *against* ϕ if and only if and because S is disvaluable. For instance, the fact that testing the promising new drug on rats would cause them considerable pain is a reason against doing so, since the rats experiencing considerable pain is disvaluable. (In every thesis that follows, I will take the relevant 'against' thesis as assumed.)

The VALUE-BASED THEORY OF REASONS purports to provide an *analysis* of what it is to be a reason, not merely a substantive account of which canonical facts are reasons.⁸ The central claim is that to be a reason for an option is to be a fact about that option's promoting some state of affairs, on the condition that that state of affairs is valuable. In keeping with this ambition, I will develop the account in a way that is maximally neutral about substantive issues. In particular, the Value-Based Theory of Reasons is consistent with a variety of more particular ways of thinking about options, promotion, and value.

The key feature of *value* for our theoretical purposes is that it is a gradable monadic property of states of affairs.⁹ This gradable property imposes at least a partial order on states of affairs.¹⁰ In keeping with traditional Act Consequentialism, by 'value' I mean *neutral* value, rather than value-relative-to-me-or-you. On this way of thinking, the disvalue of the drowning of the small child in Peter Singer's pond has nothing to do with you or me, the passers-by,

⁸ There are many ways to understand 'to be F is to be G' claims. It will not be important for current argumentative purposes to pick a side. For some helpful distinctions, see Kment 2014.

⁹ I assume that the value-bearers are states of affairs rather than objects. I don't think much turns on this; cf. Zimmerman 2001. I'll occasionally talk as though non-obtaining state of affairs are valuable. This is shorthand for saying that they would be valuable if they obtained. For convenience, I will not distinguish between facts, true propositions, or obtaining states of affairs.

¹⁰ This order corresponds to some kind of betterness ordering. We needn't take a stand on the explanatory relationship between value and betterness. Which properties of the betterness ordering we need will depend upon our theory of the relation between reasons and strict deontic properties like ought. If ought is a maximizing notion, and there are no other strict deontic notions, then an ordinal ranking will suffice for our later purposes. If we want some account of weight differentials, e.g. between what you did and what you ought to have done, e.g. for purposes of assessing degrees of criticism, then we'll need a cardinal ranking. Some analyses of *sufficient* reason and even *decisive* reason involve cardinal quantities. For these we'll also need some kind of 'absolute zero' on any betterness ranking.

and *a fortiori* not to any reasons we might have to care or want to help. It has entirely to do with the child, his suffering, his loss of life prospects, and so on.¹¹ I restrict 'value' to final (non-instrumental) value; this will extend to any cases of extrinsic or conditional or organic final value.

It is common to define Act Consequentialism – by contrast with Utilitarianism – in such a way that it takes no stand on the question of what is valuable.¹² It just maintains that *whatever* is valuable, what you ought to do at any time is maximize it. Similarly, the Value-Based Theory of Reasons takes no stand on the question of what is valuable.¹³ A fact about an option promoting a state of affairs is a reason if and only if the substantive axiology assigns value to the state of affairs in question.

I also characterise the *promotion* relation neutrally to minimize substantive commitments in our explanatory framework. To have something to work with, I suggest, with Niko Kolodny (2015), that an option *promotes* a state of affairs by instantiating it, causing or partially causing it, constituting or partially constituting it, preventing the preventing of it, or non-superfluously probabilifying it. The important detail here is that an option may 'promote' a state of affairs by constituting or instantiating it, for instance when some option constitutes an expression of love, or the coveting of your neighbour's goat.¹⁴

¹¹ For the notion of *relative* value, specifically understood in terms of reasons for desires, cf. Michael Smith 2003 and Doug Portmore 2012. Stephen Finlay's end-relational semantics (2014) is best understood as working with relative value as well. (For an important criticism of relative value see Mark Schroeder 2007.) It seems to me that relative value accounts relinquish the deepest consequentialist insight – exemplified in Singer's example – which is that axiological properties the instantiation of which has nothing to do with an agent, are, as it were, normative for that agent. Still, with adjustments here and there the account I offer is consistent with these accounts of value.

¹² Cf. Smith 2009.

¹³ *A fortiori*, I certainly don't assume value monism: that only one kind of state of affairs is valuable or disvaluable, for instance pleasure or desire-satisfaction, or that only states of affairs involving agents are valuable. I leave it open that a variety of kinds of states of affairs are valuable: love, beauty, achievement, understanding, ecological diversity, etc.

¹⁴ What about Philip Pettit's distinction between promoting and honouring (e.g. in his 1991)? He uses this distinction to distinguish consequentialism from deontology. It is very important to point out that – to the extent that I have a clear idea of how Pettit is using the term 'promotion' – I think I am using 'promotion' differently from Pettit. I shied away from the word 'promotion' for a while, in favour of 'realisation,' precisely to avoid connoting Pettit's particular theory. However, in the last few years a burgeoning literature attempts to develop a rigorous account of promotion, independently of any commitment to a value-first position (see Lin (forthcoming) for a nice survey). In the context of this new literature I think one can reclaim the word 'promotion' to play the role that I have articulated. What about 'honouring'? There are two things to say here. If honouring is just an act-type, then according to the Value-Based Theory of Reasons there is a reason to for you to honour *x* in situation *S* if and only if your honouring *x* in *S* would be valuable. But I suspect that Pettit has in mind that honouring has fittingness conditions, like other attitudes. Consequently an account of the normativity of honouring will fall outside the restricted scope of this essay. Thanks to a referee for encouraging me to clarify this contrast.

Neutrality about value and promotion gives the Value-Based Theory of Reasons significant neutrality concerning debates between Consequentialist theories and Deontological theories. Firstly there is no restriction to states of affairs *caused* by the agent. The relevant valuable state of affairs might be instantiated by the action itself. I once saw a paraplegic drive her wheelchair right down the middle of Broadway, slowing traffic to a crawl. It was magnificent! The fact that performing such an action would itself be valuable was a reason for her to do so, irrespective of its causal consequences (the many irate New Yorkers). Secondly there is no restriction on what states of affairs are valuable. The view is consistent with the ‘consequentialising’ point that there is an axiology that assigns disvalues to intuitively wrong-making action-types or types of outcomes, in such a way as to yield, together with suitable linking principles, the facts about reasons and their weights advocated by any Deontological theory.¹⁵ Indeed, the Value-Based Theory of Reasons is consistent with the substantive thesis that *all* reasons are facts about properties instantiated by options (e.g. promise-keepings, love-expressings, goat-covetings).

The Value-Based Theory of Reasons is compatible with a variety of ways of thinking about *options*. As a rule of thumb, our options are the things we choose between or rationally control – whether these are actions, omissions, activities, or plans. (I say a little more about options in section five.)

I say nothing about reasons for *attitudes* as such. The Value-Based Theory of Reasons will capture a significant class of reasons for attitudes. Consider some examples: perhaps anxiety about the rickety bridge would increase your chance of falling; believing you will die miserable and alone would make it harder for you to find true love; an overly strong desire for your own happiness might ensure its own frustration. These so-called ‘wrong kinds’ of reasons for attitudes could easily be explained by the Value-Based Theory of Reasons.¹⁶ It is plausibly the disvalue of falling, dying alone, being unhappy, that underwrites these reasons. It is far less plausible that so-called ‘right kinds’ of reasons are value-based reasons. When your reliable but inconsiderate friend tells you the Butler did it, you have a reason to believe that he did even if no value would be promoted by your doing so (it will spoil your evening).¹⁷ Similarly, the fact that the bridge is rickety might be the ‘right kind of reason’ for you to be anxious about crossing, but no value would be promoted by your having that attitude in that circumstance.¹⁸ In general, so-called ‘right kinds’ of reasons are facts that

¹⁵ On ‘consequentialising’ see Dreier 1993, Smith 2009.

¹⁶ Apart from the ‘rational control’ heuristic. For a helpful introduction to the burgeoning right/wrong reasons literature, see Daniel Jacobson (2013). I favour an account along the lines sketched in D’Arms and Jacobson (2000).

¹⁷ Cf. Tom Kelly (2003); see also Selim Berker (2013). The challenge here is to account for the plausibility of evidentialism and its relationship with a broadly value-first approach to normativity. I discuss this and similar cases at great length in ‘Explaining Epistemic Normativity’ (with Jack Woods).

¹⁸ Assuming there isn’t value in having the attitudes supported by ‘right kind of reason’ as such.

make attitudes *fitting*. In many cases the fact that some fact makes some attitude fitting will not be explained by value-promotion but by constitutive facts about the attitude itself. This raises questions about the relationships between reasons for options and fittingness facts in general, and perhaps between reasons for options and reasons for belief in particular. I take up these questions in detail elsewhere. The account in this paper will be restricted to reasons for options. This may well exclude so-called right kinds of reasons for attitudes. There is a hostage to fortune here, since many philosophers assume that the reason-for-an-option relation and the 'right kind of reason' for an attitude relation are instances of the very same relation.¹⁹ These philosophers will consider it a cost if an analysis of the reason relation can only account for some putative instances of it. However, whether we should accept any such 'unity' argument, or rather reject the relevant 'unity' premise, will depend in part on the strength of the case for the Value-Based Theory of Reasons itself. At least it seems worthwhile to lay out the view carefully and consider its advantages, in order to properly address any putative shortcomings.

Section Three: Some Substantive Motivation, All The Same

The Value-Based Theory of Reasons is a definitional thesis, not a substantive thesis. I haven't placed any restrictions on which states of affairs are valuable. But as a matter of fact, which states of affairs are valuable is explained by a substantive axiology, and we all have strong intuitions about which substantive axiology is true. We can appeal to these intuitions to test the *substantive* plausibility of the Value-Based Theory of Reasons. I believe that it does an excellent job of capturing our intuitions about what reasons we have.

Consider the fact F that wading into the pond would promote the saving of the life of a child. Wading into the pond is a way to save the child, and saving the life of a child is valuable. These facts explain why F is a reason. Now consider the fact G that wading into the pond will ruin your shoes. Wading into the pond would ruin your shoes, and the ruin of your shoes is disvaluable. These facts explain why G is a reason not to wade into the pond. There will be other reasons for and against as well – perhaps the fact that going into the pond would ruin my shoes will disincline me towards benevolent acts in the future.

The Value-Based Theory of Reasons provides a natural explanation of your reasons to give to charity, to read good books, to attend a festival. These options all promote valuable states of affairs – bearing in mind one way to promote a valuable state of affairs is to constitute or instantiate it. This principle also

¹⁹ I challenge this assumption in my 'There are no reasons for emotions or desires'. I argue that putative right-kinds of reasons are strict and non-gradable, and hence are not reasons (which are non-strict and gradable). On the 'unity of reason' see Michael Smith (ms). For a unified and broadly value-first semantics, see Stephen Finlay (2012). For an important objection to 'unity' theses, see Ralph Wedgwood 2015.

explains central examples used to motivate alternative theories of reasons. “The fact that a resort is pleasant is a reason to visit it” (Scanlon), since visiting a pleasant resort is a valuable state of affairs. “While the fact that there will be dancing at the party is a reason for Ronnie to go, it is *not* a reason for Bradley [who doesn’t enjoy dancing] to go” (Schroeder). Dancing at the party would be enjoyable for Ronnie, but not for Bradley, and Ronnie’s dancing enjoyably at the party is a valuable state of affairs.²⁰ But perhaps Bradley’s going to the party would promote Ronnie’s enjoyable dancing at the party. That fact would, presumably, be a reason for Bradley to go. Alternatively, the fact that there will be dancing at the party might *not* be a reason to go in other contexts, in which dancing at the party would *not* be valuable: if the dancers dance on burning coals, or naked, or in contemptuous disregard for the delicate sensibilities of their host (a disabled ex-dancer).

As we can see, it is a virtue of the Value-Based Theory of Reasons that some versions will have substantive bite. Plausibly, you have no reason to stop at a stop sign in an empty desert. No disvalue would be promoted by your failing to do so. Plausibly, you have no reason not to engage in fully consensual and safe sexual acts, if none of the many different susceptibilities to physical or psychological harm are present. Indeed, plausibly it is by looking around to see which values underlie our workaday moral principles (e.g. obey the law, don’t lie, don’t chain your sexual partner to the ceiling) that one can explain why certain cases constitute exceptions to those principles.

The target notion of ‘counting in favour of’ is not a simple relation. It has two separate elements, both of which need to be accounted for by a theory of reasons. There is the general notion of favouring or disfavouring an option, which involves both the idea of support and also some idea of the direction of this support, whether for or against.²¹ Then there is the more specific notion of *counting* towards the option, which brings in the idea of the weight of the reason. The Value-Based Theory of Reasons comes equipped with a complementary theory of the weight of reasons. It maintains that the weight of a reason is a non-monotonically increasing function of the *value* of the state of affairs in question:

VALUE-BASED THEORY OF WEIGHT: the weight of a reason of the form [ϕ would promote S] is a non-monotonically increasing function of the value of the object state S.

The more important the charity work, or the better the book, or the more fabulous the festival, the weightier your reasons to participate. To take a

²⁰ These are central examples of reasons respectively from Scanlon 1998, p.97 and Schroeder, 2007, p.1. These reasons are not rendered into the canonical form in the quotations, but easily could be: e.g. the fact that going to the party would promote my dancing, the fact that going to this resort would promote my having an enjoyable holiday, etc.

²¹ I assume that we can explain direction in terms of weight (though cf. Greenspan 2005).

contemporary example, the more valuable it is to rank philosophy departments, the weightier your reason to accept an invitation to contribute to the ranking committee. The value of saving the child's life is vastly more (absolutely) valuable than the ruining of the shoes is disvaluable. Hence the value-based theory of weight yields the intuitive result that the fact that wading into the pond is a way to save the child is a weightier reason than the fact that wading into the pond will ruin your shoes.²²

This theory of weight also equips the Value-Based Theory of Reasons to respond to a familiar objection, namely that some reasons seem not to have their source in value, but in something primitively deontic – for instance reasons arising directly from promises and/or from people's rights to make bad choices. Promises and rights not only seem to be normatively significant, but also – so the objection runs – they seem to generate reasons with a kind of stringency that seems not to be weighted in the same way as value-based reasons, nor to be easily compared with value.

The idea that 'deontic reasons' are not weighted, or have infinite weight, is a little implausible. It is clear from cases of minimal rights and extreme values that one can have more reason to promote the value than to respect the right, and similarly with promises. Take your right to personal space and the fact that I'll die unless I jump into your lap. Or suppose you promised to be home at seven for dinner but were suddenly called upon to save a child from a pond. As Walter Sinnott-Armstrong (2009) has argued, when it comes to the question of what explains the *weights* of reasons to keep different promises, the most plausible answer is: the various values at stake in the keeping or breaking of that particular promise. Compare promises to do two things that differ dramatically in importance: to pull you out of a rain shower or a raging river. Clearly I have a weightier reason to pull you out of the raging river. In many cases it is clearly the value of the thing promised, and not the mere promise itself, that explains the weight of the reason. In other cases, it is clearly the value of the expectations, or trust, or future cooperation that explain the weight of the reason. This is also plausible with important rights such as the right to life. Take the question whether you have more reason to save five people than one. The Value-Based Theory of Reasons will insist, as against John Taurek (1977), that you do.²³

Section Four: Structural Features, Part one: Modification

²² For an early value-based account of weight, see Michael Philips 1987. For an excellent recent discussion see Sinnott-Armstrong, 2009. For some unpersuasive criticism, see Joseph Raz, 2016. On whether we should think weight explains weightier than or vice versa, see Maguire & Lord (2016).

²³ To be clear: I am not arguing that this is a fully satisfactory theory of promising. It is not. I'm just showing that the value-based theory of weight has the resources to account for a significant range of putatively deontic intuitions. Whether you have more reason to save five than one will turn on an axiological question – but one with a simple answer.

I have provided some initial motivation for the idea that the weights of reasons vary systematically with the value of the state of affairs picked out by a reason in the canonical form. Call this state of affairs the *object state*. So far I have only maintained that the weight of the reason is a non-monotonically increasing function of the value of the object state. If the value of the object state increases, the weight of the reason increases. But other things may also affect the weight of a reason.

Let's start with a remark from *The Port-Royal Logic* from 1662:

...to judge what one ought to do to obtain a good or avoid an evil, one must not only consider the good and the evil in itself, but also the probability that it will or will not happen.

Suppose that holidaying in Turkey and Greece would be equally enjoyable. Suppose all the other reasons bearing on which country to visit are equally weighted, and that you have no significant alternatives to either. However, the two differ as follows: you have a 75% chance of getting a visa in time to visit Turkey, and a 25% chance of getting a visa in time to visit Greece. In that case, you have more reason to apply for a Turkish visa than to apply for a Greek visa. The weight of reasons to apply for either of these visas varies with the value they promote (holidays in either country) but also with its probability.²⁴

Now consider 'instrumental reasons.' The fact that going to the flower shop would promote giving flowers to Frank is a reason to go to the shop, if there is a reason to give flowers to Frank. Moreover, the weight of this reason to go to the shop is no greater than the weight of the reason to give flowers to Frank. The principles behind these two claims concern the transmission of reasons and their weight. They are roughly:

INSTRUMENTAL REASONS: If there is a reason to A, then the fact that B-ing promotes A-ing is a reason to B.²⁵

INSTRUMENTAL WEIGHT: the weight of an instrumental reason is not greater than the weight of the relevant non-instrumental reason.

Notice that INSTRUMENTAL REASONS might be challenged on the grounds that various facts might defeat the instrumental reason, perhaps if B-ing is evil, or if B-ing undermines Z-ing, and there was reason to A only because there was a reason to Z. It isn't necessary to refine these principles further here. Clearly

²⁴ At some point we will need to be more careful about the probability base, whether we are talking about the probability conditional on the option relative to non-performance, or performance of some other options, etc. See Eden Lin (forthcoming) for further details.

²⁵ What if A and B are identical? Would this lead to an explosion of reasons? As we'll see, this is fully accounted for the theory of overlap developed in the following section. Thanks to Ralf Bader for discussion.

something in the ballpark of the first principle is plausible, and the second principle appears to be unimpeachable. My point for now is that something needs to be explained. We have a necessary connection between two distinct relations. It is surely not a *coincidence* that every time there is a reason to A, the fact that B-ing promotes A-ing is a reason to B. It is certainly not a coincidence that the weight of the instrumental reason does not exceed the weight of the non-instrumental reason.

It is more plausible that these two principles – INSTRUMENTAL REASONS and INSTRUMENTAL WEIGHT – obtain necessarily. To account for this necessity, the defender of *reasons fundamentalism*, the thesis that facts about reasons admit of no further explanation, will feel some pressure to build some suitably refined version of INSTRUMENTAL REASONS into their account of what it is to be a reason. This is not quite in the spirit of the famous opening passage in *What We Owe To Each Other* (Scanlon 1998):

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.

By contrast the Value-Based Theory of Reasons handily explains INSTRUMENTAL REASONS. If there is a reason to A, then by the Value-Based Theory of Reasons, A promotes some valuable state of affairs. Promoting something that promotes S is, in general and in all the cases we are concerned with, a way of promoting S. So if B promotes A, B promotes the thing promoted by A, which is valuable. So if B promotes A, B promotes a valuable state of affairs. So by the Value-Based Theory of Reasons, there is a reason to do B. For want of a nail the shoe was lost, for want of a shoe the horse was lost, ...for want of a battle the kingdom was lost, and all for the want of a horseshoe nail. The reason, presumably, transmits all the way down from the value of the kingdom to the blacksmith’s reason to shoe the horse.

According to the Value-Based Theory of Reasons, the weight of a reason is *modified* by the extent to which the option promotes the valuable outcome. Since the probability in question is some real number between 0 and 1, and since probability is (presumably) a multiplier, this can only attenuate the weight of the reason to promote the outcome; it cannot intensify its weight. Consequently, the Value-Based Theory of Reasons is perfectly suited to explain the weight of instrumental reasons, and *a fortiori* to explain the condition specified in INSTRUMENTAL WEIGHT.

In fact, according to the Value-Based Theory of Reasons, there is no non-superficial distinction between ‘instrumental’ and ‘non-instrumental’ reasons. Really, *all* reasons are instrumental reasons.

Take the fact that playing golf is a way of having an enjoyable afternoon. Your having an enjoyable afternoon is finally valuable. Being 'a way of doing something' is a one kind of promotion relation. So, according to the Value-Based Theory of Reasons, this fact is a reason to play golf. Quite generally the best candidates for the reasons labelled 'non-instrumental' reasons are reasons for options which instantiate value, for instance the fact that ϕ -ing is a way of having a really good time, or φ -ing is the keeping of a promise, or λ -ing is an expression of kindness.

Discussions of 'non-instrumental reasons' rarely specify what the non-instrumental reason *is* – just as I didn't in my statement of INSTRUMENTAL REASONS. This is not just unclear, but can lead to indeterminacies, e.g. when a variety of different reasons support some option(s) that is promoted by one 'instrumental' option. I can now explain this omission. Talk about 'non-instrumental reasons' – if not just talk about regular instrumental reasons – is often really just talk about value. We are implicitly assuming that when some state of affairs in question is valuable, the fact that some option would promote it is a reason to perform the option. But this is just an instance of the Value-Based Theory of Reasons. Just like in the case of 'objective reasons,' with 'non-instrumental' reasons one can be led to say implausible or strictly incorrect things by a failure to distinguish and properly relate facts about values and reasons.

These different cases – modification of weight by objective probability and instrumentality – are naturally accounted for by appealing to some gradable property of the object state of a reason and some gradable property of the relation between the option and the object state in the reason. This motivates our characterization of basic reasons as canonical facts of the form [ϕ would promote S], and provides defeasible abductive support for appealing to the value of S. To be clear, I am not denying that other accounts can explain these data. Clearly, S could also instantiate the gradable property of *promoting a desire of yours to some degree*, for instance. I think reasons fundamentalism is rendered less attractive by the comparison, but I do not wish to press this point here. Rather, I wish to draw attention to how well suited the value-based theory is to explaining central structural facts about reasons.

This discussion of modification also provides us with the materials to reply to another common substantive objection to the Value-Based Theory of Reasons, namely that it cannot account for reasons of *partiality*.²⁶ Compare your reasons to keep your own promises with your reasons to promote general promise keeping, or your reasons to save the lives of people you love rather than strangers, or your reasons to promote your own projects rather than those of others. The Value-Based Theory of Reasons can fully vindicate these intuitions. There are two key

²⁶ The seminal statement of this objection is Scheffler 1994.

moves. The first is to point out that in determination of what you ought to do, your friends' interests, the needs of your family, the promotion of your own projects and ideals, etc. have *more weight* than the interests and needs of strangers. Your friends' interests are not (thereby) more neutrally valuable than the interests of others. So, again, we have a difference between the value of an outcome, and the weight of associated reasons. It is a theoretical advantage of the Value-Based Theory of Reasons that it can easily address more complex relationships between values and reasons. The second point is that reasons are modifiable. Plausibly, partial reasons are reasons the weights of which are a function of the value of the relevant outcome modified by facts about the value of *caring* about the outcome in question. It is independently plausible, as a matter of first-order axiology, that it is more neutrally valuable to care about your family and friends than about strangers. Thus it is plausible that your reasons to promote the well-being of your family are weightier than those to promote the well-being of your friends. I have argued for this position in detail elsewhere.²⁷ This quick sketch is not likely to convince you to accept this account, of course. I simply wish to advertise some of the flexibility available to a value-first theorist who begins by attending carefully to reasons and their weights.

Section Five: Structural Features, Part Two: Overlap

There is often more than one way to promote some state of affairs. Suppose that you have a remote controlled dog treat dispenser. It dispenses a single treat to a single dog. You can give the dog a treat *tonight* either by pressing a large red button this morning or pressing the same button this afternoon. If you fail to press the button, nothing bad will happen, and nothing good would come of the unused treat. You have a reason to press the button now and you have a reason to press the button later on. But there is only one treat, so there is no point in pressing it twice. (Ignore the risk of mistakes, etc.) You cannot give both of these reasons separate weight. If you also have to go to the cobbler either this morning or this afternoon, but the reason to press the button is weightier than your reason to go to the cobbler at either time, it would be silly to stay at home on both occasions to press the button. I shall say: the reason to press the button now *overlaps* with the reason to press the button later.

Not all similar reasons overlap. The fact that it would be extremely pleasant to walk my dog through the woods is a reason to do so this morning, and a reason to do so this afternoon. But these reasons don't overlap. The more often I walk my dog through the woods, the better. This suggests a simple explanation: the two treat-related reasons overlap because they are relevantly related to the same valuable state of affairs (i.e. the one treat enjoyed). It is no better to press the button twice than once. The dog-walking reasons do not overlap because two

²⁷ 'Love in the Time of Consequentialism'

valuable states of affairs are at stake (since the more walks, the merrier the dog). It is twice as good to walk my dog in the woods twice today as just once.

Some slightly subtler cases show that this simple explanation will not quite suffice to explain all cases of overlap. Suppose that by pressing some button once you would make it the case that two dogs, in separate rooms, would each enjoy a delicious treat. Plausibly the happiness of each dog would give you a reason to press the button. You have two reasons, we can say, to press the button. But do you have a third reason, arising from the state of affairs consisting in the two dogs being satisfied, in addition to each reason arising from each of the dogs' individually being satisfied? Clearly not. But the state of affairs of two dogs being satisfied *is* a valuable state of affairs. Perhaps we can *say* that this state of affairs constitutes a reason to press the button. (More on linguistic felicity in Section Five.) But this wouldn't count in favour of pressing the button in addition to the two reasons respectively provided by the each of the would-be happy dogs. The reason to make two dogs happy *overlaps* with the reason to make the one dog happy and the reason to make the other dog happy. By way of contrast, none of these reasons overlaps with the reason to make the dogs' owners happy, or the reason to make yourself happy by seeing the happy dogs.

What about the fact that you would make an intelligent sentient creature happy? Or the fact that you would make a dog *very* happy indeed? Or the fact that you would give a dog at least one unit of happiness? Or the fact that you would make a dog happy on a Tuesday afternoon? These are all perfectly good reasons to press the red button. But they do not count *separately* in favour of pressing the red button. They overlap. They could not combine to outweigh some other reason slightly weightier than each of them. Some evidence of this is provided by the following fact about our reason-giving practice. Suppose I ask you why you pressed the red button, when there was seemingly weightier reason to do something else. You couldn't *list* these considerations as though they were independent – except as a kind of joke.

Our question is: in virtue of what do two reasons overlap? This is not the same as the earlier question of how to distinguish facts that are reasons from those that are not. Here we have the task of distinguishing from within the class of reasons those that *overlap* from those that do not. This will serve an important normative purpose, since it amounts to an account of which reasons contribute separate weight towards an option.

To answer this, I begin with the distinction – familiar to axiologists – between *basically* and *non-basically* valuable states of affairs. Intuitively, the basically valuable states of affairs are the states of affairs that are the *source* of an (positive

or negative) evaluative contribution to the world on their own.²⁸ For the sake of definiteness, I'll offer a version of Fred Feldman's seminal account (2000):

BASIC VALUE: a state of affairs is *basically* valuable if and only if it is a pure attribution of a core valuable property or relation to a bare individual.²⁹

The *core* properties or relations are identified by a substantive axiology. Hedonism identifies *being happy to some degree* as the only core property. The debate about whether equality of opportunity for welfare in some community is valuable for its own sake is a debate about whether such equality is a core relation in this sense. An attribution of such a property or relation F is *pure* if the attribution is precisely restricted to F. The *bare* individual restriction is designed similarly to ensure that we pick out an individual directly rather than by any exclude evaluatively irrelevant detail. It is {Fred is happy to *n* degrees} that is basically valuable rather than {the amaro drinker is happy to *n* degrees}.

Many valuable states of affairs fail to meet these conditions, usually by including insufficient or superfluous detail.

Some valuable states of affairs are not basically valuable because they include irrelevant information. For instance the state of affairs of [one person's being happy to degree *n* for duration *d* within one million miles of *m* other planets] contains superfluous detail; the fact about planets is evaluatively irrelevant. It is not a core valuable property according to our axiological theory (we can reasonably suppose).

Some valuable states of affairs are not basically valuable because they contain insufficient information – they are too abstract. The state of affairs of [Jones being happy] is not basically valuable, since it doesn't involve an attribution of a core valuable property. It obtains in virtue of some more detailed state of affairs which does, for instance that [Jones is happy to degree *n* for duration *d*]. For the same reason, the state of affairs of [two utterly unconnected dogs being happy to degree *n* for duration *d*] is not basically valuable.

Every (non-instrumentally) valuable state of affairs that is not basically valuable is valuable in virtue of standing in some particular relation to a basically valuable state of affairs. Which particular relation? Clearly some relations won't work, e.g. the relation *obtaining at the same time as*. Lots of basically disvaluable states of affairs obtain at the same time as basically valuable states of affairs. *Is necessitated*

²⁸ I say they are the 'source' of an evaluative difference rather than that they themselves make an evaluative difference, because some basically valuable states of affairs are *conditionally* valuable. (On the source/condition distinction, see Ralf Bader forthcoming.) The basic/non-basic distinction crosscuts the intrinsic-extrinsic distinction. This principle also applies, *mutatis mutandis*, to disvalue.

²⁹ Feldman adds a spatio/temporal/worldly index to ensure unrepeatability. This is not needed for the analysis of reasons.

by won't quite work, since the value of a singleton doesn't seem related to the value of its member. The relation *is entailed by* won't quite work either. "A is mildly happy or B is utterly miserable" is entailed by "A is mildly happy" but the disjunctive sentence doesn't thereby refer to a state of affairs that is valuable overall: the disjunctive sentence might be true in virtue of the miserable disjunct. This leads us to the relation *partially grounds or is partially grounded by*.³⁰ This is the most plausible candidate so far.³¹ Both directions work. Suppose that the fact that Frank is extremely happy is basically valuable. This fact grounds the fact that Frank is happy. The fact that Frank's c-fibers are firing in such-and-such a way grounds the fact that Frank is extremely happy. Plausibly these other two facts are non-basically valuable in virtue of their relation to the fact that Frank is extremely happy, together with the fact that this fact is basically valuable. Let us indent this principle for ease of reference:

NON-BASIC: a state of affairs is non-basically valuable in virtue of the fact that it (partially) grounds or is (partially) grounded by some one basically valuable state of affairs.³²

The Value-Based Theory of Reasons allows that any fact of the form [ϕ would promote S] is a reason, so long as S is valuable, irrespective of whether or not S is *basically* valuable. I can now use this relationship between basic and non-basically valuable states of affairs to give a sufficient condition for the overlapping of two reasons that are regimented in the canonical form. I can say the following:

CANONICAL OVERLAP: a reason of the form [ϕ would promote S1] overlaps with a reason of the form [ϕ would promote S2] if there is some basically valuable state of affairs S3 and S1 is valuable in virtue of being identical to or standing in a grounding relation to S3 and S2 is valuable in virtue of being identical to or standing in a grounding relation to S3.

To show this principle at work, consider the following three reasons to ϕ :

³⁰ Ground is an irreflexive, asymmetric, transitive, non-monotonic, metaphysical explanatory relation between facts or states of affairs (cf. Gideon Rosen 2010). As mentioned in footnote 8, for presentational convenience I assume that facts are identical with states of affairs.

³¹ The relation *is non-vacuously grounded by* is an even better candidate – see Jack Woods (ms). This account will also need to be further refined to account for conditionality, e.g. cases of preventive value; c.f. Bader *op. cit.* This account is sufficient for current purposes.

³² What if a state of affairs is grounded by a basically valuable state of affairs and a basically disvaluable state of affairs? Take the fact that [A is mildly happy and B is utterly miserable]. If this situation arises, then I distinguish two distinct reasons both corresponding to the promotion of the same state of affairs: a reason in favour and a reason against. This supports the approach to overlap under consideration. Consequently I don't need to take a stand on whether the 'basic value' of a state of affairs is an overall or aggregative notion, or whether I should instead talk about the various distinct 'basic values' inherent in the same state of affairs.

- R1: [ϕ -ing would promote the happiness in degree n for duration d of two separate dogs, Tink and Tonk]
- R2: [ϕ -ing would promote Tink's being happy to degree n for duration d at 4pm]
- R3: [ϕ -ing would promote Tink's being happy to degree n for duration d at 4pm on the 1st January 2000 by chewing a rubber toy stuffed with peanut butter]

The object state in R1 is grounded by the object state in R2. (It is also grounded by the analogous fact about Tonk; see Rosen *op. cit.*) The object state in R2 is basically valuable, I shall assume. By NON-BASIC, the object state in R1 is valuable in virtue of these two facts. Hence, by CANONICAL OVERLAP, R1 and R2 overlap. The object state in R2 obtains in virtue of the more determinate object state in R3. Since the object state in R2 is basically valuable, it follows from NON-BASIC that the object state in R3 is also valuable. In virtue of these facts and CANONICAL OVERLAP, R3 and R2 overlap.

With this account of basic value in hand, we can take a first pass at an account of *basic reasons*. These are analogues in the realm of reasons of the basically valuable states of affairs in the realm of axiology. These are reasons that contribute weight on their own, and with which all other reasons overlap:

FIRST PASS: a fact of the form [ϕ would promote S] is a *basic reason* to ϕ if and only if and if so in virtue of the facts that [ϕ would promote S] and that [S is basically valuable].

This is promising but not quite right. Consider an important challenge, which I owe to an anonymous referee. Suppose that by pressing some button you would guarantee that there will be a pure attribution of a core valuable property to an individual (be happy to degree n , say) – but that there is, at the time of action, no fact of the matter about *which* individual. This isn't a far-fetched case. By donating to charity you will benefit people, but there may be no fact of the matter about which individuals you will benefit.³³ Hence there is, at time of action, no basically valuable state of affairs that will be promoted by the action. (There will be a basically valuable state of affairs later, of course, but not at the time of the action.) Hence, by the FIRST PASS, there is no basic reason to perform the action, and since only basic reasons contribute weight, there is no reason of any weight to perform the action. Clearly this is a *reductio* of any characterization of a value-based theory of reasons.

A reply to this objection proceeds from the following thought. The weight-providing reason to press the button in this case is that *there is some individual*

³³ There are two versions of the case, assuming full information: one in which you will benefit one of a determinate disjunction of people, and another, more 'non-identity problem-ish' in which there is no fact of the matter even about a disjunction of individuals one of which you will benefit. I will work with the latter, stronger version of the challenge.

who will be happy to degree n . For the purposes of metaphysically determining what you ought to do it doesn't matter *which* individual will be happy.³⁴ But for a theory of reasons and their weights it only matters *that* some individual will be happy, and *how* happy they will be. (Also the probability of this object state given the action, etc.)

Let's say that the existential generalization of some basically valuable state of affairs of the form { a is F } (where a is a bare referent and F is a core valuable property) is an *existentially basically valuable* state of affairs of the form {there is an x such that x is F }. Then we have:

SECOND PASS: a fact of the form [ϕ would promote S] is a *basic reason* to ϕ if and only if, and if so in virtue of the facts that ϕ would promote S and that S is existentially basically valuable. S is of the form {there is an x such that x is F } where F is a core valuable property.

This entails that the fact that there is some individual who will be happy to degree n if you press the button is a basic reason to do so. Lots of reasons will overlap with this one: the fact that Larry will be happy to degree n , the fact that someone will be happy to some degree, the fact that Larry will be happy to some degree.

The problem with SECOND PASS is that too many reasons will overlap with it, which intuitively shouldn't. Suppose that giving flowers to Albert would make Albert happy and would also make Egbert, his husband, happy. It is true that there is an x such that x 's being happy would be promoted by your ϕ -ing. But it is also true that there is an x and there is a y such that x does not equal y and x 's being happy and y 's being happy would be promoted by your ϕ -ing. Consistently with the first existential generalization, there could be millions of happy people. The simple existential generalization does not suffice to pick out the *basic* reasons, where these are defined as the most fundamental weight-bearers, with which all other reasons overlap. The problem is that, intuitively, the weight of the reason should increase as the number of people made happy increases, but existentially basically valuable states of affairs are insensitive to how many such people there will be.

In order to give a unified theory of basic reasons we need something in between these two: something between the outcomes, which may unfortunately be (epistemically or metaphysically) unavailable and the time of action, and a simple existential generalization over outcomes, which would lack relevant detail. What we need is a *model* of all the different token basic valuable states of

³⁴ Since we are working with neutral value rather than relative value the natures of individuals will not be core valuable properties. But even if they were, we could build reference to this property into the property attribution rather than the individual. Perhaps it is important that one's *motivating* reason refer to an individual? This is an issue to be addressed in a theory of moral worth not a theory of reasons.

affairs that stand in promotion relations to all the different options.³⁵ For the purposes of determining whether you ought to ϕ or φ , we just need to know the weights of the reasons to ϕ and to φ . For this we just need to know the *number* of distinct basic values they promote and *how* valuable they are. We don't – yet – need to know any *descriptive* features of the individuals to whom the core valuable properties are attributed – just that there are some such individuals.³⁶ There are various different ways to develop this account. I will use the language of 'representing some distinguishable state of affairs' as a placeholder for your preferred account. This leads us to the:

VALUE-BASED THEORY OF BASIC REASONS: a fact of the form [ϕ would promote S] is a *basic reason* to ϕ if and only if and if so in virtue of the facts that ϕ would promote S and that S represents some distinguishable state of affairs that would be basically valuable.

Before moving on, I'll address two objections to this account of the basic reasons.

Objection One: Some values seem to be *mass-y* rather than *count-y*. Here's what I mean by that. Suppose that you have a pain that lasts for thirty-three minutes. Its intensity fluctuates as follows: for the first three minutes it is 3/10 (on the Numeric Pain Scale), then for twenty minutes it rises to 7/10, then for the remaining 10 minutes it reduces to 5/10. How many disvaluable states of affairs are there here? One state thirty-three minutes long? Three states, corresponding to the three stable intensities of pain? An infinity of states, one corresponding to each infinitesimal moment of pain? Some finite number of states, one for each phenomenological 'quantum' during which you experience pain of any intensity? This is an objection because I have been assuming that 'reason' is a count noun, and that reasons can be individuated by appealing to basic value. But, the objection runs, this is like saying one can distinguish waves by units of water. How many reasons do you have to prevent thirty-three minutes of pain?

Reply: This is not an easy question to answer, but fortunately I don't think I need to do so. My thesis is simply that facts about reasons are parasitic on facts about the promotion of values. Whatever I say about the individuation of basic values, I'll say the correlative thing about basic reasons. The crucial thing is not to count overlapping reasons separately. So long as we keep the relevant disvaluable states of affairs clear in mind, we can regiment our reasons language to suit our needs. We can talk about one reason to avoid the thirty-three minutes of pain, when we are weighing the full amount against alternatives (perhaps sedation). We can talk about a reason to avoid the increase in pain after three minutes when

³⁵ For one such proposal, see McDaniel 2014.

³⁶ Of course there are 'separability of person' worries and 'utility monster' worries and 'replaceability' worries about this set-up. But these are issues for another day that can plausibly be addressed either by a suitably nuanced axiology or a suitably nuanced theory of weight. The point for now is just to provide a theoretical framework within which these issues can be addressed.

the alternative is something that begins after those three minutes. Perhaps some less painful but slightly less effective instrument is available for that middle section of the procedure. What we cannot do is add your reason to avoid the full thirty-three minutes of pain as characterized above together with your reason to avoid the increase in pain after three minutes. We explain this by attending to the basic values.

Objection Two: Sometimes a single reason will correspond to many basic values, and sometimes a single basic value will correspond to many reasons. The fact that giving \$100 to GiveWell would promote an improvement in child health in Sub-Saharan Africa is a reason to do so. An improvement in child health is not a basically valuable state of affairs. Indeed, even individual states of *improvements in health* are themselves non-basically valuable, being made up of various more specific improvements. But we would not say that there are *as many reasons* as there are *improvements in health*, far less as many reasons as there are specific improvements. On the other hand, sometimes it takes many people to promote a single value (e.g. to find evidence of the Higgs boson). But there is more than one reason here, surely. If only the basic reasons really contribute weight, and basic reasons are individuated by basic values, are these unobjectionable ways of talking strictly speaking incorrect?

Reply: These ways of talking are perfectly appropriate, even if they do not draw distinctions quite as finely as might be desirable for various philosophical purposes. It is perfectly appropriate to say that the fact that ϕ promotes S1 & S2 is a reason to ϕ , where each is a basically valuable state of affairs, or each an existentially basically valuable state of affairs. It is perfectly appropriate to *talk* about a reason the weight of which corresponds to some function of these basic values. Strictly speaking this talk would be incorrect if you were using the concept of *basic reasons*. But this concept is rarely needed in everyday reasoning contexts. We are only likely to get into trouble when we are counting different reasons. It *would*, however, be infelicitous to say “the fact that ϕ promotes S1 is a reason and the fact that ϕ promotes S1 & S2 is another reason” – unless the conjunction of S1 and S2 had some ‘organic’ value above and beyond the value of S1 and the value of S2. Again, though, we are usually rather proficient at avoiding this sort of double counting. The theory needs to catch up with our practice, not the other way around.

Section Six: Promiscuity, Indicating Reasons, and The Pragmatics of Explanation

I was a little cavalier about reasons *talk* in my replies to those objections. I’ll now explain why. It is important to distinguish between two projects. One is to give an account of the normative considerations that play a particular metaphysical role in the determination of facts about what you ought to do. The Value-Based Theory of Reasons is, in the first instance, an account of this metaphysical role.

This account starts by distinguishing canonical facts, by distinguishing canonical facts that are reasons from those that are not reasons, and then distinguishing canonical facts that are *basic* reasons from those that are not. These basic reasons are the weight-bearers. Strictly speaking, only the basic reasons play a grounding role in determining the total weight of reason favouring each option, and hence (by way of further principles not defended here) in turn in determining what you ought to do.

A separate (related) project is to give an account of the conditions under which it is felicitous to *give* some consideration *as* a reason to do something. This will often not be canonical facts. Quite generally, the considerations we actually offer as reasons are those that are saliently needed, in the specific conversational context, to pick out a larger explanatory structure.³⁷ You should generally make your "...conversational contribution such as is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which you are engaged" (Grice, 1975). If it is obvious that Frank needs cheering up, and that flowers would do so the best, and that flowers are easily attained, and the only uncertainty is whether Frank would be able to see the flowers, and that if he can't some alternative option (a shoulder massage, for instance) would be optimal – in that context telling us that Frank has good eyesight sounds precisely like giving us a reason to bring flowers. But this consideration's playing this conversational role is a separate matter from the question of what role it plays in the metaphysical explanation of the normative status of the act of bringing Frank flowers.

The appropriateness of giving some consideration as a reason is sensitive to lots of things that are not sufficiently relevant for our purposes: conversational pragmatics, different background interests, prevailing standards of thoroughness. I cannot provide a full-dress account of these appropriateness conditions here.³⁸ I will instead develop in some further detail the metaphysics of reasons yielded by the account of basic reasons. My central conjecture is that different parts of this structure are often *good representatives* (to borrow Fogal's term (2016)) either for the basic reasons themselves or for other significant chunks of this overall structure. I will say that considerations playing this role are *indicating reasons*.

Value-Makers: Let's start with facts that are part of the explanation of some basic reason (i.e. the fact that is the reason, not the fact that it is a reason). Since we are dealing with the usefulness of citing one fact as a way of picking out another, I swap the theoretical relation of ground here for a loose folksy notion of

³⁷ See Joseph Raz (1975:22), John Broome (2013:49), and Daniel Fogal (2016:86). This is an instance of a more general distinction between metaphysical explanations and considerations that one can give as an answer to a 'why' question in a specific context (see David Lewis (1986:215).

³⁸ Since these non-canonical reasons overlap with basic reasons, I will not offer a *necessary* condition for two reasons to overlap.

explanation. Let the following be a basic reason: [giving flowers to Frank would give him pleasure of n intensity for d duration at 12:00 on 1 Jan 2000]. Suppose that Frank enjoys flowers because they remind him fondly of his grandmother, or because he likes to marvel at the fragile beauty of nature, or because he likes their smell. Any one of these facts – the fact that flowers remind him of his grandmother, for instance – may be said to be a reason to give him flowers. Each is part of an explanation of why giving him flowers would make him happy with some intensity, for some duration – which is the basic reason.³⁹ These facts are among the most common reasons.

Conditions: The most obvious condition in our explanatory structure is the fact that the object state is valuable. This is not a *basic* reason to do anything. It is a condition on some canonical fact of the form [ϕ promotes S], in virtue of which that fact is a reason. The defender of the Value-Based Theory of Reasons does not insist that facts about value *are* reasons, far less that all reasons are facts about value. Consider this famous remark from T.M. Scanlon (1998:97):

...when I consider particular cases it seems that these reasons are provided by the natural properties that make a thing good or valuable. So, for example, the fact that a resort is pleasant is a reason to visit it or to recommend it to a friend, and the fact that a discovery casts light on the causes of cancer is a reason to applaud it and to support further research of that kind. These natural properties provide a complete explanation of the reasons we have for reacting in these ways to things that are good or valuable. It is not clear what further work could be done by special reason-providing properties of goodness and value, and even less clear how these properties could provide reasons.

Hopefully it is now clear that the appeal to value yields various advantages: distinguishing canonical facts that are reasons from those that aren't; explaining our two instrumental principles; providing a criterion of overlap. By distinguishing the reason from its condition we can avoid Scanlon's objection (for a similar gambit, see Schroeder 2007). I would also say, however, that Scanlon overplayed his hand. We often appeal to facts about the values of outcomes in reason-giving contexts. If you know that this option would save that vase, but assume it is worthless, telling you it is valuable (a priceless Ming) would very effectively indicate the relevant reason.

Grounded: First there are simple existential generalisations. In some contexts, it doesn't matter *what* the reasons are, just that there are some. Consider the fact that there is reason to ϕ , or that there is *a* reason to ϕ , or that there is at least as much reason to ϕ as to φ . It is plausible that these facts obtain in virtue of basic

³⁹ These facts are among the most common reasons. Take the fact that Teresa's tyre is flat. Since Teresa drives to work, the fact that her tyre is flat partially explains the fact that that ing her a lift will realise the state of affairs of her being in work on time. That's a reason to give her a lift to work.

reasons. It is in virtue of the fact that [ϕ promotes S1] and the fact that S1 is valuable, that there is reason to ϕ . Of course this may harmlessly be overdetermined. It may *also* obtain in virtue of the fact that [ϕ promotes S2] and that S2 is valuable. Any of these facts may be given as reasons in an appropriate context.

Alternatives: Consider the fact that you have a reason R1 to ϕ and that you cannot possibly both ϕ and φ . Intuitively, the conjunction of those facts yields a reason R2 not to φ . But when we are distinguishing the reasons that count separately in favour of our various options, we should not count both R1 and R2, e.g. counting the 5 units (or whatever) in favour of ϕ -ing and also the 5 units against φ -ing. For that would double the advantage of ϕ -ing over φ -ing. It may be perfectly felicitous to cite one of these facts in a reason giving context. However, such reasons as these are explanatorily downstream from the relevant basic reasons. It is an advantage of the Value-Based Theory of Reasons that it explains why such facts can play the role of indicating reasons without themselves contributing non-weight independently.

Combinations: Putting Grounded and Alternatives together we get another set of facts: about the net weight of reason supporting any given option, and more pertinently involving comparisons between the net weight of reason supporting any two options (whether distinct actions, or an action and its negation). Similarly, the fact that you ought to ϕ may perhaps serve as a reason to ϕ in some context. But it clearly doesn't add weight to ϕ -ing in addition to the basic reasons to ϕ . On the contrary, it obtains, in part, in virtue of all those reasons.

False but Felicitous: Once these various structural features have been picked out, we can consider some contrast cases. It can be felicitous and appropriate to say things in contexts even though they are literally false. For instance it can be appropriate to say "there is no reason" to do something in a case when it is assumed that all reasons to do that thing would be easily outweighed by some salient reasons not to do the thing. In such a context the utterance "there is no reason to ϕ " is literally false, but may be perfectly acceptable (compare Schroeder 2007). Similarly, it may be appropriate in a context to give as a reason for ϕ -ing that 'there is no reason not to ϕ '. This is strictly speaking not a reason, nor does it indicate a reason – except perhaps obliquely by indicating the fact that you *wanted* to ϕ . Generally in such contexts, there is a clear reason for ϕ -ing and an expectation that there would be an even weightier reason not to ϕ . Hence 'there was no reason not to' plays the role of indicating that there reasons that are part of the background knowledge were not in fact outweighed in this instance.

Finally and importantly, **Evidence:** Suppose that Trustworthy Joe tells you that flowers give Frank intense long-lasting pleasure. The fact that Trustworthy Joe told you that flowers give Frank intense long-lasting pleasure often seems to be a reason to give Frank flowers. But again this reason does not count in favour of giving flowers in addition to the fact that Frank will experience pleasure. Strictly

speaking the fact about Trustworthy Joe is *evidence* that giving Frank flowers would give him pleasure, which in turn is a reason to give him flowers.

There is a complication here. Suppose you know that giving flowers to Frank will give Frank intense long-lasting pleasure, and you come to learn that flowers give Frank such pleasure because they remind him of Amsterdam. It would be felicitous to say that this further fact is a reason to give Frank flowers, in some relevant context. It would be incorrect to say that this is an *additional* reason, or that this consideration adds any weight to the basic reason. It overlaps with the basic reason. Now suppose that Unreliable Rita had told you that Frank enjoys flowers, and then Trustworthy Joe came along and backed this up. These two facts pick out the same basic reason. They overlap, in a certain sense. Clearly they don't amount to a *greater* weight than the weight of the relevant basic reason. If there were some more valuable state of affairs you could equally promote with an alternative option, this additional testimony wouldn't change which of these basic reasons is weightier. However, extra testimony plausibly *does* add weight to the option. You have more reason to act on a good thing you have more evidence about. So we need to distinguish between two sets of overlapping evidence-related considerations. Firstly we have facts that stand in a variety of metaphysical relations to basic reasons, e.g. existential generalization, or by *pars pro toto* reference (Broome, *op. cit.*). These facts can be appealed to in lieu of the basic reasons, but they don't affect the weight of basic reasons and they don't contribute weight in addition to the basic reasons. They are mere "representatives" of basic reasons. Secondly we have facts that provide evidence for the basic reasons. These facts affect the weight of the reason in proportion to the degree of confidence they jointly provide in the basic reason itself. They cannot contribute weight *independently* of the basic reason, and they certainly cannot contribute *more* weight than full potential weight of the basic reason. But they can intensify or attenuate the weight of that same basic reason.

So we have basic reasons. These are the weight-bearers. These are the facts that play the key role in the explanation of what you ought to do. Then quite separately we have indicating reasons. An indicating reason is any consideration that plays the pragmatic role of picking out some basic reason in a context. In fact, even more broadly, an indicating reason is any consideration that plays the pragmatic role of picking out some significant chunk of normative explanation. The fact that 'there is no reason [or not much] not to ϕ ' doesn't pick out any basic reason; it assures you that there isn't any such basic reason! In fact, it is rarely important in a reason giving context to precisely distinguish the basic reason itself. Quite generally, the Value-Based Theory of Reasons can explain our promiscuous, or you might say, relaxed reason giving practices, based on a systematic account of which considerations contribute separate weight to options.

Conclusion

The Value-Based Theory of Reasons draws on ethical theory and ethical intuitions with deep heritage. Peter Singer's Pond-style reasons are built into the framework from the outset, so the question of whether there are reasons to benefit others never arises.

The Value-Based Theory of Reasons is also explanatorily powerful. It offers the prospect of explaining a wider range of phenomena than simple consequentialism, for instance modification to deal with partiality, different principles relating reasons and other deontic facts, for instance obligations, and complexity in the relationships between different values. The Value-Based Theory of Reasons can also accommodate a number of deontological insights, by acknowledging the relevance of values instantiated by options, or by allowing for a range of different accounts of promotion. The Value-Based Theory of Reasons also lends itself nicely to normative theories of reasoning, action, and moral worth, but those details will have to wait.

I have not argued against other theories of reasons. The considerations adduced here provide the beginnings of an objection to Scanlon-style reasons fundamentalism, for they all point in the same direction: that a theory with a gradable monadic primitive, together with a variety of non-normative relations, is more explanatorily nimble than a theory with a three- or four-place relation as its primitive. There is an advantage over Broome's (2014) view, since the Value-Based Theory of Reasons can easily pick out the 'for A' and 'against A' roles in a 'weighing explanation' (for the classic statement of this challenge, see Kearns & Star 2008). The distinction between basic reasons and indicating reasons also invites an exciting hypothesis concerning the REASONS AS EVIDENCE approach (Kearns & Star *op.cit.*; see also Thomson 2009) according to which, necessarily, a fact F is a reason for an agent A to ϕ if and only if F is evidence that A ought to ϕ . This biconditional thesis may well be consistent with the Value-Based Theory of Reasons.

The closest competitors to the value-based theory on these issues will be those which can analyze basic reasons in terms of a promotion relation and an alternative gradable property, for instance degree of desire. The dispute between the value-based and desire-based views will turn on substantive challenges to desire-based theories of reasons and their weight and metaethical challenges to value-based theories. It is noteworthy that the most significant recent defender of a desire-based theory (Schroeder 2007) argued that the desire-theorist's best hope of responding to these substantive challenges is to relinquish the attempt to analyse weight in terms of degrees of desire. There is much more to be said about each one of these structural features in particular, and for and against various theories of reasons, including the value-based theory, in general. My hope is just

to bring the Value-Based Theory of Reasons into greater prominence in these discussions.⁴⁰

⁴⁰ Special thanks to Gideon Rosen, Jack Woods, and Eden Lin. Many thanks to Geoffrey Sayre-McCord, Meghan Flaherty, Ralf Bader, Mark Johnston, Anthony Appiah, Samuel Scheffler, Carla Merino-Rajme, Michael Smith, Daniel Fogal, Whitney Schwab, Adam Lerner, and two exceptionally helpful anonymous referees.

Works Cited

- Bader, R. forthcoming. 'Conditions, Modifiers, and Holism,' in Lord & Maguire (eds), *Weighing Reasons*, OUP
- Berker, S. 2007. 'Particular Reasons,' in *Ethics*
- Berker, S. 2013. 'Epistemic Teleology and the Separateness of Propositions.' *The Philosophical Review*, vol. 122, no. 3.
- Broome, J. 2013. *Rationality Through Reasoning*. London: Wiley Blackwell.
- Brunero, J. 2013. 'Reasons as Explanations.' *Philosophical Studies*.
- Bykvist, K. 2009. 'No Good Fit: Why the Fitting Attitude Analysis of Value Fails.' *Mind*, vol.118, 469.
- Cohen, G.A. 2008. *Rescuing Justice, Rescuing Equality*. Cambridge: Harvard University Press.
- Dreier, J. 1993. 'Structures of Normative Theories.' *The Monist*, 76 (10): 22-40.
- Enoch, D. 2011. 'On Mark Schroeder's Hypotheticalism.' *The Philosophical Review* 120, 423-446.
- Feldman, F. 2000. 'Basic Intrinsic Value.' *Philosophical Studies* 99: 319-346, 2000.
- Fine, K. 2012. 'A Guide To Ground.' In Correia, F. & Schneider, B. (eds). *Because*. Cambridge: Cambridge University Press.
- Finlay, S. 2012. 'Explaining Reasons.' in the *Deutsches Jahrbuch für Philosophie* vol 3.
- Fogal, D. 2016 'Reasons and Reason: Count and Mass,' in Lord & Maguire (eds) *Weighing Reasons*, OUP.
- Greenspan, P. 2005. 'Asymmetrical Practical Reasons.' In M. E. Reicher, J. C. Marek (Eds.), *Experience and Analysis*. 387-394.
- Grice, P. 1975, 'Logic and conversation.'. In Cole, P.; Morgan, J. *Syntax and semantics*. 3: Speech acts. New York: Academic Press
- Jackson, F. 1991. 'Decision-Theoretic Consequentialism and the Nearest and Dearest Objection,' in *Ethics*
- Jacobson, D. 2013. 'Wrong Kind of Reasons Problem.' *The International Encyclopedia of Ethics*.
- Jacobson, D. & D'Arms, J. 2000. 'The Moralistic Fallacy: On the 'Appropriateness' of Emotions.' *Philosophy & Phenomenological Research*. Vol. 61, No. 1.
- Kagan, S. 1988. 'The Additive Fallacy,' in *Ethics*
- Kearns, S. 2016. 'Bearing the Weight of Reasons.' In *Weighing Reasons*, Maguire & Lord (eds), Oxford: OUP.
- Kearns, S. & Star, D. 2009. 'Reasons as Evidence.' In *Oxford Studies in Metaethics*, vol. 4.
- Kelly, T. 2003. 'Epistemic Rationality as Instrumental Rationality: A Critique.' In *Philosophy & Phenomenological Research*, Vol. LXVI, No.3.

- Kment, B. 2014. *Modality and Explanatory Reasoning*. Oxford: Oxford University Press.
- Kolodny, N. forthcoming. 'Instrumental Reasons,' in Star (ed), *Oxford Handbook of Reasons and Normativity*
- Lewis, D. 1986. 'Causal Explanation.' In *Philosophical Papers*, volume 2, 214-240. Oxford: Oxford University Press
- Lin, E. forthcoming. 'Simple Probabilistic Promotion.' *Philosophy & Phenomenological Research*.
- Maguire, B. & Lord, E. forthcoming, 'An Opinionated Guide to the Weight of Reasons.' In Maguire & Lord (eds), *Weighing Reasons*, Oxford: OUP.
- Maguire, B. (ms) 'Love in the Time of Consequentialism'
- Maguire, B. (ms) 'There are no reasons for emotions or desires'
- McDaniel, K. 2014. 'A Moorean View of the Value of Lives,' *Pacific Philosophical Quarterly* 95, 23-46.
- Manne, K. Forthcoming. 'Democratizing Reasons.' In Maguire & Lord (eds), *Weighing Reasons*.
- Maguire, B. (forthcoming) 'Extrinsic Value and the Separability of Reasons' in *Oxford Studies in Normative Ethics*, volume 6.
- Parfit, D. 2011. *On What Matters*. Oxford: Oxford University Press.
- Philips, M. 1987. 'Weighing Moral Reasons' in *Mind*
- Portmore, D. 2011. *Commonsense Consequentialism*. Oxford: OUP.
- Raz, J. 1999. *Practical Reason and Norms*, Oxford: OUP
- Raz, J. 2011. *From Normativity to Responsibility*. Oxford: OUP.
- Raz, J. 2016. 'Value and the Weight of Practical Reasons.' In *Weighing Reasons*, Maguire & Lord (eds). Oxford: OUP.
- Rosen, G. 2010. 'Metaphysical Dependence: Grounding and Reduction.' In Hale, B. (ed.) *Modality*. Oxford: OUP.
- Rosen, G. ms. 'Normative Necessity.'
- Scanlon, T. 1998. *What We Owe To Each Other*. Harvard: Harvard University Press.
- Scanlon, T. 2014. *Being Realistic about Reasons*. Oxford: OUP.
- Schaffer, J. forthcoming. 'Grounding in the Image of Causation' in *Philosophical Studies*.
- Scheffler, S. 1994. *The Rejection of Consequentialism*. Oxford: Clarendon Press.
- Schroeder, M. 2007. *Slaves of the Passions*. Oxford: OUP.
- Schroeder, M. 2007. 'Teleology, Agent-Relative Value, and the "Good".'
- Singer, P. 1971. 'Famine, Affluence, and Morality.'
- Sinnott-Armstrong, W. 2009. 'How Strong is This Obligation...' in *Analysis*

- Smith, M. 2003. 'Neutral and Relative Value after Moore'
- Smith, M. 2009. 'Consequentialism and the Nearest and Dearest Objection.' In Ian Ravenscroft (ed.), *Minds, Ethics, and Conditionals: Themes from the Philosophy of Frank Jackson*. Oxford: Oxford University Press
- Smith, M. (manuscript). Unity in the Theory of Reasons.
- Snedegar, J. 2014. 'Contrastive Reasons and Promotion.' *Ethics* 125(1):39-63.
- Thomson, J.J. 2008. *Normativity*.
- Väyrynen, P. 2006. 'Resisting the Buck-Passing Account of Value.' In Russ Shafer-Landau (ed.), *Oxford Studies in Metaethics, Volume 1*. Oxford University Press.
- Way, J. 2013. 'Value and Reasons to Favour.' In Shafer-Landau, ed, *Oxford Studies in Metaethics, volume 8*.
- Wedgwood, R. 2009. 'Intrinsic Values and Reasons for Action.' In Sosa, E., & Villanueva, E. (eds). *Metaethics*. Wiley Periodicals, Inc.
- Wedgwood, R. 2014. 'Two Grades of Non-Consequentialism.' *Criminal Law and Philosophy*.
- Wedgwood, R. 2015. 'The Pitfalls of Reasons.' [*Philosophical Issues* 25 \(1\):123-143.](#)
- Wiland, E. 2012. *Reasons*. New York: Continuum Press.
- Williams, B. 1981. *Moral Luck*. Cambridge: Cambridge University Press.
- Woods, J. Manuscript. 'How to get an 'Is' from an 'Ought.''
- Zimmerman, M. 2001. *The Nature of Intrinsic Value*. Rowman & Littlefield: Maryland.