

Reasons and Oughts: An Explanation
and Defence of Deontic Buck-
Passing



Euan Karl Hans Metz
Department of Philosophy
University of Reading

A thesis submitted to the University of Reading for the degree of

Doctor of Philosophy

April 2018

Declaration

I confirm that this is my own work and the use of all material from other sources has been properly and fully acknowledged.

Euan Metz

Abstract

This thesis is about what a normative reason is and how reasons relate to oughts. I argue that normative reasons are to be understood as relational properties of favouring or disfavouring. I then examine the question: What is the relation between reasons, so understood, and what we ought to do, believe, or feel? I argue that the relation is an explanatory one. We should explain what we ought to do in terms of reasons, and not the other way around. This view faces a number of difficulties, in particular in accounting for supererogatory acts and the distinction between an action being required and an action being recommended. The analysis that I provide explains how we can solve these problems. In providing such an analysis, this thesis aims to be a contribution to the discussion of how we might elucidate the structure of what is sometimes called *normativity*.

Acknowledgements

I would like to thank my supervisors, Philip Stratton-Lake, Brad Hooker, and Jonathan Way for their invaluable support and feedback. I would also like to thank my family and friends for their enduring support and wish to dedicate this thesis to Caitlin, whose advice, encouragement, and companionship lightened the burden and made my task that much easier.

TABLE OF CONTENTS

Introduction	1
---------------------------	---

PART I

1. The Structure of Normative Reasons

1.1. Reasons as Considerations which Count in Favour	4
1.1.1. Reasons as Three-Place Relations	6
1.1.2. Reasons as Four-Place Relations	10
1.2. Complete Reasons	13
1.3. How do Reasons Interact with each other?	22
1.4. The Distinction between Reasons For and Reasons Against	24
1.5. The Weight of Reasons.....	29
1.5.1. Weight-First or Weightier-Than First?.....	31
1.5.2. Reductive Views.....	32
1.5.3. Value-Based Views	33
1.5.4. Reductive Views, again.....	38
1.5.5. Remarks on a Non-Reductive Account of the Weight of Reasons.....	40

2. Responding Correctly to Normative Reasons

2.1. <i>Compliance</i> and <i>Conformity</i>	43
2.2.1. The Argument from Accomplishment.....	46
2.2.2. The Argument from Response-Guiding	50
2.3. Objections to <i>Compliance</i>	55
2.3.1. The Huckleberry Finn Counterexample	55
2.3.2. Omissions	56
2.3.3. The Surprise Party	59

PART II

3. Analyses of Normative Reasons

3.1.	Introduction	62
3.2.	Reasons as Explanation	63
3.3.	Reasons as Evidence.....	67
3.3.1.	Objections to RE.....	69
3.4.	Parfit's Reasons Fundamentalism	76
3.4.1.	Smith's Objection.....	76

4. Deontic Buck-Passing: The Positive Thesis

4.1.	Introduction	81
4.2.	The Enticing/Peremptory Distinction.....	83
4.3.	The Justifying/Requiring Distinction	88
4.4.1.	Second-Order Reasons	92
4.4.2.	Against Second-Order Reasons	93
4.5.	Towards a Different Analysis.....	97
4.6.	Clarifying Some Concepts.....	98
4.7.1.	Requirements to Φ and Requirements not to Φ	100
4.7.2.	Requirements to Φ	101
4.7.3.	Requirements not to Φ	106
4.8.	Recommended to Φ and Recommended not to Φ	111
4.9.	Further Remarks on Requirements and Moral Obligations.....	112
4.10.	Resolving the Supererogation Problem.....	116

5. Deontic Buck-Passing: The Negative Thesis

5.1.	Some Objection to the <i>Argument</i> for NBT.....	119
5.1.1.	Heuer on Wrongness	119
5.1.2.	Schroeder on Buck-Passing Facts.....	123
5.1.3.	Derivative Reasons	125
5.2.	Scanlon on Wrongness	128

5.3.	Darwall on Wrongness.....	132
5.3.1.	The Objection from Morally Conscientious Agency.....	135
5.3.2.	Objections to Premise 2	138
5.4.	Conclusion	140
Overall Conclusions		141
References		142

0. Introduction

This thesis falls into two parts, centred on the question of what normative reasons are, and how they relate to deontic properties. In Part I, I discuss and explain structural questions about normative reasons by examining (a) what a normative reason is – and (b) what it is to respond to a normative reason.

In chapter 1 I try to bring some of this work together and sketch a plausible account of what a normative reason is. In this chapter I distinguish a normative reason from other uses of the term ‘reason’ and outline the most widely-accepted account of a normative as a relation. I discuss whether we should accept that is a two, three, or four-place relation and try to support the claim that it is best to construe it as a three-place relation. I then consider whether *context* should form part of the content of a reason, or not. The issue has divided some over the issue of the notion of ‘complete’ reasons, which I argue we should replace with Dancy’s distinction between a reason and an enabler.

I then argue that there is an asymmetry between the reasons which count in favour, and which count against. Building on work by Greenspan and Snedegar I claim that this overlooked distinction is important in that it allows us to formulate the right analysis of deontic properties in terms of reasons.

Finally in this chapter, I discuss the weight of reasons, and give reasons for thinking that the prospects for a value-based account of reasons are not good, and give reason to doubt whether any reductive account of the weight of reasons will be successful, given the context-sensitivity of reasons in general.

In chapter 2, I consider the question of what it is to respond to a reason. Following Raz I distinguish between two different possible views on the subject: a *Compliance* view and a *Conformity* view. I argue that *Conformity* is not sufficient to satisfy responding to a reason, and give reason to think that *Compliance* is better suited to this job.

In Part II of this thesis, I examine the question whether normative reasons are more fundamental than deontic properties (what I call *Weak Reasons Fundamentalism*), such as the property ‘ought’. To begin, in chapter 3 I argue against various putative analyses of reasons in terms of the property ‘ought’. I argue that

both (a) Broome's analysis of reasons in terms of 'ought' and explanation (b) and Kearns & Star's analysis of reasons in terms of 'ought' and evidence are mistaken on various grounds. I then consider an argument due to Smith, building on work by Thomson, which attempts to reduce reasons to the notion of 'correctness' and argue that it, too, fails.

Chapter 4 considers a number of candidates for the 'positive deontic buck-passing thesis' which states that deontic properties are analysed in terms of normative reasons. I reject analyses defended by Dancy, Gert, and Raz, before formulating my own version of the thesis (§4.5). I go on to show how my analyses allow us to solve the most pressing problems of the positive deontic buck-passing thesis: the problem of entailment between 'recommended' and 'required', and the problem of supererogation.

The final chapter examines what I call the 'negative deontic buck-passing thesis' which states that deontic properties do not provide normative reasons. I consider and object to two arguments which deny this claim, and then examine two arguments in favour of the claim that deontic properties *do* provide reasons (in particular, the property of an act's being *wrong*). I argue that neither of these views is persuasive.

PART I

1.1. The Structure of Normative Reasons

I understand *Reasons Fundamentalism* (RF) as the conjunction of two claims. First, that the property of being a normative reason cannot be reduced to any other (normative or non-normative) property (call this view *Reasons Primitivism*). Second, that all other normative properties can be reduced to the property of being a normative reason (call this the *Buck-Passing View*).

In this thesis I defend what I call *Weak Reasons Fundamentalism* (WRF). WRF is the conjunction of two claims. First, that the property of being a normative reason cannot be reduced to oughts (call this *Weak Reasons Primitivism* which I defend in chapter 3) and second that ought *can* be reduced to the property of being a normative reason (call this the *Deontic Buck-Passing View* which I defend in chapters 4 and 5). Since the property of being a normative reason is of central importance to this thesis, before discussing WRF as such, I consider in this chapter some structural questions about normative reasons.

What do I mean by ‘structural’ questions? We should distinguish between two kinds of explanations of what a normative reason is. First, there are reductive explanations that aim to explain what a normative reason is by reducing it to some other (normative or non-normative) property. It is this kind of explanation I discuss in later chapters. Second, there are non-reductive explanations that aim to explain what a normative reason is without aiming to reduce it to other properties. This second kind of explanation is what I call a ‘structural’ explanation of a normative reason.

In this chapter I set the scene for the thesis by reviewing various structural issues about normative reasons. By the end of this chapter it should be clearer what the target notion of a normative reason is, which will inform subsequent discussion in the later chapters.

1.1. Reasons as Considerations which Count in Favour

Many philosophers agree that a normative reason is a relation that holds between a consideration¹ and a response (an action or an attitude). When we say that Susan's loving devotion to her children is a reason to respect her, we seem to be asserting that the consideration 'Susan's loving devotion to her children' stands in the relation of 'being a reason' towards 'respecting Susan'. Susan's loving devotion to her children counts in favour of, or lends weight to adopting, the attitude of respect toward Susan. That seems to accord with a central sense of what people mean when they talk about normative reasons for responses.

How does this consideration get to stand in this relation toward adopting this attitude? Perhaps it does so in virtue of the fact that having a loving devotion to one's children is necessary for being a good mother and that good mothers are worthy of respect. Or perhaps it does so in another way. This question points to an ambiguity. We can ask what *formal* features the relation has, such as how many relata the relation has, and what is the ontological status of its relata. Separately, we can ask the *substantive* question: what reasons do we actually have? For the purposes of this chapter, I will be concerned with the formal features of the relation, and leave the substantive question to one side.

It can be seen that the use of the term 'reason' in the example of Susan is to be distinguished from another use of the term 'reason' such as in the sentence 'The presence of high levels of leaf beetles in this area is a reason why the crops are depleted'. The relation of being a reason referred to in that sentence does not indicate that the presence of leaf beetles counts in favour of, or lends weight to the fact that, the crops are depleted. Rather, the relation in this sentence is that of being an explanation why something is the case. In English we usually distinguish the two senses of the word 'reason' by using the locution 'reason-for' and 'reason-why'. In the literature, the latter sort of reason is often named an 'explanatory' reason.

¹ Though there is disagreement about whether to understand this as a fact, proposition, state of affairs, or mental state. For discussion see, for example, Dancy (2000a: 112-120), Mantel (2014), Tanyi & Matteo (2017), Turri (2009).

We also use the term ‘reason’ to denote a further distinct relation, where we say for instance that ‘his hatred of children is the reason for which he hits them’. In this case the reason cites the consideration that explains what motivated some agent to act in the way that they did, even though that consideration may not count in favour of their acting in that way. Since reasons in this sense can figure in explanations of what motivated an agent to respond in some way they can also feature in explanations of what reasoning led the agent to so respond.

I think there are at least three different things that could be meant by a reason *in addition to* the idea of a normative reason. I’ll give these readings different labels in order to distinguish them here:

Explanatory reason	A consideration that contributed to explaining why A Φ -ed.
Motivating reason	A consideration that (partially or wholly) motivated A to Φ .
Normative-motive reason	A consideration that A took to call for A’s Φ -ing <i>and</i> that (partially or wholly) motivated A to Φ .

It is not hard to see how these definitions can come apart. A consideration can *explain* why A Φ -ed which did not motivate A to Φ , such as a sociological explanation of his Φ -ing, or a physical explanation of his Φ -ing. A consideration can *motivate* A to Φ without A recognising that consideration as favouring his Φ -ing.² I use the term *normative-motive reason* to distinguish a motivating reason which necessarily makes reference to the concept of a normative reason in citing what motivated A to Φ . In order to understand why A Φ -ed in this case we need to be in possession of the concept of a normative reason. However, that doesn’t mean that A

² Some deny that there could be such a reason as they hold that motivation by a consideration necessarily involves taking that consideration to be a normative reason. See, for example, Raz (2010).

himself must be in possession of this concept in order to satisfy *normative-motive reason*. It is an open question just how an agent can take a consideration to call for his Φ -ing when he Φ 's. I won't address this question further here.³

1.1.1. Reasons as Three-Place Relations

In the example of Susan, I have taken a reason to be a two-place relation of counting in favour which takes considerations and attitudes as relata. By attitudes I have here in mind such things as beliefs, desires, wishes, intentions, and feelings. Besides attitudes, it is commonly claimed that actions can figure in reason relations, such as 'That the incumbent Prime Minister is corrupt is a reason not to re-elect her' (where I take it that refraining from doing something as well as doing something are regarded as actions).

In what follows I will examine how many relata a reason relation has. Although there may not be a canonical form of the reason relation and it may be permissible to use the term 'reason' to denote *either* a two-place relation, a three-place relation, or a four-place-relation and so on, it is worth asking what implications higher-order relations have, and whether it is possible that there is a canonical form of the property of being a reason.

I introduce first the claim that a normative reason is a three-place relation 'counting in favour or against' that holds between (i) a consideration, (ii) a response, and (iii) an *agent*. This is a particular way of understanding reasons that is committed to the claim that the two-place reason relation is to be understood, ultimately, in terms of this three-place reason relation. Schroeder suggests that the most obvious way to explain how claims such as 'That there is dancing at the party is a reason to go' are elliptical of the form 'That there will be dancing at the party is a reason for A to go to the party' is to claim that the two-place reason relation universally quantifies into the agent-place of the three-place reason relation.⁴ Two-place reason relations, fully spelled out, are to be read: R is a reason *for all of us* to Φ . Contextual differences that rule out reasons being for certain specific sets of agents are to be

³ For more on these distinctions see Alvarez (2010) especially chapter 2.

⁴ Schroeder (2007: 280ff).

explained in terms of the scope of ‘us’ in the definition. For example, in moral cases ‘us’ likely refers to every agent, while in hedonic cases ‘us’ may refer to a set of agents defined in terms of their hedonic preferences.

This account of reasons having three places is a natural one. One simple reason for this is that if reasons favour responses, since responses require responders, so too reasons require responders. Differences can arise from the ways in which responders (agents) figure in the reason-relation, but they seem to figure in an irreducible way into the relation.⁵ So it is natural to claim that reasons quantify over agents.

Here we should note a distinction between the idea that reasons quantify over agents and the idea that agents have epistemic access to reasons that apply to them anyway. It is possible that there is a reason for some agent, A, to act in some way and that A not know that he has this reason.⁶ Since John doesn’t know about Susan’s loving devotion to her children, he does not have the reason that is implied by that fact, and hence it is not reasonable to expect him to respond to it. Nevertheless, it is true to say that there *is* a reason for John to act, it is just one that he doesn’t *have*.⁷

In addition to being ignorant of non-normative facts that imply normative reasons, an agent may be ignorant of the normative fact that some descriptive fact is, or provides, a reason to respond in some way. Here too the reason may quantify over that agent, but is one to which he fails to have epistemic access.⁸

There are alternative views we could hold. We could hold that the two-place relation formulation is the more basic, and seek to define the three-place relation in

⁵ Broome is agnostic about the possibility that reasons are always reasons for somebody, and that there may be such things as (in his terminology) ‘unowned’ reasons that ‘... are not normative properties of people’. See Broome (2013: 69) and Broome (2015: 84). However, it is difficult to see just how reasons could fail to be a property of agents, if they are (as he admits) a property of responses.

⁶ Gibbons (2010) disputes the cogency of this distinction as he conceives of reasons as ‘things that make things reasonable’ and hence cannot be independent of an agent’s epistemic access.

⁷ On the being/having distinction, see Skorupski (2010) ch.2.

⁸ On possession see Lord (2015: 27).

terms of it.⁹ On the other hand we might think that there is no need to decide between which of these relations is more basic and that there are two serviceable ways of thinking about reasons that we can use.

Note that the distinction between a reason as a two-place relation without an agent-place, and as a three-place relation with an agent-place, is not the distinction commonly referred to by philosophers as the ‘agent-relative’ and ‘agent-neutral’ distinction.¹⁰ Thinking in terms of reasons, we can define an agent-relative reason as a reason in which the designation which picks out the agent to whom that reason applies features in the ‘consideration-place’ of the reason relation (e.g. ‘the fact that John Smith’s children are in need is a reason for John Smith to help them’) and an agent-neutral reason as a reason which does not feature the designation of the agent for whom that is the reason in the ‘consideration-place’.¹¹

The three-place formulation of the reason relation captures the natural thought that reasons are for somebody without collapsing the agent relative/neutral distinction as understood above. This is an important constraint on reasons. If reasons are for somebody, it is tempting to be led to believe that whatever reasons there are, they must be such that they can be responded to appropriately by whomsoever they are for.

Though this may be true, it does not follow from what has been said about reasons so far. I have only said that reasons are considerations that favour, or lend weight to, attitudes or actions that we supposed must be for somebody. Suppose that the particular attitude favoured by some consideration, R, was extremely difficult to adopt by A. Would that imply that R is not a reason for A after all? Suppose that A has cancer, and that curing cancer would thus greatly improve A’s wellbeing. It is apparent that the fact that curing cancer would improve A’s wellbeing is a reason for A to cure cancer. However, the fact that it is extremely difficult for anybody, and for

⁹ For a discussion of the options on this question see Schroeder (2007).

¹⁰ Somewhat confusingly, Schroeder marks the distinction between two-place and three-place reason predicates with the labels ‘agent-neutral’ and ‘agent-relational’.

¹¹ For more on this version of the distinction see Parfit (1984: 27ff) and Nagel (1970) ch.7.

A in particular, to cure cancer diminishes the strength of A's reason. We can thus felicitously say that A has a reason to cure cancer, but a very weak one.¹²

Secondly, it might be thought that the agent-place constrains what reasons there are because of a specific function that reasons have, namely that reasons not only favour but also *guide* agents. If that were the case, then we have reason to think that the three-place predicate formulation in conjunction with the guidance claim would constrain what reasons there are. The idea that normative reasons are the sort of thing that guide agents to respond in certain ways is captured by some writers in the following way:

Response Constraint (RC) If R is a normative reason for A to Φ , then A is able to Φ for the reason that R¹³

Some have objected to RC by citing counterexamples. Most prominent among these are self-effacing reasons, such as those cited in 'surprise party' cases.¹⁴ If A loves surprise parties, then the fact that there is a surprise party waiting for A at home cannot be a reason that A acts for, since if he were to act on the basis of that consideration it would disappear as a reason (since the party would no longer be a surprise).

Way & Whiting argue that we should distinguish between difference senses of ability in RC.¹⁵ A *general* ability is the competence a person has to do something, such as play tennis well. But you may have the competence to play tennis well, while not having a racket at hand. In that case you have a general ability to play tennis well, but lack a *specific* ability to play tennis (well) here and now. They claim that the ability in RC is a general ability to follow certain patterns or rules of reasoning.

¹² Compare Schroeder's claim that anybody has a very weak reason to eat their car because they will gain a small amount of iron. That case however seems less clear since the value of acquiring the iron is offset by the disvalue of the damage done to one's body by eating the rest of the car. See Schroeder (2007: 95-7).

¹³ See Kiesewetter (2016), Lord (2015), Whiting & Way (2017), Williams (1979).

¹⁴ The case is from Schroeder (2007: 165).

¹⁵ Way & Whiting (2017: 214).

For example, A's wearing boots for the reason that it is snowing is a general ability to follow the rule: 'if it is snowing: wear boots'.

So is A able in the *general* sense to go home for the reason that there is a surprise party there? If that ability takes the form of the rule: 'There is a surprise party at home: go home' then it seems perfectly possible for an agent to follow that rule because the surprise party mentioned in the rule need not be understood by A as a party for A. Of course A cannot follow the rule: 'There is a surprise party for me at home: go home,' but Way & Whiting claim that it is implausible to require abilities to be so fine-grained. Consider another fine-grained ability example: suppose that B does not possess the concept of iron; it nevertheless seems plausible to say that B has a reason to eat spinach because it contains iron. They claim that it is sensible, for example, for C to advise B to eat spinach because it contains iron. The response is to make plausible a weakening of the ability condition in RC.¹⁶

I think that this is the right sort of response to these cases, and in chapter 2 of this thesis I build on this thought by arguing that correctly responding to a reason implies manifesting a certain kind of ability.

1.1.2. Reasons as Four-Place Relations

Another feature that is claimed to form part of the reason-relation is the *context* in which a reason holds. Consider Susan's reason again, but this time let us artificially suppose that a loving devotion is an appropriate attitude for Susan to take towards her children before the age of 14, and that after the age of 14 such an attitude is no longer appropriate (since her attitude will foster an unhealthy attachment in her children). So, whether or not Susan has a reason in this case is time-dependent. Reasons also seem to be dependent on other situational factors, such as place: whether you have a reason to tell the joke depends on whether you are at the pub or at the funeral.

¹⁶ Others reply to self-effacing cases by arguing that it is more plausible that if A is not able to go to the party on the basis of the fact that there is a surprise party for A there, then that fact cannot be a reason for A to Φ . See Kiesewetter (2016: 10-2).

The apparent context-sensitivity of reasons has led some philosophers to expand the number of relata of the reason relation, with some claiming that we ought to include relata which denote context or circumstance.¹⁷ Thus a four-place reason relation could then be stated as follows: R is a reason for A to Φ in context C.

One source of difficulty with this formulation of the reason relation is that since context and circumstance can be obscure or indeterminate, it makes it correspondingly difficult to understand what actually counts as a reason. In my artificial example of Susan I assumed that the link was reasonably clear between a mother's loving devotion and respect for that mother. But it may be otherwise, and whether it is may depend on facts about Susan herself, Susan considered as a mother, or as a wife, or as a daughter; it may depend on facts about her children, facts which may vary considerably over time; it may depend on instrumental considerations such as whether Susan in acting at a time, exemplifying her loving devotion to her children, causes other children in the vicinity to become murderously jealous, in which case she ought not to have exemplified that attitude, and we have reason not to respect her for so exemplifying it. This is a problem because, as noted above, reasons are plausibly *for* somebody, and if reasons are for somebody then they should be capable of being understood by that somebody. But if the four-place formulation is right, it looks like it will often be exceedingly difficult for anybody to grasp what reason they have. Since it is plausible that many of us do in fact understand many of the reasons we have, this formulation of the property of a reason is thrown into doubt.

Others, notably Dancy, have cautioned against building relata of context into the reason relation.¹⁸ According to Dancy, when we think about cases involving reasons we need to pay close attention to what counts as part of the reason, and what forms part of the context in which that consideration is a reason. Dancy distinguishes between reasons, whose job it is to favour attitudes or actions, and what he calls

¹⁷ Such as Scanlon (2014: 31) and Väyrynen (2011: 186).

¹⁸ See Dancy (2004a: 38ff). See also Schroeder (2007a) and (2007b). He tells us that he finds it difficult to understand the reason relation as a four-place relation, and allows the distinction between favouring and enabling (what he calls 'background conditions').

enabling and disabling conditions, which form part of the context which allow that reason to be a reason.¹⁹ For instance, my promising to do it favours my doing it, while that my promise was not made under duress does not itself favour doing it but is a condition without which my promise would not favour doing it.²⁰ Of course, the agent-place and the attitude/action-place do not favour either; rather they are what is favoured by the consideration. Can circumstance play *both* the role of enabling *and* being favoured?²¹ Surely not, since that would render C's role as an enabling condition redundant.

This view tells us that C is a condition that enables R to count in favour of [A's Φ -ing in C]: that I am able to do it is a condition that enables my promising to do it to count in favour of my doing it if I am able to do it. But the condition is idle since my ability to do it is already contained in the reason-relation. So, we have a choice between accepting the existence of enabling conditions as a condition of favouring which plays no role in the favouring relation, or denying the existence of enabling conditions (at least as Dancy understands them) and understanding circumstance as part of what is favoured by a consideration.²²

There are at least two more ways that a sceptic of the four-place relational analysis of a reason could proceed. He might understand the reason relation as a three-place predicate, with consideration, attitude/action, and agent as relata, and argue that we can understand circumstance by sole reference to the agent, thus building circumstance into the agent-place. Alternatively we could replace 'agent'

¹⁹ See Dancy (2004a: 38-40).

²⁰ *Ibid*: 38.

²¹ By virtue of standing on the right hand side of the favouring relation such that what is favoured is that A Φ s in C. This must be the case since if C is a relata of the favouring relation it must either favour or be favoured.

²² Note that Dancy is still concerned to capture the context-sensitivity of reasons; indeed a central claim of Dancy's view of the structure of reasons (his so-called 'holism in the theory reasons') is that a reason which favours Φ -ing in one context can be no reason, or a reason against Φ -ing in another context. However, holism about reasons depends on the denial of the claim that context figures as part of the *consideration*, R, that favours. I discuss holism below.

with ‘circumstance’ and understand the agent as a feature of the circumstance. If these strategies are plausible, we have reason to believe that the extra relata in the reason relation are redundant.

Take the first strategy first. This view tells us that we can retain the constraint of circumstance on a reason by reference to the agent-place, that is, with reference to facts about the agent. In the promising example, that I can do it is a circumstantial constraint on my reason, which is also a fact about me. However, clearly not all circumstantial facts relating to a reason for A are about A. Suppose that I no longer have a reason to keep my promise if the person to whom I promise dies. That the person to whom the promise was made died is not a fact about me but it is nevertheless relevant for the assessment of the reason. So this view, on its own, does not give us the right result.

Alternatively, we might seek to understand whom the reason is for by reference to the circumstances where the reason holds, the agent simply being mentioned as part of the circumstances in which the reason holds when the reason applies to them. However, I find it difficult to understand how to formulate this version of the reason relation without making the concept of a reason more obscure.

1.2. Complete Reasons

A closely related question about the relation between a reason and its context is whether we should understand context as forming part of the content of a reason, i.e. as figuring in R. Here is an example of what I have in mind: suppose that the fact that conversing with your friend is pleasurable is a reason to do so. Pleasure, in this context, is reason-providing. Now suppose that you are torturing somebody, from which you derive pleasure. Given the shift in context, from conversing to torturing, pleasure is no longer reason-providing.²³

²³ This example assumes the sensation account over the attitude account of pleasure. The sensation accounts tells us that what counts as pleasure are the sensations felt, while the attitude accounts tells us that pleasure is the attitude of enjoyment *of something*. On the attitude account the sadistic and non-sadistic pleasures would

In other contexts pleasure even seems to become a reason *against* responding in a certain way. Suppose that it is the job of a prison guard to discipline his wards, which in practice can occasionally involve physically harming them. In this context it seems plausible that your deriving pleasure from this aspect of the job is a reason *against* taking the job in the first place. Not only does it seem wrong for a guard to take pleasure in these unfortunate but necessary parts of the job, we might think that his being so renders him less able to do his job well (his enjoying it might lead him to be over-zealous in meting out punishment).

These examples show a few striking ways in which not only the weight, but also the valency of a reason can shift given a shift in context. Dancy calls this phenomenon ‘holism in the theory of reasons’.²⁴ Stronger and weaker versions of the claim are usually distinguished:

Strong holism in the theory of reasons	A feature that is a reason in one case <i>must</i> be able to be no reason at all, or an opposite reason, in another. ²⁵
---	---

Weak holism in the theory of reasons	A feature that is a reason in one case <i>can</i> be no reason at all, or an opposite reason, in another.
---	---

The strong version of holism rules out the possibility of what are sometimes referred to as *invariant* reasons. An invariant reason is a reason that never changes its valency. If it is a reason to Φ it is always a reason to Φ . The weak version of holism does not rule out the possibility of invariant reasons understood as considerations that do not *as a matter of fact* change their valency. That fact is simply due to

simply count as different kinds of pleasures. See Feldman (1997) for a review of the accounts.

²⁴ Dancy (2004a: 7)

²⁵ *Loc. cit.*

contingent features of the world. What is referred to by Dancy as *atomism* about reasons is the denial of either strong or weak holism about reasons.²⁶

Whether we will be persuaded by examples like these and others that Dancy adduces may depend on whether we think that context is a proper part of the reason relation, or whether context must be left out of a complete specification of the reason. Take the fact that the joke is funny as a reason to tell it, and suppose that context is a proper part of the reason relation. We can include the relevant contextual information in the reason as follows:

- (i) [The fact that the joke is funny, and that A is at a pub] is a reason for A to tell the joke
- (ii) [The fact that the joke is funny, and that A is at a funeral] is a reason for A not to tell the joke

The bracketed parts indicate the content of the reason, the fact which stands in the favouring relation. In light of examples like these, I claim that *weak* holism in the theory of reasons is plausible. I defend it against some objections in this section.²⁷

As mentioned above, debates over holism versus atomism centre on whether one accepts that context figures ineliminably in the consideration-place (R) of the reason relation. There are different views about what this conception of a normative reason amounts to. A prominent explicit account of one such conception is given by Raz, who calls this a ‘complete reason’:

A complete reason consists of all the facts stated by the non-redundant premises of a sound, deductive argument entailing as its conclusion a proposition of the form ‘There is a reason for P to V’ (where P stands for an expression referring to an agent or a group of agents, and V for a description of an action, omission, or a mode of conduct) [...] I refer to any fact stated by

²⁶ *Loc. cit.*

²⁷ See for example Crisp in Hooker & Little eds. (2000: 36-40), Hooker in Hooker & Little eds. (2000: 13-5), and Hooker in Lance & Strahovnik eds. (2008: 12-30).

any proposition which can be a non-redundant premise in a sound argument of the kind just described as a reason or as part of a (complete) reason.²⁸

What counts as a non-redundant premise is, of course, an issue of contention between Raz and defenders of holism about reasons. Building on the idea of a complete reason, Raz presses the following objection to the claim that context is not part of the content of a reason: if context is part of what determines the rightness or wrongness of acting in some way (what Raz calls its ‘evaluation function’) and that context is *not* part of the content of a reason for action, then reasons for action and ‘evaluation’ will come apart.²⁹ Why is this problematic? Suppose that you have a reason R that favours in C but disfavors in C*, then it ought to be intelligible just why this is the case. Moreover, the explanation should be a justifying one, i.e. one that should appeal to normative reasons.

But if we accept holism about reasons that explanation necessarily will not appeal to normative reasons. The explanation will figure as part of the enabling conditions. So holism is false because it has a false implication. It precludes us from explaining why a consideration changes valency in terms of normative reasons.

A second way of cashing out the idea of a complete reason comes from Cullity:

Surely a fact can only be a reason for a given action if the obtaining of that fact is sufficient to make it the case that there is a reason to perform that action. However, the fact that an action would be enjoyable cannot be sufficient to make it the case that I have a reason to do it, if there are kinds of enjoyment that are not reason-giving. Therefore, when there is a reason to do something enjoyable, the reason cannot simply be *that it is enjoyable*. In order to mention the whole of the reason for doing it, we need to mention the

²⁸ Raz in Hooker & Little eds. (2000: 59 fn.22). Schroeder (2007: 23, fn.1) suggests that the deductive claim should be replaced with an explanatory one. I take no stand on the issue here.

²⁹ Raz (2006: 110-1).

kind of enjoyment that it involves, specifying that it is one of the kinds that is reason-giving, rather than one that is not.³⁰

The lesson that we should take from the fact that pleasure can be reason giving in some contexts and not others is not that reasons are holistic in nature, but that the fact that something is pleasurable is not, by itself, a reason. The reason must, for example, be *non-sadistic pleasure*. So where the holist would happily admit that pleasure is reason-giving in the right context, the atomist would deny the claim.

We may also appeal to the idea of *ultimate* reasons to support an argument against the claim that context resides outside the content of the reason. Crisp claims that in order to be a substantive claim holism must be true of what he calls ‘ultimate’ reasons.³¹ An ultimate reason is a reason that cannot be further explained in terms of other normative reasons. For example, we might think that A has a reason to take the bus because he has a reason to go to the dentist, because he has a reason to avoid suffering, because he has a reason to promote his own well-being.³² A’s well-being is an ultimate reason for A to do it, because there is no *further* normative explanation why A has a reason to do it. The argument then is that for any putative reason that one gives, one can provide a further, ultimate reason which explains that first reason. Given the existence of ultimate reasons, we are to see that the appearance of reasons’ context-sensitivity is merely an appearance. We can in principle always appeal to an ultimate reason at bottom and ultimate reasons – like non-sadistic pleasure – do not plausibly change their valency.

We can reply to this last objection immediately. As noted above, we should distinguish between a weaker and stronger version of holism about reasons. We can define holism as the claim that a consideration’s status as a reason *can* depend on context, or that a consideration’s status as a reason *must* depend on context. If we accept the former claim we can accept that there are invariant reasons ‘... not because they are reasons but because of their specific content’.³³ Given the possible

³⁰ Cullity (2002: 174). See also Nagel (1970: 90-5).

³¹ Crisp in Hooker & Little eds. (2000: 35ff).

³² Crisp (2000: 8).

³³ Dancy (2004b: 136). See also McKeever and Ridge (2006: 30ff).

existence of invariant reasons, ultimate reasons are commensurate with holism. Thus we need not see the possibility of an ultimate reason as a threat to *weak* holism.

In addition, I do not find the above arguments persuasive. As noted above (§1.1.1) a plausible constraint (RC) on a reason is an agent's ability to respond to it. RC was understood in the *de dicto* sense that in order for R to be a reason for A to Φ , it must be *possible* for A to Φ on the basis of R. As we have seen, the defender of the idea of a complete reason adheres to the claim that potential contextual defeat always forms part of the content of the reason. We have seen some disputable cases of this, for example where only pleasure that is *non-sadistic* is reason-giving (the content of the reason being: 'X is pleasurable and X is non-sadistic'). If the action were sadistic, that would defeat the reason.

How many potential defeaters are there for a reason? In the case of pleasure it is difficult to imagine many other potential defeaters, other than sadism. But other reasons may be quite different and may admit of many potential defeaters. There may be an indefinitely long number of potential defeaters for a reason, and this possibility raises a serious problem for complete reasons given the truth of RC. Clearly, since an agent cannot act on the basis of an indefinitely long disjunction of potential defeating considerations, such a complete reason would be ruled out by the constraint imposed by RC.

The argument runs:

1. In order for R to be a reason for A to Φ , it must be possible for A to Φ on the basis of R (RC).
2. Complete reasons are compatible with the possibility that there are reasons whose content (R) contains an indefinitely long disjunction of considerations.
3. It is not possible that an agent act on the basis of an indefinitely long disjunction of considerations.
4. So: complete reasons are compatible with a contradiction, i.e. complete reasons entail the possibility of a contradiction.
5. Contradictions are not possible.

6. So: either it is not possible that there are reasons whose content contains an indefinitely long disjunction of considerations, or there are no complete reasons.

There are two moves open to the defender of complete reasons here. They can argue that RC is false, or they can argue that it is not possible that a complete reason can contain an indefinitely long disjunction of considerations. I have already independently motivated RC and argued that it is a plausible principle. Moreover, it commands a fairly wide acceptance among theorists. So perhaps the best strategy for the defender of complete reasons is to argue that it is not possible that there are complete reasons that contain an indefinitely long disjunction of defeaters.

Admittedly, as in the case of pleasure, it is difficult to enumerate more than a few defeaters, such as cases in which the pleasure taken is sadistic or malicious. On the other hand, that it is not easy to discover which potential defeaters exist for reasons for at least one reason does not itself undermine the possibility that there are such indefinitely long conjunctions of defeaters. The fact that we have not yet found more defeaters for a reason may simply point to the difficulty of discovery, not that there is nothing to be discovered. In order for the argument to work, we only need to defend the claim that it is *possible* that there are indefinitely long disjunctions of defeaters for reasons, and it is hard to see on what basis we could sensibly rule this out.

Cullity, who ultimately rejects the idea of a complete reason, provides a slightly different argument for weak holism beginning from RC.³⁴ He claims that what it is to be motivated to act on the basis of a reason is a form of ‘self-explanation’. It is an explanation of why an agent acted as they did which takes the form of that agent’s belief about what their normative reason is, and is primarily an explanation directed towards that agent. However, where an agent cites their belief about what normative reason they had in explaining to themselves why they acted in some way, their explanation, Cullity argues, is subject to a ‘background expectation of normality’.³⁵ That is, there are constraints on what an agent should sensibly ask of themselves in providing this explanation of why they acted. In particular, an agent

³⁴ Cullity (2002).

³⁵ Cullity (2002:178).

who has no evidence to the contrary is entitled to assume the absence of normally-absent defeaters.

Thus, the rational explanation of why an agent acted for some reason plausibly should not contain defeaters. Compare this kind of explanation with other kinds of explanation. Suppose that we want an explanation why the train was delayed. Some considerations are clearly relevant to that claim, for example that there was a fault with the engine that caused the train to slow down. Other considerations do not seem relevant to this explanation at all, such as that if the train was struck by lightning it would have caused an electrical surge on the line that would have sped the train up by 20mph, causing the train to arrive at the station on time. The absence of defeating considerations is not *generally* part of an explanation for something. The point generalises to the sort of rational explanation invoked above. Thus, given RC, the constraint on rational explanation carries over as a constraint on how we should individuate a normative reason. Normative reasons should not *generally* include as part of their content absent defeaters. So we should give up the idea of a complete reason and adopt an alternative conception of a reason: reasons are *pro tanto* unless they are undermined. This is what Cullity refers to as a *presumptive* reason and it preserves the distinction between a reason and its context.³⁶

Against this, Väyrynen argues that there is no necessary relation between a presumptive reason and background expectations of normality.³⁷ He objects that there is nothing about a presumptive reason itself that implies normal background conditions. According to Cullity, an agent is entitled to assume that a consideration that is generally reason-giving is a (presumptive) reason. Väyrynen claims that in one sense Cullity is thinking that such reasons function specifically as *epistemic defaults*, in that they entitle an agent to *believe* that this a reason to act, by default.³⁸ That is part of the ‘self-explanation’ invoked by Cullity. In providing that self-explanation an agent is (a) entitled to assume that R is a reason *if* it is normally a reason, and (b) entitled to exclude absent defeaters.

³⁶ *Ibid*: 188.

³⁷ See Väyrynen (2004: 53-79).

³⁸ *Ibid*: 65.

But suppose we are in a ‘nasty world’ in which, say, promising is usually defeated as a reason. In this world an agent is not entitled to assume that promising is reason-giving, but just as in the actual world it may be a reason unless undermined. That is, in our world and the nasty world, promising is a reason unless it is undermined, but in the nasty world it is consistently undermined. If normality has nothing to do with the conditions for some consideration being a (presumptive) reason, then Cullity cannot appeal to normality to explain how we individuate such reasons.³⁹

To put the point another way, we shouldn’t (according to Väyrynen) claim that some consideration, R, should not form part of A’s reason to act for the reason that R isn’t part of the background expectation of normality, because (i) A may not be entitled to assume that conditions are normal and (ii) in cases where A is not in a ‘normal’ world, and thus not entitled to assume normality, the consideration which was reason-giving in the ‘normal’ world may *still* give A a reason to act.

I disagree with Väyrynen. In particular, I think that some normality constraint is plausibly applied to explanation in general, and rational explanation is thus also subject to some such condition. In the nasty world expectations of normality shift, so that the relevant disabler for promising becomes part of what is normally expected in that world. Similarly, in the train example, if you live in a world in which trains are typically delayed because they are struck by lightning, that consideration, in that world, becomes relevant to the explanation of the event.

If that is right, then given our adherence to RC, this constraint on a rational explanation will carry over to the content of a normative reason. According to RC, normative reasons are the sorts of things that an agent can act on the basis of, where that is understood as a kind of explanation by an agent, to themselves (in terms of normative reasons) of why they acted as they did. If, in the nasty world, the kind of explanation that is plausibly demanded is different to the normal world, in that it will include a defeater, then that defeater will be included in the content of the normative reason.

So I have conceded that, if we accept Cullity’s background condition as a way of constraining reasons, defeaters may *sometimes* form part of the content of a

³⁹ *Ibid*: 61-6.

reason. But that is not to concede that the form that a normative reason must take is of a complete reason. Given this way of understanding things we can still accept the distinction between reason and context (whether defeater or enabling condition) and still have a principled way to individuate the content of a reason. In many cases, considerations of context will not form part of the content of the reason because of the normality constraint and, if that is right, it suffices to falsify the complete reasons view, since that view tells us that context *always* forms part of the content of the reason. Väyrynen's argument does not put that conclusion in jeopardy.

If these arguments are right, then the notion of a complete reason should be discarded because it either entails the possibility of an indefinitely long disjunction of considerations such that within the content of a reason it implies a contradiction, or it entails a mistaken view of rational explanation.

1.3. How do Reasons Interact with Each Other?

One confusing aspect of the literature about normative reasons is the sheer number of disparate distinctions made about reasons. Some of the confusion stems from a lack of a uniform terminology for these distinctions, but it is mainly due to the variety of conflicting views about the structure of reasons. In this section I will briefly survey some of these views.

So far I have been discussing the idea of a consideration which speaks in favour of a response, what is often referred to as a *pro tanto* reason. We noted that this is a consideration that counts in favour of a response for some agent. We can now further specify that it counts in favour of that action/attitude so long as it is not outweighed by some other reason, or cancelled by some consideration or circumstance. That I will be late for the meeting is a *pro tanto* reason for me to leave now, and remains a *pro tanto* reason to leave now, even if I have stronger reason to stay in and look after my mother. Thus a distinctive feature of *pro tanto* reasons is that they each have a *weight*. To appeal to a familiar metaphor, reasons are like weights that go into the pans of a scale, and the way in which the scale tips tells us the strongest reason we have to respond in some way.

Pro tanto reasons can retain their force, relative to their weight in the face of competing considerations. This can be shown by considering what attitudes it would be appropriate to adopt in cases of outweighed *pro tanto* reasons. My reason to stay in and look after my mother outweighs my reason to leave, but it is, nevertheless, appropriate that I feel regret that I cannot go to the meeting.⁴⁰ That regret is appropriate suggests that my reason to go to the meeting still retains some force for me. On the other hand, there seem to be cases in which *pro tanto* reasons are cancelled by some other reason, in the sense that they completely lose their force as reasons. Suppose I have a reason to give back the book, since I promised my friend I would do so. If my friend freely releases me from this promise, my reason seems to be cancelled: I no longer have *any* reason to give back the book. We can, again, consider whether it is appropriate to feel regret in this case to judge whether the promise retains any residual normative force. Since it does not seem appropriate to feel regret, that gives some substance to the distinction between outweighing and cancelling.

Another distinction often cited with respect to reasons is *modification*, which can affect the weights of reasons directly without themselves constituting reasons.⁴¹ Certain considerations can modify the weight of a reason by ‘intensifying’ it, i.e. by increasing the weight of the reason, or by ‘attenuating’ it, i.e. by decreasing the weight of the reason.⁴² For example, suppose that I ought to buy you a present because it is your birthday. Now suppose that you are my daughter. Plausibly, this further consideration intensifies the reason that I already have, without adding an additional reason to the stock already apparent to me.

⁴⁰ I discuss this ‘regret principle’ further in §4.9.

⁴¹ See Dancy (2004a: 41-2).

⁴² *Loc. cit.*

1.4. The Distinction between Reasons For and Reasons Against

I here provide arguments to support a different distinction among reasons that I have implicitly relied upon thus far, between reasons for, and reasons against. I deny that that reasons against are a kind of reasons for. Defences of the distinction have been mounted on two fronts. First, that the distinction is required in order to properly account for choices among fine-grained alternatives, and second that the distinction is required in order to respect the role that criticism plays in thinking about reasons.⁴³ I discuss both approaches here.

There are a number of ways of understanding reasons against in terms of reasons for.

1. A reason against Φ is a reason for some alternative to Φ
2. A reason against Φ is a reason for all alternatives to Φ
3. A reason against Φ is a reason for not- Φ ⁴⁴

Let us take each interpretation in turn, focusing on reasons for action. According to (1) reasons against acting in some way are just reasons for acting in some other way. Suppose that the fact that the film is terrible is a reason not to watch it. What alternative is the fact that the film is terrible a reason for? Perhaps it is a reason to dissuade others from watching the film, or to try to suppress it, or to write a damning article about it in a popular newspaper. (1) tells us that the reason against is equivalent to a single alternative, but gives us no principled way to decide which of these alternatives that should be.

Second, it is plausible that, in general, a reason, whether for or against Φ , should be understood as taking Φ as its object. The fact that you wouldn't enjoy skiing isn't thereby a reason for you to play cricket. That the film is terrible is a reason not to go to the cinema, but isn't thereby also a reason to go to the restaurant. So I doubt whether (1) or (2) are plausible putative analyses of reasons against.

⁴³ On the former strategy see Snedegar (2017); on the latter strategy see Greenspan (2005).

⁴⁴ See Snedegar (2017: 729)

Now suppose that as (3) tells us reasons against Φ are reasons for not- Φ . This is a more popular view. It comports with the idea that a normative reason to Φ takes only Φ as its object, and seems to get some cases right. The fact that you wouldn't enjoy skiing is a reason against going skiing, and also, one might think, a reason for not going skiing. It is easy to believe that these claims are equivalent.

I think, however, that we have good reason to reject (3), that if we have reason to reject (3) we have no plausible analysis left of reasons against in terms of reasons for, and thus that reasons for and reasons against are non-equivalent relations.

First, reasons against Φ are not equivalent to reasons for not- Φ because justifying an omission can be different from undermining an action. Suppose that the fact that it shows your steely resolve is a reason for refraining from taking the drug. Here we have a reason which supports the case for the omission. Now suppose that the fact that the drug is addictive is a reason against taking the drug. Here we have a reason which undermines the case for the action. According to (3) these reasons *both* function as reasons for the omission, but that doesn't seem right. The first tells you something specifically about the omission, the second something specifically about the action, where these are two quite different objects. For example, if you are addicted to the drug, it is difficult to refrain from taking it, and easy to take it. So the point mentioned above that normative reasons take a single object is appropriate here. The omission is a different object with different properties than the action, and thus should be treated differently.

Second, if reasons against Φ are equivalent to reasons for not- Φ , then treating omissions and actions as interchangeable should lead us to claim that reasons against not- Φ are equivalent to reasons for *not-not- Φ* , i.e. reasons *for- Φ* . But that is false. Suppose that the fact that the drug is pleasurable is a reason for taking it, and that the fact that it will contribute to your boredom is a reason against refraining from taking the drug. Claiming that your reason against not- Φ is equivalent to a reason for Φ is clearly bizarre. It makes no sense to suggest that the fact that it will contribute to your boredom is a reason for taking it. Again, I take this as evidence that a reason against an omission is not equivalent to a reason for the action.

Let me turn to a further problem for (1) articulated by Bedke.⁴⁵ Bedke argues that our rejection of (1) and (2) implies a further very serious problem for (3). In rejecting (1) and (2) we rejected the idea that in general a reason against Φ is also thereby a reason for Ψ (some or all alternative(s) to Φ). Now consider the following case. You have three options open to you: either kill the innocent, refrain from killing the innocent, or eat the ice-cream. Since, according to (3), reasons against Φ are reasons for not- Φ , you have weightiest reason to refrain from killing the innocent. Your reason to eat the ice-cream is comparatively less weighty. So according to (3) you have most reason to refrain from killing the innocent *rather than* eating the ice-cream. But that seems absurd.⁴⁶

The obvious response to this problem is to reply that you have most reason to eat the ice-cream because *that is a way of refraining from killing the innocent* and, in addition, it would be enjoyable to eat the ice-cream. How should we amend (1)? Snedegar suggests that we appeal to the following principle:

Facilitative Principle (FP) If R is a non-derivative reason for A, and doing B facilitates doing A, then R is a reason for B.⁴⁷

You have a non-derivative reason to refrain from killing the innocent, and eating the ice-cream is a way of refraining from killing the innocent, so your reason for refraining from killing the innocent transmits to your reason for eating the ice-cream.

Unfortunately we have good reason to reject FP. That is because it is false in conjunction with (1). (1) tells us that a reason for not Φ -ing is equivalent to a reason against Φ -ing. But *any* action is a way of not Φ -ing, so FP and (1) imply that for any reason you have for not Φ -ing also implies reasons for any alternative to Φ -ing. And, given some range of alternatives, an alternative to Φ may in fact be *worse* with respect to the objective of that reason than the reason against Φ . Suppose that you could visit café A, café B, or café C. The fact that there are long queues at café A is a reason against going to café A. A way to refrain from going to café A is to go to café

⁴⁵ Bedke (2017).

⁴⁶ *Ibid*: 23.

⁴⁷ Snedegar (2017) who also argues that this creates a problem for (1).

B. So according to FP the crowdedness of café A gives you reason to go to café B. But suppose that café B has *even longer* queues than café A. Then the mere fact that it is a way of refraining from going to café A is not enough to give you reason to go there, on the basis of that consideration. Given that café B is more_crowded than café A, it can't sensibly be true that the fact that café A is crowded is a reason to go to café B.

Bedke's problem forces the defender of (1) to appeal to the idea that reasons against can somehow be reasons for an alternative. But for the reason just given, and for the reasons given above, we should not accept this equivalence. Thus we have good reason to reject all three analyses of reasons against in terms of reasons for.

The second sort of defence of the distinction reflects intuitions about the appropriateness of criticism with respect to reason for and reasons against. Greenspan considers the following cases:

I recall an occasion when a University administrator, trying to offer positive motivation for faculty to serve on committees, appealed to the possibility of attaining power in the University. Now, I would grant that this *is* a reason to serve — and a reason for me, a reason I 'have', at any rate once it is brought to my attention — though I think I would still be within my rights, rationally speaking (as well as otherwise), to turn it down. [...]

[However] maybe it is a serious reason against failing to serve on a committee that some degree of power in the University, more than I would otherwise have, is needed to keep others from instituting policies that would undermine my academic or other goals. In that case — if I were persuaded that I had better serve on a committee to prevent some bad outcome — it *would* seem irrational not to serve.⁴⁸

Greenspan here distinguishes between a reason that offers positive motivation for acting in some way, but is waivable by the agent who has that reason. That serving on the committee will gain you some modicum of power is a feature of that action

⁴⁸ Greenspan (2005: 388).

that explains why serving is worth doing. But just because something is worth doing it doesn't follow by necessity that your failure to do it highlights a rational failure in you. You are within your rights to decide not to act on this basis.

On the other hand there are negative features which undermine the case for failing to serve on the committee, such as that if you fail to serve, a ruthless administrator will take your place and cause widespread misery. This results as a consequence of your omission and does not seem to be rationally waivable by the agent. You are open to criticism if you fail to take this fact into account, in a way that you are not if you fail to take the fact that your serving on the committee will gain you a modicum of power.

As Greenspan suggests, it is not helpful to claim that these are both reasons for. There is clearly *some kind of* asymmetry between the two kinds of reasons that is reflected in our judgements about whether the agent can permissibly waive each reason. That again suggests that we should reject analyses of reasons against in terms of reasons for as in (1)-(3). The difficult question is what exactly characterises the distinction between reasons for and reasons against?

Most theorists recognise a distinction between positive, negative, and neutral reactive attitudes, and that these are non-interchangeable, distinct attitudes that can be held towards the same object. One way to shed light on the distinction between reasons for and reasons against is to appeal to the attitudes that it is appropriate to hold towards an agent who has reasons. A plausible generalisation might be: where A has undefeated reason *against* Φ and, nevertheless, Φ 's, A is appropriately subject to others' holding negative reactive attitudes towards A, as well as it being appropriate for A to hold negative reactive attitudes towards themselves. On the other hand, where A has undefeated reason *for* Φ and they fail to Φ , it is generally at most appropriate only for A to hold negative reactive attitudes (such as regret) towards themselves.

In the cases above, your failing to serve on the committee on the basis of the fact that you will gain a modicum of power may make appropriate at most an attitude of regret towards your decision. In the second case, your failing to serve on the committee given that you know that a ruthless administrator will take your place

plausibly also makes appropriate your co-workers' holding of negative reactive attitudes towards you.

Although this seems a decent analysis of what is happening in the committee example, aren't there examples of reasons for that *don't* imply reasons against the alternatives, but which you are criticisable for failing to act on? Suppose that the fact that it will improve your piano-playing is a reason for practising, and that you decide not to practice. Isn't your piano teacher justified in criticising you for failing to practise? I don't think so. I think we should say that they are justified in being *disappointed* that you did not practice, but it is appropriate for them to hold critical attitudes towards you only if, e.g. you promised them that you would practice. Your promising that you will practice gives you a reason for practising and a (*pro tanto*) reason against acting on alternatives (such as: 'The fact that I promised my teacher that I would practice is a reason against going to the cinema tonight').

Appealing to positive, negative, and neutral reactive attitudes to identify reasons for and reasons against may not be decisive but nevertheless seems to be a potentially fruitful way to understand the asymmetry. I explore other ways to understand the differences between these reasons in chapter 4.

1.5, The Weight of Reasons

As we have seen, normative reasons are *pro tanto* considerations. They count in favour or against some response *to some extent* and some are stronger than others. That the butter is tasty is a less weighty reason to eat it than the unhealthiness of eating the butter is a reason against. A discussion of the structure of reasons should thus have something to say about the weight of reasons.

According to some, there is a pressing issue for a defender of Reasons Fundamentalism (RF).⁴⁹ Correctly understanding the weight of normative reasons, they claim, constitutes a potentially decisive objection to RF, drawing on the following argument:

⁴⁹ Recall that this is the disjunctive claim that normative properties can be reduced to reasons, and that reasons cannot be reduced to other normative properties. I defend a weaker version (WRF) of this claim that is restricted to only deontic properties.

1. All normative reasons have a weight.
 2. To say that a reason has a weight is to make a normative claim.
 3. Normative claims are, fundamentally, to be explained in terms of reasons (RF).
- C. So: The weight of reasons is to be explained in terms of reasons.

But presumably, given RF, we cannot non-circularly understand the weight of a reason, so perhaps we should give up RF. What reason do we have to believe premise 1? According to Mark Schroeder:

... [T]o say that one reason is weightier than another is to make a normative claim akin to the claim that one person is more admirable than another. One person is not more admirable than another because she is more widely admired, or because it is easier for people to admire her, but because it is *appropriate* or *correct* to admire her more. Similarly, one reason is not weightier than another because it carries more weight in deliberation, but because it is *correct* for it to carry more weight in deliberation.⁵⁰

This looks initially plausible. For a reason to have a certain weight in deliberation doesn't look like a descriptive fact about that reason. When we ask whether a consideration favours some response we are making a paradigmatically normative claim. Similarly when we ask whether a consideration favours *to a certain extent* (that is, has a certain weight) we also seem to be making a normative claim.

So, at the very least, a defence of RF (and WRF) must be able to explain either what is wrong with one of these claims, or why circularity is not a vice of the view. In chapter 5 (§5.1.4) I defuse this objection by arguing, following Stratton-Lake, that in order to make sense of premise 2 we must sharply distinguish between a reason and its weight. But in order to do so on the conception of a normative reason that I have been sketching we must appeal to an implausible theory of weight. So I don't think that the objection is particularly strong against RF.

⁵⁰ Schroeder (2007: 129).

Nevertheless, I will examine what we might say about the weight of normative reasons here in the service of explaining the structure of reasons. I'll first consider the most prominent view that the weight of reasons is to be explained in terms other than reasons, namely that this explanation should be given in terms of promoting value. I'll argue that this view is not plausible, and that in fact any such reductive view of the weights of reasons is likely to fail. I then go on to sketch some ways we might better understand the weight of reasons.

1.5.1. Weight-First or Weightier-Than-First?

An issue that arises when thinking about the weight of reasons is whether the weight of a reason can be determined prior to any determination of its being *weightier than* some other reason. To paraphrase Scanlon: does the weight of a reason have any significance for us apart from the considerations of potential conflict among our reasons?⁵¹ It certainly seems to have a significance that goes further than potential *conflicts* among reasons, since we also talk about the weight of reasons in the context of *combining* reasons. It makes good sense to say that of two reasons, A and B, both of which favour Φ -ing, that A favours Φ -ing more strongly than B so that were A not to obtain, B would not be a sufficient reason to Φ , while were B not to obtain A would remain a sufficient reason to Φ .⁵²

Let us call talk of the conflict and combination of reasons the *interactions* between reasons. Following Lord & Maguire, we can call 'weight-first' those views that understand the weight of a reason in isolation from interactions among reasons, and 'weightier-than-first' those views that deny that weight can be understood in isolation from the interactions among reasons.⁵³ Does the weight of a reason, then, have any significance for us in isolation from interactions among reasons?

⁵¹ Scanlon (2014: 112).

⁵² Normative reasons are *sufficient* to Φ where those reasons are at least as strong as reasons to do otherwise. This is an example of a verdictive property that involves normative reasons. I discuss verdictive properties of reasons further in §4.7.1.

⁵³ See Lord & Maguire in Lord & Maguire eds. (2016: 14-22).

It seems reasonably easy to understand the weight of a reason in isolation from any actual interaction that it has with other reasons. That it would be a sadistic act of torture seems intuitively to be a weighty reason not to do it even where there are no other reasons around. On the other hand, it seems more difficult to understand the nature of the weight of a reason in isolation from any *possible* interaction with other reasons. That seems to me a difficult question to answer because it implies that it is possible that we can understand normative reasons in isolation from practical reasoning. Abstracting from possible interactions among reasons also courts the danger of neglecting potential modifying interactions among reasons (see §1.3). This *suggests* that ‘weightier-than’ is fundamental. However, in what follows I try not to let much hang on the ‘weightier-than’/ ‘weight’ distinction and focus on specific views.

1.5.2. Reductive Views

The simplest of these views identify the weight of a reason with a function F , of some other property P , which can be expressed as a magnitude. The weight of any reason will be determined by this function as it increases or decreases. The greater F is, the weightier your reason for responding is.⁵⁴

Suppose that P is value. Then according to this view, we can determine the weight of any reason in terms of some function of P (suppose that P is a state of affairs). A simple ‘value-based’ view will claim that the more valuable the state of affairs promoted by the reason, the weightier that reason is. For example, the fact that keeping your health is more valuable than experiencing some gustatory pleasure explains why you have a stronger reason to refrain from eating the chocolate bar.

An objection to these kinds of views is that they are too simple in a certain respect. According to these views the larger the magnitude of P the weightier the reason you have to respond in that way. Suppose that P has a very high order of

⁵⁴ We also might mean by strength not the degree to which a reason favours, but that the reason favour an attitude that is itself stronger than another reason for another attitude. For example R may be a stronger reason than R' because R favours outrage, while R' favours mere disapproval.

magnitude but is almost impossible to achieve. Then you have a very strong reason to respond in such a way as to promote or bring about P. But it is plausible that the likelihood of being able to bring about P affects the amount of weight that you should place on trying to achieve P. So these views must supplement the simple thesis of explaining weight in terms of function of some quantifiable property with a further claim that the weight of a reason can also be affected by the probability that what it is a reason for will succeed.⁵⁵

With this qualification in mind we arrive at the following generic reductive view:

Weightier-than-first reductive view

For some set of reasons S to Φ to be weightier than another set of reasons S' to Ψ is an increasing function of (a) the promotion of some quantifiable property P, and (b) the probability that A will Φ .

In fact, a popular candidate for P is value. One version of this view holds that the weight of a reason is determined by some increasing function of the value promoted by the state of affairs which that reason favours.

1.5.3. Value Based Views

The value based view tells us that the weight of a reason is determined by an increasing function of the value of the state of affairs promoted by the response which that reason recommends.⁵⁶ The simplest such function tells us that the more

⁵⁵ See Maguire (2016: 236), Schroeder (2007) ch.7.

⁵⁶ I use this formulation in part to avoid evil demon style counterexamples which will arise for value-based views of weight that take the object of the response as the valuable state of affairs. Suppose that you have two options A and B, where B is somewhat more valuable than A. Now suppose that an evil demon will torture you

valuable the state of affairs promoted by responding in the way that the reason recommends, the weightier the reason one has to respond in that way. What is meant by ‘promotion’ here? I want to remain as neutral as possible about what it is to promote some state of affairs, so A’s promoting S (where Φ -ing promotes S and A has a reason to Φ) can be understood either as A’s causing S, constituting S, or as making S more likely.⁵⁷ I also take it for granted that value is a property of states of affairs. Finally, I am neutral on whether value is monistic or pluralistic.

The value based account of the weight of a reason nicely explains certain cases. In my example above it looks like an attractive explanation of why I have a stronger reason to eat the fruit than I do to eat the chocolate, which is because by eating the fruit I promote a more valuable state of affairs since healthiness is more valuable (assume for the sake of argument) than gustatory pleasure.

The value based account gains plausibility not only by providing a nice explanation in cases like the one above, but for other reasons. Ordinarily we refer to the weight of reasons in evaluative terms, in terms of whether the reason is good or bad. Since language suggests that we use specifically evaluative terms, we should understand the weight of reason in evaluative terms, and a promising way to cash that out is the value-based account just sketched. Further, when we are trying to figure out what we ought to do, one way we do so is to figure out what we have *best* reason to do, i.e. we try to find the best reason to respond to.⁵⁸ The value-based picture of the weight of reasons also makes immediate sense of a reason’s being better or worse.

However I don’t think that the value based account of the weight of reasons can be right. The value-based theory of weight is not general enough because it fails to explain the weight of all reasons. My example concerned reasons for action, and it seemed plausible, at least in that case to explain the weight of the reasons to perform certain actions as a function of the value promoted by those actions. Other reasons do not seem to gain from the *prima facie* plausibility of this treatment, in particular,

unless you choose A. You thereby have weightier reason to perform the less valuable option.

⁵⁷ See Finlay (2012), Kolodny in Star ed. (forthcoming), Maguire (2016).

⁵⁸ Cf. Gregory (2014).

reasons for belief. Suppose that I have stronger reason to believe that the package will arrive on Monday than that it will arrive today. Whatever story that we provide as to what explains why I have stronger reason to believe that the package will arrive on Monday it doesn't look plausible to explain why it is stronger in terms of the value that will be promoted by my believing that it will. That explanation seems to look in the wrong place.

We have to be careful not to make the value-based theory a straw man. As I said, we can think of value pluralistically, so that having a view on which the strength of an epistemic reason depends on value need not entail a crude pragmatism. We can suppose, reasonably, that the relevant value in this domain is, for instance, knowledge. You have stronger reason to believe that p than that q where believing that p promotes a world which contains more knowledge than believing that q. Still, although this allows us to avoid a crude pragmatism, the view remains controversial.

First, the view is revisionary without justification. Common sense often suggests that whether an epistemic reason p is weighty depends on the degree to which that reason raises the likelihood that p. The value-based view tells us that this claim could only be true in a derivative sense. The fact that a reason to believe that p raises the likelihood that p is true, on this view increases the weight of the reason *because* believing that p promotes a world in which there is more knowledge (or justified beliefs). But we usually find acceptable the claim that a reason is weightier if the evidence provided by that reason is weightier, without being committed to the stronger claim that the value-based view implies.

Second, the view is implausible as a view about what determines weight of epistemic reasons in particular, as opposed to a view about the value of believing for reasons. The value-based view may provide a plausible answer to the second question. But we are not here interested in answering the second question, we are interested in the first question.

Turning to reasons for action, I note that a popular argument against the value-based theory runs as follows. Suppose that the promotion relation for the value-based theorist is a maximising one, and that it is the case that performing some morally abhorrent act would be a means to producing a very small increase in overall value. Then according to the value-based view you have stronger reason to perform

the morally abhorrent act. But it seems deeply implausible to think that it is rational to perform acts because of such tiny gains of value.⁵⁹

An obvious response on behalf of the value-based theorist is to simply deny that the promotion relation takes a maximising form. Perhaps instead we should understand the promotion relation in a satisficing way.⁶⁰ Nevertheless, this revision is not trivial. I take this to be a problem for the view since any view of the weight of reasons ought to be as parsimonious as possible. That is to say that in formulating the structure of normative properties, we should be as neutral as possible on first-order normative views, such as moral theories. The value based view of weight threatens to do that in the cases mentioned above. That is because on some moral theories the reasons that arise from promises and other more ‘categorical’ considerations may not lose their force, or have their weights attenuated merely by the fact that they do not promote valuable states of affairs.

A second objection is that the value-based view finds it difficult to make sense of *partiality*. If the strength of a reason to act is determined solely by the amount of value that is thereby promoted by that act, it may imply that you have no more reason to save your wife than to save a stranger. The value-based theory misses this distinction as it overlooks the fact that the strength of reasons can also vary in light of facts about one’s *relation* to the valuable state of affairs.

A popular response to this objection is to draw a distinction between agent-relative and agent-neutral value, and to dispense with the latter. You have stronger reason to save your wife because you have more value, *relative to you*, to promote this state of affairs, than to save the stranger.⁶¹ A less controversial response claims that we should not reject or replace the idea of agent-neutral value, but make room for partiality through adding a modifying factor into our function from values (understood agent-neutrally) to weights.⁶² One suggestion is that the strength of a

⁵⁹ For the objection as applied to act utilitarianism see Hooker (2000: 151-2), Portmore (2011: 26).

⁶⁰ See Slote (1985).

⁶¹ Lord & Maguire in Lord & Maguire eds. (2016: 17).

⁶² *Ibid.*

reason is modified by the agent-neutral value of ‘moral distance’. What is moral distance? Bader tells us that

[Partiality] can be determined by centring the evaluation on the particular perspective of the agent, allowing one to take into consideration where values are located and in what relations they stand to the point of view on which the evaluation is centred, in particular what distance relations obtain in moral space.⁶³

I interpret the idea as follows. Moral distance is determined by the degree to which you care or ought to care about another person. Standing in a relation of friendship or marriage to another person narrows the ‘moral’ distance between you and them. In assessing the value of that state of affairs we must take into account the moral distance between A and B: the narrower the distance, the higher the value. So, the fact that A and B are married modifies the agent-neutral value of the state of affairs ‘A’s saving B’s life’. Since there is more moral distance between A and C (the stranger) than that between A and B, ‘A’s saving B’s life’ is more valuable than ‘A’s saving C’s life’.

If this is the right characterisation of the view, then I am not convinced that the view is really distinct from the previous solution. It does not seem correct to say that the state of affairs ‘A’s saving B’s life’ has more agent-neutral value than ‘A’s saving C’s life’ *unless* we also agree that the only agent to whom we are ascribing the promotion of value is A himself. So, again, it looks like this state of affairs is to be promoted (is more valuable) only *relative to A*. The defender of the view at least needs to say more here to distinguish the two views, and to persuade us that he is really identifying agent-neutral value.

A third objection draws on Scanlon’s influential discussion of value in *What We Owe to Each Other*.⁶⁴ Scanlon persuasively draws the distinction between promoting value and respecting value. To show this he cites the example of friendship. If value was all about promotion, your response to the value of friendship

⁶³ Bader in Lord & Maguire eds. (2016: 42).

⁶⁴ Scanlon (1998) ch. 2.

would be to promote states of affairs in which friendship inheres. But this misunderstands the nature of friendship. You could satisfy the promotion relation by trying to make as many friends as you possibly can. What would be more appropriate to the value of friendship would be your acting in certain specific ways towards them, such as calling them every now and again, giving moral support, etc. That is, your being a friend involves your recognising that you have important reasons to act (and adopt attitudes) towards your friend, and that you act on those reasons.

The objection, then, is that the value-based view cannot implement the distinction between promoting and respecting value because the best way to understand respecting value is in terms of having normative reasons to act in various ways. Revising the value-based view might run as follows. The strength of a reason is determined either by the amount of value that is promoted by acting in that way, or by the amount of value that is respected in that way. But respect is not measured by amounts of value. Since a chief virtue of the monistic value-based view of the weight of reasons was that they could identify the weight of a reason with an amount of value, the view falters on this point.

1.5.4. Reductive Views, Again

Any reductive view of the weight of reasons faces the difficulty of accounting for the context sensitivity of reasons. As I argued in §1.3 reasons can be sensitive to the context in which they obtain. Monistic reductive theories of weight will find it difficult to account for the context-sensitivity of reasons, even if we deny the weaker version of holism in the theory of reasons. All we need is the claim that the *weight* of a reason can also be changed given a change in context.⁶⁵ Given the above examples, if a change in context can change the valency of the reason, it is plausible that context can also affect the weight of reason.

Given that reasons and their weights are context-sensitive in this sense any plausible reductive theory of weight must take this into account. Simple reductive theories which claim that the weight of a reason increases as P increases will clearly fail, as the example of reasons provided by pleasure illustrates. One way that

⁶⁵ On this point see Berker (2007: 113-122).

reductive theories will be able to do this is to adopt a similar thesis of context sensitivity with respect to P. We will be forced to accept holism about desires or values or whatever we take P to be *and* we will have to show that there is a one-to-one correspondence between P conceived holistically and the intuitive shifts in weights determined immediately by context. But the problem is that even if such a theory can be given, without its looking gerrymandered, it will simply look like a restatement of what we already know about the weights of reasons, given their contexts. Reductive theories will fail to be any more informative about weight than non-reductive theories.

Second, as we have seen for one of the most popular theories of the weight of reasons – the value-based account – even on the pluralistic understanding of value it looks very plausible that the weights of at least *some* kinds of reasons are not value-based. Reductive accounts thus find it difficult to satisfy the requirement of unification that the theory should be able to tell us what determines the weight of a reason regardless of what sort of reason it is.

Third, as I stated the problem above I tried to make it salient that while the notion of a consideration favouring a response looks, intuitively, like a primitive normative relation, the notion of the weight of a reason looks more like it stands in need of a normative explanation. If reasons are fundamental then it looks just a little mysterious why it is correct to place a certain amount of weight on some reason. Having said this, I should add that a view which (a) analyses the weight of a reason in other normative terms but (b) does not analyse the favouring relation would look odd. It would be odd that there is nothing informative that we can say about ‘favouring’ itself, though we can say a great deal about favouring *to some extent*. In fact, most analyses of the weight of reasons in normative terms (that I am aware of) also offer analyses of favouring. In the end, then, whether a reductive account of the weight of reasons is right may also depend on whether the further analysis of favouring is right. In subsequent chapters I will try to show that there is no plausible analysis of favouring at least in terms of deontic properties. If I am right about that, the prospects for a reductive theory of the weight of reasons look worse.

1.5.5. Remarks on a Non-Reductive Account of the Weight of Reasons

If the reductive accounts fail then in order to preserve unification and explanatory power, what can we say about the weight of reasons? We can begin to sketch this by noting a few ways in which the weight of reasons can be determined or modified. I have argued that the canonical form of a normative reason is a consideration that counts in favour of an agent's responding in some way. These considerations can help to determine the weight of the reason by providing an explanation of the extent to which the response that the reason recommends is favoured. As noted above reasons can have multifarious 'sources' some of which appear to be captured broadly by the idea that they are valuable and others that do not seem to be 'value-based'. For instance conventions, rules, norms, and evidence may give rise to normative reasons, but these sorts of sources are not, without revision, best thought of as considerations of value. The non-reductive view is silent on what sources give rise to normative support so it does not run into problems of the limitations of the reductive views.

Another important point to make is that it is plausible that often whatever it is in virtue of which some consideration favours some response is also what explains why we should favour it to a certain extent. Once we are satisfied why a consideration favours, we are well on the way to understanding the weight of the reason. But as mentioned above, the sorts of things (the 'sources') in virtue of which considerations favour are multifarious. The reasons fundamentalist thesis is silent on what considerations in fact favour.

Note also that we can consistently accept that the weight of some reasons can be at least partly explained in terms of value. How so? According to the buck-passing account of value for X to be valuable is for X to have the higher-order property of having other properties that give us reason(s) to hold pro-attitude towards X.⁶⁶

According to one version of the value-based view of weight, for a reason, R, to Φ to have a weight W is determined by the degree to which the state of affairs promoted by Φ -ing is valuable. This view need not be circular. It can analyse the weight of reasons for action in terms of that action having the property of having

⁶⁶ On the buck-passing theory of value see Dancy (2000b), Rabinowicz & Rasmussen (2004), Scanlon (1998: 96-100), Skorupski (2007), Stratton-Lake (2005).

reasons to hold pro-attitudes towards those states of affairs. Thus we can, in some cases, explain the weight of reasons for action in terms of reasons for pro-attitudes.

Take the case of friendship. That Smith is your friend is a reason to phone him this week. Given that it will presumably defeat the other reasons you have, what explains why this is a weighty reason? One explanation would be to advert to the value of friendship. That is just to say that friendship is reason-providing, but it might still shed some light on how reasons' weights are determined. We can say more about *in what way* being a friend has properties that are reason-providing. Being a friend gives one reasons of respect, reasons to commit yourself to various things, reasons to hold attitudes of love and care towards your friend, for example. These reasons are provided by the various properties that are implied by what it is to be a friend (which properties might be understood in terms of the various relations which one person stands to another).

The buck-passing account of value so understood does not imply a circular account of the weight of reasons. However we can presumably ask a further question about the reasons given by one's being a friend. What determines the weight of *those* reasons you have to hold certain pro-attitudes towards your friend? I think that a full answer to this question will involve appealing to first-order normative theories. However, one helpful way to think about the weight of reasons without invoking the results of first-order normative theories is to appeal to constitutive facts about the activity from which those reasons arise.

Consider certain constitutive facts about friendship. There are certain reasons that are such that if we don't respond to them we fail to count as being a friend in that regard. For instance, your failing to place more weight on honouring certain commitments to your friend than on some egoistic reason that you also have implies that you are failing to live up to your role at a friend. These normative considerations, in the case of friendship, also seem to imply silencing and are determinative of sufficiency. It is plausible that you are not really a friend if you don't take certain reasons as decisively outweighing, at certain appropriate times. There is, of course, a further question about the justifiability of the activity of friendship in the first place. However these sorts of constitutive considerations can help to provide concrete guidance to how the weights of reasons can be determined.

Summing up the foregoing remarks, we might say that the non-reductive view of the weight of reasons has a few distinctive virtues. It respects the fact that reasons are context-sensitive and that reasons can have various sources whether in value-based, rule-based, evidence-based or multifarious other considerations. As I have argued, given the context-sensitivity of reasons any ‘theory’ of the weight of reasons is going to be difficult to formulate in meta-normative terms. In order to give a complete account of weight we need to appeal to the resources of first-order normative theory. Second, the non-reductive view can account for the relevance of the agent for whom the reason applies in determining the weight of a reason since the reason relation itself already contains an agent-place. Third, the non-reductive view can accept that the weight of some reasons is explained in terms of value in a way that is consistent with the buck-passing account of value. Finally, on this view we can provide a partial explanation of the weight of certain reasons by adverting to certain constitutive facts about the nature of the activity from which that reason arises.

Given that it is plausible that reductive views will not capture the weight of reasons, these tentative remarks should serve to highlight a few ways in which a non-reductive account of the weight of reasons can be illuminating. Clearly much more work needs to be done to provide a fuller picture of the non-reductive account, but what I take myself to have begun to do here is to point to some ways in which such an account might proceed.

2. Responding Correctly to Normative Reasons

2. 1. Compliance and Conformity

I turn in this chapter from questions about the structure of a normative reason to the question of what is to respond correctly to a normative reason. In other words, what does a normative reason, in general, require that we do in order to satisfy it?

In moral philosophy there is a familiar distinction between two kinds of morally-assessable action: action that is done ‘in accordance with duty’ and action that is done ‘for the sake of duty’.⁶⁷ One way to express the idea that action must be done for the sake of duty is that correct action requires an internal connection between motive and justification, one which excludes luck or accident.

Others have made similar claims about believing correctly, in drawing a distinction between believing in accordance with one’s epistemic reasons, and believing in light of one’s epistemic reasons, the latter claim implying that *believing* correctly excludes luck or accident.⁶⁸ If the Kantian claim is right with respect to moral action and right with respect to belief, we might wonder whether the claim is an instance of a more general truth about the normative. That general claim is about what it is to respond correctly to the reasons one has. It tells us that a necessary condition on what it is to correctly to respond to the reasons one has is that one responds *on the basis of that reason*.

I follow Raz in distinguishing between two views, one which denies that this is a necessary condition on correct reasons-responsiveness, namely *Conformity*, and the other which affirms the necessary condition, which Raz calls *Compliance*.⁶⁹ In this chapter I give reasons to accept *Compliance* by adducing two arguments against *Conformity*, an argument from accomplishment and an argument from response-

⁶⁷ See Kant (1966: 4:400).

⁶⁸ It is a commonplace among epistemologists post-Gettier that knowledge excludes certain kinds of luck. For an extended discussion of which kinds of luck knowledge should exclude see Pritchard (2005). For an explicit avowal of the parallel between the Kantian claim in moral theory and epistemology see Russell in Steup ed. (2001).

⁶⁹ Raz draws the distinction in the postscript to Raz (1990).

guiding. I then examine some objections to *Compliance*, which I argue do not succeed in convincing us of its falsity.

I here spell out a little more precisely what *Compliance* is. *Compliance* can be stated as follows:

Compliance An agent responds correctly to his reason(s) to Φ if, and only if, (i) they Φ , and (ii) the consideration(s) that figure in the explanation why they Φ -ed are also the considerations which favour their Φ -ing.

The consideration (a fact, a proposition, a state of affairs, etc) will pick out the *content* of a reason in the following way: suppose that the consideration that it has begun raining is a reason for Jim to take the washing in. The content of Jim's reason (granted his epistemic access to that content) is 'that it has begun raining' and the response that this consideration calls for is that Jim takes the washing in. When it is the case that (a) Jim is moved to take the washing in by the consideration that it has begun raining, and (b) its having begun raining *calls for* Jim to take the washing in (recall that this is just what it is for some consideration to stand in the relation of being a reason for responding in some way), and Jim in fact takes the washing in, he has acted on the basis of his reason, and according to *Compliance* he has done all that his reasons require of him.⁷⁰

We can then define *Conformity* as the acceptance of (i) and the denial of (ii). Since it may not seem obvious why someone might deny (ii) I present some considerations that motivate philosophers such as Raz to deny (ii) in order to show that *Compliance* is substantive, and needs argumentation to support it.⁷¹

First, *Compliance* appears too intellectualised. It tells us that responses are only correct when the reasons that agents have figure in their motivation for so responding. But don't we do all sorts of things without thinking about the reasons

⁷⁰ This is sometimes referred to in the literature as the *Identity Thesis* in that it claims an identity between normative and motivating reasons. For discussion see Dancy (2000a) ch.5, Hieronymi (2011), Lord (2008), Mantel (2014), Wiland (2012) ch.5.

⁷¹ Raz (1990: 178-82). Gardner also defends the view in Gardner (2007: 100ff).

which support our actions, which are nevertheless not failures in us? Must I have in mind the reason there is to pick up my car keys on the way to work lest I be rationally failing in some way?

Second, the view implies the wrong kind of ‘mental background’ in cases of omissions. I have undefeated reason not to kill you, and I don’t kill you, but not because I was moved by the reason I had not to kill, but because the thought of killing you did not even occur to me! According to *Compliance* I am guilty of failing to comply with my reason not to kill you. But, Raz claims, not only am I *not* guilty of any failure here, I would be a better person if the thought doesn’t even cross my mind (I conform to my reason), than if I act on the basis of my reason not to kill (I comply with my reason).⁷²

Third, there may be lots of cases of simply conforming to one’s reasons that are perfectly good responses. Raz has an example of Lucy, who is generally sensitive to her friend’s needs, dropping by her friend’s house because she is bored.⁷³ It turns out that unbeknownst to Lucy her friend has had some bad news and is in need of cheering up. Lucy’s dropping by the house gets her friend exactly what she wanted, but Lucy did not comply with her reason. Is there any rational failing in Lucy’s response? Raz claims that there is nothing amiss in this case whatsoever. As long as Lucy’s friend got what she wanted, and Lucy is *generally* sensitive to her needs, why insist further that she must act on the basis of the reason there is for her to go to the house? It is not clear.

One final piece of clarification before I discuss the arguments. So far I have talked about an agent’s responding correctly to their reasons, but many of these reasons will be outweighed. Surely an agent shouldn’t respond to *these* reasons by responding in the way they favour? It is precisely the reverse. They should respond to outweighed or defeated reasons by *not* responding in the way they recommend, or by responding in some other way. But these reasons should not always simply be ignored. Sometimes deliberation requires agonising over the options. Sometimes our decisions can be difficult. It can be the case that it is appropriate to regret that a reason is outweighed by others, such as when you must renege on your promise

⁷² Raz (1990: 181).

⁷³ *Loc. cit.*

because of an immediate emergency. It is an interesting question what it is to correctly respond to defeated or outweighed reasons, but one that I won't address here. So this chapter idealises somewhat by asking only what it is to respond correctly to undefeated reason(s).⁷⁴ With that in mind, I turn to present two arguments against *Conformity*.

2.2.1. The Argument from Accomplishment

This argument works by spelling out the intuition that whatever correct response to reasons consists in, it is an *accomplishment* of some sort. However, *Conformity* entails that responding correctly to reasons cannot be an accomplishment. So the *Conformity* view must fail:

1. Responding correctly to a reason is an accomplishment.
2. Conformity to a reason is not an accomplishment.
- C. Therefore, correctly responding to a reason requires more than conformity.

In order to assess the first premise we need to know first what it is for something to be an accomplishment. Let's take as a toy example of an accomplishment an archer hitting a bullseye on his target. Two conditions are clearly present in this case: (1) A *success* condition. The aim of the archer, which is to hit the bullseye on the target, sets a standard of success which the archer can meet or fail to meet. If he fails to hit the target, though he may accomplish a different, related success such as hitting very close to the bullseye, or overcoming an arm injury, he does not succeed in achieving hitting the bullseye.

(2) A *competence* condition. In order to hit the target the archer must have had some relevant competence or skill, without which he would be unable to, reliably, hit the target. Moreover, this competence must be related in some way to the

⁷⁴ Broome idealises in a similar way by narrowing his focus to what your reasons require of you in Broome (2004). My focus is not as narrow as Broome's since an undefeated reason may permit me to respond in some way, but not require me to so respond.

success of the action. If the archer was competent in hitting bullseyes, but a freak gust of wind blew his arrow off course, then another gust back on course again, the hitting of the target would not be an accomplishment, but a fluke. Thus, in addition to a success condition and a competence condition, there must be a further condition linking these two conditions together. I will say that the two are linked together when the success is one which *manifests* the competence of the agent.⁷⁵ That the arrow hit the bullseye because the arrow was directed toward the bullseye by a gust of wind does not manifest the competence that the archer has of hitting the bullseye. It fails to manifest any competence; it is simply by the wind caused to hit the bullseye. Accomplishment is then defined as follows:

Accomplishment X is an accomplishment of A's if, and only if, (i) X is a success, (ii) A has a competence without which he would be unable to reliably perform X, and (iii) A manifests that competence when he performs X.⁷⁶

Now that we have a putative definition of accomplishment on the table, we should ask whether we have pre-theoretical reason to believe that responding correctly to one's reasons is an accomplishment. Sometimes correct responses to one's reasons seem like paradigmatic kinds of accomplishments. When you manage to correctly figure out the quickest way to get somewhere (and get there), or when you act against your addictions because you recognise there are strong reasons not to take the drug, your actions seem to be accomplishments at least in part because they are correct responses to reasons. But other kinds of response may look less like accomplishments – scratching an itch or flicking on a lightswitch, for instance, don't seem to be the sort of thing that we call accomplishments. One reason why we might

⁷⁵ On the epistemic version of this claim see Greco (2010), Sosa (2007). On the practical version of this claim see Mantel (2013).

⁷⁶ I distinguish here between *accomplishment* and *achievement*. Bradford argues that achievement also requires a difficulty condition. In distinguishing between what she calls 'mundane accomplishment' and 'capital-A Achievement', she accepts that the former does not require a difficulty condition. See Bradford (2015: 12-18).

refrain from calling scratching one's arm an accomplishment is that scratching one's arm is easy to do.

Accomplishment, one might think, involves *difficulty* or *overcoming obstacles*. But not all genuine accomplishments involve difficulties or the overcoming of obstacles. Think of the ease with which Ronnie O'Sullivan often wins frames in snooker. We don't take the fact that O'Sullivan is immensely talented at the game to imply that his winning frames and matches are not genuine accomplishments.⁷⁷

Another reason we might have for refraining from thinking that responding to my reason to scratch my arm is an accomplishment is that doing so has relatively little value. Apart from the slightly pleasurable sensation I get from scratching, there is nothing else to be said for the action. But value is a red herring in understanding accomplishment. Counting every blade of grass in the garden no doubt has no value, but it is nevertheless an accomplishment on the part of the counter.

If the bare idea of accomplishment does not imply claims about difficulty, overcoming obstacles, or value, I see no reason at the outset to think that responding to one's reasons by performing an everyday action such as scratching one's arm should not count as an accomplishment. Responding to a reason is a *rational* accomplishment which we can understand divorced from the particular content of that reason. Furthermore, response implies success conditions in virtue of the nature of a reason. Since reasons are normative, they imply a standard of success that issues from their nature. Response also presumably requires the competence of being able to discern one's undefeated reasons (i.e. it requires good reasoning), and the ability to respond in the way that those reasons recommend (putting aside considerations about *akrasia* and physical obstacles).

Must that response be a manifestation of that competence? Accomplishment requires that one's competence is manifested in the successful action because flukes are not accomplishments. But accomplishment may not exclude all kinds of luck. Suppose that there are 10 targets, 9 of which are protected by an invisible force-field that will prevent any arrow from hitting the target. The archer happens to fire at the

⁷⁷ Though as Bradford argues, for this reason O'Sullivan's winning easily may diminish the degree to which this is an *achievement* of O'Sullivan. See *Loc. cit.*

only target that is not protected by a forcefield, and hits the bullseye.⁷⁸ This case satisfies the definition of accomplishment given above, and intuitively the archer's hitting the bullseye seems to be an accomplishment. Yet it was partly down to luck that the archer hit a bullseye. We can distinguish this kind of luck (what Pritchard calls 'environmental luck') from the sort of luck involved where the arrow is blown on and off course ('non-environmental luck'). Cases of non-environmental luck are typified by breaking the (either necessary or contingent) connection between competence and success. Accomplishment seems to exclude only non-environmental luck. So in order to answer the question above, we can ask whether responding correctly to one's reason excludes non-environmental luck.

It is, I submit, plausible that responding to one's reasons excludes non-environmental luck. To see this, notice that responding to reasons is partly what it is to be rational,⁷⁹ and that rationality requires that there must be *some* connection between competence and success, even if only a contingent one. Suppose that A knows that he has undefeated reason to believe that p (competence) but A's believing that p depends on whether A feels like it today. Then it is plausible that A is irrational. Here rationality requires the connection between success and competence to be a necessary one. Now consider a practical case. A knows that he has undefeated reason to Φ , and A is able to Φ (competence) but A only Φ s because A feels like it today. The connection between success and competence is contingent, but here A need not be irrational. Now consider the following case: A knows that he has undefeated reason to post the letter today, and A is able to post the letter today (competence) but A posts the letter in the afternoon by sheer fluke while sleepwalking. Did A act rationally? In this case it seems that your action is not even up for assessment as rational or irrational. This suggests that where we talk about rationality, we are required to presuppose *some* connection between competence and success.

⁷⁸ The example is from Pritchard (2008: 30-1).

⁷⁹ Though this is a substantive claim, denied for instance by Broome (2007), it is a plausible one. On the view, see Lord (Forthcoming), Parfit (2011).

I think, therefore, that we have good reason to believe that correctly responding to a reason is an accomplishment both pre-theoretically, and in light of our definition of accomplishment.

Premise 2 tells us that merely conforming to a reason is not an accomplishment. According to *Conformity* you may satisfy all that your reasons favour just by responding in the way they recommend, without having to respond on the basis of those reasons. We can immediately see that *Conformity* does not exclude non-environmental luck. Since you can do what your reasons recommend by sheer fluke or accident, non-environmental luck may be involved in your conforming to your reasons. For this reason alone, on the *Conformity* view it can be the case that you respond correctly to your reasons, but doing so is not an accomplishment.

Further, we might think that complying with one's reasons *is* an accomplishment. *Compliance* captures the fact that an accomplishment is a *manifestation* of a skill in some successful response. This view does look like it excludes non-environmental luck, while being compatible with environmental luck. *Compliance* tells us that in the good cases the considerations that favour are also the considerations which motivate. This connection between favouring and motivation implies some sort of connection between competence and success. Let me illustrate with an example. I have undefeated reason to post the letter today, and I post the letter on the basis of this consideration, complying with my reason. This implies (i)-(iii) of my definition of accomplishment. It implies success conditions simply in virtue of the presence of reasons, implies that I am competent, since in order to post the letter on the basis of my reason, I must know that I have that reason (or at least have some kind of epistemic access) and implies that my posting the letter manifests my competence because I do so on the basis of my reason. Since complying with my reason implies (i)-(iii) it also secures the claim that compliance with reasons excludes non-environmental luck.

2.2.2. The Argument from Response-Guiding

As we have already seen in examining RC it is plausible that normative reasons are the sort of thing that can guide agents. If R is a reason for A to act, it must be

possible that A can act for that reason.⁸⁰ So, anything that counts as a reason is the sort of thing that an agent could respond to.

But beyond this claim about normative reasons, there is no consensus about what it means for reasons to be ‘guides’. Raz provides material with which to understand what it is for reasons to guide on the *Compliance* view and the *Conformity* view. The following argument claims that *Conformity* cannot make sense of this feature of normative reasons. Here is the argument:

1. Normative reasons are response-guiding.
 2. According to *Conformity*, reasons are guides in the sense that a Michelin guide is a guide.
 3. Reasons are not guides in the sense that a Michelin guide is a guide
- C. Therefore, *Conformity* is false.

The first premise of this argument expresses the claim mentioned above. The second premise expresses Raz’s view about guidance.⁸¹ Recall that *Conformity* tells us that there are cases in which all that matters is that the agent respond in the way that his reason favours, regardless of whether he does so on the basis of that reason or not. If sometimes it fails to matter whether we respond on the basis of our reasons, then it is natural to think that we need not always be guided by reasons in our responses. Raz introduces a metaphor to explain the sense in which reasons are guides on the *Conformity* view:

They [reasons] are guides in the sense that the Michelin guide to Paris is a guide. I may use it, but I do not have to. I do not even have to be aware of its existence. There is absolutely nothing wrong in using another guide, if it is as

⁸⁰ That *Compliance* is possible depends on the truth of this claim, since if correctly responding to reasons implies responding on the basis of those reasons, it must be possible that we can respond on the basis of those reasons.

⁸¹ Raz (1990: 180).

good. The important thing is that I get to see the things which are worth seeing in Paris.⁸²

On the *Conformity* view, if there is an equally good way to respond in the way that my reasons favour which does not involve my responding on the basis of those reasons, I am permitted to do so. Take Lucy again. If Lucy can cheer her friend up by visiting her but does so simply because Lucy is bored, then she is permitted to act on this basis. There is nothing wrong with using her reason to alleviate her boredom as a guide to cheering up her friend, if it does in fact lead to her cheering up her friend. Note that the alternative guides that I am permitted to use may not be particularly reliable guides. It is enough, on this view, that I end up responding in the way I am supposed to. That one's guide need not be reliable is implied by *Conformity* itself, not merely Raz's particular characterisation of what it is to be guided by reasons.

Raz also tells us that in being guided one does not even need to be aware of the existence of the reason which favours one's response. Again, this seems to be implied by the *Conformity* view itself, and in particular with the view's compatibility with non-environmental luck. Lucy is unaware that her friend needs cheering up, and nevertheless manages to cheer her up. Further, the idea that one could respond to one's reasons by conforming to them may get mileage from the idea of conformity in other areas. For instance, it may be justifiable to conform to a law by doing what the law says not because it is a law (perhaps you in fact think it is an unjust law) but because you think that things would be worse if you set a precedent by breaking it. It can often be sensible and permissible to merely respond to laws, conventions, rules, commands, etc on the basis of *other* kinds of considerations. So why not think that it may often be sensible and permissible to respond to your reasons on the basis of other sorts of considerations?

On the *Conformity* view you may act on a basis other than the reasons which call for Φ -ing so long as you respond in the way that your reasons favour (by Φ -ing). According to Raz, this is a permissible way to be guided in your responses. Here, I consider three kinds of cases of response that are compatible with *Conformity* but which are plausibly not cases of being guided by reasons in the relevant sense.

⁸² *Loc. cit.*

First, if the basis of your response is not a reason, such that we cannot explain your response by citing any reasons you took to favour responding in that way then being guided is a matter of subverting your own rationality. An example of this sort of case is getting somebody to hypnotise you in order to get you to Φ . We might call these sorts of cases ‘external’ guidance, or ‘arational’ guidance.

Second, if the basis of your response is a reason which plays no role in favouring that response, then your conforming to that reason happens as a result of non-environmental luck. The mere fact that your response happens accidentally itself impugns the idea that you were guided to that response.

Third, and most obviously, there are cases in which it is morally suspect that you act on certain motivating reasons that do not contribute to favouring the action, but are good at getting you to do the favoured thing. Suppose that A is a malicious person and is motivated to do good deeds only if he also gets to harm others along the way (as a sort of perverse compensation). According to the conception of guidance by reasons on the *Conformity* view, A’s doing the good deeds on the basis of the fact that he will also get to harm others is as fine a basis for doing so as any other. Since this seems deeply morally suspect, we should admit that there must be some relevant constraints on what can permissibly count as a good basis for action. It is *prima facie* plausible that those constraints be grounded in the reasons that favour the action.⁸³

There are, however, less obvious cases. There are cases in which it seems that it would be *better* if you were guided to Φ on some other basis than the reasons that favour Φ -ing. For example, suppose that you want to impress somebody. If you try to impress them on the basis of your reasons for impressing them (say, because you badly need the job that they can give you), say, by thinking about how badly you need the job, you will in virtue of those thoughts appear less impressive. Or suppose that you must walk carefully across the bridge because it is unstable and will collapse if you don’t tread carefully. We can imagine a case where, if you think about this fact, it will cause you to become so nervous that you stumble and cause the bridge to

⁸³ I note that Raz could reply by amending the *Conformity* view to include some constraints that rule out such cases without accepting the *Compliance* view. However this would still necessitate his rejecting the *Conformity* view as stated above.

collapse. Here it looks like you *ought* to Φ on a basis other than the reasons that favour Φ -ing.

Other cases arise from indirect moral theories. For instance, rational egoism tells us that although the reasons why we ought to act are grounded in self-interest, we should not act on the basis of these grounds because they will frustrate or defeat the ends which they give us. Suppose that you treat your friend kindly only because you know you will get more kindness in return. Your being motivated by this consideration would prevent you from being properly able to treat your friend kindly, preventing you from getting what is ultimately in your self-interest. So in order to best promote your self-interest you should act on some other basis – say that you should take the fact that friends deserve kindness *qua* friendship as your basis for treating him kindly.

However, we should point out that *Compliance* would be implausible if it suggested that agents must have as an *occurrent* thought at the time of their responding the content of the reasons that favour that response. It can be enough that they are *disposed* such that their response can be explained in part in terms of the content of those reasons. When I flick the lightswitch I need not be thinking about the fact that I want the room to be illuminated. It suffices that we can use the fact that I want the room to be illuminated to explain my behaviour, say that if you asked me why I flicked the switch, I would tell you that I wanted the room to be illuminated. Recognising this already goes some way to addressing the cases above. In the case of the unstable bridge, for instance, it may be enough to avoid your becoming nervous that you do not have the fact that the bridge is unstable in mind as an *occurrent* thought when you walk across the bridge.

I have provided good reason to accept premise 3 of our argument. Reasons are not guides in the sense that a Michelin guide is a guide. Reasons are not simply interchangeable, but are plausibly subject to constraints on the basis of their contribution to favouring the action. Thus I submit that *Conformity* is false.

2.3. Objections to *Compliance*

In §2.1 I mentioned some motivations for rejecting *Compliance*. Now that we have some reason to accept *Compliance* I turn to examining some of these objections.

2.3.1. The Huckleberry Finn Counterexample

Compliance may face a counterexample in the example of Huckleberry Finn. Finn believes that he ought to turn Jim (a slave) in to the authorities, but decides to act against what he believes he ought to do, and helps Jim to freedom, despite feeling guilty about his decision. This case introduces a complication that I have not dealt with yet – of what to say about cases in which an agent, through no fault of their own, is misled as to what they have reason to do. Perhaps it is plausible that we should in fact *give credit* to agents who respond in a way that is entirely reasonable (more, that given their beliefs or their evidence, it may appear obligatory to them) but do not respond to undefeated reasons, because the considerations that the agent responds on the basis of do not *in fact* favour responding in that way.

That issue aside, the problem that this case seems to raise for *Compliance* is that Finn seems to do the right thing, but not do so on the basis of his undefeated reason to do so. There are different ways to interpret the case: either Finn does not believe he has any reason to help Jim, or he believes he has some reason to help Jim which he believes is defeated or outweighed by other considerations, chiefly that he believes Jim to be the property of Miss Watson, his owner.

So this seems to be a case in which Finn conforms to his reason to help Jim, which he believes is wrong. Does Finn correctly respond to his reason? Presumably he is sensitive at least somewhat to a reason that he recognises (if only dimly) that Jim's request is warranted. If the case is one in which Finn doesn't balance his reasons correctly, we can still insist that there is a partial overlap between motivating and normative reasons. But to make the case as strong as possible, let us suppose that Finn does not recognise (even dimly) any reason to help Jim, i.e. that there is no overlap between motivation and favouring. Then he would be merely conforming to his reason without in any way acting on the basis of that reason. But this case now

makes it difficult to understand why Finn did what he did. He believes that he ought not help Jim, and yet he does so for no reason apparent to himself. Acting for no reason makes it impossible to explain from the agent's point of view why he did what he did. Granted we could provide an 'external' explanation of Finn's action, but forcing the case to comport with the *Conformity* view makes at least makes Finn's action unintelligible from his point of view. But I doubt whether readers of the passage would think that Finn's action was unintelligible. Therefore, this can't be the right way to interpret the case. So I don't think Finn serves as a convincing counterexample to *Compliance*.⁸⁴

2.3.2. Omissions

The second objection to *Compliance* centres on certain cases of omissions. Recall that I defined the responses that reasons call for as activities broadly construed to include not only actions, beliefs, feelings etc, but also *omissions* of actions, beliefs, feelings etc. Raz provides an example in which an agent is not only not irrational in merely conforming (explicitly not complying) with his reason but is in fact more worthy of criticism *if* he complies with his reason. Suppose, plausibly, that you have reason not to murder me. According to Raz's interpretation of *Compliance*, in order to correctly respond to the reason you have not to murder me, you do not murder on the basis of the reason you have not to murder me (whatever that in fact is, it doesn't matter for our purposes). But it looks like it would be better if, instead, thoughts about murdering me didn't even cross your mind.⁸⁵

This objection centres on which view of reasoning – complying or conforming – is better in certain cases. There may also be a more general objection. Presumably, at any given time an agent has many reasons not to do all kinds of things. I have (undefeated) reason not to harm John, who is sitting next to me, not to tickle Pam, who is in the next room, not to scratch the car on the street outside, not to scream at the top of my lungs, and so on. Depending on what view we have about the

⁸⁴ For further discussion of the Huckleberry Finn case see Arpaly (2002), Arpaly (2015).

⁸⁵ Raz (1990: 183).

individuation of reasons, we may have *uncountably* many reasons not to respond in certain ways (and add to this stock all the reasons one has not to believe, feel, etc). The more general worry about reasons for omissions then is that it would be *impossible* to comply with our undefeated reasons for omissions.

On the first version of the objection, that it is better that an agent does not think about killing, here we need to say more about what I am taking to be going on psychologically on the *Compliance* view. I've said that on the *Compliance* view, you respond to your reason through the explanation of why you so responded also being why the response is justified. That leaves a lot to be said about how intellectualised the view is. In explaining why somebody acted in some way, our explanation can (and if it is to be informative should) include the dispositions of the agent, as well as his occurrent beliefs. If those dispositions help to explain why the agent responded in some way, and they also comport with the justification of the response, then we satisfy *Compliance*.

Once we see this, we have available an explanation why Raz's case is not a counterexample to *Compliance*. First ask what the reason is that he has not to murder his friend. Let's say that the mere fact that he is his friend is an undefeated reason not to murder him (among other reasons). It is compatible with *Compliance* that the agent is disposed to take the fact that he is his friend as a reason not to murder him, without that fact being an occurrent thought. All we need is some matching between the explanation of action and the justification (normative reason) of response. That gives us a way to explain this case in a way that is compatible with *Compliance*. In fact, it is a good explanation, since otherwise *Compliance* would be susceptible to a further charge that it was far too intellectualised a view. Nobody who defends *Compliance* would go in for such an implausible version of the view.

It is not much of a stretch to apply these remarks to the more general objection. It comports with why I am not scratching the car outside that it would be vandalism, and that I am disposed not to do so, on this basis. As a test, we can imagine asking me: 'Why *aren't* you scratching that car outside?' I answer: 'Because to do so would be vandalism.' So it is a good explanation of my not scratching the car that I am disposed to think it would be vandalism to do so. This explanation of my omission comports with the justification for the omission, and in a non-accidental

way (if in fact I thought that vandalism of cars was a fine thing to do, it would explain why I was not scratching the car that I was ignorant of the car outside, or that perhaps I had better things to do, but if not, would be scratching it). Since this matching occurs here, we can make sense of the fact that I am not doing what is wrong because it is wrong in the sense adduced.

A possible problem with this view is that it may muddy the distinction between *mere* omission and *intentional* omission, if what it is to intentionally omit something is to omit it on the basis of a reason. I think we can get around this problem by distinguishing between explanation, motivation, and justification. As long as what explains why you omitted matches with the justification in a non-accidental way you comply with your reason. To distinguish mere from intentional omission we require further that you omitted to act on the basis of a motivating reason.

So, in the case in which you don't scratch the car, your omission is at least partly explained by your being disposed not to do so, on the basis of the reason that forbids doing so. But I don't think that your being disposed to omit scratching the car, where that disposition is founded on a reason, implies that you intended not to scratch the car. You can felicitously say *both* that you had a reason not to scratch the car *and* that you did not intend not to scratch the car (since it didn't enter your mind at the time). That is, I think that failing to do something for a reason and intentionally failing to do something are not equivalent claims.

Does this answer conflict with my argument from accomplishment? I claimed that responding to a reason requires that you respond via the manifestation of a competence that you have. When you do not attack the person sitting next to you in the café (understood as a mere omission), was not doing so the manifestation of a competence that you have? I think we can understand cases of mere omissions as manifestations of competences. I claimed that a competence is an ability without which an agent would not reliably perform (or omit) some action. In particular, to respond correctly to a reason an agent must manifest a competence to so respond. Crucially, this excluded non-environmental luck from the agent's response to their reason and implied that there was a non-accidental relation between competence and success.

The relevant distinction is being disposed to refrain from acting in certain ways, where that disposition is had on the basis of a reason, and simply failing to act in some way. In the first case, an agent complies with their reason not to act, where the explanation of their omission implies the manifestation of a competence. For example, they know that they have undefeated reason not to cause damage to others' property, and are thereby disposed to refrain from scratching the car. In the second case, an agent may conform with their reason not to scratch the car, but the connection between competence and success is broken since, for example, they are not aware that they have undefeated reason not to cause damage to others' property.

2.3.3. The Surprise Party

The third objection is another counterexample, which I will read as a counterexample to *Compliance* (rather than to the claim that reasons are essentially response-guiding). The counterexample presents a case in which an agent's reason will be defeated if he acts on its basis. Thus we seem to have a case in which an agent has a reason, but that they cannot respond on the basis of that reason. This is not immediately a counterexample to *Compliance*, but only to that view in conjunction with RC. Although some deny the latter claim, I think that it is plausible.⁸⁶

I have briefly discussed this counterexample above (§1.1.1). Nate loves surprise parties, and his friends are going to throw one for him tonight.⁸⁷ Nate finishes work, and wonders whether he should now go shopping, or to the cinema, or just go straight home. Because Nate loves surprise parties so much he should go straight home, where there is a surprise party waiting for him. So it seems that the fact that there is a surprise party at home is a reason for Nate to go straight home after work. But if this *is* a reason for Nate, it looks like a reason that it is *impossible* for him to respond to, at least on the *Compliance* view. For, on the *Compliance* view, an agent correctly responds to their reasons by responding on the basis of those reasons. But as soon as the reason that Nate has enters his deliberation or reasoning, that reason is defeated – Nate can't go straight home on the basis of the fact that

⁸⁶ See Crisp (2006: 43), Schroeder (2007: 92-7).

⁸⁷ The example is from Schroeder. See Schroeder (2007: 33).

there will be a surprise party for him waiting there without defeating the point of going home (since Nate is a bit odd – he only really cares about the surprise, not so much the party).

So the conjunction of the surprise party case and *Compliance* entails that there is at least one reason to which agents cannot correctly respond. But it seems plausible that Nate should go home. In the terms of our debate, it looks like Nate should conform to, but cannot comply with, his reason.

One way to respond to this case is to argue that in fact Nate does not have a reason to go home, by appealing e.g. to premise 1 of the argument from response-guiding. Many philosophers agree on premise 1, and even Raz, whom I have taken as opponent in this chapter and who denies *Compliance*, believes that reasons are essentially action-guiding in the sense that in order for X to be a normative reason for Y, Y must be *able* to respond on the basis of X. But the surprise party case presents a reason that entails the falsity of this claim. We can build an argument that since response-guiding claim is plausible, if we have independent reason to doubt that Nate in fact has an undefeated reason then we can argue from the response-guiding claim to the falsity of the claim that Nate has a reason.

One way to elicit doubt that Nate has an undefeated reason to go home is to draw a distinction between different sorts of *overall* claims. We can agree that it would be good for Nate if Nate went home, and we can agree that there is no reason that he can use as a guide to decide to go home. So we might distinguish between an evaluative and a deliberative sense of what Nate ought to do.⁸⁸ According to Schroeder, the evaluative sense picks out the option which A *in fact* has most reason to perform. The deliberative sense picks out what you are rationally required to decide to do *if* you believe you ought to. That the distinction matters can be seen from cases in which an agent must act under limited information, such as the Miners case.⁸⁹ Here, it is agreed, what the agent ought to do will not (as he knows) produce the best state of affairs possible. With this distinction in hand we can better explore our intuitions about the surprise party case. Clearly it is not the case that Nate ought, in the deliberative sense, to go home, since he could not deliberate on that basis (nor

⁸⁸ See Kiesewetter (2016: 11), Schroeder (2011: 1).

⁸⁹ On the Miners case see Parfit (2011: 159-60).

could he be advised to do so). But Nate ought, in the evaluative sense, to go home – it would produce the best state of affairs were he to do so. However, the thought goes, the evaluative sense does not pick out a reason for an agent to respond in any way it simply picks out the best state of affairs, or in Schroeder's terms:

[...] an evaluative sense, on which it means, roughly, that were things ideal, some proposition would be the case.⁹⁰

Since Nate ought not, in the deliberative sense, to go home, and this is the sense picked out by his reasons, we can defuse the Surprise Party in this way.

⁹⁰ Schroeder (2011: 1).

PART II

3. Analyses of Normative Reasons

3.1. Introduction

In this chapter I begin my defence of WRF. I begin by distinguishing what this thesis can mean, and what version of it I intend to defend in this thesis. First, WRF is an analytical intra-normative view. It provides an analysis of normative properties in terms of other normative properties. It is therefore to be distinguished from analytical *inter*-normative views which seek to provide an analysis of normative properties in terms of non-normative properties, or vice versa.

Second, we can distinguish between different forms that WRF might take. As an analytic view it can be concerned with properties (metaphysical reduction), concepts (conceptual analysis), or language (semantic analysis). In this thesis I understand *WRF* as a metaphysical reduction of oughts to reasons. Therefore I understand my thesis as providing not only necessary and sufficient conditions for oughts in terms of reasons, but as the thesis that reasons are explanatorily prior to oughts.⁹¹ Recall that the thesis I intend to defend is expressed as follows:

Weak Reasons Fundamentalism (WRF) (i) Reasons cannot be reduced to deontic properties⁹² (*Weak Reasons Primitivism*) and (ii) reasons are explanatorily prior to oughts (*Deontic Buck-Passing*).

⁹¹ We could instead understand WRF as the view that provides biconditionals with oughts on one side, and reasons on the other, but as silent on explanatory priority

⁹² I understand, for instance, ought, obligation, right, wrong, and permission to be deontic properties.

Note that (ii) can be broken down into two further claims: first, the *positive* deontic buck-passing claim that deontic properties are to be reduced to reasons (which I discuss in chapter 4), second, the *negative* deontic buck-passing claim that deontic properties do not themselves provide reasons (which I discuss in chapter 5).

In this chapter I begin my defence of WRF by examining (i). Here I address various putative analyses of normative reasons in terms of deontic properties. I argue that none of these analyses succeed. If my claims are plausible then we have good *prima facie* reason to believe that reasons are explanatorily more basic than oughts. I then consider a more general challenge to the thesis that reasons are fundamental in the normative domain and argue that it is not persuasive.

3.2. Reasons as Explanations

Broome claims that the property of being a reason is the property of explaining the obtaining of *other* normative relations; in particular he thinks that reasons are to be understood in terms of ought and explanation.⁹³ His strategy is to distinguish between what he takes to be two different kinds of reasons, which are defined in terms of the role they play in explaining facts about what an agent ought to do. Since reasons can be adequately explained in terms of oughts, it cannot be the case that reasons are explanatorily prior to oughts. Thus, if Broome's view is right, (i) is false. Broome first defines a '*pro tanto*' reason in the following technical sense:

Pro Tanto R is a *pro tanto* reason to Φ if, and only if, R is a fact that plays the for-role in a weighing explanation of why A ought to Φ , or in a weighing explanation of why A ought not to Φ , or in a weighing explanation of why it is not the case that A ought to Φ and not the case that A ought not to Φ .⁹⁴

⁹³ See Broome (2004), Broome (2014), and Broome (forthcoming),

⁹⁴ Broome (2004: 36). Note that this differs from my usage of *pro tanto* reason above.

Note how this definition of a *pro tanto* reason differs from the sense in which I have used it above. The crucial difference between my use of that term and Broome's use is his introduction of the notion of a weighing explanation. What is a weighing explanation? According to Broome a consideration figures in a weighing explanation of an ought-fact by playing a 'for Φ ' or an 'against Φ ' role in that explanation, and by having a weight. Broome does not commit himself to a particular theory of weight, or tell us how to understand these weights.

Broome also introduces a different kind of reason, a putative kind that I have not yet considered, as follows:

Pro Toto R is a *pro toto* reason to Φ if, and only if, R is a consideration which (all by itself) explains why A ought to Φ .⁹⁵

A *pro toto* reason, like a *pro tanto* reason, explains why A ought to do something. But a *pro toto* reason does not do this by playing the role of weighing explanation. It explains why A ought to do something without adverting to weighing considerations against one another. Thus, if Broome's definitions are correct we have at hand an explanation of normative reasons in terms of ought and explanation, which would falsify (i). Note that Broome goes further than these definitions in providing a reductive explanation of what a normative reason is. He claims that to say that 'A ought to Φ ' and 'There is a *pro toto* reason for A to Φ ' are equivalent in the sense that there is *always* an explanation of some ought-fact, and this explanation is a *pro toto* reason.⁹⁶

Are Broome's definitions plausible? Brunero offers the following objection to Broome's view.⁹⁷ Suppose that there are two considerations (call these R1 and R2) which, taken individually, are insufficient to explain why A ought to Φ , but are sufficient when taken together. R1 counts as a *pro tanto* reason, and R2 counts as a *pro tanto* reason. However, since there is an explanation why A ought to Φ , A also

⁹⁵ In earlier work Broome calls a *pro toto* reason a 'perfect reason'. See Broome (2004: 34).

⁹⁶ *Ibid*: 35.

⁹⁷ Brunero (2013: 810).

has a *pro toto* reason in addition to these *pro tanto* reasons. But what is the *pro toto* reason in this case? It can't be R1, nor can it be R2 since neither is itself sufficient to explain why A ought to Φ . Suppose Broome claims that the *conjunctive* fact (R1 & R2) is the *pro toto* reason. Since this conjunctive fact explains why A ought to Φ , perhaps this is the *pro toto* reason. So, according to Broome, A has three distinct normative reasons to Φ . Can that be right? Surely if A has three reasons, each of these reasons must have some weight. But how can (R1 & R2) have weight in addition to R1 and R2 individually? The conjunctive fact seems to add no weight to the case for A's Φ -ing. But since, as we have seen above, it is plausible that a normative reason is a weighted notion, Broome's account of a *pro toto* reason is undermined by the fact that it is a weightless notion.⁹⁸

Later (§5.1.3) I consider some ways in which it might be thought that some R could be a reason but add no weight. If there is some way to make sense of that thought, then we can save Broome's view by claiming, for example, that although (R1 & R2) is a reason distinct from R1 and from R2, it does not add any additional weight to the case. However, I argue there that we should in the end reject the idea of a reason which adds no weight to the case.

Turning to Broome's definition of a *pro tanto* reason, I note that the most popular objection to this claim is that the analysis is circular.⁹⁹ This is because Broome defines *pro tanto* reasons in terms of playing either the 'for Φ ' or 'against Φ ' role in a weighing explanation. But it does not seem plausible that the concepts of 'for Φ ' and 'against Φ ' can be thought without the concept of a consideration which counts in favour, and the concept of a consideration which counts against. Since the for/against Φ distinction must be captured on any view (as I claim, this distinction is normatively irreducible) Broome cannot sidestep the complaint. Thus his account of *pro tanto* reasons too fails.

Broome has a reply to this objection. He claims that we need not make reference to the 'for Φ ' role in defining a *pro tanto* reason. It can be defined more simply as something which (i) has a weight, and (ii) combines in some way to

⁹⁸ See Brunero (2013: 810).

⁹⁹ See for example Brunero (2013) and Kearns & Star (2008).

determine whether or not A is to Φ .¹⁰⁰ But this just pushes the question back a step. How are we to conceive of the weight of a reason prior to the idea of counting in favour? As claimed in chapter 1, the idea of a weight involves the idea of some response being favoured *to an extent*. In chapter 5 I will argue that the idea of the weight of a reason, and the idea of a reason are not separable in this way.

Lastly, there are counterexamples to Broome's putative definitions. Recall above that we distinguished between a normative-motive reason and an explanatory reason. An explanatory reason can provide an explanation for why an agent acted in some way that does not justify their so acting. An explanation does not have to be a justificatory one. But if it is not then it is not a normative reason, which justifies acting in some way. This fact suggests a difficulty for Broome's analysis. Since he defines a normative reason partly in terms of explanation, and explanation need not be justificatory, his definition may go awry. Could an explanation why you *ought* to do something fail to be a normative reason to do it?

I think there are two kinds of cases that we could cite. Consider first indirect moral theories. Suppose that act utilitarianism is true, but the best decision procedure is to act on the basis of commonsense moral considerations, rather than the principle of utility. Then it is true that at least one explanation why you ought not Φ is, for example, that Φ -ing is ruled out by a commonsense moral consideration, such as that it would be a promise-breaking. However, what makes your Φ -ing the wrong thing to do (i.e. the normative reason that makes that act wrong) is not that it would be a promise-breaking, but that it would fail to maximise net utility. So the explanation why you ought not Φ and the normative reason which provides the justification for your not Φ -ing can be different.

Second, consider thin deontic and evaluative facts. Thin evaluative facts seem to be able to play the role of explaining why you ought to do something.¹⁰¹ That it would be a bad thing to do can provide *one* explanation why you ought not do it. But if, for example, we accept Scanlon's claim that evaluative facts do not themselves provide us with normative reasons, then again we have a case in which some fact

¹⁰⁰ Broome (2013: 54).

¹⁰¹ See also Schroeter & Schroeter (2009: 291).

explains why A ought to Φ which is not itself a normative reason.¹⁰² I argue in chapter 4 that we should accept the parallel claim with respect to deontic facts. Deontic facts do not themselves provide reasons. Again, this would constitute a counterexample to Broome's definitions.

3.3. Reasons as Evidence

In a series of papers, Kearns & Star (K&S) outline and defend a novel analysis of a normative reason.¹⁰³ They share Broome's view that reasons are to be analysed in terms of the deontic property ought. However they disagree with Broome's reduction of reasons to ought and *explanation*. They claim instead that reasons for some agent A to Φ are considerations that provide *evidence* that A ought to Φ . In this section I offer a number of objections to this account.

First, I briefly motivate and explain the account. K&S give the following analysis of a normative reason:

Reasons as Evidence (RE) Necessarily, a fact, F, is a reason for A to Φ if, and only if, F is evidence that A ought to Φ (where Φ is an action or a belief).¹⁰⁴

As we have already seen, an essential feature of a normative reason is that it justifies by favouring to some extent the response that it supports. Does evidence play a justifying role? One very widespread source of evidence is testimony. That a reliable friend has told you that the match kicks off at 12pm should, plausibly, count as evidence for that proposition. One reason why it should count as evidence is because, given the reliability of your friend about these matters, it raises the probability that the match kicks off at 12pm. So evidence seems to be a sort of justification. If I have evidence, E, that supports p, then I can cite E in defence of my belief that p.

¹⁰² See Scanlon (1998: 97).

¹⁰³ Kearns & Star (2008), (2009), (2013).

¹⁰⁴ Kearns and Star (2009: 216)

Unsurprisingly, since evidence is closely aligned with justification, it is also closely related to reasons to believe. Reasons are the sorts of things that we can cite in justifying the beliefs that we have. That my very reliable friend told me that the match would kick off at 12pm also seems to be a good reason to believe that the match will indeed kick off at 12pm.

A reasonable claim we could make about the relation between evidence and reasons to believe, then, is that

A fact is a reason to believe that p iff that fact is evidence that p.

In other words, facts that are reason-providing are also evidence-providing, and vice versa. This simple analysis may fail, however, according to K&S, because there may exist practical reasons to believe. If there are practical reasons to believe, the facts that provide those reasons are not evidence that p. Rather, they are evidence that you should believe that p, for some practical reason. K&S therefore revise the simple claim in the following way:

A fact is a reason for A to believe that p iff that fact is evidence that A *ought to believe* that p.¹⁰⁵

As K&S claim, evidence has a dual role to play. It may be evidence that p, which supports the ought-claim, or it may be evidence of practical considerations that support the ought-claim.

Analysing epistemic reasons in terms of evidence, then, has some plausibility. But K&S extend their analysis to include practical reasons. Why think practical reasons can be analysed in terms of evidence? To see why, consider the following example. Suppose that you pass somebody who is in severe pain and in need of your help. However, you promised to meet your friend across town in 5 minutes. What ought you to do? As K&S see things, these considerations each count as evidence that you ought to perform different actions.¹⁰⁶ That someone is in great pain is

¹⁰⁵ *Ibid*: 217.

¹⁰⁶ *Ibid*: 222-3.

stronger evidence that I ought to help them, and outweighs the evidence that supports my meeting my friend instead. In this case, the concept of evidence seems to track reasons well.

Treating both considerations as evidence respects the fact that there is a ‘pull’ towards both actions in this case. You see that you ought to help the distressed person, but feel badly for failing to keep your promise. As we have seen, one puzzling feature of normative reasons is their weight. One way to think about practical deliberation is by counting up reasons for and against an action, and making a decision based on which side the reasons fall. But how do they weigh against one another? Do moral reasons invariably have greater weight than prudential or other sorts of reasons? A seeming advantage of K&S’s analysis is that understanding reasons in terms of evidence allows us to also understand *weight* in terms of evidence.

However, despite the initial plausibility of this analysis, and its promise to shed light on some obscure features of normative reasons, I will argue that we have reason to believe that RE is false.

3.3.1. Objections to RE

The first objection to RE focuses on the differences between reasons and ought. There are two cases which seem to provide counterexamples to RE. First, consider *enticing reasons*. Enticing reasons are reasons that do not contribute towards oughts. These reasons therefore cannot be defined in terms of evidence that you ought to do something. So, RE rules out the category of enticing reasons.

K&S could reply to this objection in at least two ways. They could restrict their analysis of reasons to exclude enticing reasons, or they could deny that there are enticing reasons. The first approach is unsatisfactory by K&S’s own lights. Since they want to provide a ‘unified and informative’ analysis of reasons, it would be *ad hoc* to deny a component of the concept they wish to analyse.

Even were they to give an argument to this effect, notice how strong K&S’s claim is here. According to RE enticing reasons are ruled out by definition. This seems overly strong as a conceptual claim. To rule out enticing reasons as a

component of the concept of a reason forestalls the possibility of substantive debate about them. But surely such debates are not just instances of talking past one another.

Second consider *supererogatory* actions. Supererogatory actions are actions that would be good to perform (in some sense) but of which it is false that you ought to perform. One can have reason to perform a supererogatory action, but such a reason cannot, by definition, be evidence that you ought to do that thing.

How could K&S respond to the second counterexample? One way to go would be to distinguish between different kinds of ought-claims. It is not clear what distinctions would be most helpful to use, however. We might distinguish, for example, a prudential sense of ‘ought’ with a moral sense of ‘ought’ but neither seems to capture any sense in which it may be that you ‘ought’ to perform a supererogatory actions.

In a recent article Kearns revises RE in a way that may help to answer this criticism.¹⁰⁷ He claims that we should widen the analysis of reasons to include evidence for other kinds of deontic and evaluative verdicts. In particular, we should accept an analysis of normative reasons in terms of evidence that it is valuable that A Φ 's, that A ought to Φ , or that A is permitted to Φ .¹⁰⁸ As long as one of those disjuncts is satisfied, A has a normative reason to Φ .

This revision seems to get around the problem of ruling out enticing and supererogatory reasons if we understand these in terms of evidence that acting in that way was valuable, as opposed to being evidence that A ought to Φ .¹⁰⁹ So let us assume that this revision applies to RE.

Broome objects to RE on the grounds that it implausibly implies reasons to act where there are none.¹¹⁰ He cites the example of the fact that a reliable book tells you that you should eat cabbage which, in fact, would make you very ill. Let us

¹⁰⁷ Kearns (2016).

¹⁰⁸ Kearns (2016: 180).

¹⁰⁹ One minor problem with this analysis is that it misconstrues value on many accounts. Many have found it plausible that value attaches not to an action, but either to the agent who authors that action, to the state of affairs produced by that action, or to the aim of that action.

¹¹⁰ Broome (2008: 101).

suppose that in addition to this bad-making property of the cabbage, it has no good-making properties, such as for example its having no nutritional value whatsoever.¹¹¹ Since K&S understand probability in epistemic terms the fact that the book is reliable increases the probability that it is rational for you to assign to the proposition that you ought to eat the cabbage. On any plausible account of reasons you have a reason to *believe* that you ought to eat the cabbage. Nevertheless, since there is nothing to be said for eating the cabbage, and it would be quite bad for you if you did eat the cabbage, it is plausible that you have no reason to eat the cabbage. Thus, according to RE, you have evidence that you ought to do something which, by hypothesis, you have no reason to do. So reasons to act cannot be equivalent to evidence that you ought to act.

K&S reply that if we deny that evidence that you ought to do something necessarily gives you reason to do that thing this will implausibly *restrict* what we should count as reasons.¹¹² In the cabbage example they will claim that the fact that the reliable book tells you to eat cabbage is one that we would ordinarily respond to in practical deliberation. Since it is something that we would ordinarily rationally respond to, we should agree that it plays the role of a normative reason. The thought is something like this. Since you have a reason to believe that you ought to eat the cabbage, it is *prima facie* rational for you to eat the cabbage. But if it is rational for you to eat the cabbage, you must have a normative reason to eat the cabbage. So we should accept that the reliable book gives you reason to eat the cabbage.

However, I am not persuaded by the inference from what it is rational for you to do to what you have reason to do. That is because eliding this distinction entails the implausible result that the very same consideration can be both a reason for and a reason against for the same agent at the same time. Consider an example from Parfit, slightly modified.¹¹³ Suppose that you believe on the basis of reliable testimony that vipers will attack stationary targets. You therefore take the fact that there is a

¹¹¹ I use the terms ‘good-making property’ and ‘bad-making property’ here to avoid explaining the example in terms of reasons, which might be question-begging against K&S in this context.

¹¹² Kearns & Star (2008: 47-9).

¹¹³ Parfit (2011: 34).

rattlesnake in front of you as a reason to run away. The fact that it is a rattlesnake is, given your background beliefs, evidence that you ought to run away. So it is, by RE, a normative reason to run away. However, it turns out that although *most* vipers will attack stationary targets, rattlesnakes are the exception. They will, in fact, only attack moving targets. So it is also true that the fact that it is a rattlesnake is a normative reason *against* running away.

I think it comports better with our common-sense usage of the concepts ‘reason’ and ‘rationality’ to claim that in this case you acted rationally in running away, even though there was no reason for you to do so (and decisive reason against doing so) rather than claiming that the fact that it is a rattlesnake is both a reason to run away and a reason to stay put. So I don’t think we should collapse reasons and rationality as K&S’s reasoning suggests that we do.

It is an important part of K&S’s analysis of a reason that we understand evidence in terms of what *probability* it would be rational to assign to the proposition that A ought to Φ . Brunero picks up on this feature of RE to press further counterexamples. He argues that it can be true (a) that a fact could be a reason to Φ without raising the probability that A ought to Φ , and (b) that a fact can raise the probability that A ought to Φ and not be a reason for A to Φ .¹¹⁴ I will focus solely on (a) in what follows.

To show that a fact can be a reason to Φ without raising the probability that A ought to Φ , Brunero considers the following case. Suppose that

E1 Dad would be happy were I to get Mom some specific gift he found featured in the Sears Catalogue.¹¹⁵

I have no reason to doubt that E1 is true, but I know, as a piece of background information, that whenever Dad is happy getting Mom a gift, there is a stronger reason that outweighs getting that gift for Mom, because he usually has pretty bad taste.

¹¹⁴ Brunero (2009: 540-1).

¹¹⁵ *Loc. cit.*

Brunero claims that E1 is a reason for me to get the gift because it would make Dad happy, but not only does it not raise the probability that I ought to buy the gift, given my piece of background information it actually *lowers* the probability that I ought to buy the gift, because I know that if E1 is true, the reason to buy the gift will be outweighed by competing reasons. Hence RE is false read left-to-right.

K&S offer some replies to this objection.¹¹⁶ First, they claim that since the fact that it would make Dad happy to get Mom the gift is *always* outweighed suggests that this is not, in fact, a reason for me to get the gift. They say that ‘... it seems that it might be best to say he *never* has a reason to get gifts for Mom that he would be happy with.’¹¹⁷ But this seems wrong. What is true is that I never have *sufficient* reason to get the gift *when* it will make him happy. But why deny that the fact that Dad would be happy gives me *pro tanto* reason to buy the gift? Of course Dad’s happiness is outweighed by more important considerations, namely that Mom will hate it. As we have seen, getting happiness can be undercut as a reason when it is borne of sadism or some other vicious motive. But nothing in the case suggests that these undercutting considerations are present.

Their second reply is to admit that E1 is a reason and thus that it *does* raise the probability that A ought to Φ . The difficulty now is to avoid contradictorily admitting that the same consideration both raises and lowers the probability that A ought to Φ . K&S get around this difficulty by claiming that considerations increase the probability that A ought to Φ relative to some salient subset of our evidence. Relative to the piece of background information that Mom hates the gifts that Dad likes, E1 lowers the probability that I ought to buy the gift. Without having that piece of information, E1 raises the probability that I ought to buy the gift.

However this move seems to just run into the same problem as above. Where I know the piece of background information the probability that I ought to buy the gift is lowered but I still have the reason. The reason doesn’t disappear just because I am privy to the information. If K&S admit that the fact that it would make Dad happy *is* a reason then I see no reason why it should not count as a reason with or without the background information. Again, being privy to that information I merely

¹¹⁶ Kearns & Star (2013).

¹¹⁷ Kearns & Star (2013: 82).

understand the reason stemming from Dad's happiness is outweighed. The case does not suggest that the reason is undercut.

Their third response is to deny that E1 lowers the probability that I ought to buy the gift. Since it would make Dad happy, that makes it more likely that I ought to buy the gift. Of course this seems implausible since the case is set up precisely such that when Dad is happy in his choice of gift, *decisive* reasons appear on the other side.

K&S argue for this by first noting that the proposition 'That it is not the case that I ought to buy the gift' (P1) is a different proposition from the proposition 'That I ought not to buy the gift' (P2). K&S claim that Dad's happiness can increase the probability of the proposition 'That I ought to buy the gift' *and* the probability that P2 but not the probability that P1. Because the propositions 'I ought to buy the gift' and 'It is not the case that I ought to buy the gift' are strictly speaking contradictory, they cannot admit that Dad's happiness increases the probability of both of these propositions. But because 'I ought to buy the gift' and 'I ought not to buy the gift' and *not* contradictory propositions, the same consideration can raise the probability of both.

However, even if we admit that it is logically possible that Dad's happiness does not lower the probability that I ought to buy the gift, that is not a reason to believe that happiness in fact does not lower the probability that I ought to buy the gift. This also runs into the same problem mentioned above with respect to the distinction between reasons and rationality. It does not comport with commonsense to claim that the same consideration counts both in favour of and against acting in some way at the same time. If this is an implication of RE that is a cost of the view.¹¹⁸

Further, this response seems to conflate what we might call indicative and non-indicative considerations. K&S will claim that the fact that it would Dad happy is a reason not to buy the gift. But that 'reason' is merely *indicative* of the genuine reason-giving fact, which is that Mom will hate the gift (whatever it is). Granted, I

¹¹⁸ See also Snedegar (2018: 31) who claims that the objective in virtue of which some fact is a reason cannot both explain why the same fact is a reason for Φ -ing and a reason against Φ -ing. He calls this thesis 'Restricted Exclusivity'.

am opposing my intuition about reasons against K&S but I think that the distinction between indicative considerations and normative reasons is one reflected in common-sense and again RE's denial of this looks like a cost of their view.

Finally, consider a recent revision to RE, suggested by Whiting. He amends the analysis as follows:

R = ERR The fact that p is a reason for a person to Φ iff R is a respect in which it is right for her to Φ ; and the fact that p is evidence of R.¹¹⁹

What does it mean for something to be right (or wrong) in some respect? Whiting does not intend by this *pro tanto* rightness or wrongness, and that is anyway unhelpful since it is doubtful that there is any way that we can understand *pro tanto* rightness or wrongness without appealing to the idea of counting in favour. He also cannot intend by 'right in some respect' that there is some respect in which A ought to respond. That is because, if it is false that A ought to respond, then there is therefore no respect in which A ought to respond. Nevertheless, A may have a weak reason to respond in that way.

The thought is that if right/wrong are overall properties then it makes no sense to say, of a particular response, that it is right/wrong in a respect because opposing overall verdictive properties exclude each other. We could say something like: Φ -ing is wrong, though there is a reason to Φ which, in the absence of some wrong-making feature(s) of Φ -ing, would make it right to Φ . But that is a respect in which Φ -ing *would be* right if it lacked certain wrong-making features. It is not a respect in which Φ -ing is, in fact, right.

So rightness cannot be understood as an overall property on this view, and is not plausibly understood as contributory in the sense of being equivalent to the idea of *pro tanto* rightness. Since rightness must be understood as either an overall or a contributory property, I do not think we should follow Whiting's suggestion.

¹¹⁹ Whiting (2017a: 7).

3.4. Parfit's Reasons Fundamentalism

Finally, I consider a slightly more complex objection to the claim that reasons are not analysable in terms of other properties. In chapter 1 of *On What Matters*, Parfit claims that

[I]t is difficult to explain what the concept of a reason or the phrase 'a reason' means.¹²⁰

He compares the concept of a reason to the concepts of time, consciousness and possibility in that, he claims, they are all indefinable concepts. He concludes that reasons are fundamental in this sense. What kind of justification would be required to support this claim? Since the claim is that the concept of a reason cannot be defined, one way to justify this would be proof by exhaustion. If it turns out that the best analyses of reasons fail, we have a good reason to accept Parfit's claim.

3.4.1. Smith's objection

Here I will discuss whether the relation of being a reason is fundamental by addressing a challenge to this crucial stage-setting part of Parfit's book. I examine Michael Smith's recent attempt to reduce the concept of a reason to the concepts of evidence and truth.¹²¹ Smith's argument runs as follows. Parfit claims that the concept of a reason is unanalysable ('fundamental' in Smith's terminology). Reasons come in broadly two sorts: epistemic reasons and practical reasons. The concept of an *epistemic* reason, Smith claims, can be reduced to the concept of evidence. But it is implausible that the concept of an epistemic reason is not fundamental, while the concept of a *practical* reason is fundamental. Since we have good reason to think that epistemic reasons are non-fundamental, we have reason to think that practical reasons are non-fundamental. One way in which this could be would be for practical

¹²⁰ Parfit (2011: 31).

¹²¹ Smith in Singer ed. (2016) .

reasons to reduce to the concept of an epistemic reason, which then reduces as above. Thus the concept of a reason is fully reducible, *pace* Parfit.

Smith claims that the concept of a reason for belief can be analysed in terms of entailment in the following way:

Epistemic Reason R is a reason for A to believe that p in C if, and only if,
R entails p.¹²²

This definition is controversial. Do we not often have excellent reason to believe that p when R fails to entail p? The fact that the weather report tells you that it will rain later is generally a good reason to believe that it will rain later, but it does not *entail* that it will do so.

It is clear that Smith does not intend to defend this analysis. His point is that since there are *putative* analyses in the offering, Parfit needs an argument to defend his claim that all reasons are fundamental. Smith's objection runs deeper than merely pointing out a possible analysis. Once he has established that it is possible to analyse the concept of an epistemic reason, he then claims that it would be implausible to think that practical reasons are unanalysable. It would be so since the concept of a reason would then be disjunctive, being one part analysable, and another part unanalysable. Since it is implausible that the concept of a reason would be disjunctive in this way, Parfit should give up his reasons fundamentalism.

However, I doubt whether this conceptual claim is plausible, for the following reason. If the concept of a reason for belief were to be analysed in terms of evidence, then that would rule out pragmatic reasons for belief as a conceptual claim, where a pragmatic reason for belief is a reason to believe something because it would be good for you to believe it. We should note how strong this claim is. It is not the claim that there are no pragmatic reasons for belief, but that the concept of a reason for belief itself excludes such pragmatic reasons. Thus those who claim that there *are* pragmatic reasons for belief are not only mistaken, but are misusing their concepts when they make their arguments. But surely intelligent, informed philosophers are

¹²² *Ibid*: 101.

not just talking past one another when they conduct debates about whether there are such pragmatic reasons. So Smith's conceptual claim is very strong.

Smith claims that Parfit contradicts his own claim that reasons are unanalysable in at least one instance: he claims that reasons for acting can be analysed into reasons for having the aim for which one acts. Suppose the reason that I have for writing these words is that I have the aim of writing about Parfit's view. That reason, he claims, can be analysed in terms of a reason to have that aim in the first place, since if I did not have reason to have the aim, I would have no reason to perform the action (to type the words). Since every action is performed with some aim, the claim is generally true. All practical reasons are analysable in terms of reasons for A to have the aim for which that action is done.

However, I doubt whether this means that Parfit contradicts his earlier claim, since it does not violate his reasons fundamentalism that we explain reasons *in terms of reasons*. As long as the explanation includes reasons, Parfit's claims are consistent.

Smith then claims that we should reduce reasons for aims to reasons for intending, reasons for intending in terms of reasons to desire, and finally reasons to desire in terms of reasons to believe desirable. I won't here consider whether all of these steps in the argument are plausible, but will focus solely on the last step: reasons to desire to reasons to believe desirable. How does this argument run? He first argues, following Thomson, that many mental states have correctness conditions.¹²³ What is a correctness condition? According to Thomson, they are conditions under which a mental state is *deserved*. For example, on this account, beliefs are correct in virtue of the truth of their content, and desires are correct in virtue of desirable objects and states of affairs, since desirable objects and states of affairs are deserving of desire.

Given this notion of correctness conditions, Smith describes Thomson's account of what it is for there to be reasons for being in a mental state:

A consideration is a reason for being in a mental state that has a correctness condition just in case that consideration is evidence for, or makes probable, or

¹²³ See Thomson (2008: 83ff).

lends weight to, the truth of the proposition that is that mental state's correctness condition.¹²⁴

So that,

[R]easons for desire are considerations that lend weight to the truth of the proposition that the desired object is desirable, as the truth of the proposition that the object of desire is desirable is the correctness condition of desire.¹²⁵

In other words, we have a reason for a desire in virtue of whether the object of that desire is desirable. Finally, with this general analysis of reasons in hand, we can reduce reasons for desire ‘... to reasons for believing that acting in the relevant way is desirable.’¹²⁶ I take it that Smith merely identifies the latter reason to believe with a reason to desire.

Just as we saw above that it is at least a *conceptual* possibility that there are pragmatic reasons for belief, so too there may be pragmatic reasons for desires, i.e. reasons to desire or not desire some object which are not grounded in claims about that object’s desirability. They may be grounded in claims about what would be good for you, were you to desire it. Again, claiming that pragmatic reasons for desire are ruled out conceptually seems a very strong claim. Parfit denies that such ‘state-based’ reasons are reasons to believe, but a defender of reasons fundamentalism need not follow him in this.

Another difficulty with this claim is that the satisfaction conditions for belief and desire differ. Having sufficient reason for belief is satisfied by your coming to adopt the belief that is supported. If reasons for desire are reducible to reasons for belief, then you have done everything that you ought rationally to do in responding to your reason by forming the belief that the object is desirable. Now it may seem plausible that if you believe that p is desirable then you ought, rationally, to desire it, but that forms no part of what makes it rational to respond to your reason. That

¹²⁴ Smith in Singer ed. (2016: 109).

¹²⁵ *Loc. cit.*

¹²⁶ *Loc. cit.*

reason is satisfied simply by your coming to adopt a belief. But we might think it is part of the satisfaction conditions of a reason for desire that if you have sufficient reason to desire some object you ought, rationally, to desire it . That truth is not captured on this account.

Finally, we should note that the account still needs to tell us what is to count as desirable. But Smith's account of 'desirable' in terms of what we would desire after rational and well-informed deliberation falls short of the mark. I note here merely that it is this is a controversial account of what it is to be desirable.

In the end then I do not think we should be persuaded that WRF (i) is false on the grounds that Smith outlines. I turn in the next chapter to discuss the thesis that *other normative* properties, in particular deontic properties, can themselves be analysed in terms of reasons.

4. Deontic Buck-Passing: The Positive Thesis

4.1. Introduction

In this chapter I outline an argument for WRF (ii), the claim that normative reasons are a more fundamental normative property than deontic properties. I argue that, in a structurally similar way to Scanlon's 'buck-passing' account of the normative property of goodness, which says that X's goodness is the property of X's having other properties (other than the property that it is good) that give us reason to care about X, deontic properties are properties of an action's having other properties (other than those deontic properties) that combine to give us decisive reason to perform, or not to perform that act.¹²⁷

This buck-passing thesis combines two claims: a negative claim, that deontic properties are formal properties which do not themselves provide reasons for action, and a positive claim, that deontic properties are to be understood in terms of reasons. Here I examine the positive claim, and attempt to fill out the details of the proposal. It turns out that there are several obstacles facing the positive claim, which must be addressed and overcome if the analysis is to be viable.

Let us first distinguish a first-shot analysis of the positive claim, focusing on the deontic property ought:

Simple Analysis (SA)	A ought to Φ if, and only if, A has decisive reason to Φ .
-----------------------------	--

There are two main sets of difficulties with SA, as it stands. The first set of difficulties centres around the idea that SA is unable to sufficiently discriminate between different kinds of practical conclusions. Suppose that to have decisive reason is for the reasons which favour acting in that way to be stronger than the

¹²⁷ Scanlon (1998: 96-100). See chapter 1 for my account of what a normative reason is.

reasons against, in some intuitive sense.¹²⁸ We can compare two cases in which this might be true.

Consider the following case. You are deciding whether to go to the cinema or the theatre tonight. You have already seen every film currently showing and there is a particularly interesting play that is only in town tonight. In some intuitive sense you have stronger reason to go to the theatre. According to SA, you ought to go to the theatre. Let us further suppose that you decide not to go to theatre. It seems to be an open question whether you are appropriately subject to serious criticism for not going to the theatre: on the one hand you had decisive reason (in the sense just defined) to do so but on the other it seems to be merely an option that it is up to you to flout.

Now consider a different case. You have promised your friend that you will meet him for lunch at 12pm tomorrow, but you would rather spend the afternoon watching television. Let us suppose that your promise is a stronger reason to meet your friend than the satisfaction of your desire to watch television is a reason not to do so. So, according to SA, you ought to meet your friend. Suppose that you decide not to meet your friend. In this case it does *not* seem to be an open question whether you are appropriately subject to criticism for failing to meet your friend. You are clearly at fault for your failing. In this case meeting your friend does not seem to be merely an option that it is up to you to flout.

The problem for SA is that it seems to be unable to discriminate between these cases, between cases in which reasons support a practical conclusion that *recommends* acting in some way, and cases in which reasons support a practical conclusion that *requires* acting in some way.

Not only does SA fail to discriminate between different practical conclusions, it is objected that any buck-passing account of ought is unable to capture what is distinctive about, in particular, *moral* requirements. The argument depends on a simple thought: normative reasons are the sorts of things that merely make actions eligible or attractive, telling us what makes some action desirable. They count in favour of acting in some way, and counting in favour of an action is just

¹²⁸ I discuss and try to precisify these concepts below; for the moment I just want to use them to get the issues on the table.

categorically different from being under a moral requirement, to act in some way, or its being morally wrong to do something. *No matter how many reasons we add into the picture* they will never add up to give us these moral properties.

Naturally, many theorists have been attracted to introducing further normative distinctions in order to explain the difference between recommendation and requirement. A further problem that arises for these theorists, once they have captured each of these practical conclusions, is to capture intuitively plausible entailments between them. In particular, that an action is required seems to imply that it is recommended, and that it is recommended does not entail that it is required. As we will see, introducing further normative distinctions makes guaranteeing these entailments tricky.

The second set of difficulties for SA is that it appears to make supererogation impossible. Suppose that it is supererogatory for you to give half of your annual income to charities providing aid to starving people. That is, although it would be very good if you did that, it is false that you ought to do it. Doing so would be to go ‘above and beyond the call of duty’. We might think that it is possible that the fact that those people are starving and in great pain, which you can alleviate, gives you overwhelmingly strong reason to give half your income. If SA is true, the fact that you have decisive reason to give half your income entails that you ought to give half your income. Thus either your giving half your income to charity is not a supererogatory act, or SA is false. Since it seems overwhelmingly plausible that if anything is a supererogatory act, your giving half your income to charity is a supererogatory act, the truth of SA would seem to make supererogation impossible.

In what follows I will discuss and reject various solutions to the first set of problems. I develop my account in contrast to these failed putative solutions.

4.2. The Enticing/Peremptory Distinction

The first solution to our difficulties is raised by Raz and defended by Dancy.¹²⁹ We should revise SA because, as stated, it glosses over a distinction between two different kinds of normative reasons.

¹²⁹ See Dancy (2004a), (2004b), and Raz (1999: 100).

What is the distinction that Dancy has in mind? He tells us that

[As well as] the sort of ‘peremptory’ reasons I have so far been discussing, which certainly do stand in some close relation to oughts (even if we are finding it hard to characterize them in terms of that relation), there are reasons of another style, which I call enticing reasons, and these do not stand in the same relation to oughts at all. Enticing reasons are to do with what would be fun, amusing, attractive, exciting, pleasant, and so on. They can be stronger and weaker, and they are often strong enough for action. But (as I understand the matter) they never take us to an ought; it is not true of an enticing reason that if one has one of them and no reason of any other sort, one ought to do what the reason entices one to do. One can do that; but one has the right not to. With peremptory reasons we could not say any such thing.¹³⁰

The distinction Dancy has in mind is between considerations that favour by ‘enticing’ and considerations that favour in a more peremptory fashion. Enticing reasons ‘never take us to an “ought”’ but peremptory reasons can.¹³¹ This suggests a revision to SA:

SA* A ought to Φ iff A has most *peremptory* reason to Φ .

We can overcome our first set of difficulties by claiming that those acts which are merely recommended involve enticing reasons, while those that we are required to perform, and thus ought to perform, involve peremptory reasons.

Let’s try to get clearer on what Dancy’s distinction amounts to. He provides us with a negative and positive characterisation of an enticing reason. First, an enticing reason can be typified with respect to its content. If a favouring consideration’s content is related to fun, amusement, excitement or pleasantness, that consideration is an enticing reason. On the other hand we have a negative

¹³⁰ Dancy (2004a: 21).

¹³¹ *Loc. cit.*

characterisation: enticing reasons are considerations which favour acting in some way but they ‘never take us to an ought’ in cases in which we have an enticing reason to Φ and no reason ‘of any other sort’.

Consider the negative characterisation first. What does it mean to say that a reason *never* takes us to an ‘ought’? Dancy provides us with an example. If A has a reason to Φ and no reason to Ψ (i.e. to do otherwise) and it is not the case that A ought to Φ , A’s reason to Φ is an enticing reason.¹³² It would be pleasant for Smith to remain sitting on the couch, and since the football match has been postponed, he has no reason to do anything else. Still, it is implausible to claim that Smith is not permitted not to leave the couch. He may do so, but he may do otherwise.

The example, however, is an odd one. Suppose that Smith fails to Φ . Then he Ψ ’s (I understand action here to include omission) but has, by hypothesis, no reason to Ψ . Smith’s failing to Φ (given his knowledge that there is no reason for him to Ψ) does look plausibly like a case of irrationality.¹³³ I also find it impossible to conceive of having only one reason – to Φ – which makes it difficult to assess the example. Perhaps we should say that Smith has many available actions, each supported by a *pro tanto* reason, but has *more* reason to go out and play football than to do otherwise. Yet, it is not the case, according to Dancy, that Smith ought to play football.

We can revise this definition of an enticing reason as follows. If A has more reason to Φ than to Ψ and it is not the case that A ought to Φ , A’s reason to Φ is an enticing reason. That avoids the oddity of the previous example, but clearly this will not do in the present context of investigation. We are trying to amend our analysis of ought in terms of reasons and are now appealing to a distinction among reasons. In that context it is clearly circular to appeal to ought in order to draw that very distinction. If enticing reasons are considerations which favour but never add up to ought (as in the previous case) while preemptory reasons are those that *do* add up to ought we will have done little more than to re-label our first problem for SA above. In order to make progress on this front then, we need to find an independent characterisation for the enticing/preemptory distinction.

¹³² *Loc. cit.*

¹³³ See for example Davidson (1963).

Consider now Dancy's positive characterisation of an enticing reason. Whether a reason is enticing depends on its content: that it is fun, pleasant, amusing, or exciting. If Smith's options above turn out to be supported by enticing reasons, then we have ready to hand an explanation why it is not the case that Smith is permitted not to act on his best option. Smith's best option is to play football because it will allow him to catch up with his friends, and will be good for his health (among other benefits) while sitting on the couch watching TV will be mildly entertaining and relaxing. Some of these considerations naturally fall under these categories: catching up with friends can be pleasant and fun. It is not as clear whether the promotion of health falls under any of these categories, but perhaps the list is not intended to be exhaustive.

In fact, I think it is clear that Dancy does not take the list of the contents of potential enticers to be exhaustive. He tells us that those considerations that count as enticing '... have an *evaluative* focus' while those that count as peremptory '... have a *deontic* focus'¹³⁴ For Dancy's purposes, a consideration has an evaluative focus insofar as it contributes to making it the case that Φ -ing is the *best* thing to do. Pleasantness, fun, amusement and excitement are grounds for an assessment of a potential action as the *best* or *better* alternative.¹³⁵

Appealing to the content of reasons in this way to draw the distinction may avoid the problem of vicious circularity. However, there are two important objections to this characterisation which put pressure on us to rule out this approach. First, we should be very wary about a content-based distinction as a general strategy. Why should we expect that the structure of any concept, and a normative reason in particular, could be determined by its content? Contrast the concept of a reason with the concept of a belief. We would not expect to construct a theory of belief on the basis of the content of particular beliefs. Why should we expect it to be otherwise for reasons?

Second, as Robertson points out, enticing reasons introduce an asymmetry that is not easy to explain on Dancy's model.¹³⁶ Reasons can count for or against an

¹³⁴ *Ibid*: 116, my emphasis.

¹³⁵ See *ibid*: 95, 110, and Dancy (2004a: 24-5).

¹³⁶ Robertson (2008: 269).

action. The fact that it would be pleasant to eat your favourite pasta dish for lunch counts in favour of eating it; the fact that it would be downright unpleasant to drink paint counts against drinking it. Suppose we agree that if you have no more reason to do otherwise than eat the pasta, you are nevertheless permitted not to eat it. There may be nothing irrational in your failure to do so.

On the other hand, deciding to drink the paint looks like a pretty irrational decision. If that's right then it doesn't look implausible to claim that A (rationally) ought not to drink the paint. After all there is nothing to be said for drinking it, and it would be terribly unpleasant to do so. If pleasant feelings paradigmatically constitute enticing reasons, then painful feelings must constitute *dis*-enticing reasons. But surely the fact that some action would be painful can 'add up to an ought' especially given the absence of any defeaters? What about the fact that the pain would be inflicted on another agent? Couldn't this dis-enticing reason make it the case that I ought not to do so?

Appealing to the evaluative/deontic distinction makes the asymmetry look stark. Clearly being in *pain* is, all else being equal, a bad state of affairs, and therefore falls on the evaluative side of the distinction. If enticing reasons take us to 'bests' while preemptory reasons take us to 'oughts' then just as preemptory reasons against take us to 'ought nots', dis-enticing reasons take us to 'worsts' or what it would be overall worst to do. While it seemed somewhat plausible that A is permitted not to do what is best for him, it strikes me as far less plausible that A is (rationally) permitted to do what is *worst* for him. There thus seems to be an asymmetry within the evaluative that leaves an explanatory gap in Dancy's theory.

Notice that the claim that I am arguing against, that reasons with evaluative content can never (all by themselves at least) generate an 'ought' is very strong. I have argued that the existence of dis-enticing reasons shifts the burden of proof onto the defender of enticing reasons, at least on the present definition. It is fair to demand at the very least an explanation why in a number of cases it is not true that I ought to refrain from inflicting pain on A (for example).

A possible move here would be to weaken the positive definition of enticing reasons. Whether a reason is enticing does not depend on whether the content of that reason is evaluative, but on whether that content is some subset within the set of

evaluative considerations. But on what basis can we now distinguish between those evaluative reasons which are enticing, and those that are peremptory? (or of some other ‘flavour’ of favouring?) Perhaps we should appeal to intuition at this point, and ask ourselves, of some considerations, whether it seems to be such that they could not, all by themselves, generate an ought. We should however be cautious in relying too heavily on such intuitions, since for example arriving at correct answer may involve, instead, carrying out first-order moral or practical theorising.

For these reasons, I do not find either the positive or the negative characterisation of the enticing/peremptory distinction particularly convincing.

4.3. The Justifying/Requiring Distinction

In a series of papers, Gert argues for a distinction between two kinds of strength of practical reasons: practical reasons can have requiring or justifying strength. Gert tells us that

Justifying strength is a matter of making it rationally permissible to do something that would otherwise be irrational.¹³⁷

Here, for A to be irrational in Φ -ing is for it be rationally *required* that A not Φ . On the other hand

Requiring strength is a matter of making it rationally impermissible to do something that would otherwise be rationally permissible (optional or required).¹³⁸

Gert claims that this is a distinction which many theorists about practical reasons have overlooked. In light of this distinction he claims that *maximising* (or satisficing) views about reasons are false, that it makes no sense to claim that what you ought to do is to act on what you have most reason to do. He thinks that the question ‘What

¹³⁷ Gert (2000: 230).

¹³⁸ *Loc. cit.*

do I have most reason to do?’ often does not make sense because it is systematically ambiguous.

His *non-maximising* view is that because reasons have two dimensions of strength we should instead ask how strong your *justifying* reasons are, and how strong your *requiring* reasons are, or what do your reasons justify, and what do they require? Gert cashes this out in terms of rationality (though Gert also discusses morality, which he thinks is a separate issue) in the sense that, just as Broome claims that reasons are defined by the role they play in contributing to oughts, Gert claims that reasons are defined by the role they play in contributing to the *rational status of an action*.¹³⁹

Justifying and requiring strength are independent of each other. Justifying strength can offset requiring strength and vice versa. Gert claims that these dimensions of strength are independent in that they do not necessarily co-vary.¹⁴⁰

To illustrate the view, compare A’s saving some strangers from malnutrition with B’s saving £100. A can *justify* much more than B. A may be able to justify something as rationally impermissible as incurring serious harms upon herself, while B is not able to justify to this degree. We don’t think that it would mitigate the rational impermissibility of incurring great harms that by doing so you would save £100. But, conversely, B seems to have greater *requiring* strength than A. We think that it would be irrational to part with £100 without good reason but that it would not be irrational to fail to save the strangers from malnutrition, even in the absence of good reasons not to. Simply put, A can justify more than B, while B can require more than A.¹⁴¹

How can we utilise Gert’s distinction between justifying and requiring to solve the problems that arise for SA? There were two main problems: first, that the analysis cannot distinguish between an action that comes recommended and an action that is required, and second that the analysis rules out, as a conceptual matter, supererogation. In accounting for the difference between recommendation and

¹³⁹ Note here again the worry that the analysis looks circular for reasons fundamentalism applies here.

¹⁴⁰ Gert (2003: 26).

¹⁴¹ See Gert (2005: 356).

requirement, we might directly map Gert's distinction onto this difference. Recommendation gets cashed out in terms of what can justify, requirement gets cashed out in terms of what can require.

Suppose that you have two reasons to eat a banana: that it is tasty and that it would make your friend laugh. On Gert's schema, these reasons are clearly justificatory rather than requiring reasons. However it is not clear to me that Gert's characterisation of justificatory strength can account for cases like this. Recall that for a reason to have justificatory strength is for it to make rationally permissible an action, which would not otherwise be (not without the existence of that reason). I don't think that correctly characterises why A's eating the banana is recommended (an overall verdict).

One difficulty here is that permissibility comes cheap. In order for it to be permissible to Φ , all we need is an absence of an undefeated reason against Φ -ing and an absence of requiring reasons to do otherwise. This fits with Gert's conception of justifying reasons, which are considerations which *answer* criticisms, that answer to the reasons there are against acting in that way. In the banana case, however, we need not make the assumption that there are any reasons against eating the banana, and can still perfectly well make sense of it being recommended (it is perfectly permissible for me to eat the banana just because I feel like it!). But in order to explain recommendation in terms of justifying strength we need to make this assumption, and it is just not clear why. So justifying strength may not map onto recommendation, and be able to provide the needed analysis. Clearly we cannot appeal to requiring strength – the tastiness of the banana doesn't render the permissibility of not eating it somehow impermissible. So the account does not fare well in cases in which it is not *generally* rationally impermissible to do something.

One way of putting the point is that practical reason cannot be wholly characterised in terms of either levelling or answering criticisms. A criticism-based picture might be useful when we think about morality (though as I argue in chapter 5 a criticism-based account of moral obligation such as that defended by Darwall is implausible) but not when thinking about bananas.

A second problem with the justifying/requiring distinction is broached by Snedegar.¹⁴² As we saw above, a plausible constraint on a deontic buck-passing view is that it should guarantee that, in my terms, if you are required to Φ , then Φ -ing is recommended (what Snedegar calls ‘ought’). That is, that there is an entailment from requirement to recommendation (or ‘ought’) but no converse entailment. This seems right. It is not correct to say that Φ -ing is what A is required to do, but Φ -ing is not what is recommended. On the other hand, we should be able to say that Φ -ing is recommended, but A is not required to Φ (say, eat the banana).

The problem is that Gert’s distinction does not guarantee the entailment from requirement to recommendation. Clearly, if requirement maps onto requiring strength, and recommendation maps onto justifying strength, and requiring and justifying strength are independent, then requirement and recommendation are independent. Thus, on this view, it can be the case that A is required to do something which is not recommended. Suppose that A has two options: Φ or Ψ . A has more requiring reason to Φ than to Ψ , which means that the justifying reasons for Ψ -ing are not strong enough to compensate for not Φ -ing. But it can be the case nevertheless that A has strongest justifying reason to Ψ – it is just that they are not strong *enough* to compensate for not Φ -ing.

The crucial distinction that Gert wants to capture is the difference between a consideration answering criticism(s) of acting in some way, and a consideration which tends to rule out acting in certain ways, thus forming the basis for criticism of acting in other ways (and since requiring is a way of making irrational for Gert, and irrationality is what it would be stupid, crazy, etc to do, these are the criticisms he has in mind).¹⁴³ What Gert wants to emphasise is that these roles are conceptually distinct, and respecting their being distinct entails denying a maximising conception of practical reason. In outlining my account I will agree that there is a distinction to be made between reasons that count in favour and reasons that count against, but that recognising this does not force us to give up a maximising picture of practical reasoning.

¹⁴² See Snedegar (2016: 167-72).

¹⁴³ See also Greenspan (2005) and (2007).

4.4.1. Second-Order Reasons

In *Practical Reason and Norms* Raz introduces a distinction between first-order and second-order normative reasons in the service of understanding particular forms of practical conflicts.¹⁴⁴ First-order reasons are, as indicated above, considerations which count in favour of, or against acting in certain ways, or in favour of, or against holding certain attitudes. Second-order reasons are considerations which count in favour of or against acting *on the basis of* first-order reasons. Raz offers the following definition of a second-order reason:

A second-order reason is any reason to act for a reason or refrain from acting for a reason. An exclusionary reason is a second-order reason to refrain from acting for some reason.¹⁴⁵

What does it mean to have a reason to refrain from acting on the basis of a first-order reason? Suppose that it would serve to further your career were you to send your daughter to a local school. That it would further your career gives you a reason to send your daughter to that school. However, you have promised your wife not to base your decision about where to send your daughter to school on your career needs, but instead only on her educational needs. According to Raz, the promise that you made to your wife is not a first-order reason to simply act, but a second-order reason not to act on the basis of a first-order reason.

Although second-order reasons can compete with each other, according to Raz they cannot compete with first-order reasons. He tells us that ‘... the very point of exclusionary reasons is to bypass issues of weight by excluding consideration of the excluded reasons regardless of weight.’¹⁴⁶ Indeed, it is necessary to explain the case above that the fact sending your daughter to a local school would further your career is not to be weighed against the second-order reason. That reason persists (and so makes appropriate regret) but is excluded by your promise.

¹⁴⁴ See Raz (1978) and (1999) ch.1.

¹⁴⁵ Raz (1978: 132).

¹⁴⁶ Raz (1999: 190).

Note that second-order reasons should be distinguished from reasons to refrain or engage in deliberating in certain ways. We may have reason to exclude some consideration from deliberation, say on the grounds that it is unimportant and not worth taking into account in coming to a decision. Such reasons are ordinary first-order reasons which favour or disfavour deliberating in some way.¹⁴⁷ According to Raz, it is not that having made the promise to your wife you have reason not to take your career into consideration when thinking about where to send your child. Perhaps you do. But you have reason not to send your child to the local school *on that basis*.

Introducing second-order reasons gives us an attractive way to account for requirements while remaining consistent with the buck-passing analysis. It does by giving us the tools to construct a ‘two dimensional theory of requirement’ as follows:

Razian Requirement	A is required to Φ iff A has first-order reason to Φ and second-order reason not to act on reasons that favour not Φ -ing. ¹⁴⁸
---------------------------	--

What the Razian Requirement analysis allows us to do is to hold the claim that there can be prudential reasons against acting in the way required by morality but that these prudential reasons should not be acted upon because *excluded* by those moral reasons, *even if these prudential reasons outweigh the moral reasons*.

The view allows us to distinguish between requirement and recommendation by defining requirement as above, and recommendation, by contrast, purely in terms of first-order reasons to act.

4.4.2. Against Second-Order Reasons

We should note that one of Raz’s main motivations for appealing to second-order reasons is his *service conception of authority*, which roughly says that some authority

¹⁴⁷ See Whiting (2017b) for discussion on this point.

¹⁴⁸ What I’ve called a requirement here is sometimes referred to as a ‘protected reason’. See for example Essert (2012).

has legitimacy in virtue of the fact that in following its demands you better conform to the (first-order) reasons that you have anyway, than you would on your own.¹⁴⁹ To illustrate, take the case of Jeremy the soldier. In discussing Jeremy's order to requisition a citizen's vehicle, Raz claims that '[Jeremy] admits that if he were ordered to commit an atrocity he should refuse. But this is an ordinary case, he thinks, and the order should prevail.'¹⁵⁰ This suggests that whether a reason is exclusionary may depend on the strength of first-order reasons, which seems to conflict with Raz's principle that second-order reasons do not weigh against first-order reasons. The resulting view is one in which second-order reasons don't conflict with first-order reasons except where the first-order reasons are sufficiently weighty.

Can we render Raz's view consistent by insisting that in cases in which orders should not be followed, there is no conflict between a first-order and second-order reasons? Suppose instead that the first-order reason serves as a disabler for the second-order reason, removing its normative force, rather than outweighing it. But disabling is not the same thing as excluding. To disable some reason is to remove entirely its normative force. This is not what Raz has in mind by exclusion. The soldier should be able to recognise that where the non-authority-backed excluded reasons are sufficient, he is torn in two directions. That makes sense only if the excluded reason retains its normative force. So we should not collapse the disabling/excluding distinction. Thus Raz's principle that first and second-order reasons do not weigh against each other admits of exceptions.

I want now to focus on three objections to Raz's solution, which imply that we have reason to reject the category of second-order reasons, and that we have reason to reject the non-maximising view that this category commits us to.

First, we should reject second-order reasons because they conflict with RC.¹⁵¹ RC told us that if R is a reason for A to Φ then it is possible that A Φ 's on the basis of R. Strengthening this claim, in defending the *Compliance* view I tried to make plausible that correctly responding to a reason requires that A Φ 's on the basis of R. Thus if RC is true, in order for R to be a second-order reason for A to Φ /not to Φ , it

¹⁴⁹ On the service conception see Raz (1986: 56-67).

¹⁵⁰ Raz (1990: 38).

¹⁵¹ See Whiting (2017b: 403-7).

must be possible that A Φ 's on the basis of that second-order reason. However, as Whiting makes plausible, although it is possible to act on the basis of a reason, it is not possible to act on the basis of a reason, *on the basis of a reason*.¹⁵²

Why can't A act on the basis of a reason, on the basis of a reason? To see why recall the distinction between motivating and explanatory reason. An explanatory reason why A Φ -ed explains why A Φ -ed without necessarily adverting to A's intentions. A may not recognise a (nevertheless true) explanatory reason why they Φ -ed. On the other hand A must recognise their motivating reason to Φ . That reason explains why they Φ -ed in part by adverting to A's intention in Φ -ing. In fact, some think that A's motivating reason is just *equivalent* to A's intention.¹⁵³ So the idea of acting for a second-order reason would be the idea of acting on the basis of a second-order motivating reason. But that doesn't seem right.

Take the example of excluding career-based reasons to send your child to a school. You in fact send your child to the local school for the (unexcluded) reason that it is a good school. Complying with RC would require you to send your child to school on the basis of the (motivating) reason that it is a good school, on the basis of the (motivating) reason that you will not act for career-based reasons. But is it plausible that the latter reason is a motivating, as opposed to an explanatory, reason? Can you form the intention to send your child to school because it is a good school because you will not act for career-based reasons?

I believe that it is more plausible to claim that the reason that *motivated* you to send your child to school was that it was a good school. Not acting on career-based reasons is an *explanatory* reason which offers a partial explanation why you sent your child to that school. To support this intuition Whiting appeals to the claim that acting for the right reasons makes you creditworthy.¹⁵⁴ Plausibly, if you do the right thing for the right reason you are deserving of credit for doing so. But this principle conflicts with second-order reasons. Take again the case above about sending your daughter to school. Suppose that you send her to the school on the basis of educational reasons, which are the reasons which make that action right. Your

¹⁵² *Loc. cit.*

¹⁵³ See Davidson (1963).

¹⁵⁴ *Ibid*: 407-8.

promise to your wife not to act on the basis of career-based reasons serves as your second-order reason. Now suppose that were it not for that promise you certainly would have sent her to a different school for career-based reasons. You only do not do so because of your promise, not because you care about her education.

However, in this case you have acted on the basis of the right reasons, so by the plausible principle you ought to be creditworthy. But surely, given that you don't care about her education, you are *not* creditworthy. So we have a choice. Either reject the plausible principle that if you act for the reasons that make that act right you are creditworthy, or reject second-order reasons. Since the principle is independently plausible that gives us reason to reject second-order reasons.

The second objection is that exclusionary reasons are unable to explain obligations because positive second-order reasons can hypothetically always come in and outweigh exclusionary considerations. So if we can find a suitable basis on which to do something other than that which we are morally obligated to do, it is rational to weigh up these (exclusionary and inclusionary) second-order reasons. In this sense then, the problem is merely shifted up one level. Raz may reply that just as we cannot balance first-order reasons against second-order reasons, we cannot balance inclusionary and exclusionary reasons. However, exclusionary reasons then appear to fail to be a weighted notion. As we have seen in chapter 1, reasons are paradigmatically weighted properties, so this would imply that exclusionary reasons are not reasons at all.¹⁵⁵

The third objection is that we should reject the Razian picture as it commits us to denying that you should always act on the balance of reasons. Exclusionary reasons are reasons not to act on the basis of some first-order reasons. But sometimes one's first-order reasons can be sufficient or decisive reason to act in some way. If exclusionary reasons can exclude acting for some reason, they can exclude acting for some *decisive* reason. But this seems *prima facie* problematic. Suppose that you reason that you have decisive reason to act in some way in the sense that the balance of reasons favours acting in that way. Now suppose that you have an exclusionary reason not to act on the basis of this reason (or for these reasons). From this possibility it follows that it is not rational to act on the balance of reasons. But how

¹⁵⁵ On this point see also Owens (2008: 416-20).

can it not be rational to act on the balance of reasons?¹⁵⁶ The two notions are so closely related that it does violence to the idea of rationality.

Consider again the example of the soldier. The soldier has been instructed to appropriate the car, but if he does so he will render the person who owns the car unable to travel to his place of work, resulting in his losing his job, with deleterious effects on his family. We can suppose in this case that there is in fact most first-order reason for the soldier not to steal the car. However the soldier also has second-order reason not to act for this reason, derived from the authority of his superior. Thus, according to Raz, what the soldier ought to do is to steal the car. But claiming that practical reason is two-tiered in this way seems to imply a problematic kind of fragmentation. How can I respond to the fact that I am in a position to know what I have most (first-order) reason to do, and know that I have (second-order, decisive) reason not to do that thing? According to Raz I should hope that the thing gets done, even though I shouldn't do it. That claim can make sense for, say, agent-relative reasons (I hope he gets the birthday present even though I shouldn't give it to him because he would then reject it) but absent such reasons the case looks paradoxical.

4.5. Towards a Different Analysis

So far I have argued that we should reject various ways of amending SA. We should not appeal to a content-based distinction between different kinds of normative reasons because doing so is either *ad hoc* or circular. We have seen two kinds of solutions that appeal to a non-maximising strategy that reject the claim that an agent always ought to act on the balance of reasons. Both of these views face *fragmentation* objections. Gert denies the maximising view by claiming that there are two irreducibly distinct practical verdicts: what you are required to do, and what you

¹⁵⁶ Gert claims that if second-order reasons can have as their content decisive first-order reasons, they introduce a problematic paradox. In order to avoid the difficulties involved in the *ad hoc* strategy of choosing a level of sufficiency of strength beyond which second-order reasons cannot apply, Gert advocates making use of his requiring/justifying distinction. Only justifying reasons can be excluded by second-order reasons. See Gert (Forthcoming: 20-1).

are justified in doing. Drawing this distinction gives us the tools to distinguish between requirement and recommendation, but at the cost of severing any necessary connection between these verdicts. On Raz's view the introduction of second-order reasons created a more immediate difficulty in raising the possibility that agents can be in a position in which he knows both that he has most (first-order) reason to do something, but nevertheless has reason not to do it. That seems highly counterintuitive.

In formulating my own analysis I try to learn from these difficulties. My analysis does not rely on content-based distinctions, accepts the intuitively plausible maximising structure, and ensures that requirement and recommendation are conceptually related in the desired way.

4.6. Clarifying Some Concepts

Before outlining my own view I need to discuss some of the concepts that I will make use of in what follows. First I must say something about normative reasons and how they interact with each other in practical reasoning. I take normative reasons for acting as facts that either count in favour or count against A's acting in some way. Most normative theorists claim that the following claims are true:

Negative Reasons Equivalence	'R is a reason against A's Φ -ing' is equivalent to 'R is a reason for A's not Φ -ing'.
Positive Reasons Equivalence	'R is a reason for A's Φ -ing' is equivalent to 'R is a reason against A's not Φ -ing'.

As I argued above in §1.4 I deny these equivalences. I do so because I claim that to be a reason for is to provide support for acting in some way, while to be a reason against is to count against acting in some way. Even at an intuitive level these seem to be quite different relations.

Let me briefly recap the motivation for this denial. Consider the following example. The fact that you will become addicted to the drug is a reason against taking the drug; this fact directly counts against doing so. But that the drug is addictive does not tell us what supports the case for refraining from acting in that way. Why? As I understand it, to support the case for either acting or refraining from acting in some way that consideration must make some positive contribution towards that case ('favouring'). Here we have only a negative contribution ('disfavouring') which attaches to the action, and no positive contribution which attaches to the omission.

We can add more detail to the case: suppose that the fact that it shows your steely resolve is a reason for refraining from taking the drug. Here we have a positive reason which supports the case for the omission. However the fact that it shows your steely resolve does nothing to undermine the case for taking the drug. That consideration only provides positive support for refraining from doing so, but tells us nothing about the negative undermining of the case against.¹⁵⁷

Second, I use the terms 'strong reason', 'greater reason', and 'most reason' to pick out different uses of overall verdicts about reasons as follows:

Strong reason

A has strong reason to Φ if, and only if, the reasons that favour A's Φ -ing outweigh the reasons against A's Φ -ing.

A has strong reason not to Φ if, and only if, the reasons that are against A's Φ -ing outweigh the reasons for A's Φ -ing.

Greater reason

A has greater reason to Φ if, and only if, there is more net reason for A to Φ than to Ψ (where Ψ is a set of relevant alternatives to Φ -ing).

¹⁵⁷ See also Snedegar (2017) and Bedke (2009) for arguments that deny the inference from 'A has reason against Φ -ing' to 'A has reason for not- Φ -ing'.

Most reason A has most reason to Φ if, and only if, A has strong reason and greater reason to Φ .

What I aim to do here is to distinguish between two components of practical reasoning: in particular, between what there is to be said for and against acting in some way, and the balancing of this net verdict against other net verdicts. In formulating my solution to the difficulties above I exploit the above two sets of claims. I will argue that we need to recognise the distinction between reasons as counting against and reasons as counting for, and the balancing of these reasons against reasons for/against an alternative, in order to get to the right view.

4.7.1. Requirements to Φ and Requirements not to Φ

In what follows I outline my analyses of some important deontic verdictive properties. In particular, I distinguish between two senses of the verdictive property of an action picked out by the term ‘ought’. The first sense I call a *requirement*, and the second sense I call a *recommendation*. In an intuitive sense, a requirement functions to rule out courses of action (as Gert points out). On the other hand, where it is recommended that you do something, this roughly expresses the idea that it would be best to do it or it that would be a good course of action.

We can be required either to perform some action, or to refrain from performing some action. Where we are required to perform an action, I claim that this expresses the idea that that action is to be done, and that all other actions are ruled out. Where we are required to refrain from performing an action, I claim that this expresses the idea that that action is ruled out and any other permissible action is to be done. Thus, whether requirements take actions or omissions as their object implies a different idea, and I will discuss these ideas separately.

I here lay out the first pass at the analyses, and will comment on and refine these claims in what follows.

Requirement to Φ What it is for A to be required to Φ is for there to be some set of considerations $\{R\}$ that give

A strong reason to Φ , and which also give A strong reason against doing whatever conflicts with Φ -ing.

Requirement not to Φ What it is for A to be required not to Φ is for there to be some set of considerations $\{R\}$ which give A strong reason against Φ -ing and $\{R\}$ has a significant degree of strength.

Recommended to Φ What it is for it to be the case that A is recommended to Φ is for there to be some set of considerations $\{R\}$ that give A greatest reason to Φ .

Recommended not to Φ What it is for it to be the case that A is recommended not to Φ is for A to have greatest reason to Ψ .

We can note and set aside that there are other definitions of putative verdictive claims that can be built out of the distinctions I have drawn above. I focus on the kinds of verdictive claims that I hope will allow us to solve the difficulties in discrimination and the problems of supererogation outlined in §4.1. I will discuss each of these analyses in turn, beginning with requirements to Φ .

4.7.2. Requirements to Φ

I claim above that requirements express the idea either of ruling out any alternative course of action, or of ruling out a single action. Thus, I take a distinctive mark of requirement, in contradistinction with recommendation, to be the idea that a requirement is typically concerned with, or takes as its object, *both an act and its*

alternatives. Recommendation, on the other hand, is typically concerned only with, or takes as its object, a single act, and bears on alternatives only *indirectly*.¹⁵⁸

So, one controversial claim that I need to defend is that the consideration that favours Φ -ing in a requirement itself *also* counts against those actions that conflict with Φ -ing, such as that A promised B that you would return his fiver this afternoon counts against going to the cinema if, and only if, going to the cinema is a way of reneging on that promise. On the other hand, that the restaurant's food is tasty does not count against going to the cinema, *even where going to the cinema is a way of not going to the restaurant*. In those cases, the explanation why you should not go to the cinema is simply an expression of the *balance* of reasons. On the other hand, promise-giving reasons count against doing otherwise not just because there is something better to do, but because they bear directly on that course of action. So the tastiness of the food at the restaurant doesn't give you reason not to go to the cinema, but the fact that you promised to return the fiver this afternoon gives you reason not to go to the cinema (say, this afternoon). These sorts of considerations do more than just 'make eligible' some option: they directly count against other options.

Does the fact that you promised to return the fiver give you reason not to go to the cinema? Only if going to the cinema conflicts with your being able, within a reasonable amount of time, to return the fiver. But isn't that also true of any consideration, for example the fact that going to the cinema is enjoyable is a reason to go, and not to do whatever would conflict with going?

Here is one way of spelling out the difference. We could say that the fact that the cinema is *more* enjoyable than going to the restaurant is a reason not to go to the restaurant. But, again, that is just to express the way that the balance of reasons lies. There is more reason to go to the cinema than to the restaurant, so you should go to the cinema, not the restaurant. We are expressing the fact that there is reason of some strength to go to the cinema, and reason of weaker strength to go to the restaurant, and so conclude, on balance, that we should go to the cinema rather than the

¹⁵⁸ As will be obvious this already distinguishes requirements from normative reasons, given my adherence to the asymmetry of the for/against reason-relations. Normative reasons bring to bear a single valence on a single act, while requirements bring to bear different valences to multiple acts.

restaurant. We should not say that the fact that going to the cinema (now) is enjoyable is a reason not to go to the restaurant (now) because that reason does not bear directly on that alternative.

This is not the way we express a requirement. It is not that there are reasons favouring Φ -ing and reasons favouring Ψ -ing and what we should do is to (simply) act on the balance of those reasons. We should say that the fact that you promised to return the fiver (now) *is* a (strong) reason not to do whatever conflicts with giving back the fiver. Here the *same* consideration both favours Φ -ing *and* disfavors Ψ -ing. Requirements, in this sense, don't simply express the balance of reasons in the ordinary way. They express the claim that you should do something not primarily because the balance of reasons falls on one side, but because the reasons that fall on one side themselves also count against acting otherwise.

We can make the point clearer by applying the distinction between reasons for and reasons against to these cases. As I claimed above we should accept that practical reasons display a kind of asymmetry. Where A has reason against acting in some way, all else being equal he needs some justification for acting in that way or is appropriately subject to negative reactive attitudes (from others) for acting in that way, given his recognition of this reason. On the other hand, where A has reason for acting in some way, and he decides not to act in that way, A may not be appropriately subject to such attitudes for failing to do so (except for self-appropriate attitudes). Note that this is not a point specifically about morality but about practical reason generally. For instance, that you would find the food disgusting implies that you are less than fully rational for eating it without some explanation, but that you would find the food tasty, and decide not to eat it does not entail that you are less than fully rational. Failing to eat the food is simply optional for you. In this sense, reasons for tend to make options eligible, while reasons against tend to undermine such options.

Now let us apply this to our pair of contrasting cases: deciding whether to go to the restaurant or the cinema, and deciding whether to keep your promise.

In the first case, your decision to act on the basis of the weaker reasons makes you criticisable, if at all, only because you fail to do what you had most reason to do. Here you are subject to criticism of a different kind to that which is involved in the

promise-keeping case, of failing to act on the strongest reason that you have. You are not subject to criticism in the above sense, where considerations count (decisively) against your acting in that way, thus demanding an explanation for your decision to act in that way. Those considerations, in the promise-keeping case, serve to effectively rule out alternatives not simply because your promise generates overwhelmingly weightier reason to keep it, but because it also puts an unanswerable onus on the agent who decides to do otherwise to justify himself.

There is, however, a problem for this analysis. Suppose that there is nothing to be said for doing X, but it is not as bad as Y or Z. Given the distinction between reasons for and reasons against, A could not be required to X. But surely you can be required to perform your least bad option? Suppose that the only difference between your three options is the degree of harm that they will cause, with X causing the least harm, Y more, and Z still more. There is nothing to be said for doing any of these things. However, A has no further choices –these are the only options open to him.

However, we can get around this objection by slightly modifying the principle by dropping the condition that where you are required to do something, there must be something to be said for doing that thing:

Requirement-to-Φ*	What it is for A to be required to Φ is for there to be some set of considerations {R} that give A strong reason not to do whatever conflicts with Φ -ing
--	---

Thus although requirements are typically concerned with both an act and its alternatives, there can be cases in which you must perform your least bad option. This modification reflects this.

Consider now the following objection. Suppose that the fact that it is sunny outside gives you strong reason to go to the plaza, strong reason against going to the cinema, and strong reason against not doing either but instead staying at home. And suppose that these are the only options available to you now. According to my definition of requirement, we have a consideration that possibly gives you strong reason to Φ which also gives you strong reason not to act in any of the ways that

conflicts with Φ -ing, where the set of options is reduced to these three. Therefore you are required to go to the restaurant, which seems false.

The first thing to ask about this case is whether the fact that it is sunny does in fact give you strong reason against going to the plaza and against staying home, in the sense that the consideration against *outweighs* any reasons for performing each of these options. According to my analysis, to generate a requirement it is not enough that this consideration simply gives you *some* reason against these options. So in order to be a counterexample to my analysis of requirement the fact that it is sunny must *decisively* count against the other two options by constituting a strong reason.

In addition, the case appears to me to be mis-described. That it is sunny outside is a reason to go outside, not a reason against staying in, at least according to my schema of reasons. To count against acting in some way, that consideration must make appropriate negative reactive attitudes for acting in that way, i.e. it must highlight a negative quality of the action. But that it is sunny outside just doesn't bear on staying inside in the sense that missing out on something there is reason for is not itself a reason against.

We need to specify more precisely what the relevant ground is in this case. Suppose that it is that your body requires a certain amount of vitamin D per day and the only way to meet your quota today is to go outside. Let us understand that failing to meet your quota damages your health. So, one way of looking at things is to say that you have reason to go outside because doing so is a way of avoiding damaging your health. You have reason to go outside because of the bad consequences of not going outside. So according to me this is really a reason against not going outside.

Suppose, on the other hand, that the ground is the pleasure you get from being in the sun. Then according to me that pleasure gives you a reason to go to the plaza, but does not count as a reason against staying in or going to the cinema. So the claim 'It is sunny outside' really conceals a positive and negative aspect which allows us to truthfully say that it gives you a reason to go outside and a reason against staying in. Once we separate the negative and positive aspects we see that it is not really true that the very same consideration counts both for acting in some way and against acting on conflicting alternatives.

4.7.3. Requirements not to Φ

As I say above, requirements to do something differ from requirements not to do something, in the sense that the former rules out all alternative courses of action, while the latter rules out acting in one way. Requirements not to do something in effect say simply ‘Do not do this’. According to my analysis such requirements are expressed through the conjunction of two claims: first, that A has strong reason not to Φ , i.e. that the reasons against A’s Φ -ing win out against the reasons for A’s Φ -ing. Second, those considerations are of a significant strength. That is, not only is A criticisable for Φ -ing but the criticism is significant enough to warrant it being the case that A is required not to Φ .¹⁵⁹

The definition of a requirement not to do something cannot be right as it stands however. Consider again a scenario in which you must perform your least bad option. Assuming that you have strong reason not to perform each of these options, and that there are no reasons favouring any of the options, according to my definition of requirement not to Φ , as it stands, you are required not to perform *any* of these options. Furthermore, according to my definition of requirement to Φ^* you are required to perform your least bad option. That is an outright contradiction. Again, however, we can easily amend the definition to take account of these cases as follows:

¹⁵⁹ Note that this view is not inconsistent with absolute prohibitions. To say that X is absolutely prohibited is to say that there sufficient reason for A not to X and there is no consideration that *could* justify A’s X-ing. It is helpful here to appeal to the distinction between requirements considered synchronically versus requirements considered diachronically. Even though it may be true at t_1 that you are required not to Φ , it may be false at t_2 that you are required not to Φ because new information comes to light or a new reason is created. But in saying that a consideration could not justify A’s Φ -ing (absolute prohibition) we mean that at no time will it be false that A is required not to Φ . We should also say that in a situation in which the only options open to A are to Φ or to Ψ , where there is sufficient reason against Ψ -ing, unless Ψ is also absolutely prohibited, there is (by hypothesis) stronger reason not to Φ .

Requirement not to Φ * What it is for A to be required not to Φ is for there to be some set of considerations $\{R\}$ which give A strong reason not to Φ and $\{R\}$ has a significant degree of strength, *unless* Φ -ing is the only way that A can avoid Ψ -ing, and there is strong reason not to Ψ of greater strength.

With this modification in mind, we should note some further differences between requirements not to Φ and requirements to Φ .

We can ask two quite different questions of ourselves and others: what am I required to do now, and what am I required not to do now? In the first case an answer places you under a demand to perform a single action (or series of actions, but I want to simplify things here) and only a single action.¹⁶⁰ In the second case, an answer may place you under a demand not to perform many actions, or many demands not to perform different actions. So where you are required to do something, there is one unique action that it is required of you that you do. You cannot, at t_1 , be required to do more than one thing. On the other hand you can, at t_1 , be required not to do more than one thing.

Another difference between acts and omissions with respect to their normative support is that when considering whether to perform some act, comparison classes are often important. Often you are deciding whether to X *rather than* Y or Z. When we deliberate over omissions, on the other hand, comparison classes are often of less significance. Usually when deliberation involves omissions it is by implicit or explicit reference to an action that is incompatible with that omission, as when we

¹⁶⁰ I assume here that an agent can only perform a single action at a time, where a single action can be ‘complex’ in the sense that it can include separable parts. For example, your drinking the water and raising your finger at the same time count as a single action though you could drink the water without raising your finger and vice versa. We should also be careful to distinguish actions and omissions here.

ask ‘should I, or should I not go to the cinema?’ or ‘should I not go to the cinema (and rather stay in and watch TV) or should I go out to meet my friend?’.

In general then, when considering *acts*, practical reasoning may include both considerations for and against acting in that way, *and* that value balanced against the considered value of acting in other ways. Being required to Φ , or Φ -ing being recommended, involve weighing competing reasons for competing actions (this is why there is conceptual pressure to accept that there is an ‘ought, all-things-considered’). On the other hand, having strong reason *not* to act in some way involves (a) only competing reasons for and against that omission and (b) there can be multiple non-competing ‘most reason’ verdicts supporting different omissions.

Are there cases in which negative verdicts involve comparison classes? Suppose I am deciding between the green jacket and the red jacket. My friend recommends that I don’t wear the red jacket. However cases like this are just the converse of asking what I should do: if I am asking which I should wear I am trying primarily to decide what I should *do*. Although reasons against may be the only considerations relevant here (my choice is a matter of damage-limitation – picking the least worst rather than the best) where there are less strong reasons against wearing the green as opposed to the red, we are doing so in the service of answering the question what I should do rather than what I should not do.

I claim that to be required not to act in some way is for there to be strong reason not to act in that way of significant strength. To have strong reason not to act in some way is for there to be more reason against Φ than for Φ and to have reason against Φ is to explain why Φ -ing would be criticisable. Criticism, however, can take different forms, and the way in which criticism differs can in part determine what counts as ‘significant’ for the purposes of requirement. For example, consider the difference between prudential and moral reasons. I have prudential reason to avoid acting in some way that specifies some harm that will befall me if I do it. I have moral reason to avoid acting in some way that specifies, for example, some harm that will befall you if I do it.

In the first case we might construe this in broadly teleological terms: I have reason not to cause myself to be in pain because of the badness of pain and I am subject to criticism on the basis of my failure simply to care about the badness of

pain. In the second case it is less plausible that this can be entirely captured in teleological terms. It is not my failure to recognise the badness of pain that primarily form the basis of criticism against me. It is, in addition, the fact that it is *your body* that gives me reason to care about avoiding causing you harm. So I am also criticisable on two fronts. One front is associated with the teleological reason against promoting pain, and another is associated with a respect-based reason against harming agents.¹⁶¹

We can also note the difference between cases in which you have strong reason against for which there are no reasons for, and cases in which you have strong reason against for which there are reasons for. In the latter case one may be susceptible to misplacing or misconstruing the respective weights involved in the reasons bearing on that action. You may be more likely to mistakenly suppose that you can discount criticism in acting in that way because there is something to be said for doing so. In cases of the first sort, however, such mistakes seem more serious. That there is nothing whatsoever to be said for acting in that way makes acting in that way less understandably discountable by the agent.

By definition, if you have strong reason not to act in some way there is no consideration that you can appeal to in order to justify your doing it anyway. But our intuitions arising from the example of having strong reason not to eat the food because it is disgusting may not fully comport with this idea. We might be inclined there to say ‘Yes, he is criticisable for eating the food, but it is *up to him* whether he eats it’. Is that a way of appealing to some consideration to justify what he is doing?

Notice first that claiming that acting in this way is ‘up to him’ is consistent with the claim that he is nevertheless criticisable for acting in that way. It seems very plausible to say that he is in fact so criticisable, but that criticism is commensurate with the strength of the reasons which count against his so acting. We can admit that

¹⁶¹ Greenspan also emphasises the importance of distinguishing between cases in which criticism is primarily directed from within, such as where you prudentially ought not act in some way, from cases in which criticism is primarily directed from without, especially where your acting in some way directly and adversely affects that other agent. See Greenspan (2007).

the strength of such reasons against is relatively weak; so too is our criticism relatively mild.

We might appeal to the distinction between *justification* and *excuse* to help here. Although such an agent is necessarily not justified in acting in the way he does, and thus is subject to mild criticism, we might think that he may nevertheless be excused in acting in that way due to the relative insignificance both of the infraction, and of the fact that his action affects only himself. Plausibly, agents have a certain degree of leeway over acting in ways that harm (in the broadest sense) themselves, which, as mentioned above, is absent where the act affects others.

Obviously, when considering requirements not to Φ , our criticisms will be strong. But is the difference between a moral requirement not to kill and a prudential requirement not to cause myself some small bodily discomfort merely reflected in terms of the strength of criticism?

To answer this question let us examine whether we should appeal to the idea that reasons also play a dual-role where they figure in requirements not to Φ . We could claim that moral reasons which figure in negative requirements also play a role of shaping an agent's dispositions in the following way: that it would harm an autonomous agent gives you reason not just not to Φ but in addition gives you reason to aim to act in morally laudable ways. However, I don't think we should be committed to this view. We should leave it open whether we ought to construe morality in minimal terms, for instance as a number of constraints on action. This minimal understanding of morality simply has nothing to say about what, positively, you should do, but merely rules out certain ways of acting.

Now consider reasons for demands. Suppose that you try to harm me. I thus have reason to demand that you refrain from doing so. This is a reason which may even be waivable by me. If I decide to relinquish my demand I may have no complaint against you for harming me. In other cases it may be true (a) that I have reason to demand that you refrain from Φ -ing and (b) that others have reason to demand that you refrain from Φ -ing. So the fact that it would harm me is both a

strong reason not to do it and a strong reason for others to (legitimately) demand that you not do it.¹⁶²

So we could say this: reasons for are generally waivable, reasons against are not. But, further, an agent can only waive his *own* reasons, he cannot waive the reasons that *others* have. For example, that some act would cause you harm gives me strong reason not to do it, and reason for you to demand that I refrain from doing it, I cannot waive either reason: because (i) it is a negative reason (and thus makes criticism appropriate) and (ii) because an agent cannot waive another agent's reason. That is a further way of distinguishing between morally forbidden acts and (generally) rationally 'forbidden' acts.

So if we appeal to intuition about what would count as appropriate criticism it seems plausible that you are not appropriately subject to criticism for failing to do what you had less-than-most reason to do (as long as you do what you had most reason to do). However, I don't think that this is the full story about how to capture negative requirements in terms of reasons, and make some further remarks below (in §4.9).

4.8. Recommended to Φ and Recommended not to Φ

What the claims that A is recommended to Φ and that A is recommended not to Φ are intended to express are that some course of action is either the best option from some set of relevant alternatives (recommended to Φ) or that, of some set of actions, you should do otherwise than Φ -ing (recommended not to Φ). Note that one way in which recommended not to Φ differs from having a requirement not to perform some action is that if you are recommended not to Φ there may be some reasons in favour of Φ -ing but there is still more to be said for every relevant alternative. On the other hand, where A is recommended to Φ he is not required to Φ because those considerations that give him reason to Φ do not also rule out those considerations that conflict with A's Φ -ing.

¹⁶² However, see Wallace in Bakhurst, Hooker, & Little (2013: 150-8) on how demands are too narrow to capture moral obligation.

It is important to note that my definition of the claim that A is recommended not to do something does not mention reasons against. That is because, as I have tried to make clear, the presence of reasons against usually signals the existence of a requirement, and implies criticism (however weak and potentially excusable). So I have avoided defining recommended not to Φ in terms of reasons against in order to avoid collapsing this claim into requirements not to Φ .

Now recall our first desiderata of a good deontic buck-passing theory. It should preserve the correct entailments between verdicts of requirement and verdicts of recommendation. If A is required to Φ , this entails that it is recommended that he Φ , but not vice versa. This was a problem, especially for Gert's view, whose strict distinction between justifying and requiring ensured that he could not guarantee those entailments. Can my theory preserve the entailments?

First, it should be clear that requirement to Φ entails recommended to Φ since if A has strong reason not to do otherwise than Φ -ing, then A has more net reason to Φ than to Ψ since his reasons against defeat any reasons for any conflicting action. Conversely, recommended to Φ clearly does not entail requirement to Φ . So the desired entailments are preserved here.

What about the entailment from requirement not to Φ to recommended not to Φ ? I defined recommended not to Φ as having reason to do otherwise. If it can be the case that A is required not to do something and it can be the case that A does not have reason to do otherwise, then the desired entailments between recommended not to Φ and required not to Φ are not preserved. But, as we have seen above given the modified version of requirement not to Φ it could not be the case both that you are required not to do something and that there is not greater reason to do otherwise. So the entailments seem to be preserved here too.

4.9. Further Remarks on Requirements and Moral Obligations

I have presented a content-neutral way of expressing the distinction between requirements and recommendations without appealing to a non-maximising theory of practical reason, but I do not think that requirements are equivalent to moral obligations. I think that you can be required (in my sense and in an intuitive sense) to

do what you are morally obligated to do, and that it can fail to be the case that you are required to do what you are morally obligated to do. In this section I outline what more we need to say to capture moral obligation within my framework. Scanlon helpfully suggests two relevant cases:

First, positively, since others could reasonably refuse to license us to decide what to do in a way that gave concrete factors such as [‘he needs me help’ or ‘doing that would put them in danger’] no weight, the aim of justifiability to others gives us reason to recognize these considerations as ones that are generally relevant, and are in some circumstances compelling reasons to act.

Second, negatively, ‘being moral’ involves seeing certain considerations as providing no justification for action in some situations even though they involve elements which, in other contexts, would be relevant. The fact that it would be slightly inconvenient for me to keep a promise should be excluded as a reason for not doing so. Even if I am in great need of money to complete my life project, this gives me no reason to hasten the death of my rich uncle or even to hope that, flourishing and happy at seventy-three, he will soon be felled by a heart attack. [...]

It does not seem true even of most of us, let alone of a person who was fully moved by moral reasons, that the moral motivation not to act wrongly has to hold in check, by outweighing, all these opposing considerations. [...] Being moral involves seeing reason to exclude some considerations from the realm of relevant reasons (under certain conditions) just as it involves reasons for including others. The contractualist account can explain this fact, since these considerations are ones that others could reasonably refuse to license us to count as reasons.¹⁶³

I add to the promising case and the inheritance case a further case: that you must gain permission to act in an optional way, e.g. where I must ask your permission to eat your cheese.

¹⁶³ Scanlon (1998: 156-7)

I distinguish the promising and inheritance cases because they have a different structure (which is not reflected by Scanlon). Though it may be right that the fact that you stand to inherit money from your uncle is disabled as a reason to hasten his death, when I promise to meet you later, and then regret the fact that I made the promise when I discover that in order to meet you I must get on an overly-packed train, that I feel regret is a reasonable response to my situation. But if it can be appropriate to feel regret in promising cases, then the reason on the other side cannot be simply disabled. It still retains some normative force.

The common factor in all three cases is that the incorrect way of describing these normative situations is as one in which one course of action wins out by outweighing a competing consideration. In the cheese case, I do not seek permission because the tastiness of the cheese is not a sufficient reason alone to eat it. In the promising case, even though I may regret making the promise, I do not weigh the irritation likely caused by the packed train against my making the promise and decide to keep the promise because it is a strong reason. Finally, in the inheritance case, I do not weigh up my duty not to murder innocents against the value of gaining the inheritance and decide that my duty is stronger. Because outweighing does not reflect these normative situations, these cases present a *prima facie* problem for any deontic buck-passing analysis.

Here is how I handle these cases. The cheese case is mis-described. The tastiness of the cheese gives you reason to eat it but the fact that the cheese is tasty (and that you want it) also gives you reason to seek permission from the owner to eat it. So this reason is not in competition with the reason you have to seek permission to eat it; it just *is* this reason.

In the inheritance case I agree with Scanlon's diagnosis. As I mentioned above, normative reasons can interact not only by outweighing, but also by undercutting. Where a normative reason undercuts it removes entirely the justification that may have existed on the other side.¹⁶⁴ In this case I believe that the e.g. the fact that your uncle is an autonomous agent (which implies constraints on what can be done to him) *undercuts* what would otherwise be a perfectly good

¹⁶⁴ Undercutting can be either 'content' undermining or 'context' undermining. I do not go into the details here. See Cullity (2013).

reason: that you would gain the inheritance. That you would gain the inheritance is not outweighed by the fact that you would be causing the death of your uncle, but is undercut by that fact. So the fact that you would gain the inheritance provides, I claim, no reason at all for murdering your uncle. We can bolster this claim by appealing to a plausible principle:

Regret Principle If there is reason for A to Φ then (all else equal) it is appropriate either (a) for A to regret not Φ -ing, or (b) for A to regret not being able to Φ .

Suppose that the fact that the chocolate is tasty is a weak reason to eat it, but that you have more reason to instead to eat the healthy cabbage. In choosing to eat the cabbage it seems appropriate for you to have some regretful feelings towards being (in this case, rationally) unable to eat the tasty chocolate. By the same token, it is plausible that it is *not* appropriate for you to regret not causing the death of your uncle. So according to the *Regret Principle* it is not plausible that you have any (outweighed) reason to do so.

In the promising case, I don't think that the fact that you made the promise undercuts the fact that the train would be uncomfortable to ride on is reason not to get on it (and thus rationalises a reason to regret making the promise). That reason persists, but is not simply outweighed. But we need not revert to Razian exclusionary reasons to understand the status of this reason. We can distinguish between second-order reasons not to act on the basis of some first-order reason (an exclusionary reason) and first-order reasons not to include in deliberation some other first-order reasons. These reasons are first-order, not second-order, because they take as their object deliberation rather than acting-on-the-basis-of-a-reason.¹⁶⁵ So we do not fall foul of the general objections to second-order reasons.

¹⁶⁵ It could be true that if you have a second-order reason not to act for some first-order reason then you also have reason not take that first-order reason into account in deliberation (give it weight in your deliberation). But that is not implied by the very idea of an (exclusionary) second-order reason. All that an (exclusionary) second-

In fact first-order reasons not to deliberate with other first-order reasons are uncontroversial, and we respond to them all the time. Most significant are reasons not to include in one's deliberation other reasons that bear on a deliberative question, but are just too weak to be worth deliberating about. Time-based constraints provide us with perfectly legitimate reasons not to take into consideration many reasons that nevertheless bear on a practical conclusion.

To explain the promising case I appeal to such first-order reasons to exclude *from deliberation*. When you make a promise to meet your friend, the fact that you made that promise gives you reason not to include in your deliberations about what to do putatively competing reasons which recommend acting in a way that conflicts with keeping your promise (and of course *these* reasons may be outweighed, where for example there is a significant personal or other-regarding cost to keeping your promise). I take this to be a general truth about the nature of promising.¹⁶⁶

So it can make sense to regret making your promise where reasons that support conflicting actions persist, but are excluded from deliberation in my sense, rather than excluded in Raz's sense. Such reasons persist, but it is not appropriate to weigh them against your promise-keeping reasons.

Correctly describing cases, paying attention to the undercutting/outweighing distinction, and recognising deliberative constraints such as implied by promises help us to capture cases of moral obligations on the deontic buck-passing view. As I have said, moral obligations are not synonymous with what I call requirements, and as we have seen, do not in any case have a uniform normative structure. Nevertheless, I have argued that that normative structure can be understood in a way that is consistent with the deontic buck-passing view.

4.10. Resolving the Supererogation Problem

The second difficulty with SA is that it appeared to conceptually rule out supererogation. Suppose that you have strongest reason to sacrifice your life to save

order reason requires of you is that you do not act on the basis of some consideration, not that you do not take that consideration into account.

¹⁶⁶ For more on the nature of promising see Heuer (2012).

5 people in the burning building. It seems plausible to suppose that such a reason could outweigh opposing reasons you have, i.e. it seems possible that you can have most reason to perform a supererogatory act, in this sense. According to SA you therefore *ought* to sacrifice yourself to save the 5 (where I am again using ‘ought’ in an intuitive sense). But it is a conceptual truth that, even though it would be morally laudable to do so, you cannot be required to perform a supererogatory action. Thus SA rules out the conceptual possibility of supererogation.

I have claimed that when we use the word ‘ought’ we can be talking (amongst other things) either about what I call ‘requirement’ or about what I call ‘recommendation’. Further, according to my analysis of recommended to Φ , for A’s Φ -ing to be recommended is just for A to have most reason to Φ . That removes the problematic inference on the picture drawn above. If you have most reason to perform a supererogatory action, it is false that you are *required* to perform that action, though it may be recommended in the sense that it is your best option.

It may be objected that my analysis only defers the objection. According to my analysis if it true both that R gives A strong reason to Φ and R gives A strong reason not to act in any way that conflicts with Φ -ing, then A is required to Φ . If it is possible that Φ -ing counts as a supererogatory act, my analysis similarly rules out supererogation.

Suppose that A can save the lives of 5 people only by throwing himself in front of the runaway trolley. If the fact that it would save 5 lives counts decisively in favour of throwing himself in the path of the trolley, and that fact also decisively counts against any alternative, then according to my analysis A is required to throw himself in front of the trolley. I think, however, that the crucial difference between a supererogatory act and a required act lies in whether or not that fact counts decisively against A’s alternatives. In this case, although the fact that it would save 5 lives *does* plausibly count against A’s refraining from throwing himself in front of the trolley that fact does not plausibly count *decisively* against that alternative. The same intuition that motivates our belief that this is a supererogatory act should motivate the denial of the claim that the fact that A will save 5 lives counts decisively against refraining from jumping.

However, we can allow that A has *most* reason to save the 5 in that he has *strongest* reason to save the 5 (A's reasons for saving the 5 outweigh the reasons against saving the 5) and *greater* reason to save the 5 (A has more reason to save the 5 than to perform an alternative). Still, according to my analysis, A is not required to save the 5. A is only required to save the 5 if the consideration that gives him reason to save the 5 *also* counts decisively against alternative courses of action. That is, it can still be true, on my analysis that A has most reason to perform a supererogatory act, but that A is not required to act in that way.

Further, I find it difficult to imagine what content there could be to claiming that such an act would be required. As I have already argued, to act 'in the teeth' of reasons against is different from acting on the basis of reasons which favour acting in some way. Decisive reasons against imply a strong form of criticism for flouting what they disfavour. But we are to suppose that in this case it is not reasonable to criticise you for failing to sacrifice yourself. Given that these two claims seem inconsistent together, we have reason to reject this possibility: supererogatory acts which represent your best option can come recommended in the sense above, but they cannot be required in my sense.

5. Deontic Buck-Passing: The Negative Thesis

If the deontic properties of acts such as the property of being required or recommended should be analysed in terms of normative reasons in the way I have argued, that lends support to the negative claim that those deontic properties themselves are not reason-providing. Suppose that I am required not to do something. According to my analysis, that is to say that I have strong reason not to do it. But if what it is to be required not to do it is that I have strong reason, then the fact that I am required not to do it isn't plausibly a *further* reason not to do it. Compare: having sufficient reason not to do it is having more reason not to do it than to do it, and in addition, that I have sufficient reason not to do it is an extra reason not to do it. That looks like a clear example of double-counting reasons.

There are, however, various objections to the negative buck-passing thesis (NBT). Some of these object to the arguments for the claim, while others object to the claim itself. In this final chapter I examine both sorts of objections.

5.1. Some Objections to the *Argument* for NBT

5.1.1. Heuer on Wrongness

Here is one version of an argument for NBT that Heuer gleans from Dancy.¹⁶⁷ I have modified the argument slightly to avoid some ambiguities and unnecessary difficulties. Define V as 'A is required not to Φ ' and let $(r_1, r_2, \dots r_n)$ represent the normative reasons that make it the case that A is required not to Φ :

1. V is a verdict based on all the relevant reasons $(r_1, r_2, \dots r_n)$.
2. If V were itself a reason, it would have to be a further reason (in addition to $(r_1, r_2, \dots r_n)$).
3. If V were a further reason, this reason would affect the balance of reasons.

¹⁶⁷ Heuer (2010). See Dancy (2005) ch.2 for the inspiration for this argument.

4. If the fact that makes V true would affect the balance of reasons, (1) would be false, i.e. V would not be a verdict, based on all the relevant reasons ($r_1, r_2, \dots r_n$).
- C. Therefore, V is not a reason.¹⁶⁸

What is it for a consideration to affect the balance of reasons? As I have claimed, to affect the balance of reasons for or against acting in some way, a consideration must affect the strength of the case for or against Φ -ing. Although by invoking the idea of strength I am appealing to a metaphor, the basic thought should not be difficult to grasp. To say that a consideration does not affect the balance of reasons is, then, to say that that consideration does not add its weight, for or against, to the case for or against Φ -ing.¹⁶⁹

As I understand it, a normative reason that bears on A's Φ -ing could fail to add its weight to the case for or against Φ -ing either (a) because it doesn't have weight or (b) because, although it has weight it is somehow not transferred to the case, even though the consideration itself forms part of the case. We might understand (b) in terms of Scanlon's idea of a reason not to take another, further reason into account in deliberation. R1 is a reason for A to Φ but A also has a further reason R2 not to place any weight on R1.

However, there is a difference between an agent placing a certain amount of weight on a consideration in deliberation and that consideration actually having that weight. Appealing to the idea of not taking a reason into account in deliberation may not help in this respect. There the idea is that even though that consideration has a certain weight, that consideration should be excluded from deliberation. In any case, it is clear that Heuer is interested in considerations failing to add weight in the sense of (a).¹⁷⁰ So we can focus in what follows on that version of the claim.

According to Heuer the above argument fails because it relies on a mistaken assumption, which is a generalisation of premise (3):

¹⁶⁸ See Heuer (2010:141).

¹⁶⁹ See e.g. Berker (2007: 113-18) and Lord & Maguire in Lord & Maguire ed. (2016) for more on this conception of the balance of reasons.

¹⁷⁰ See Heuer (2010: 143-4).

(Δ) A consideration which does not affect the balance of reasons is not a reason.

Heuer mentions two examples of reasons that do not affect the *balance* of reasons, and thus provide counterexamples to Δ . First, consider the following example: in addition to the undefeated reason you have to pick up your child from school, the fact that you are looking forward to it is an additional reason. But this reason does not alter the balance of reasons since the verdict remains the same regardless of whether this additional reason holds or not. Second, reasons can be undercut in that some other reason completely removes the normative force of that reason. But, Heuer claims, such undercut reasons are still reasons.¹⁷¹

In the first case Heuer apparently understands the idea of ‘affecting the balance of reasons’ in terms of the verdict that those reasons produce. But I doubt that Dancy and others understand the balance of reasons in this narrower way. Rather, a consideration can affect the balance of reasons by strengthening or weakening the case. Such strengthenings and weakenings may not change the verdict of the overall case. Once we allow for this more natural and wider understanding of ‘the balance of reasons’ we can see that the fact that you are looking forward to picking up your child is a reason, and affects the balance of reasons by strengthening the case for going, even though its absence would not tip the scales the other way.

The second case that Heuer appeals to is not compelling either. To say of a reason that it is *undercut* (as opposed to outweighed) is precisely to say that it is no longer a reason for a response. Compare undercutting defeaters in epistemology. Suppose that the reason provided by Smith’s testimony is defeated by the fact that, as you later found out, Smith was not, in fact, a witness at the scene. That undercutting defeater makes it the case that Smith’s testimony no longer gives you that reason for believing.

Heuer has a further argument that Δ is false that appeals to what she calls ‘specifiable properties’.¹⁷²

¹⁷¹ *Loc. cit.*

¹⁷² Heuer (2010: 145-6).

Some G is a specifiable property if, and only if

1. Necessarily all Fs are Gs, but
2. possibly some Gs are not Fs, and
3. if a G is not an F then necessarily there is some E such that necessarily all Es are Gs, but possibly some Gs are not Es, and E is different from G & $\neg F$.¹⁷³

Thus, goodness-in-a-respect is a specifiable property since it can be further specified in terms of, say, pleasure. Specifiable properties do not provide further reasons, in addition to the reasons provided by their specifications. That is, specifiable properties necessarily never affect the balance of reasons (only their specifications do that). But, Heuer claims, the properties of being entertaining, pleasant, or tasty are *also* themselves specifiable properties. Now by the definition of specifiable properties, these considerations too do not affect the balance of reasons, and so if Δ is true such considerations are not themselves reasons. But this is at odds with our intuitive understanding of what counts as a reason (insofar as we think that tastiness, pleasantness, and so on provide reasons). So Δ is false. Specifiable properties do not affect the balance of reasons; yet intuitively we think that they are reasons.

In reply to this argument, first note that it is controversial whether the three conditions on what makes a specifiable property holds even for goodness. It is a familiar thought that it is false that all pleasant things are good things, given the case of sadistic pleasure.¹⁷⁴ So it is not clear that the view holds even for its most plausible target. Second, I doubt whether the view can plausibly be re-applied to the ‘good-making’ properties in the way Heuer thinks. What F satisfies condition 1 for pleasure or tastiness for example? Is there some F such that all Fs are tasty? Heuer does not provide us with an answer to this question, and it seems difficult to see what answer is forthcoming.

Now suppose that we can further specify these kinds of properties. That leaves open the question of priority: suppose that all Richard Prior concerts are

¹⁷³ Heuer (2010: 145).

¹⁷⁴ Again, I am assuming the attitude account of pleasure here.

entertaining (and *necessarily* entertaining?). We can agree to that and still claim that there is reason to see Prior *because* he is entertaining, not that we have reason to watch something entertaining because it is Richard Prior. Heuer's definition of specification doesn't settle the question of priority. According to the view adduced by Heuer you have here three different reasons stemming from the concert's goodness, its being entertaining, and its being a Richard Prior concert. This looks in fact like a cumbersome and unintuitive view.

If a consideration necessarily never affects the balance of reasons, then it is a reason with no weight. If it had weight then it would, necessarily, affect the balance of reasons if by that we mean not that it would necessarily change the conclusion arrived at, but that it would somewhat weaken or somewhat strengthen the case for or against. But I doubt whether a reason that has no weight is really a normative reason at all. As claimed above, a reason is a consideration that counts in favour or against some response. But to count in favour of some response it must do so to a certain degree. That the cake tastes nice counts in favour to a certain degree (perhaps, amongst other things, to the degree to which you will derive pleasure from eating it) of eating the cake. To say that a reason has no weight is just to say that it counts in favour of an action to no degree. But I don't think it is coherent to claim that R counts in favour of a response, but not to *any* degree. That just sounds like it is not a reason at all.

In the next section I consider the role that derivative reasons might play in objecting to the argument above.

5.1.2. Schroeder on Buck-Passing Facts

Recall again the argument above:

1. V is a verdict based on all the relevant reasons (r_1, r_2, \dots, r_n).
2. If V were itself a reason, it would have to be a further reason (in addition to (r_1, r_2, \dots, r_n)).
3. If V were a further reason, this reason would affect the balance of reasons.

4. If the fact that makes V true would affect the balance of reasons, (i) would be false, i.e. V would not be a verdict, based on all the relevant reasons (r_1, r_2, \dots, r_n).
- C. Therefore, V is not a reason.¹⁷⁵

Mark Schroeder denies the inference from the positive buck-passing claim to the negative claim. He does so by denying premise 3 of the argument above. Some R can both be a normative reason to act and not affect the balance of reasons.¹⁷⁶

He makes room for the denial of premise 3 by claiming that there are reasons of a special sort, which we can call ‘derivative’ reasons. These derivative reasons (somehow) depend upon, but are nevertheless distinct from the other (derived) reasons which in fact make them the case. Furthermore we should think that these derivative reasons are fully-fledged reasons, due to this argument:

5. If A can act rationally on the basis of motivating reason, M, then M can be a normative reason.
6. A can act rationally in acting on the basis of a buck-passing fact.¹⁷⁷
- C. Therefore, buck-passing facts can provide normative reasons.

So not only does the existence of derivative reasons undermine the argument for the negative claim, Schroeder also gives us an argument for the claim that buck-passing facts *do* give us reasons.

¹⁷⁵ See Heuer (2010:141).

¹⁷⁶ Schroeder (2009).

¹⁷⁷ Where a buck-passing fact is a fact about a property that passes the normative buck, such as (according to Scanlon) the fact X is valuable.

5.1.3. Derivative Reasons

What role does a derivative reason play? Again, we can go one of two ways here: either (a) a derivative reason is distinct from the non-derivative reasons from which it derives, but has no normative force or weight, or (b) a derivative reason is distinct from the non-derivative reason from which it derives, and has a normative force or weight, but does not transfer that weight to the case for/against Φ -ing. In both cases a derivative reason is a distinct reason but does not affect the balance of reasons. Thus derivative reasons falsify premise 3.

Stratton-Lake argues that if one understands normative reasons as considerations that count in favour then the concept of a derivative reason is either incoherent or very unclear.¹⁷⁸ A consideration that counts in favour, counts in favour to a certain degree, which is difficult not to hear as corresponding to the idea of the weight of a reason. So on this understanding of a normative reason, the ‘favouring’ and the ‘weight’ do not look like they can be teased apart. A derivative reason in the sense of (a) looks incoherent on this conception – a consideration which counts in favour to no extent just sounds like it is not a reason.

The (b) interpretation also seems problematic. How are we to understand a consideration that has weight but it does not transfer it to the case for/against Φ -ing? Perhaps the derivative reason ‘inherits’ its weight from the non-derivative reasons and thus does not ‘add’ its weight to the case already made by the non-derivative reasons? But we can understand this in two ways: either the non-derivative reasons lose the weight they have, which is absurd, or the derivative reasons inherit the full powers of the non-derivative reason, which seems wrong.

Perhaps a different way to think about the relation between derivative and non-derivative reasons on the ‘counting in favour’ model is to think about them in mereological terms. Some defend the view there can be two distinct objects that are located in the same space-time region (captured by the slogan that ‘constitution is not identity’). For example, an object may be located in the same space as its constituent parts but those parts may not be identical to the object due to their having different properties (such as different persistence conditions). We might then think of

¹⁷⁸ Stratton-Lake (2017: 87-8).

derivative reasons as like the object and the non-derivative reasons as the parts that constitute that object. That would allow us to conceive of the idea that they are distinct reasons, but that their weights are not distinct. We don't 'add' the weight of the derivative reason to the weight of the non-derivative reason (though the derivative reason has weight) in the same way that we do not add the weight of the object to the weight of its constituent parts.¹⁷⁹

Note that I am not endorsing this model of reasons, but merely pointing out that there may be other ways that we can understand derivative reasons given this alternative mode.

As Stratton-Lake points out, Schroeder makes logical room for the idea of a derivative reason by appealing to his theory of weight. Schroeder distinguishes between a normative reason and its weight, and claims that the weight of a set of reasons is that which it is correct to place weight on. What is it to be correct to place weight? It is just to have reason to place that weight on the set. A regress looms. Schroeder claims that we stop the regress by (i) recursive analysis: the weight of a set S_1 is determined by whether it is weightier than some other set of reasons $S_2 \dots S_n$ (ii) that there will always be some definite point at which there is *some* reason to place weight on S_n and *no* reason to place weight on S_{n+1} .¹⁸⁰

This theory of weight appears to allow us to make sense of derivative reasons. Suppose that we have some set of reasons that includes non-derivative reasons. We can then define a derivative reason as a reason which is such that, when it is added to a set of reasons, it adds no weight to that set.

Stratton-Lake argues that this theory is rendered implausible because it cannot make sense of reasons which are excluded in the sense that we should not place weight on them in deliberation, but which nevertheless have a weight.¹⁸¹ Since Schroeder defines weight in terms of having reasons to place weight on reasons in deliberation, he cannot keep these claims apart. So his theory cannot countenance these reasons.

¹⁷⁹ See Brown (2014) for a development of a mereological theory of reasons.

¹⁸⁰ See Schroeder (2007b) ch.7.

¹⁸¹ Stratton-Lake (2017: 91).

Although exclusionary reasons of the Razian variety are controversial (as I argued in §4.4.2, we have strong reason to doubt the existence of such reasons) the Scanlonian exclusionary reasons that Schroeder cannot countenance are not of this controversial sort. It is very plausible that there are reasons not to place weight on other reasons in deliberation. By contrast, there are many who deny that there can be a reason not to act for another reason (in particular in cases in which you first-order reason is decisive). So Stratton-Lake's argument is quite persuasive here. Schroeder must alter his theory of weight somehow to accommodate exclusionary reasons of this sort.

Are there *other* theories of weight that can make sense of derivative reasons? One might appeal to a view defended by Way and Setiya, that reasons are the premises of good reasoning. The view can be stated as follows:

Reasoning View For the fact that R to be a reason for A to Φ is for there to be a good basic pattern of reasoning from the belief that R, perhaps together with other correct attitudes which A has, to Φ -ing.¹⁸²

This is intended to be a constitutive account of what a normative reason is. So it may be possible to avoid the conceptual difficulties of formulating derivative reasons in terms of considerations that favour by instead understanding derivative reasons in terms of good reasoning and correct attitudes. However it is not at all clear how we might do this. The view, framed as it is in terms of good (or correct) reasoning, also looks vulnerable to the complaint that it cannot make sense of exclusionary reasons. In particular, since a reason is defined in terms of good reasoning, if we have an excluded reason then it would presumably not be a pattern of good reasoning to deliberate on its basis (since it is excluded) but then according to the *Reasoning View* such an excluded reason is not a reason at all. So I don't think that this view advances us much further.

¹⁸² Way (2017: 254). See also Setiya (2007).

5.2. Scanlon on Wrongness

Let us turn from objections to the argument for the negative buck-passing thesis, to objections to the thesis itself. Scanlon argues that the deontic property of *wrongness* provides agents with normative reasons.¹⁸³ So far I have focused my discussion on ought and have distinguished two properties that ‘ought’ can refer to, namely recommendation and requirement. In what follows I understand wrongness to be equivalent to what I called ‘moral obligation’ in §4.9. Roughly, it is wrong that you Φ if, and only if, you are required not to Φ , and vice versa. I proceed to discuss this property under the label ‘wrongness’ because that is the term used by Scanlon (as well as Darwall, whom I discuss in the next section).

Scanlon claims that there are two roles that wrongness plays in providing reasons. The first he calls the *backstop* role and the second the *shaping* role. The shaping role that wrongness plays is (i) to shape the way an agent should think about the decision they face, and (ii) to determine which other considerations an agent should take as reasons to act. Suppose you are hired as a guard to protect A and then notice that B is about to be injured. The fact that you have undertaken to protect A and that A should reasonably expect you not to leave your post can undermine or remove the decisive-reason status of the fact that B will be harmed without your intervention.¹⁸⁴ Wrongness shapes deliberation here by providing an agent with a guide to weighing reasons. Appealing to the shaping role of wrongness allows us to answer the question: how much weight should be given to acting on some reason given that doing this would be wrong?

The backstop role, according to Scanlon, is that wrongness can also provide reasons to exclude other reasons at an earlier stage of deliberation. Here, Scanlon claims that wrongness provides a reason which serves to answer the question: why should I take this consideration as having decisive force? He remarks that

[...] when I feel convinced by Peter Singer’s article on famine, and find myself crushed by the recognition of what seems a clear moral requirement,

¹⁸³ See Scanlon (1998: 155-8) and (2007).

¹⁸⁴ Scanlon in Shafer-Landau ed. (2007: 8-10).

there is something else at work. In addition to the thought of how much good I could do for people in drought-stricken lands, I am overwhelmed by the further, seemingly distinct thought that it would be wrong for me to fail to aid them when I could do so at so little cost to myself.¹⁸⁵

In order to make sense of this ‘distinct thought’ Scanlon distinguishes between higher-order reasons (reasons that bear on other reasons) and lower-order reasons (reasons that don’t bear on other reasons).¹⁸⁶ He claims that wrongness provides a higher-order reason which tells us why we should take certain lower-order reasons as decisive and certain standards to be authoritative. Scanlon arrives at his particular conception of this one way in which actions can have the property of being wrong:

The strategy of my argument was [...] based on what might be called the *remorse test*: that is, the idea that an account of wrongness and its normative significance ought to fit with our sense of the kind of self-reproach that is occasioned by having done something wrong.¹⁸⁷

According to Scanlon, this higher-order reason is the

[...] reason [we have] to care about whether our actions could be justifiable to others on grounds they could not reasonably reject. [...] [It is a] reason to think in a particular way about what to do and to accept as reasons the first-order considerations that this mode of thinking directs us to.¹⁸⁸

This reason-providing property is not the property of being wrong in the most general sense (that it mustn’t be done) but *one way* in which an action can have that property. Scanlon answers the question of why it makes sense to ask why I should give weight to the fact that not giving money to the famine charity would be wrong

¹⁸⁵ Scanlon (1982: 138).

¹⁸⁶ Note that these differ from Razian first and second-order reasons.

¹⁸⁷ Scanlon in Shafer-Landau ed. (2007: 8).

¹⁸⁸ *Ibid*: 17.

in the following way. It is a decisive reason to act because it is wrong not to give in this case, and the way in which this act is wrong is that we have reason to care about whether our actions could be justifiable to others on grounds they could not reasonably reject. In not giving, not only are you doing something you have reason not to do, you would also be demonstrating that you do not care about whether your actions could be justifiable to others. But you have reason to so care. Here, the way in which you show that you do not care is that you do not take suffering as a decisive reason to give.

The result is a two-tiered view: agents have lower-order reasons for and against acting in certain ways, and higher-order reasons to place decisive weight on those reasons.¹⁸⁹ The higher-order reason is provided by (one way an action has) the property of wrongness. It also shows that the shaping and backstop roles are in fact playing the same role insofar as they are suited to provide an answer to the question of why an agent should place a certain weight on a reason that they have.

However, I think that this view encounters difficulties. First notice that it is plausible that if Scanlon's view is right, the view generalises. If we can sensibly ask why we should give weight to the consideration that it would cause suffering, or that it would break a promise, it seems sensible to ask this question of *any* lower-order reason.

According to Scanlon, a higher-order reason justifies treating a lower-order reason as decisive. But if it is sensible to ask 'Why should I treat promise-breaking as a decisive reason?' then surely it is sensible to ask 'Why should I treat *justifiability to others* as a decisive reason to treat promise-breaking as a decisive reason?' One way to stop a potential regress is to claim that you must treat justifiability to others as a reason to treat promise-breaking as a decisive reason, because justifiability to others provides the *ultimate* or *fundamental* reason not to act in the way you would by breaking the promise. But stopping the regress in that way collapses the higher/lower order distinction. Justifiability to others becomes the ultimate reason not to act in that way, which is simply another lower order reason.

¹⁸⁹ The view is thus similar to Schroeder's, but restricted to placing decisive weight on other reasons.

In addition, returning to the shaping and backstop roles, I think that we can account for Scanlon's 'seemingly distinct thought' in a number of ways which do not commit us to the two-tiered view. Take again the example of promising to protect your client. Here, assuming the role of guarding your client determines which other considerations you should take as decisive reasons to act. But that can be explained without making use of the idea of wrongness. We could invoke a simple balancing explanation. We might think that the fact that you are to guard your client gives you very strong reason to protect him, since that is implied in your undertaking to guard him. Leaving your client open to attack, given this fact, is less justifiable and hence outweighs the weaker reason to assist the other person about to be injured. So the fact that you have undertaken to guard your client determines which other considerations you should take as a decisive reason to act by outweighing reasons that in other contexts would prevail.

Or we might think that the fact that you have undertaken to guard your client not only gives you decisive reason to protect him, but also itself (that fact) gives you decisive reason not to act in any other way. Here your assuming that role determines which other consideration you should take as a decisive reason to act by itself counting against acting in that way. What is the difference between these two explanations? As I argued in chapter 4, reasons against acting in some way imply a rational demand to justify yourself if you choose to act in that way, while reasons for generally imply at most self-criticism if you fail to choose to act in that way. Thus on this second explanation it is not simply the case that your assuming the role of guard outweighs the reason you have to intervene to prevent injury to the other person. That fact also directly bears on, by providing a reason which counts against, your acting so as to intervene.

A further reply is to distinguish between the *ground* of a reason (the fact that is reason-providing) and the *way* in which that reason counts in favour or counts against. We can understand favouring in at least two ways: in terms of promotion, and in terms of respect.¹⁹⁰ We can understand 'R counts in favour of A's Φ -ing' as ' Φ -ing promotes R'; e.g. that it would give A pleasure is a reason for A to eat the ice-

¹⁹⁰ See Pettit (1989: 116 – 26) who draws the distinction in discussing the different ways in which one can value something.

cream as A's eating the ice-cream promotes a pleasant objective/state of affairs. On the other hand we can understand 'R counts in favour of A's Φ -ing' as ' Φ -ing is a way of respecting R' e.g. that B is A's friend counts in favour of A's returning B's call as: because B is A's friend, A's returning B's call is a way of respecting B as a friend.

With this distinction in mind we can interpret the Singer case above as a consideration playing both the promotion and the respect roles. The fact that the group are subject to great suffering gives you decisive reason to intervene (by giving money to Oxfam, for example) but it does so in both of the ways adumbrated above. The reason-providing fact that they are suffering can be understood both as an injunction to diminish the objective/state of affairs in which there is such suffering, and is also, were you not to act on its basis, a way of manifesting disrespect towards that group.

Thus, given an plausible alternative way to make room for Scanlon's 'distinct thought' that does not appeal to the controversial two-tiered view, I do not think we should accept the two-tiered view as a reason to believe that wrongness provides reasons.

5.3. Darwall on Wrongness

According to Darwall, wrongness provides reasons for action. If this view is right then NBT is false. We should thus pay close attention to Darwall's argument.

First note that Darwall appeals to what he calls second-personal reasons.¹⁹¹ How might B 'give' A a reason to stop standing on his foot? B could point out that there reason against causing the pain of standing on a foot (that happens to be B's) and get A to see this. This is to make an epistemological demand on A, to ask him to believe that there is a reason to act. On the other hand, B could demand of A that he stop standing on his foot by giving A a claim-based agent relative reason not to stand on his foot. The second sort of agent-relative reason is one that, necessarily, is

¹⁹¹ See Darwall (2006: 3-4).

backed-up by some particular authority and makes a demand directly on A's will. It is reasons of this second sort that Darwall calls 'second-personal'.¹⁹²

Darwall is a buck-passer from the deontic property of wrongness to the property of having certain Strawsonian 'reactive attitudes', most importantly blame.¹⁹³ He argues that wrongness provides reason for certain reactive attitudes as follows:

... although a buck-passing theory of a normative concept (or similarly, a 'warranted attitude' or 'fitting-attitude' theory of the concept) entails that the fact that something instantiates the concept (say, is good, right, or wrong) is not a further reason for the specific attitudes that are conceptually tied to the normative concept, the theory may not entail that that fact is not a reason for some relevant choice, intention, or *action*, since the conceptually implicated attitude may not entail any relevant action attitude (like intention or choice).¹⁹⁴

The most promising account of buck-passing according to Darwall is a 'fittingness' approach. For example, to be desirable is to be the fitting object of desire. One way to construe buck-passing about value thus is that to be valuable is to be the fitting object of some valuing attitude (such as desire). So on this way of understanding the view with respect to wrongness we need to know which attitudes are conceptually implicated by the concept of wrongness. Darwall suggests that Strawsonian reactive attitudes like blame, guilt, and indignation are implied by wrongness, i.e. the second-personal relation of holding someone accountable. What is morally wrong is the fitting object of blame. However, this tells us nothing about reasons for action, as opposed to reasons for having attitudes.

On the other hand, *to be desirable*, does seem to have a conceptual connection not only with desire, but also with action. Darwall claims that '... some disposition to bring the desired state about seems to be part of the very concept of

¹⁹² *Loc. cit.*

¹⁹³ See Strawson (2008: 6ff).

¹⁹⁴ Darwall (2010: 137-8).

desire'.¹⁹⁵ But is blame like desire in this respect? Darwall claims that it is not. It is possible that you have reason to blame somebody for acting wrongly, and that you fail to have any reason to act in any way. Compare with the estimable. That something is estimable is for there to be reasons to esteem it, but that you have a reason for acting given that X is estimable is not conceptually tied to esteem. So it could be that something being estimable is itself a reason to act. Similarly, the mere fact that somebody is blameworthy, according to Darwall, does not necessarily give us any reason to do anything.

If we accept that wrongness passes the buck to reasons for blame, this opens up conceptual space to claim that wrongness can itself provide reasons to act. What is this reason? Darwall claims that the fact that an act *violates a demand that is legitimately made* itself provides a reason not to perform the act that is additional to its wrong-making features. So, for example, the fact that you broke a promise gives me reason to blame you and *because* promise-breaking violates a legitimate demand it gives me a reason not to do so.¹⁹⁶

Thus Darwall's argument proceeds in two stages. In the first stage Darwall argues for the particular positive deontic buck-passing view that he believes should be accepted. Adopting this version of the positive deontic buck-passing view gives us logical room to be a buck-passer about wrongness, but *also* claim that wrongness provides reasons for action:

1. The positive deontic buck-passing claim should take the form of a fitting-attitude account in which to say that Φ -ing is wrong is to say that it is appropriate to blame A for Φ -ing.¹⁹⁷
2. Some fitting attitudes do not have conceptual implications for action. Blame is one such attitude with no implications for action.
- C. So , you can be a buck-passer about wrongness (in terms of fitting attitudes) *and* claim that wrongness provides reasons for action.

¹⁹⁵ *Loc. cit.*

¹⁹⁶ *Ibid.*: 153.

¹⁹⁷ This view takes its inspiration from Mill. See Mill (1979: 60-10).

So if wrongness passed the buck to reasons for action, then wrongness could not itself provide reasons for action without circularity. But, says Darwall, wrongness instead passes the buck to the attitude of blame, which is not a reason for action. So there is nothing circular about claiming that wrongness provides reasons for action.

In the second stage Darwall argues that not only is it *possible*, given his positive deontic buck-passing view, that you can hold that wrongness provides reasons for action, you *should* agree that wrongness provides reasons for action. In what follows I will focus on the first stage of the argument and on the plausibility of premises 1 and 2.

5.3.1. The objection from morally conscientious agency

I first examine premise 1, that the best form of the positive deontic buck-passing view is a ‘fitting-attitude’ account that analyses the deontic property of wrongness in terms of the reactive attitude of *blame*. In effect, the thesis claims that for X to be wrong is for it to be appropriate for B to blame A for X-ing.

According to the fitting-attitude view Darwall holds, for an action to be wrong is for it to be appropriate to blame the agent who performs that act.¹⁹⁸ Now consider the view that an agent who is moved to act by the thought that acting in some way would be right, and moved to avoid acting in some way by the thought that it would be wrong, is praiseworthy. This need not be true as a conceptual claim, but is found plausible on a number of views.¹⁹⁹

Combining this view about praiseworthiness in light of moral motivation with Darwall’s view that wrongness is blameworthiness, an agent is praiseworthy when they are motivated not to act in some way because they would thereby be appropriately subject to blame. The resulting view is questionable. In order to assess these claims we need to know in what way we should characterise blame. In particular must our view of the nature of blame imply that it is a ‘sanction’ in a sense incompatible with conscientious moral motivation?

¹⁹⁸ *Ibid*: 143.

¹⁹⁹ The classic example would be Kant. See Kant (1966), and Herman (1981), Stratton-Lake (2000) for discussion.

According to Darwall, blame is to be understood as a reactive attitude in Strawson's sense. For Strawson a 'reactive attitude' is the sort of attitude that we feel are importantly attached to those to whom we stand in interpersonal relations, such as friendship, kinship, or marriage.²⁰⁰ On views of this kind, blame is claimed to be the exercise of an emotional response such as resentment, indignation, guilt, and anger.²⁰¹ Darwall goes further in his analysis of blame, telling us that reactive attitudes must involve demands that are addressed to the object of those attitudes interpersonally.²⁰²

We should note and set aside the claim that blameworthiness and wrongness are conceptually distinct on the basis of the fact that it is plausible that an agent can act wrongly, without being blameworthy in cases in which he had an excuse. I accept for the purposes of this discussion that it is understood here that blameworthiness refers to blameworthiness without excuse.

Let us proceed on the assumption that Darwall accepts the view that blame is a primarily emotional reaction to wrongdoing that is a way of holding agents accountable through (a) censuring them for acting in that way, and (b) making appropriate that the wrongdoer should feel guilt for what have done.

I believe that the avoidance of being the fitting object of blame is not the best explanation of why we should praise agents for avoiding wrongdoing. Consider an agent who is moved to act in some way towards somebody simply because they care about that person. Why does this agent fail to be a morally conscientious person? Often it can seem that being motivated in this way is a paradigmatic form of moral motivation. Suppose that, having done me a good turn, I decide to return the favour. On Darwall's view, my returning the favour because I care about our friendship (say) is less morally praiseworthy than my returning the favour because I want to avoid it being appropriate to blame me for not returning the favour. Not only does this view counterintuitive, it seems wrong as a description of the phenomenology. So

²⁰⁰ See Strawson (2008: 6-7).

²⁰¹ See Toganazzini & Coates (2016) 1.2. There are a number of objections to the reactive-attitude theory of blame, and a number of rival views. I will not consider these here. See e.g. Hieronymi (2004) for discussion.

²⁰² Darwall (2010: 142).

avoidance of being the fitting object of blame does not plausibly look like a necessary condition of morally conscientious agency.

One reply open to Darwall is to follow Scanlon in claiming that there are many different ways in which actions can be wrong, and that the analyses of wrongness in terms of blame is an analysis of only one way in which an act can be wrong. In particular it is true only of what Scanlon calls ‘what we owe to each other’ and Darwall sometimes calls ‘bipolar obligations’. An agent has a bipolar obligation where the action that they are obligated to perform is, in some sense, owed to some other particular agent (again, these obligations involve Darwall’s notion of a second-personal, or agent-relative reason). Appealing to this restriction may allow us to avoid problematic cases such as the charity case above, in which it is not true that the agent owes his acting in some way to some other person.

I do not think, however, that adopting a more restricted usage of rightness and wrongness will entirely get around the objection. We can adduce cases of bipolar obligations that still seem problematic when interpreted as Darwall suggests. Suppose that I ought today to give you back the £5 that you lent to me last week. My being motivated by the thought that unless I give you back the money I will be appropriately subject to blame does not necessarily look like the thought that a morally conscientious person would have. What is morally suspect about my being moved by my recognition of the fact that I owe you the money is a decisive moral reason?²⁰³ Since this does not seem morally suspect, the view is rendered implausible.

²⁰³ Bedke defends a similar Millian-style buck-passing view in Bedke (2011). The difference between Bedke’s view and Darwall’s is that Bedke claims that deontic properties pass the buck to reasons for *speech acts* not reactive attitudes. For example, for some act to be required of A is for B to have most reason to say ‘no’ to the question ‘may I Φ ?’. I find this view even less plausible than Darwall’s as it is just not clear what the performance or non-performance of speech acts has to do with what you ought to do. Since Bedke strictly distinguishes between reasons to act and reasons to perform speech acts, he also faces the problem of ensuring the entailments between ‘should’ and ‘must’ that besets Gert’s theory (see above §4.3).

5.3.2. Objections to Premise 2

It is crucial to Darwall's argument that blame is more like amusement than desire, in the sense that fitting attitudes like amusement plausibly do not imply reasons to act in any way towards the object of amusement, while fitting attitudes like desires plausibly do imply reasons to act in some way towards (the fulfilment of) the object of the desire. Recall that if we are buck-passers about wrongness, if what it is to be the fitting attitude of wrongness is to have a reason to act in some way towards the object of that attitude, claiming that wrongness itself provides reasons would again invoke a charge of double-counting reasons.

Darwall claims that being blameworthy does not imply having a reason to act because

Even if esteeming some trait generally gives rise, say, to some desire to acquire the trait, the desire does not seem to be *conceptually* implicated in the former. We might imagine beings who admire some trait while lacking any desire whatsoever to act in any particular way with respect to it.²⁰⁴

Presumably, by analogy with blame Darwall would claim that we could imagine beings that hold A blameworthy for acting in some way while lacking any desire whatsoever to act in any particular way towards A. But if we have reason to blame A don't we also have reason to act in certain ways towards him? Here are two ways in which we might support this claim. First, as Miller claims, blame implicates action because it involves compunction, which itself implicates action.²⁰⁵ Compunction is a reactive attitude of uneasiness, similar to guilt that one feels either after performing a morally wrong action, or that one feels before being tempted to perform a morally wrong action. Compunction however differs from guilt in that an agent who feels compunction is thereby disposed to avoid actually performing the morally wrong action. Thus, just as something that makes an object desirable gives

²⁰⁴ Darwall (2010: 145).

²⁰⁵ See Miller (Manuscript: 7-8).

us a practical reason to aim for it, what makes an object (action) worthy of compunction gives us a practical reason to avoid doing it.

Second, consider the role that blame plays in making appropriate sanctions. As discussed above, the best interpretation of Darwall's position is that blame is a reactive attitude that functions as making appropriate sanctions on wrongful conduct. Plausibly, however, conceiving blame as appropriate sanction has force *only if* it implies action, in particular that a blameworthy agent will be acted on in some way. To illustrate the point, consider Skorupski's views on the nature of blame. Skorupski argues that blame is an emotion that falls under what he calls the 'bridge principle' that

Whatever facts give x reason to feel Φ give x reason to do the Φ -prompted action, in virtue of being a reason to feel Φ .²⁰⁶

He claims that the action to which blame gives rise is '...withdrawal of recognition'²⁰⁷ which is to cut off one's relations with the targeted agent of blame, putting a distance between oneself and the agent. Part of what it is to withdraw recognition is to act or refrain from acting in certain ways towards the target individual. If an account like this is right, Darwall's claim looks less plausible.

Further, Darwall seems to think that the fact that A is blameworthy doesn't imply any reasons *for* A. That seems questionable, since it is plausible that A's being blameworthy implies that A has reason to refrain from acting in that way. But even if blameworthiness doesn't imply reasons for A, it implies reasons for others to act in certain ways towards me, such as to censure me for acting in that way. Those reasons to censure me may be outweighed by other reasons, but still seem to be appropriate. If that is right then blameworthiness conceptually implies reasons for action.

²⁰⁶ Skorupski (2010: 267).

²⁰⁷ *Ibid*: 294.

5.3.4. Conclusion

In this chapter I have considered the negative claim of the deontic buck-passing view. I have discussed both arguments against arguments for the negative claim and arguments against the claim itself. I do not believe that these arguments are successful, and that we can hold the negative claim consistently with the positive. This puts in overall argument in a strong position as, if the negative claim is true, it constitutes some pressure to accept some version of the positive claim (though it does not depend on it). Thus the failure of these arguments should be seen as providing indirect support for the positive, in addition to NBT.

6. Overall Conclusions

In this thesis I have explored a number of themes. I have argued that the best view of what a normative reason is holds that it is a three-place relation of either counting in favour, or counting against, which I argued should be understood as distinct relations. Such reasons plausibly demand at least more than mere conformity to them, *pace* Raz. I then argued that normative reasons are more fundamental than oughts. I considered and rejected a number of possible analyses of reasons in terms of oughts and argued that we can provide a good analysis of oughts in terms of reasons.

This thesis leaves unresolved a number of issues that I intend to take up elsewhere. I mention two here. First, a more fully worked-out theory of the distinction between reasons for and reasons against would be invaluable in helping to further specify the deontic buck-passing view. Second, although I believe that the analyses of requirement and recommendation cover the most important ground as an analysis of deontic properties, the analysis may be taken further by considering other putative deontic properties, such as norms and rules. Further work may be done to extend my existing analysis to these properties.

REFERENCES

- Anderson, Elizabeth (1993) *Value in Ethics and Economics*. Harvard University Press.
- Arpaly, Nomy (2002) 'Moral Worth', *Journal of Philosophy* 99 (5):223.
- Arpaly, Nomy (2015) 'Huckleberry Finn Revisited: Inverse Akrasia and Moral Ignorance', in Randolph Clarke Michael Mckenna & Angela M. Smith (eds.), *The Nature of Moral Responsibility*. New York, USA: Oxford University Press: 141-156.
- Bedke, Matt (2009) 'The Iffiest Oughts: A Guise of Reasons Account of End-Given Conditionals', *Ethics* 119 (4):672-698.
- Bedke, Matt (2011) 'Passing the Deontic Buck', *Oxford Studies in Metaethics: Volume 6* 6:128.
- Bedke, Matt (2017) 'Ends to Means: A Probability-Raising Account of Means and the Weight of Reasons to Take Them', *Journal of Ethics and Social Philosophy*, Vol 12: 1.
- Berker, Selim (2007) 'Particular Reasons', *Ethics* 118 (1):109-139.
- Bradford, Gwen (2015) *Achievement*. Oxford University Press.
- Broome, John (2004) 'Reasons', in R. Jay Wallace ed. *Reason and Value: Themes From the Moral Philosophy of Joseph Raz*. Oxford University Press: 28-55.
- Broome, John (2007) 'Does Rationality Consist in Responding Correctly to Reasons?', *Journal of Moral Philosophy* 4 (3):349-374.

- Broome, John (2008) 'Reply to Southwood, Kearns and Star, and Cullity', *Ethics* 119 (1):96-108.
- Broome, John (2013) *Rationality Through Reasoning*. Wiley-Blackwell.
- Broome, John (2015) 'Reason Versus Ought', *Philosophical Issues* 25 (1):80-97.
- Broome, John (Forthcoming) 'Reasons Fundamentalism and what is Wrong with it', in Star ed. *Handbook of Reasons and Normativity*. Oxford University Press.
- Brown, Campbell (2014) 'The Composition of Reasons', *Synthese* 191 (5).
- Brunero, John (2009) 'Reasons and Evidence One Ought', *Ethics* 119 (3):538-545.
- Brunero, John (2013) 'Reasons as Explanations', *Philosophical Studies* 165 (3):805-824.
- Crisp, Roger (2000) 'Particularizing Particularism', in Brad Hooker & Margaret Olivia Little (eds.), *Moral Particularism*. Oxford University Press: 23-47.
- Crisp, Roger (2006). *Reasons and the Good*. Clarendon Press.
- Cullity, Garrett (2002) 'Particularism and Presumptive Reasons', *Aristotelian Society Supplementary Volume* 76 (1):169-190.
- Cullity, Garrett (2013) 'The Context-Undermining of Practical Reasons', *Ethics* 124 (1):8-34.
- Dancy, Jonathan (2000a) *Practical Reality*. Oxford University Press.
- Dancy, Jonathan (2000b) 'Should we Pass the Buck?', *Royal Institute of Philosophy Supplement* 47:159-173.

- Dancy, Jonathan (2004a). *Ethics Without Principles*. Oxford University Press.
- Dancy, Jonathan (2004b) 'Enticing Reasons', in R. Jay Wallace, Philip Pettit, Samuel Scheffler & Michael Smith eds. *Reason and Value: Themes From the Moral Philosophy of Joseph Raz*. Clarendon Press.
- Darwall, Stephen (2006) *The Second-Person Standpoint: Morality, Respect, and Accountability*. Harvard University Press.
- Darwall, Stephen (2010) 'But it would be Wrong', *Social Philosophy and Policy* 27 (2):135-157.
- Davidson, Donald (1963) 'Actions, Reasons, and Causes', *Journal of Philosophy* 60 (23):685-700.
- Essert, Christopher (2012) 'A Dilemma for Protected Reasons', *Law and Philosophy* Vol. 31, No. 1: 49-75.
- Feldman, Fred (1997) 'Two Questions about Pleasure', in *Utilitarianism, Hedonism and Desert: Essays in Moral Philosophy*, Cambridge: Cambridge University Press.
- Finlay, Stephen (2012) 'Explaining Reasons', *Deutsches Jahrbuch Fuer Philosophie* 4:112-126.
- Gardner, John (2007). *Offences and Defences: Selected Essays in the Philosophy of Criminal Law*. Oxford University Press.
- Gert, Joshua (2000) 'Practical Rationality, Morality, and Purely Justificatory Reasons', *American Philosophical Quarterly* 37 (3):227 - 243.

- Gert, Joshua (2003) 'Requiring and Justifying: Two Dimensions of Normative Strength', *Erkenntnis* 59 (1):5 - 36.
- Gert, Joshua (2005) 'A Functional Role Analysis of Reasons', *Philosophical Studies* 124 (3):353-378.
- Gert, Joshua (Forthcoming) 'The Distinction Between Justifying and Requiring: Nothing to Fear', in Star ed. *The Oxford Handbook of Reasons and Normativity*. Oxford University Press.
- Gibbons, John (2010) 'Things That Make Things Reasonable', *Philosophy and Phenomenological Research* 81 (2):335-361.
- Greco, John (2010) *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity*. Cambridge University Press.
- Greenspan, Patricia (2005) 'Asymmetrical Practical Reasons', in M. E. Reicher, J. C. Marek eds. *Experience and Analysis*. 387-394.
- Greenspan, Patricia (2007) 'Practical Reasons and Moral "Ought"', in Russell Schafer-Landau ed. *Oxford Studies in Metaethics, vol. II*. Clarendon Press: 172-194.
- Gregory, Alex (2016). 'Normative Reasons as Good Bases', *Philosophical Studies* 173 (9):2291-2310.
- Gregory, Alex (2014) 'A Very Good Reason to Reject the Buck-Passing Account', *Australasian Journal of Philosophy* 92 (2):287-303.
- Herman, Barbara (1981) 'On the Value of Acting from the Motive of Duty', *Philosophical Review* 90 (3):359-382.

- Heuer, Ulrike (2010) 'Wrongness and Reasons', *Ethical Theory and Moral Practice* 13 (2):137 - 152.
- Heuer, Ulrike (2012) 'Promising', *Philosophy Compass* 7 (12):842-851.
- Hieronymi, Pamela (2004) 'The Force and Fairness of Blame', *Philosophical Perspectives* 18 (1):115–148.
- Hooker, Brad (2000) *Ideal Code, Real World: A Rule-Consequentialist Theory of Morality*. Oxford University Press.
- Hooker, B (2000) 'Moral Particularism: Wrong and Bad', in Brad Hooker & Margaret Olivia Little eds. *Moral Particularism*. Oxford University Press; 23-47.
- Hooker, Brad (2008) 'Moral Particularism and the Real World', in Mark Norris Lance, Matjaž Potrč & Vojko Strahovnik eds. *Challenging Moral Particularism*. Routledge: 12-30.
- Kant, Immanuel & Paton, H. J. (1966) *The Moral Law or, Kant's Groundwork of the Metaphysic of Morals. Translated and Analysed by H. J. Paton*. Hutchinson University Library.
- Kearns, Stephen & Star, Daniel (2008) 'Reasons: Explanations or Evidence?', *Ethics* 119 (1):31-56.
- Kearns, Stephen & Star, Daniel (2009) 'Reasons as Evidence', *Oxford Studies in Metaethics* 4:215-42.
- Kearns, Stephen & Star, Daniel (2013) 'Reasons, Facts-About-Evidence, and Indirect Evidence', *Analytic Philosophy* 54 (2):237-243.

- Kearns, Stephen & Star, Daniel (2013) 'Weighing Reasons', *Journal of Moral Philosophy* 10 (1):70-86.
- Kearns, Stephen (2016) 'Bearing the Weight of Reasons', in Lord & Maguire eds. *Weighing Reasons*. Oxford University Press.
- Kiesewetter, Benjamin (2016) 'You ought to ϕ only if you may believe that you ought to ϕ ', *Philosophical Quarterly* 66 (265):760-82.
- Kolodny, Niko (Forthcoming) 'Instrumental Reasons', in Star ed. *Oxford Handbook of Reasons and Normativity*.
- Lord, Errol (2015) 'Acting for the Right Reasons, Abilities, and Obligation', in Russ Shafer-Landau ed., *Oxford Studies in Metaethics, vol. 10*. Oxford University Press.
- Lord, Errol & Maguire, Barry (2016) 'An Opinionated Guide to the Weight of Reasons', in Lord & Maguire eds. *Weighing Reasons*. Oxford University Press.
- Lord, Errol (Forthcoming) *The Importance of Being Rational*. Oxford University Press.
- Maguire, Barry (2016) 'The Value-Based Theory of Reasons', *Ergo, an Open Access Journal of Philosophy* 3.
- Mantel, Susanne (2014) 'No Reason for Identity: On the Relation Between Motivating and Normative Reasons', *Philosophical Explorations* 17 (1):49-62.
- Mantel, Susanne (2013) 'Acting for Reasons, Apt Action, and Knowledge', *Synthese* 190 (17):3865-3888.

- McKeever, Sean & Ridge, Michael (2005) 'What Does Holism have to do with Moral Particularism?', *Ratio* 18 (1):93–103.
- McKeever, Sean D. & Ridge, Michael R. (2006). *Principled Ethics: Generalism as a Regulative Ideal*. Oxford University Press.
- Mill, J. S. (1979 [1861]). *Utilitarianism*. Indianapolis: Hackett Pub. Co.
- Miller, Dale (Manuscript) 'Compunction, Buck-Passing, and Moral Reasons: Reply to Darwall'.
- Nagel, Thomas (1970) *The Possibility of Altruism*. Oxford Clarendon Press.
- Owens, David (2008) 'Rationalism about Obligation', *European Journal of Philosophy* 16 (3):403-431.
- Parfit, Derek (1984) *Reasons and Persons*. Oxford University Press.
- Parfit, Derek (2011) *On What Matters: Volume One*. Oxford University Press.
- Pettit, Philip (1989) 'Consequentialism and Respect for Persons', *Ethics*, 100(1): 116-26.
- Pollock, John (1970) *Knowledge and Justification*. Princeton University Press.
- Portmore, Douglas (2011) *Commonsense Consequentialism: Wherein Morality Meets Rationality*. New York, USA: Oxford University Press USA.
- Pritchard, Duncan (2005) *Epistemic Luck*. Oxford University Press UK.
- Pritchard, Duncan (2008) 'Radical Scepticism, Epistemic Luck, and Epistemic Value', *Aristotelian Society Supplementary Volume* 82 (1):19-41.

- Rabinowicz, Wlodek & Rønnow-Rasmussen, Toni (2004). 'The Strike of the Demon: On Fitting Pro-Attitudes and Value', *Ethics* 114 (3):391-423.
- Raz, Joseph (1978) 'Reasons for Action, Decisions and Norms', in Raz, Joseph, ed. *Practical Reasoning*. Oxford University Press.
- Raz, Joseph (1986) *The Morality of Freedom*. Oxford University Press.
- Raz, Joseph (1990) *Practical Reason and Norms*. Oxford University Press.
- Raz, Joseph (1999) *Engaging Reason: On the Theory of Value and Action*. Oxford University Press.
- Raz, Joseph (2006) 'The Trouble with Particularism (Dancy's Version)', *Mind* 115 (457):99-120.
- Raz, Joseph (2010) 'On the Guise of the Good', in Tenebaum, ed., *Desire, Practical Reason, and the Good*. Oxford University Press.
- Ridge, Michael (2003) 'Contractualism and the New and Improved Redundancy Objection', *Analysis* 63 (4):337-342.
- Robertson, Simon (2008) 'Not So Enticing Reasons', *Ethical Theory and Moral Practice* 11 (3):263-277.
- Russell, B (2001) 'Epistemic and Moral Duty', in Steup, Matthias ed. *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*. Oxford University Press.

- Scanlon, Thomas (1982) 'Contractualism and Utilitarianism', in Amartya Kumar Sen & Bernard Arthur Owen Williams eds. *Utilitarianism and Beyond*. Cambridge University Press: 103-110.
- Scanlon, Thomas (1998) *What We Owe to Each Other*. Belknap Press of Harvard University Press.
- Scanlon, Thomas (2007) 'Wrongness and Reasons: A Re-examination', in Russ Shafer-Landau ed. *Oxford Studies in Metaethics*. Clarendon Press.
- Scanlon, Thomas (2014) *Being Realistic About Reasons*. Oxford University Press.
- Schroeder, Mark (2007a) 'Reasons and Agent-Neutrality', *Philosophical Studies* 135 (2):279-306.
- Schroeder, Mark (2007b) *Slaves of the Passions*. Oxford University Press.
- Schroeder, Mark (2009) 'Buck-Passers' Negative Thesis', *Philosophical Explorations* 12 (3):341-347.
- Schroeder, Mark. (2011) 'Ought, Agents, and Actions', *Philosophical Review* 120 (1):1-41.
- Schroeder, Mark (Manuscript) 'Buck-Passers' Negative Thesis'.
- Schroeter, Laura & Schroeter, Francois (2009) 'Reasons as Right-Makers', *Philosophical Explorations* 12 (3):279-296.
- Setiya, Kieran (2007) *Reasons Without Rationalism*. Princeton University Press.

- Skorupski, John (2007) 'Buck-Passing about Goodness', in J. Josefsson D. Egonsson ed. *Hommage à Wlodek. Philosophical Papers Dedicated to Wlodek Rabinowicz*. URL = <<http://www.fil.lu.se/hommageawlodek/index.htm>>.
- Skorupski, John (2010) *The Domain of Reasons*. Oxford University Press.
- Slote, Michael A. (1985). *Common-Sense Morality and Consequentialism*. Routledge & Kegan.
- Sosa, Ernest (2007) *A Virtue Epistemology: Apt Belief and Reflective Knowledge, Volume I*. Oxford University Press.
- Smith, M (2016) 'Parfit's Mistaken Metaethics', in Singer, Peter, ed. *Does Anything Really Matter? Parfit on Objectivity*. Oxford University Press.
- Snedegar, Justin (2016) 'Reasons, Oughts, and Requirements', in Russ Shafer-Landau ed. *Oxford Studies in Metaethics*. Clarendon Press.
- Snedegar, Justin (2017) 'Reasons For and Reasons Against', *Philosophical Studies* (3) 725-43.
- Snedegar, Justin (2017). *Contrastive Reasons*. Oxford University Press.
- Stratton-Lake, Philip (2000). *Kant, Duty, and Moral Worth*. Routledge.
- Stratton-Lake, Philip (2005) 'How to Deal with Evil Demons: Comment on Rabinowicz and Rønnow-Rasmussen', *Ethics* 115 (4):788-798.
- Stratton-Lake, Philip (2017) 'The Buck-Passing Theory of Value: Assessing the Negative Thesis', in Kirchin, Simon ed. *Reading Parfit: On What Matters*. Routledge.

- Stratton-Lake, Philip (Forthcoming) 'Fundamentalism about Reasons', in Star ed. *Oxford Handbook of Reasons and Normativity*
- Strawson, Peter F. (2008) *Freedom and Resentment and Other Essays*. Routledge.
- Tanyi, Attila & Matteo, Morganti (2017) 'Can Reasons be Propositions? Against Dancy's Attack on Propositionalism', *Theoria* 83 (3):185-205.
- Thomson, Judith Jarvis (2008). *Normativity*. Open Court.
- Tognazzini, Neal & Coates, D. Justin, (2016) 'Blame', *The Stanford Encyclopedia of Philosophy* Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/spr2016/entries/blame/>.
- Turri, John (2009) 'The Ontology of Epistemic Reasons', *Noûs* 43 (3):490-512.
- Väyrynen, Pekka (2004) 'Particularism and Default Reasons', *Ethical Theory and Moral Practice* 7 (1):53-79.
- Väyrynen, Pekka (2011) 'A Wrong Turn to Reasons?', in Michael Brady ed. *New Waves in Metaethics*. Palgrave-Macmillan.
- Wallace, R. Jay (2013) 'The Deontic Structure of Morality', in David Bakhurst, Margaret Olivia Little & Brad Hooker (eds.), *Thinking About Reasons: Themes From the Philosophy of Jonathan Dancy*. Oxford University Press.
- Way, Jonathan & Whiting, Daniel (2016) 'Reasons and Guidance', *Analytic Philosophy* 57 (3):214-235.
- Way, Jonathan (2017) 'Reasons as Premises of Good Reasoning', *Pacific Philosophical Quarterly* 98 (2).

Wiland, Eric (2012). *Reasons*. Continuum.

Whiting, Daniel (2017a). 'Right in Some Respects: Reasons as Evidence',
Philosophical Studies:1-18.

Whiting, Daniel (2017b) 'Against Second-Order Reasons', *Noûs* 51 (2):398-420.

Williams, Bernard (1979) 'Internal and External Reasons', in Ross Harrison ed.
Rational Action. Cambridge University Press: 101-113.