

Reasons, Oughts, and Requirements*

Justin Snedegar
University of St Andrews
js280@st-andrews.ac.uk

January 27, 2015
Not a final version

This paper raises a challenge for the recently popular *reasons first* approach to normativity, according to which all normative notions can be explained in terms of reasons. The reasons first theorist owes us an account of how these explanations go for all other normative notions. I focus here on *requirement*, and to a lesser extent, *permission*.

There is a very plausible, widely accepted account of the relationship between your reasons and what you ought to do—roughly, what you ought to do is just what you have most reason to do. But it is important to distinguish what you ought to do and what you are required to do. So we still need to give some account of the relationship between reasons and requirements, and relatedly, between reasons and permission. This is less straightforward than giving an account of ought in terms of reasons. I focus in this paper on a strategy I call the *Two Kinds of Reasons* strategy, and argue that it faces serious obstacles.

1 Reasons First

On one historically popular conception, ethics is about rules, principles, and imperatives. These rules or principles are meant to tell you in any given situation what you are morally required to do. If you break one of these rules, then you are acting wrongly. In this sense, the rules are meant to be *exceptionless*.

But this way of thinking about morality leads to a pretty obvious problem: these rules, principles, or imperatives can *conflict*. For example, sometimes avoiding causing harm requires lying. One response is to deny that there are actually such conflicts, either because there is in fact only one true principle, or because the principles are much more complicated than we realize, such that there will not be

*Thanks to Shyam Nair, Bruno Guindon, Ben Sachs, Brian McElwee, and audiences at the Institute for Philosophy in London, the 2014 St. Louis Annual Conference on Reasons and Rationality, and the 2014 Wisconsin Metaethics Workshop.

conflicts between these very carefully qualified rules—for example, “Don’t break your promises unless . . .”, where the ‘unless’ clause may be very long.¹

Another response is to abandon the idea that there are exceptionless rules that you must act in accordance with in every situation. This approach is taken by Ross (1930) in his theory of *prima facie* duties. Ross took a list of what earlier theorists would have treated as exceptionless rules or principles, like ‘Don’t cause harm’, ‘Provide help’, ‘Don’t lie’, and so on, and changed their status. Rather than thinking of them as exceptionless rules or principles, Ross thought of them as *prima facie*, or (as we’d now say) *pro tanto*, or contributory considerations—what we’d now call *reasons*.² They contribute some weight either for or against the action: that it causes harm counts against it with some weight, but that it involves telling the truth counts for it with some weight. We can then weigh up these considerations to determine the overall status of various actions.

This turn to reasons has been extended beyond morality, to normativity more generally. Just as there are moral reasons which contribute to the overall moral status of an action, there are prudential reasons, epistemic reasons, and perhaps practical reasons more generally. Recently, this approach has been taken to the extreme in the *reasons first* program. According to this program, normativity is all about reasons—all normative notions can be understood in terms of reasons. Part of the motivation for this picture comes from the difficulty in trying to explain features of normativity, and in particular morality, in terms of overall notions like rules or imperatives, and the relative success of doing so in terms of reasons, as in the (admittedly simplified) discussion above. But there are other kinds of motivations, as well.

One appeal of the reasons first program is that it promises to explain what it is that all normative notions have in common—what it is that makes the normative *normative*. What makes a notion normative is that it is to be explained in terms of, or is intimately related to, reasons.³ Relatedly, the reasons first picture may let us give a systematic description of normativity. By understanding all other normative notions in terms of reasons, we might be able to see what kinds of relationships these notions stand in to one another: reasons give us a common currency. Finally, if philosophers like Schroeder (2007) are right that reasons can be given a naturalistic reduction, a successful reasons first approach would give us a reduction of the normative.⁴

¹See Sidgwick (1907) for an early discussion of how complicated the principles would have to be in his discussion of the ‘Intuitional method’. See Scanlon (1998), Chapter 5, for a more recent discussion. See Dancy (2004b), Chapter 2, section 3 for critical discussion.

²See the discussions of Ross in Horty (2007), p. 21 and Dancy (2004b), pp. 3-12. The narrative introducing contributory reasons in the main text owes a lot to these passages from Dancy, as well.

³See Smith (1994); Scanlon (1998); Hampton (1998); Raz (1999); Dancy (2004b); Schroeder (2007); Parfit (2011). For skepticism, see Broome (2004) and Väyrynen (2010).

⁴Of course, these considerations really only motivate taking some single notion to be normatively basic. Some philosophers like Moore (1903) and Finlay (2014), for example, prefer to take the notion of goodness or value as normatively basic, and explain other notions, including reasons, in

Reasons first theorists owe us explanations of all other normative notions in terms of reasons. In particular, they owe us accounts of overall concepts like requirements—which reasons are meant to replace at the fundamental normative level—in terms of reasons. What exactly this burden involves will depend on at least two factors. First, it will depend on the extension of ‘normative’. For example, the normative is often distinguished from the evaluative, so we may hope to explain the normative in terms of reasons, but not the evaluative. The notions I discuss here fall squarely on the normative side, however. It is also an open question which kinds of requirements are normative. For example, requirements of etiquette are often cited as examples of non-normative requirements. And there has been much recent debate about whether the requirements of rationality are normative. The examples I will appeal to will be primarily from the domain of morality, which most people—or at least most reasons first theorists—take to be normative.⁵

Second, what the burden involves will depend on what the reasons first theorist takes the relevant kind of explanation to be. It may be that reasons are conceptually or explanatorily prior to other normative notions, or it may be that they are metaphysically prior, so that other normative notions are grounded in facts about reasons. I will try to remain neutral on this question here. I will assume that what the reasons first theorist needs to do, at the very least, is give us biconditionals with claims about when a normative notion (other than reasons) holds on the left and claims about reasons on the right, where these biconditionals are understood to involve some kind of priority on the right. In the next section I will give a straightforward example that will also serve to set up the rest of the paper.

2 Reasons and Ought

Perhaps the most straightforward case for the reasons first theorist is the notion of *ought*.⁶ Some principle like the following is very widely accepted:

Ought Most Reason (OMR): *s* ought to *A* iff *s* has more reason for *A* than for any other alternative.

Even those who reject the reasons first approach can accept this principle, as long as they deny that the right hand side has the relevant kind of priority. We might follow Broome (2004) and use something like **OMR** to instead provide an analysis of

terms of goodness or value. So we need other arguments—which I won’t discuss here—to really motivate a reasons first approach over a value first approach.

⁵Some philosophers—e.g., Foot (1972) and Williams (1981)—have famously denied that morality is (necessarily) normative. Whether or not these philosophers subscribed to any kind of reasons first program, then, they are not among my targets here.

⁶I will follow convention and nominalize the verb ‘ought’ to talk about oughts, as well as the notion of ought; when you say that someone ought to do something, you have ascribed an ought and used the notion of ought.

reasons in terms of ought. But again, here, my focus is on the reasons first theorist who does understand **OMR** with this kind of priority of reasons built in.

This principle captures a plausible reading of the platitude that you ought to do what's best, without the consequentialist baggage, since we can just understand 'best' in the platitude as meaning 'most reason', which need not involve (only) consequentialist reasons. This highlights the crucial feature of **OMR**, for my purposes: it involves a kind of *maximization* (again, in a not-necessarily-consequentialist sense). What you ought to do is what you have *most* reason to do. We can imagine a ranking of actions in terms of how much reason there is supporting them. It will be generated by looking at the reasons for and against each of the actions and combining them in some way, which is likely to be quite complicated.⁷ What you ought to do is the action at the top of this ranking. **OMR** gives us an example of the kind of thing the reasons first theorist needs to provide for all normative notions.

3 Ought, Requirement, and Permission

'Ought' is, or at least was, often treated as the central term in moral philosophy. Moral principles and rules were stated using 'ought', moral obligations were ascribed using 'ought', and deontic logicians use a big 'O', for 'ought', as their deontic necessity operator. In fact deontic logic—which deals with deontic necessity and possibility—is sometimes described as “the logic of *ought*”.

But more recently philosophers and linguists have emphasized that this is a mistake. 'Ought' is too weak for these purposes. Rules, requirements, and deontic necessity more generally, are much more naturally expressed using words like 'must' and 'have to'. 'Must'/'have to' and 'ought' are not synonymous, as evidenced by the felicity of claims like the following:

- (1) You ought to help with the dishes, but you don't have to.
- (2) You ought to get here on time; in fact, you must!

These sentences also show that 'ought' does not entail 'must'—it just isn't true, in general, that you must do what you ought to do. If 'ought' entailed 'must', then (1) would be contradictory, which it clearly isn't. And (2) would also be inappropriate, since the second half would be redundant. But (2) is perfectly fine.

On the other hand, consider the following sentences:

- (3) # It's not as if you ought to help with the dishes, but you must.
- (4) # You have to park in your assigned spot; in fact, you ought to!

⁷See Chapter 3 of Dancy (2004b) for a discussion of the complicated ways in which reasons can interact and combine. See also Chapter 7 of Schroeder (2007) and Horty (2012).

While (1) is not contradictory, (3) does sound contradictory. And while (2) is not redundant, (4) does sound redundant—the second half of (4) doesn't seem to be adding any new information. So whereas 'ought' does not entail 'must', the best explanation of this pattern is that 'must' does entail 'ought': you ought to do what you must do.⁸ Switching now to talking about requirements, which are ascribed using 'must': you ought to do what you are required to do, even though you aren't in general required to do what you ought to do.

Similar kinds of considerations show where 'may' (on the linguistic side) and permission (on the non-linguistic side) fit into this picture. Whatever you ought to do or are required to do, you are also permitted to do ('ought' and 'must' each entail 'may'). But in general you are not required to do what you are permitted to do, nor is it true in general that you ought to do what you are permitted to do ('may' entails neither 'must' nor 'ought'). So we have the following:

Strength: requirement \Rightarrow ought \Rightarrow permission

This is the picture we get from thinking about the logical relationships the deontic modals ('ought', 'must', and 'may', and their synonyms) stand in to one another.

The relationships captured in **Strength** mirror the relationships between some other philosophically interesting concepts, in particular, quantificational and epistemic modal concepts:

Quantificational Strength: all \Rightarrow most \Rightarrow some

Epistemic Modal Strength: must \Rightarrow probably \Rightarrow might

One important feature of these concepts is that using the terms that express the weaker notions has scalar implicatures that the stronger notions do not hold, due to Grice (1989)'s maxim of quantity.⁹ For example, if I tell you that most of the students passed, that implicates that it is not the case that all of the students passed. And if I tell you that the student probably passed, that implicates that, given my information, it isn't true that the student must have passed—there's a chance that she did not. Similarly, using 'ought' instead of something stronger, like 'must' or 'have to', implicates that the subject is not required to perform the given action. Consider the following (non-moral) case from Paul McNamara:

Imagine that we drive to work together and there are two routes we often take, respectively, involving two exits on a certain highway. I, but not you, know that a certain bridge on the route involving the second exit is closed for repairs this week. As a result, the route involving the first exit is the only acceptable route today. At the last minute, as we come

⁸See Portner (2009), pp. 32-35 and pp. 79-81 and Silk (ms) for nice discussions of the relative strengths of 'ought' and 'must'. See also McNamara (1996a,b,c), and Wertheimer (1972) for an early and especially prescient discussion.

⁹Thanks to Mark Richard for encouraging me to consider this issue.

to Exit 1, I remember and say: “You *ought* to turn here”. What I have said is true of course, but it is also inappropriate. For by saying “You *ought* to turn here”, I leave open the possibility that the second route (the one you were intending to take) is still an acceptable alternative. I have thus misleadingly suggested that the first exit is merely preferable and hence optional. In contrast, had I said “You *must* turn here”, there would have been no such suggestion of optionality. (McNamara (1996c), emphasis in original.)

The implicature is admittedly not as strong as in the quantificational and epistemic modal cases, but this is plausibly due to special features of ‘must’—namely, that it carries a kind of authoritative force, which may make it inappropriate in cases in which the speaker has no authority over the audience.¹⁰


Before moving on, I should make one important qualification to forestall worries about the linguistic data I have been appealing to. We have to read the ‘oughts’ and ‘musts’ as having the same kind of source or “flavor”; it may be, for example, that legally, you must stop at the red light, but all things considered, or even just prudentially, you ought not stop at the red light, since you have to make it to the hospital as soon as possible. But when we keep the sources the same, I think, we get the results I reported above.¹¹

Return to the idea of a ranking of actions I mentioned above, and consider the diagram below. What you ought to do is the action at the top of the ranking. Now we can complicate the picture by adding a permissibility bar: above the line are the permissible actions, below the line are the impermissible actions. As long as the top-ranked action is above the line (i.e., as long as there is some permissible action), ought will entail permission. Now suppose we have a ranking in which only one option is above the line—it is the only permissible option. Well, to say that an option is the only permissible one is to say that it is required, since all the others are impermissible. If there were more than one permissible option, then none of them would be required. So if an option is required, it is above the line, which means it is permissible: requirement entails permission. And if an option is required, it is the only option above the line, which means it must be the top-ranked option. So requirement also entails ought. But since an option can be top-ranked without being the only option above the line, ought does not entail requirement.

In addition to the linguistic evidence, this picture is also supported by some features of commonsense morality. McNamara (1996a,b,c) argues for this kind of

¹⁰See Silk (ms), sections 3 and 4.

¹¹Another interesting issue regarding the sources of the oughts and requirements is that, while distinguishing between what you ought to do and what you are required to do is very natural in the moral case, it is much more controversial in the case of epistemic and practical rationality. This is a difference that should be explored and explained, but it goes beyond the scope of this paper. See the papers in Byron, M. (editor) (2004) for discussion of rational supererogation and satisficing, which (as we will see shortly) is closely connected to the ought/requirement distinction.



structural picture by showing how easily it can accommodate various moral categories, including the supererogatory and the “least you can do”. The notion of the least you can do is a familiar one from commonsense morality. It’s what we do when we aren’t willing to go above and beyond, but rather just want to get by, morally speaking. For example, sending an e-card to your mother on her birthday may be the least you can do. On this picture, the least you can do is the lowest-ranked permissible option—it is where the bar is set. You are permitted to do this option, though you really ought to do better.

Supererogation is an even more familiar notion, at least within moral philosophy. You supererogate when you perform an action that’s higher-ranked than the least you can do. You are not required to perform supererogatory actions, by definition. But, at least for the top-ranked ones, we do want to say that, nevertheless, you really *ought* to. You don’t have to donate your bonus to charity, but you ought to.

There is good evidence from both normative discourse and commonsense morality supporting this structural picture. For the remainder of the paper, I will assume that this picture accurately captures the structural relationships between ought, requirement, and permission. What this means is that the reasons first theorist owes us some account of requirement in terms of reasons that (i) distinguishes requirement from ought, and (ii) allows us to explain these structural relationships.

4 A Simple Attempt

You ought to perform some action when you have more reason for it than you do for any other alternative. As we have seen, saying you ought to do something is weaker

than saying you are required to do it. This suggests that to explain requirements in terms of reasons, we should *strengthen* **OMR**. Here is a simple idea:

Required A Lot More Reason (RLMR): s is required to A iff s has a lot more reason for A than for any other alternative.

When you have most reason for some action, you ought to do it, though you may not be required to do it. But if you keep getting more and more reasons for it, according to this principle, eventually you will be required to do it. For example, suppose it's my mother's birthday. So I ought to go visit her, but I may not be required to—I could just give her a call, instead. But suppose that no one else will be around to spend her birthday with her, that she's had a hard time lately, and it would pull her out of her depression if I were to go visit. Suppose also that I don't really have any good reason not to go visit—I can afford it, have the time, and so on. It's plausible that, if we specify the situation in this way, I really am morally required to go visit her. And “specifying the situation in this way” just involves adding more and more reasons to go (and removing potential reasons against going).

This account also captures the structural relationships between oughts and requirements. First and most obviously, ought and requirement turn out to be distinct: having more reason to do something is not the same as having a lot more reason to do it. We also get the one-way entailment from requirement to ought: if you have a lot more reason to do something than to do any alternative, then you have more reason to do it than to do any alternative; so requirement entails ought. But you can have more reason to do something without having a lot more reason to do it; so ought does not entail requirement.

Nevertheless, I think **RLMR** is probably false. Consider first the left-to-right direction: if s is required to A , then s has a lot more reason for A than for any other alternative. Suppose that someone has left \$1000 in a bank account and set up the following mechanism. If you press button A, all \$1000 will be transferred to Oxfam; if you press button B, \$999 will be transferred to Oxfam, and the remaining \$1 will be burned.¹² I take it that you are morally required to press button A. But you do not have a lot more reason to press A than to press B, since it is only a difference of \$1. So the left-to-right direction of **RLMR** is false.

One potential response is to say that the fact that you are required to press A itself gives you a strong reason to do so, so that in fact you do have a lot more reason to press A than to press B. But this response is not open to the reasons first theorist, since this would be to explain reasons in terms of requirements. Similarly, it does not appear to be open to the reasons first theorist to claim that you have very strong reasons to press A provided by the fact that pressing B instead would display a bad character, or something like that. The notion of bad character seems to be a normative notion; if so, then the reasons first theorist is committed to explaining

¹²This case is a variation of one presented by Portmore (2011), pp. 125–126.

bad character in terms of reasons, and not vice versa.¹³

Now consider the right-to-left direction of **RLMR**: if you have a lot more reason to *A* than to do any other alternative, then you are required to *A*. Many cases of supererogation are counterexamples to this claim. By definition, you are not required to perform a supererogatory action, but you often do seem to have much more reason to do so than to do anything else—supererogatory actions often involve comparatively small self-sacrifice in order to achieve great goods (donating your disposable income to charity, risking injury to save lives, etc.). So the right-to-left direction of **RLMR** is also false.

RLMR is admittedly a naive principle, so perhaps it is no surprise that it fails. But the trouble is that it is just not clear what else the reasons first theorist should say. **RLMR** is the only obvious account of requirement, as opposed to ought, in terms of reasons, at least if we stick to the template of **OMR**. In the next section I will introduce in a general way a strategy which has been implemented in different ways by several philosophers, which is meant to explain the difference between ought and requirement, while explaining both in terms of reasons.

5 Two Kinds of Reasons

OMR appeals to ordinary, first order reasons for the action in question: you ought to do *A* when the reasons for *A* outweigh the reasons for anything else. **RLMR** was the only obvious way to extend this kind of approach to requirement, but as we saw, it is problematic. The idea behind the strategy I introduce in this section is straightforward: in order to explain requirement, as opposed to ought, in terms of reasons, we should appeal to a different kind of reason. So one kind of reason generates oughts, and another generates requirements, and also permissions, on some versions of this strategy. This is the *Two Kinds of Reasons* strategy.

To introduce the Two Kinds of Reasons strategy, I will discuss a view which can be distilled from Dancy (2004a). Dancy is interested in explaining the fact that there are often several permissible options open to us, even when one particular option has more reason in its favor than the others.¹⁴ So we can read Dancy as trying to explain why oughts do not entail requirements. A difficulty with reading Dancy this way is that he conflates oughts and requirements, and uses ‘ought’ to talk about what I’ve been calling requirement. But I think we can ignore that, and just focus on the structural features of Dancy’s strategy.

So here is the simplest and most abstract version of the Two Kinds of Reasons strategy:

¹³The reasons first theorist may deny that bad character is a normative notion, or come up with some other non-normative notion along these lines, to defend the left-to-right direction of **RLMR** in this way. I will not pursue this here, since the right-to-left direction fails, anyway.

¹⁴This is what Raz (1999) calls the “basic belief”. Dancy’s strategy is discussed and ultimately rejected by Raz. Part of Dancy’s aim is to defend the strategy from Raz’s objections.

Ought O-Reasons: s ought to A iff s has more O-reason for A than for any other alternative

Required R-Reasons: s is required to A iff s has more R-reason for A than for any other alternative¹⁵

This gives us the structure of a Two Kinds of Reasons account, but what we need is some kind of story about what O-reasons and R-reasons are. Different theorists provide different accounts; I consider two in the next two sections. Moreover, as we will see, not all advocates of the Two Kinds of Reasons strategy will accept the simple “maximizing” picture encoded in these two principles.¹⁶

Dancy calls R-reasons ‘peremptory’ and O-reasons ‘enticing’. So R-reasons tells us “you’d better do this” while O-reasons tell us “come on, I’d be nice”, on Dancy’s characterization. Peremptory reasons can generate requirements—you are required to do what you have most peremptory reason to do. Enticing reasons, on the other hand, do not generate requirements. They generate what Dancy calls “bests”, but I will just say that they generate oughts. A complication here is that Dancy’s characterization of enticing reasons is too narrow to explain the full range of cases that we’re interested in. He says that enticing reasons show why an option would be nice or pleasant. But for the purposes of this paper, reasons which generate oughts but not requirements should also be in play in cases of supererogation. These are more heavyweight than the sorts of reasons Dancy has in mind.

Nevertheless, Dancy’s account does make clear the structure behind the Two Kinds of Reasons strategy. It also illustrates its most important advantage, namely, that it distinguishes between oughts and requirements. In the next two sections I will discuss two implementations of this strategy.

6 Bedke’s Millian Inversion

First I will discuss a version of the Two Kinds of Reasons strategy proposed by Bedke (2011). After pointing out that we cannot understand requirement and permission in terms of the same kind of reasons that ground oughts, Bedke offers a “Millian inversion”, so-called because of its similarity to a view defended by John Stuart Mill (1861).

¹⁵Given **Required R-Reasons**, we could appeal to the duality of requirement and permission (compare necessity and possibility, or the universal and existential quantifier) to get an account of permission: s is permitted to A iff it is not the case that s is required not to A .

¹⁶Gert (2007), in particular, is very hard on this idea: “when one appreciates the nature of the two kinds of normative strength, it will become clear that maximizing is not really a coherent goal, that the general advice ‘Act on the stronger reasons’ is typically quite useless and confused, and that phrases such as ‘the balance of reasons’ or ‘what there is most reason to do’, even taken as metaphorical, are so misleading that they ought never to be used” (535). Gert is complaining about speaking in these ways *without* distinguishing between two kinds of reasons. But as we will see, he rejects the maximizing feature of **Required R-Reasons** even once R-reasons are distinguished from O-reasons.

In terms of the dialectic of this paper, Bedke gives us a substantive theory of what R-reasons are. O-reasons are just ordinary practical reasons for agents to perform actions. Thus, Bedke accepts **OMR** from above. R-reasons, on the other hand, are not reasons for the agent to perform actions, but rather reasons for people (including the agent herself) to perform certain speech acts, or at least speech act analogs.¹⁷

Reasons to Require: s is required to A iff there is most reason to require s to A

Reasons to Permit: s is permitted to A iff there is most reason to permit s to A

According to this proposal, requirements and permissions can indeed be understood in terms of reasons; it is just that they are not to be understood in terms of reasons for the agent to perform the action she is required or permitted to perform. Rather, they are explained in terms of reasons (for all of us, including the agent herself) to respond to the agent in certain ways.¹⁸

This proposal has a lot going for it. As Bedke illustrates, it gives us a powerful framework that lets us understand not only requirement and permission in terms of reasons, but also several other normative notions. Second, the view is supported by the observation that the notions of requirement and permission correspond to the speech acts or responses Bedke utilizes: requiring and permitting. And third, the view draws on a venerable idea from moral philosophy, that many normative notions can be understood in terms of appropriate responses to the behavior of others.

Nevertheless, this account is problematic. Bedke's account is unable to explain the relationships between ought, requirement, and permission captured in **Strength**—namely, that there is a one-way entailment between requirement and ought, and one-way entailments both between requirement and permission, and between ought and permission. The problem is that we have no reason to think that the kinds of reasons that ground the different normative notions stand in the appropriate relationships.¹⁹

First, consider the one-way entailment between requirement and ought. To capture this on Bedke's view, it would have to be true that (i) whenever there is most reason to require you to A , you have most reason to A , and (ii) there are cases in which you have most reason to A , though there is not most reason to require you to A . But we just have no reason to think this is true. For example, it is plausible that there are cases in which we may have most reason to require someone to perform

¹⁷Bedke calls them speech acts; this doesn't seem wholly satisfactory to me, since we need not actually say anything to perform the acts in question.

¹⁸It may be in one way misleading to call this a Two Kinds of Reasons view, since on one level, we really just have reasons that favor things in the same way. It is just that some of the reasons are reasons for the agent to do certain things (the action the agent herself ought to perform), and the other reasons are reasons for others to do something else (respond to the agent in certain ways). Nevertheless, I take the underlying motivation to be the same.

¹⁹Since writing this paper, I've discovered that Alex Silk points out this kind of problem with Bedke's proposal in Silk (2012).

some action, even if, perhaps because of their idiosyncratic desires or values, they have most reason to do something else. Even if we can set aside these kinds of cases, the point is just that we need some explanation for why ordinary first order reasons for the agent to perform an action will correspond in just the right ways with the reasons to require that agent to perform the action.²⁰ In fact, Bedke insists that “reasons for requiring *s*’s *A*-ing are conceptually independent of *s*’s reasons for *A*-ing”. So it is hard to see how Bedke could capture this entailment relation.

Relatedly, we have seen that that oughts entail permissions, but not vice versa. For this to be true on Bedke’s account, it must be true that (i) whenever there is most reason for *s* to *A*, there is also most reason to permit *s* to *A*, and (ii) there are some cases in which there is most reason to permit *s* to *A*, though there is not most reason for *s* to *A*. As before, we simply have no reason to think this is true. It is not obviously false, of course. But it is also not obviously true—it seems an open possibility that we could have very good reason not to permit someone to do something that she has most reason to do.

Finally, even if we just stick with reasons for speech acts, or speech act analogs—reasons to require and reasons to permit—it is not clear that Bedke’s account can capture the right entailments. In particular, we have seen that there is an entailment from requirements to permissions. Thus, it would need to be true that whenever there is most reason to require someone to *A*, there is *also* most reason to permit her to *A*. This is *prima facie* problematic, since requiring and permitting are two different speech acts; so it is puzzling how there could be *most* reason to do *both* of them.

In this case, a possible response is to provide an analysis of the speech acts of requiring and permitting such that requiring necessarily or constitutively involves permitting, but we do not have such an analysis yet. Moreover, seeing what would be required to capture this entailment makes it even harder to see how to explain the entailments involving oughts, since the reasons that generate oughts are of a totally different kind than the reasons for speech acts that generate requirements and permissions. We cannot appeal to constitutive relationships between speech acts to capture the entailments involving ought.²¹

7 Gert’s Implementation

Next I will discuss a more complicated version of the Two Kinds of Reasons strategy inspired by Gert (2004, 2007). As we’ll see, the reason this account is merely *inspired by* Gert is that he does not explicitly tell us what his account of ought, as opposed

²⁰Bedke restricts the reasons to require and permit to the *right kind* of reasons for performing these speech acts, which does rule out some cases like this. But it is far from obvious that it rules out all of them.

²¹Bedke considers objections to his account based on various entailment relations. His strategy is to argue that the entailments do not in fact hold. Even if these arguments are successful, though, Bedke does not consider the entailments I have discussed here.

to requirement, is. This is not necessarily a criticism of Gert, though, since he is not actually a reasons first theorist. Moreover, Gert is primarily interested in rational requirements, as opposed to moral requirements. Even though we should retain the conceptual distinction between what you are required to do and what you ought to do across normative domains, it is much more controversial that there is a *substantive* distinction between what you are rationally required to do and what you rationally ought to do. There is much debate about whether the notion of rational, as opposed to moral, supererogation even makes sense.²² So it is not surprising that Gert did not have the distinction between ought and requirement in view.

So I will not be discussing Gert's own view explicitly. Nevertheless, his picture does seem like a promising route for the Two Kinds of Reasons theorist, since he has a sophisticated theory that distinguishes between two kinds of reasons. Here I will attempt to pair his account of requirement with the most obvious account of ought, given the distinction between two kinds of reasons he proposes.

Gert is interested in explaining why (i) supererogatory acts are not required, often because of the degree of self-sacrifice involved, but (ii) they are (rationally) permitted, despite the self-sacrifice involved. To explain this, he appeals to two kinds of normative reasons: justifying reasons, which make an action that would otherwise be impermissible, permissible, and requiring reasons, which make an action that would otherwise not be required, required.

To be more precise, Gert is interested in distinguishing two dimensions of normative strength that a normative reason may have, rather than two kinds of normative reasons. Nevertheless, I will discuss his view here as distinguishing between two kinds of reasons. Given that the requiring strength and justifying strength of a reason are conceptually independent of one another, on Gert's view, this should not lead to any serious distortion. Since the requiring strength of a reason has nothing to do with its justifying strength, it should be fine to simply talk about it as two different reasons: a requiring reason and a justifying reason.

Gert offers an account of the comparative strength of two reasons along each of these two dimensions—requiring strength and justifying strength: a reason $r1$ has more requiring/justifying strength than a reason $r2$ if, when playing the requiring/justifying role in any circumstance, $r1$ can overcome any reasons that $r2$ can overcome, and more.²³

The most obvious view of requirement and ought appealing to this distinction is the following:

Simple Required: s is required to A iff s 's requiring reasons to A outweigh s 's requiring reasons to do anything else

²²See, for example, Slote (1989); Dreier (2004); van Roojen (2004). It is similarly unclear whether the notion of *epistemic* supererogation makes sense. Of course, there is a large literature about how to make sense of moral supererogation, but it is generally taken for granted that there must be some way of doing so.

²³Gert (2007), pp. 538-539.

Simple Ought: s ought to A iff s 's justifying reasons to A outweigh s 's justifying reasons to do anything else

This would be to stick with the simple maximizing picture, encoded in **Ought O-Reasons** and **Required R-Reasons**.

However, this is not the view Gert adopts. He gives an account of requirements like this:

Gert Required: s is required to A iff s 's justifying reasons for all of the alternatives other than A are not strong enough to compensate for acting against the requiring reasons for A ²⁴

This notion of “compensation” should be more fully explained. We can start to get the idea by thinking in terms of criticism. Requiring reasons serve to rule actions out as impermissible—requiring reasons for A tend to rule out alternatives incompatible with A . Justifying reasons, on the other hand, serve to rule actions in as permissible. Requiring reasons for an action form the basis for criticism of other actions; justifying reasons for those other actions answer those criticisms.²⁵ This means that agents may not actually be required to perform the action for which they have most requiring reason—there may be sufficient justifying reasons for the other alternatives to rule them in as permissible. So Gert rejects the maximizing account, **Simple Required**. More could be said here, but the crucial point, as Gert stresses, is that requiring and justifying reasons are conceptually distinct. Tending to require an action by providing potential criticisms of alternatives is very different than tending to answer potential criticisms leveled at the relevant action.

Since Gert rejects **Simple Required**, it seems likely that he would also reject **Simple Ought**. The trouble, as I mentioned above, is that Gert does not actually discuss ought as opposed to requirement. Thus, despite the important role that supererogation plays in motivating his theory (see especially Gert (2004)), Gert doesn't give us a way to say that you *ought* to perform the supererogatory actions, even though you are not required to do so.²⁶ All we can say about supererogatory actions, on Gert's view, is that there are strong justifying reasons in their favor. This doesn't tell us much about their overall normative status. If a reasons first theorist wants to appeal to a Gert-style view, she will need to give us some explanation of ought in terms of reasons.

²⁴This is more general than the account that Gert explicitly offers, on which requiring reasons are identified with prudential reasons and justifying reasons are identified with altruistic reasons. But as Gert stresses, the important thing for his purposes in his (2007) is the structure of the account, which I've tried to highlight.

²⁵See Greenspan (2005) for another version of this kind of view. Greenspan presents the ideas primarily in terms of giving and answering criticisms. Gert focuses on the R-reasons for an action, and holds that they tend to rule out alternatives. Greenspan instead focuses on R-reasons *against* an action, which tend to rule it out. So, some complications aside, Greenspan and Gert are coming to the same view, just from different directions.

²⁶Bedke (2011) points out this gap in Gert's theory.

Even though Gert may not accept **Simple Ought**, I propose to see how it fares, since it is a particularly natural account, given the distinction between requiring and justifying reasons, and since we lack any other account. Moreover, since supererogatory actions are meant to be supported by strong justifying reasons, but not by requiring reasons, and since these are actions that you ought but are not required to perform, this is a position worth exploring.

According to **Simple Ought**, an agent ought to *A* when she has most justifying reason to *A*. Justifying reasons for an action serve to answer potential criticisms of the action. The more there is to be said in favor of an action, the more potential criticisms it can answer. And if there is more to be said in favor of one action than any other—if it can answer any criticisms the others can answer and more—you ought to perform that action. This line of thought leads very naturally to **Simple Ought**. Moreover, adopting this principle in addition to **Gert Required** correctly distinguishes between oughts and requirements, and does not mistakenly give us an entailment from ought to requirement. Since an agent can have most justifying reason to *A* even if the requiring reasons for *A* are not sufficient to ground a requirement, since justifying and requiring strength are independent, ought does not entail requirement, just as we want.

But we also want to capture the entailment from requirement to ought. This is where we get trouble, just as we did with Bedke's account above. Gert emphasizes that the strength of the requiring reasons for *A* and the strength of justifying reasons for *A* do not covary. It is possible for there to be reasons of equal justifying strength for *A* and for *B*, but different requiring strength; conversely, it is possible for there to be reasons of equal requiring strength for *A* and for *B*, but different justifying strength. And it is also possible for there to be stronger requiring reason for *A* than for *B*, but stronger justifying reason for *B* than for *A*.²⁷ Since requiring reasons and justifying reasons can come apart in this way, we can expect cases in which the justifying reasons for one alternative are stronger than the justifying reasons for another, though still not strong enough to answer the requiring reasons.

This means that on this view, requirement does not entail ought. Suppose the options open to the agent are *A* and *B*, and suppose further that *A* is required. This means that the justifying reasons for *B* are not strong enough to compensate for acting against the requiring reasons for *A*. Still, given the sharp distinction between requiring and justifying reasons and their corresponding strengths, it is possible that there is most justifying reason for *B*; suppose this is the case. Thus, according to **Simple Ought**, you ought to do *B*. It is just that the justifying reasons for *B* are not sufficiently strong to compensate for acting against the requiring reasons in favor of *A*; so you are still required to *A*. Since you are required to *A* but ought to do *B*, requirement does not entail ought. So given **Simple Ought** along with **Gert Required**, we cannot capture the entailment from requirement to ought.

Even if we returned to the maximizing picture and adopted **Simple Required**,

²⁷Gert's headache/charity case on pp. 550-552 is meant to be such a case.

given the independence of the strengths of requiring and justifying reasons, we would not be able to capture the entailment. By distinguishing so sharply between the requiring and justifying reasons, we have undermined the entailment.

What would be needed here is for the fact that there is unanswerable requiring reason for an action to entail that there is strongest justifying reason for that action. But this violates the independence of these two dimensions of force that Gert insists on. The crucial thing to notice is that Gert's Two Kinds of Reasons picture is motivated by considering the sharp distinction between leveling criticisms and answering criticisms. Leveling a criticism at an option is a way of requiring that you do not perform it; answering that criticism is a way of justifying performing it. These are very different roles. Once we violate the independence of requiring and justifying reasons, it is not clear how we could maintain this motivation for making the distinction in the first place. For example, why should the strength of the criticisms we can level against alternatives to *A*—which correspond to the requiring reasons for *A*—have any bearing on whether we can answer potential criticisms of *A*-ing? So given his underlying criticism-based picture, it seems that Gert is right to insist on a sharp distinction between requiring and justifying reasons. The trouble is that this makes it very hard to explain the structural relationships between requirement and ought.

As I said above, if Gert were interested in explaining ought in terms of reasons, he may reject **Simple Ought**, given that he rejects **Simple Required**. But he does not tell us what he would accept, instead. What I have just argued is that the reasons first theorist cannot accept **Simple Ought**, either with **Gert Required** or with **Simple Required**. Neither combination lets us capture the entailment from requirement to ought. So we still lack an account of ought in terms of reasons on this Gert-style view. Later I will briefly consider another attempt in this vein.

8 Denying that the Entailments are Structural Features

In this section I want to consider a possible response on behalf of the Two Kinds of Reasons theorist. I have been arguing that this theorist will have trouble capturing entailments between requirements, oughts, and permissions. Throughout I have been assuming that these entailments should be captured by *structural* features of these concepts. In particular, since on the picture in question, reasons ground or explain these concepts, the relationships between them should be explained by relationships between the reasons that ground them—the entailments should be reflected in the analyses in terms of reasons. The Two Kinds of Reasons strategy makes it hard to see how to do this. But the defender of this strategy may simply reject the starting point: the so-called entailments are not actually entailments, or at least, if they are, they are due to *substantive* rather than *structural* features. So we could give a substantive theory about what the two kinds of reasons are such

that we get the entailments.²⁸

For example, Bedke could argue for the substantive claim that if an agent has most reason to *A*, then there cannot be most reason (for us) to require the agent to do something else. One potential strategy would be to argue that considerations of the agent's integrity, or concerns about alienation, would provide strong reasons not to require the agent to do something other than what she herself has most reason to do. Alternatively, we could argue that considerations that give us reasons to require the agent to *A* are always also reasons for the agent to *A*, and that these reasons are bound to outweigh any other reasons the agent has to do something else. Either way, we will get the result that whenever the agent is required to *A*, she ought to *A*.

Someone who accepts a Gert-inspired picture would have to give different arguments, since on this picture the kinds of reasons we have to work with are different. But the overall strategy would be similar.

While I think these are interesting proposals, I'm skeptical. These substantive claims will undoubtedly be controversial, and there will likely already be compelling arguments against them in the literature. This doesn't mean the claims are false, of course. But the evidence I presented for the entailments is compelling and widely accepted, and does not seem to rely on any such controversial substantive claims. For example, even if we think reasons to require an agent to *A* need not always be reasons for the agent to *A*, the linguistic evidence that 'required' or 'must' entails 'ought' is no less compelling.

Some coherent views seem to deny the entailment from requirement to ought. For example, suppose we think that moral requirements are provided by something like the Ten Commandments or the law of the land, but that reasons depend for their existence on certain motivations within the agent. And suppose we also accept **OMR**, that what you ought to do is what you have most reason to do. On this view, requirements do not seem to entail oughts.²⁹ If these views are coherent, then it looks like the so-called entailments simply cannot be structural, and instead must be substantive, if they hold at all.

But remember that I emphasized above that we need to make sure to keep the *flavor* of the oughts and requirements—moral, rational, legal, epistemic, etc.—the same in examining potential counterexamples to the entailments. While it is easy to imagine accepting a view like the one described above according to which *moral* requirements come from some external source while *rational* or *all things considered* oughts come from reasons provided by the agent's own motivations, it is harder to imagine a view according to which moral requirements come from the external source but *moral* oughts come from these internal reasons. So we don't yet have a view according to which the entailments fail.

²⁸Thanks especially to Matthew Bedke for pressing this suggestion.

²⁹This is not a Two Kinds of Reasons, or even Reasons First view, but that doesn't matter: all that matters is that it is a coherent view according to which requirement does not entail ought.

9 A General Challenge, and a Possible Way Forward

At bottom, the Gert-style view and Bedke's view face the same problem: by sharply distinguishing between reasons that explain oughts and reasons that explain requirements and permissions, these views make it very hard to see how to capture the structural relationships between these normative notions. This is not surprising, though, given that these are implementations of the Two Kinds of Reasons strategy. The thought behind this strategy is correct: we need to distinguish oughts and requirements, and the obvious way to do it—**RLMR**—does not work. But what I have been emphasizing here is that we need to be careful in distinguishing between oughts and requirements: though these are distinct normative notions, they do stand in certain structural relationships to one another. The Two Kinds of Reasons strategy can easily pry them too far apart.

This is not to say that no version of the Two Kinds of Reasons strategy can explain the structural relationships. In this section I will first introduce a very schematic Two Kinds of Reasons picture that can explain them, and then briefly discuss two possible implementations. Both of these face challenges. Nevertheless, given that they can at least explain the structural relationships between oughts, requirements, and permissions, they may be promising ways forward for the Two Kinds of Reasons theorist, and for the reasons first theorist more generally.

The general idea behind each of the views I sketch below is to appeal to *both* of the two kinds of reasons at once in the account of requirement. As we saw, the Gert-style view above did this—it appealed to both requiring and justifying reasons (or strength). The trouble was that it did so in a way that lost the one-way entailment from requirement to ought. Here is a very schematic picture which retains the entailment:

OMR: s ought to A iff s has more O-reason for A than for any other alternative

Required Both: s is required to A iff (i) s has more O-reason for A than for any other alternative, and (ii) s has more R-reason for A than for any other alternative

It is easy to see that ought will not entail requirement, but that requirement will entail ought, just as we want. We can flesh this out in different ways by giving different accounts of what the R-reasons are; I will assume that O-reasons are just ordinary practical reasons for agents to perform actions. First I will discuss a view which makes use of Gert's distinction between requiring and justifying reasons, and which is a simplified version of a view defended by Portmore (2011). Then I will discuss a view which makes use of Bedke's distinction between reasons for actions and reasons for responses.

9.1 Required Both: Portmore-inspired

One thing we can take away from discussion of the Gert-style view in section 7 is that capturing the entailment from requirement to ought on a criticism-based, rather than maximizing, picture is tricky. This is partly because, as long as we accept **OMR**, the concept of ought has maximization built in—what you ought to do is what you have *most reason* to do. This kind of concept doesn't fit very comfortably into Gert's criticism-based picture. So now I want to sketch a view which makes use of Gert's distinction between requiring and justifying reasons, but retains the maximizing structure of **Required Both**. A more sophisticated version of this view is defended by Portmore (2011).

This picture is most natural in the case of morality. I have been focusing mainly on morality throughout, but here I will make it explicit; it is an interesting question how this view would extend to other normative domains. Here is the picture:

OMR: s morally ought to A iff s has most morally requiring reason to A .

Portmore Required: s is morally required to A iff (i) s has most morally requiring reason to A , and (ii) s has most reason, all things considered, to A .

This view draws on the basic idea of **Required Both**, and builds in the entailment from required to ought straightforwardly. The basic idea behind this view is the following. Sometimes, you are not morally required to perform the action that you have most moral reason—or morally ought—to perform. The most obvious cases are cases of supererogation. According to this view, this is because you have more reason, all things considered, to do something other than what you morally ought to do. So though the non-moral reasons—which make up the set of all things considered reasons, minus the moral reasons—cannot generate moral requirements, they can *justify* you in not performing the act that you have most moral reason to perform.

It is important that the category of reasons, all things considered, which play the justifying role on this picture, includes the moral reasons. So we cannot identify the set of reasons all things considered with the set of non-moral reasons. If the moral reasons were not a subset of the reasons all things considered, we would always be morally permitted to perform the action that we have most *non-moral* reason to perform. But in many cases, this would be a bad verdict. For example, suppose there are strong moral reasons, but no non-moral reasons, in support of keeping your promise. On the other hand, there are some relatively lightweight non-moral reasons in favor of not keeping your promise. If the all things considered reasons referenced in clause (ii) of **Portmore Required** did not include the moral reasons, you would be morally permitted to break your promise. This is a bad prediction. But since the moral reasons are among the reasons all things considered, such a case will be ruled out: the strong moral reasons for keeping the promise will contribute to the strength of the reasons all things considered to do so.

This is where a potential challenge arises. Consider the following case, which seems to be a paradigm case of supererogation. You know you could save fifty lives by

donating \$500 to Oxfam, but doing so would prevent you from buying a new laptop to replace your perfectly functional, but outdated one. Clearly the moral reasons support donating the money. But for those of us who believe in supererogation, this is a case in which you are not required to donate the money, though you ought to. On the view under consideration, this means that you have more reason all things considered not to donate the money, and to buy the laptop instead. But given that the moral reasons are among the all things considered reasons, this is a surprising claim—after all, there are fifty lives at stake.³⁰ There is more to be said here, of course, but this is a challenge that the Two Kinds of Reasons theorist who takes this path must face.³¹

9.2 Required Both: Bedke-inspired

The final version of **Required Both** that I want to consider takes its cue from Bedke’s implementation of the Two Kinds of Reasons strategy:

OMR: s ought to A iff s has most reason to A

Reasons to Require*: s is required to A iff (i) s has most reason to A , and (ii) there is most reason to require s to A .³²

Once again, this view straightforwardly captures the one-way entailment from requirement to ought. It also gains other advantages of Bedke’s view. In particular, it draws on the plausible and venerable idea that some normative notions are intimately related to appropriate or fitting responses on our behalf.

A challenge for this view is to give some kind of independent motivation for the first condition of **Reasons to Require***. As it stands, the principle looks *ad hoc*. Further, unless some connection between reasons to require and reasons for the agent to perform the action in question is established—just the kind of connection Bedke’s original version needed—we will get strange results. In particular, it will be possible for there to be most reason to require an agent to perform some action, even

³⁰Portmore (2011) is willing to accept this kind of consequence, at least in a case in which what is at stake for the agent is not a new laptop but rather severe injury; see p. 139. The consequence is much less surprising in the case Portmore considers there, but if we accept a commonsense view of supererogation, where much less than severe injury can keep you from being morally required to save lives, this view commits us to consequences like the one in the main text.

³¹As I mentioned, this view is merely *inspired* by Portmore (2011). The view he defends there is actually much more sophisticated and complicated. For one thing, he distinguishes between the superperfecterogatory (actions which go beyond perfect duty) and the supererogatory (actions which also go beyond imperfect duty). Relatedly, he argues that only sets of actions, performed over an extended period of time (perhaps a whole lifetime) can properly be called supererogatory. Unfortunately, spelling out this view in any detail would take us far afield; see Portmore (2011), Chapter 7 for the details.

³²A variant of this view according to which we replace ‘reason to require s to A ’ with something like ‘reason to blame s for not A -ing’ has been suggested to me in conversation by both Walter Sinnott-Armstrong and Brian McElwee, and is a version of the venerable Fitting Attitudes theory.

though the agent is actually not required to perform the action. This can happen as long as there can be most reason to require the agent to *A*, even if the agent does not have most reason to *A*. In such a case, we ought to require the agent to *A*, even though the agent is not required to *A*. This may not be an absurd result, but it is at least counterintuitive.

10 Conclusion

The Two Kinds of Reasons strategy is motivated by the need to distinguish between oughts and requirements. I have argued here that advocates of this strategy must be very careful to recognize that, nevertheless, there are still relationships between these notions that must be captured. I have illustrated how challenging this is by arguing that two implementations of the strategy fail to do this. I also considered two versions of the Two Kinds of Reasons strategy—all of which begin from the same idea—that can explain the structural relationships. But both face significant, though not necessarily insurmountable, challenges. At the very least, the structural relationships between oughts, requirements, and permissions put serious constraints on the Two Kinds of Reasons strategy. The reasons first theorist, then, still owes us an account of requirement—as distinct from ought—and permission in terms of reasons.

References

- Bedke, M. (2011). Passing the deontic buck. In Shafer-Landau, R., editor, *Oxford Studies in Metaethics*, volume 6, pages 128–152.
- Broome, J. (2004). Reasons. In Wallace, Pettit, Scheffler, and Smith, editors, *Reason and Value: Themes from the Moral Philosophy of Joseph Raz*, pages 28–55. Oxford University Press.
- Byron, M. (editor) (2004). *Satisficing and Maximizing: Moral theorists on practical reason*. Cambridge University Press.
- Dancy, J. (2004a). Enticing reasons. In Wallace, Pettit, Scheffler, and Smith, editors, *Reason and Value: Themes from the Moral Philosophy of Joseph Raz*, pages 91–118. Oxford University Press.
- Dancy, J. (2004b). *Ethics Without Principles*. Oxford University Press.
- Dreier, J. (2004). Why ethical satisficing makes sense and rational satisficing doesn't. In Byron, M., editor, *Satisficing and Maximizing: Moral theorists on practical reason*, pages 131–154. Cambridge University Press.
- Finlay, S. (2014). *Confusion of Tongues*. Oxford University Press.

- Foot, P. (1972). Morality as a system of hypothetical imperatives. *The Philosophical Review*, 81(3):305–316.
- Gert, J. (2004). *Brute Rationality: Normativity and Human Action*. Cambridge University Press.
- Gert, J. (2007). Normative strength and the balance of reasons. *Philosophical Review*, 116(4).
- Greenspan, P. (2005). Asymmetrical practical reasons. In Reicher, M. and Marek, J., editors, *Experience and Analysis: Proceedings of the 27th International Wittgenstein Symposium*, pages 387–394. oebv&hpt.
- Grice, H. (1989). Logic and conversation. In *Studies in the Way of Words*. Harvard University Press.
- Hampton, J. (1998). *The Authority of Reason*. Cambridge University of Press.
- Horty, J. F. (2007). Reasons as defaults. *Philosophers' Imprint*, 7(3).
- Horty, J. F. (2012). *Reasons as Defaults*. Oxford University Press.
- McNamara, P. (1996a). Doing well enough: Toward a logic for common-sense morality. *Studia Logica*, 57(1):167–192.
- McNamara, P. (1996b). Making room for going beyond the call. *Mind*, 105(419):415–450.
- McNamara, P. (1996c). Must I do what I ought (or will the least I can do do)? In Brown, M. A. and Carmo, J., editors, *Deontic Logic, Agency and Normative Systems*, pages 154–173. Berlin: Springer-Verlag.
- Mill, J. S. (2001 [1861]). *Utilitarianism*. Hackett Publishing Company, Inc.
- Moore, G. (1903). *Principia Ethica*. Cambridge University Press.
- Parfit, D. (2011). *On What Matters*. Oxford University Press.
- Portmore, D. (2011). *Commonsense Consequentialism: Wherein Morality Meets Rationality*. Oxford University Press.
- Portner, P. (2009). *Modality*. Oxford University Press.
- Raz, J. (1999). *Engaging Reason: On the theory of value and action*. Oxford University Press.
- Ross, W. D. (1930). *The Right and the Good*. Oxford University Press.
- Scanlon, T. (1998). *What We Owe to Each Other*. Harvard University Press.

- Schroeder, M. (2007). *Slaves of the Passions*. Oxford University Press.
- Sidgwick, H. (1907). *The Methods of Ethics*. Hackett Publishing, 7th edition.
- Silk, A. (2012). Review of Shafer Landau, Russ (ed.), *Oxford Studies in Metaethics, Vol. 6. Ethics*, 122(3):622–627.
- Silk, A. (ms). Weak and strong necessity. University of Birmingham.
- Slote, M. (1989). *Beyond Optimizing*. Harvard University Press.
- Smith, M. (1994). *The Moral Problem*. Oxford University Press.
- Urmson, J. (1958). Saints and heroes. In Melden, editor, *Essays in Moral Philosophy*. University of Washington Press.
- van Roojen, M. (2004). The plausibility of satisficing and the role of good in ordinary thought. In Byron, M., editor, *Satisficing and Maximizing: Moral theorists on practical reason*, pages 155–175. Cambridge University Press.
- Väyrynen, P. (2010). A wrong turn to reasons? In Brady, M. S., editor, *New Waves in Metaethics*. Palgrave Macmillan.
- Wertheimer, R. (1972). *The Significance of Sense: Meaning, Modality, and Morality*. Cornell University Press.
- Williams, B. (1981). Internal and external reasons. In *Moral Luck*, pages 101–113. Cambridge University Press.