How to be a deontic buck-passer

Euan K. H. Metz¹

Published online: 31 October 2019
© The Author(s) 2019

Abstract *Deontic*, as opposed to *evaluative* buck-passing theories seem to be easier to accept, since there appears to be an intimate connection between deontic properties, such as ‘ought’, ‘requirement’, and ‘permission’ on the one hand, and normative reasons on the other. However, it is far from obvious what, precisely, the connection consists in, and this topic has suffered from a paucity of discussion. This paper seeks to address that paucity by providing a novel deontic buck-passing view, one that avoids the pitfalls both of the most straightforward view on the matter (what I call the “standard view”) as well as a recently articulated view, due to Matt Bedke. It does so by appealing first to the distinction between a reason for, and a reason against, and uses this distinction to clarify what are taken to be two fundamental, but distinct, deontic properties—*ought* and *requirement*. The resulting view allows us to capture these properties, the structural relations between them, and does so in a way that avoids making supererogation impossible.

Keywords Deontic properties · Reasons · Normative · Normativity · Buck-passing · Requirements · Oughts

1 Introduction

Recent work on normativity has seen an interest in extending the much-discussed buck-passing analysis of evaluative properties (such as good, bad, better, worse) to deontic properties (such as ought, permissible, required). According to a commonly expressed view, the all-in deontic property *ought* can be given the following analysis:

✉ Euan K. H. Metz
ekhm2@open.ac.uk

¹ Open University, Milton Keynes, UK

Standard view A ought to Φ if, and only if, A has most reason to Φ , where to have most reason to Φ is for the reasons which favour A's Φ -ing to be stronger than the reasons against A's Φ -ing

There are two important objections to the standard view. The first objection is that the standard view appears to make supererogation impossible. Suppose that it is supererogatory for A to give almost all of her annual income to charities providing aid to starving people. That is, although she would be praiseworthy if she did that, it is false that she *ought* to do it. Doing so would be 'above and beyond the call of duty'.

Given a typical way of understanding normative reasons, if the suffering of those that could be saved by A provides her with a reason to donate, then the greater the extent of that suffering, the weightier the reason she has. Thus, it is possible to imagine that this reason is strong enough to outweigh any other reasons that A has.

But if there can be a case in which A has most reason to donate almost all of her income to the charity, according to the standard view, she therefore *ought* to donate almost all of her income. But it is plausible that if anything is a supererogatory act, donating almost all of one's income to charity is a supererogatory act. Thus, the standard view is false because it falsely implies that A's donating almost all their income to charity is not a supererogatory act.

The second objection to the standard view is that it fails to capture two distinct senses of 'ought'. Sometimes, when you have most reason to Φ , that implies that Φ -ing is thereby no longer (rationally, morally, prudentially) optional for you. Consider:

Lunch A promises her friend B that she will meet B for lunch at 12 pm tomorrow, but A would rather spend the afternoon at home watching television. A knows that she has more reason to keep her promise than to satisfy her desire to watch television

According to the standard view, in *Lunch* A ought to meet B for lunch at 12 pm. That seems the right result, in the sense that, all else being equal, it is *because* A made a promise to B, that A's keeping that promise is not optional for A. It is *required* of A that, having made her promise, she keep it. But this isn't the only usage of 'ought' we are familiar with. Consider now:

Theatre A is deciding whether to go to the cinema, go to the theatre, or to stay home and read a book tonight. Although she has good reasons to choose any of these alternatives, she has a little bit more reason to go to the theatre, as she hasn't been there for months

According to the standard view, in *Theatre* A ought to go to the theatre. But it is at the very least an open question whether A is thereby *required* to go to the theatre just because she has a little bit more reason to go. We might more felicitously say that in *Theatre*, A's going to the theatre comes highly *recommended*, and is what one might *advise* A to do. The problem for the standard view is that, as it stands, it is

unable to discriminate between the *requirement* claim in *Lunch* and the *ought* claim in *Theatre*.

In this paper, I shall outline a better deontic buck-passing view than the standard view; one that avoids these two objections. To do so, I utilise the distinction between a reason *for* and a reason *against*. In Sect. 2 I shall argue for this distinction, and offer my own, positive, interpretation of that distinction, which diverges from the most prominent interpretation, due to Greenspan.¹ As I understand the distinction, we can distinguish a reason *for* from a reason *against*, not by appealing to criticism, but by appealing to *reactive attitudes*.

After having argued for and sketched this distinction, in Sect. 3 I use it to analyse the deontic properties of *ought* and *requirement* in terms of normative reasons. In Sect. 4 I argue that the resulting view can account not only for the distinction between the two deontic properties, but that it can capture the structural relations between these properties.² I argue that this gives my view an advantage over the other most prominent example of a deontic buck-passing view, due to Matt Bedke.³

In Sect. 5 I demonstrate how my view can resolve the supererogation problem, and in Sect. 6 I conclude.

2 Reasons for and reasons against

2.1 Motivating and arguing for the position

To resolve the problems outlined above for the deontic buck-passing view, I appeal to the distinction between reasons *for* and reasons *against*. Since few philosophers have explicitly discussed this distinction,⁴ and many philosophers implicitly accept an equivalence of some sort between these relations it behoves me to justify its use.

To argue that the distinction is a real one, my strategy is to identify what equivalences there could be between reasons *for* and reasons *against*, and to argue that those equivalences are untenable.⁵ Here are two tempting ways to understand the identification of reasons *against* with reasons *for*: we might (a) understand a reason *against* Φ -ing as equivalent to a reason *for* Ψ -ing (an alternative, or set of alternatives to Φ -ing); on the other hand, we might (b) understand a reason *against* Φ -ing as equivalent to a reason *for not* Ψ -ing.⁶

There are serious difficulties facing the first interpretation. First, this equivalence does not respect the plausible thought that normative reasons generally bear *only* on the response that they favour. For example, the fact that the play running at the

¹ See Greenspan (2004).

² This difficulty is articulated by Snedegar in Snedegar (2016).

³ See Bedke (2011).

⁴ For discussion see Gert (2003), Greenspan (2004), and Snedegar (2018).

⁵ Snedegar pursues a similar strategy in Snedegar (2018).

⁶ See Snedegar (2016, 2017) for discussion of these alternatives. Snedegar calls (a) the *opportunity cost principle*.

theatre on Monday has had dreadful reviews is thereby a reason not to go to see that play, but that same fact—that the play has had dreadful reviews—is not thereby a reason to, say, play golf.

To put more flesh on that idea, we tend to think that a normative reason plays a normative-cum-explanatory role. That the play has had dreadful reviews not only *disfavours* going, but also *explains why* it makes sense not to go. One fruitful way to understand how a normative reason can play that role is that, where a reason favours an action, it tells us something *about that action* that is worth doing. In order to play that justificatory-cum-explanatory role, the reason must bear on its object and not, directly, on other objects.⁷

This interpretation of the equivalence has not, however, gained much traction among normative theorists. The second interpretation of the equivalence, however, appears more intuitive. Suppose again, that the bad reviews of the play are a reason against seeing it on Monday. Is it not the case that the fact that the play has received bad reviews is a reason *for not* seeing the play on Monday? This purported equivalence seems much more secure, so we must spend more time dislodging this idea if the asymmetry of *for* and *against* is to be found acceptable.

The first doubt that I have about this version of the equivalence is that it is guilty of a failing of the first interpretation; namely, that it falls foul of the general thesis that reasons bear only on their object, not on alternatives to that object. Above, I questioned the idea that a fact *F* being against one action is generally equivalent to *F* being for another action. By the same token, it is doubtful that *F*'s being against an action is generally equivalent to that *F* being for an *omission*. Again, the driving force behind that idea is that a reason plays the justificatory-cum-explanatory role because it tells us what about *that action*, or about *that omission* is worth doing, or worth refraining from. But these are different objects, with different properties.

Consider the following illustration: 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 the second interpretation of the *for/against* equivalence, these reasons *both* function as reasons for the omission, and both function as reasons against the action. But that seems questionable at best.

Furthermore, the difference between act and omission bears on the determination of the *strength* of a reason. For example, it is a plausible thought that the strength of your normative reasons is modified by the difficulty of acting on them. Suppose that you are addicted to the drug, so that it is difficult to refrain from taking it, and easy to take it. Then a reason which bears on the omission—not taking the drug—will be harder to achieve than the action—taking the drug. The ability of the agent to perform the action differs from their ability to perform the omission, with consequences for the strengths of the reasons appropriate to each.

⁷ Further, as Snedegar and others have pointed out, a reason against Φ cannot generally be a reason for some or all alternatives to Φ because those alternatives may be *worse* with respect to the goal or aim in virtue of which A Φ 's. See Snedegar (2018: 733 and *passim*).

That normative reasons take a single object is important, then. An omission is a different object, with different properties, to an action, and thus should plausibly be treated differently.

2.2 How do I understand the distinction?

I have argued cursorily that the for/against distinction is a real one, but have as yet said little about how that distinction should be characterised. Let me now situate how I conceive of the distinction alongside another view on what the distinction amounts to.

Patricia Greenspan develops a view on which the for/against (or negative/positive in Greenspan's terminology) distinction should be understood as closely tied to the notion of rationality.⁸ According to Greenspan, a negative reason that counts against an option 'tends to rule it out' and that '[k]nowingly acting against such a [negative] reason without strong enough opposing reasons, is the paradigm of practical irrationality'.⁹ Positive reasons, on the other hand, merely rationally permit responding to options by 'rendering an option eligible for choice [...] without requiring it'.¹⁰

Furthermore, Greenspan's target notion of rationality is to be understood in terms of the offering and answering of criticisms. Negative reasons are the sorts of things that make appropriate criticism for failing to respond to them, while positive reasons are the sorts of things that can answer such criticisms.¹¹ This conception of reasons, (which she calls the 'critical conception') as being in the business of answering and offering criticisms is intended to be exhaustive: '[a]ll that rationality requires is fending off criticism'.¹²

However, I think it is a mistake to try to understand a normative reason in terms of rationality. First, the notion of rationality is more diffuse and unclear than that of the notion of a consideration counting in favour or against a response. If that is right then basing arguments for the for/against distinction on the basis of intuitions about what counts as rational or irrational in cases (as both Greenspan and Gert do in arguing for their positions) is a non-starter.

Second, I'm not convinced that we exhaust all of the sorts of things that reasons can do by appealing primarily to the notion of criticism. Even if Greenspan's distinction does pick out two ways in which reasons bear on their objects, the distinction surely does not exhaust the ways in which we understand such strengths. Consider, for

⁸ See Greenspan (2004) and Greenspan (2007). In Gert (2003), Gert develops a structurally similar view to Greenspan.

⁹ Greenspan (2004: 387–8)

¹⁰ Greenspan (2004: 389), see also Raz (2011: 5–6)

¹¹ Gert advocates a structurally similar view. He distinguishes between justifying and requiring strengths of reasons, where the strengths of reasons should be understood in terms of rationality. Positive reasons for Greenspan correspond to justifying strength for Gert: for a reason to have justifying strength is for it to contribute to making an otherwise irrational act rational. In other words, justifying strength defends against potential rational criticisms that would otherwise be applicable to an agent who acted in that way. See Gert (2003).

¹² Greenspan (2004: 389), see also Greenspan (2011: 173).

instance, simple pairwise comparisons between options that share a value base: that the 12 year old Laphroaig is tastier than the 8 year old Glenkinchie implies that you have stronger reason to drink the Laphroaig. One way to understand the strength of that reason is in terms of the underlying value of choosing that option. But the normative force of that claim doesn't seem to bear any necessary connection to defending oneself against potential criticism. Sure enough, the fact that the Laphroaig is tastier *may* provide a sensible defence of your decision to drink it, but if we are looking for an explanation of the strength of your reason, the fact you can defend your decision looks downstream from the fact that it is more valuable.

Consider another example. Jonathan Dancy identifies a subset of reasons as of an "enticing" variety.¹³ They are such that they would be pleasant to do, but no matter how strong they are, they never make it the case that you ought to act on their basis. Suppose that the fact that it is a pleasant day is a reason to go for a walk. That reason doesn't seem to be the sort of thing that rules out other options, but neither is it necessarily playing the role of a defence against potential criticism.

You are not failing to understand the normative situation if you do not take the fact that it is a pleasant day as material to defend yourself against the potential charge that you are doing a rationally criticisable thing. It seems at best like a stretch to treat this reason as if this were its primary "function". So it is questionable whether all normative questions answerable by reasons are reducible to questions about giving and answering criticisms.

So we are left with a problem unresolved: if reasons against Φ -ing are not reducible to reasons for not Φ -ing, and the critical conception is not plausible then how *should* we understand the asymmetry of normative reasons? The critical conception gains most of its plausibility with respect to its account of reasons against. If criticism enters into the picture it seems obvious that it will be here. But (1) this doesn't force us to accept that reasons for play a merely defensive role, and (2) explaining the asymmetry in terms of rationality still seems like a wrong move. Here is one alternative possibility: we pick out the for/against distinction by appealing to *reactive attitudes* in an identifying role.

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, then, is to appeal to the attitudes that it is appropriate to hold towards an agent for whom those reasons are relevant. A plausible generalisation might be: where A has undefeated reason *against* Φ -ing 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.¹⁴ On the other hand, where A has undefeated reason *for* Φ -ing and they fail to Φ , it is generally at most appropriate only for A to hold negative reactive attitudes (such as regret) towards themselves.

¹³ See Dancy (2004a).

¹⁴ In Strawson's terminology, the former would count as "personal" reactive attitudes, and the latter "self" reactive attitudes. See Strawson (1974).

This alternative view gives us the starting materials with which to distinguish reasons for and reasons against. In this way, we can avoid the complaint that we must appeal to a more imprecise notion (rationality) than reasons, to explain reasons, and it gives us an alternative to conceiving of reasons and reasoning as a primarily defensive operation. The view also does not claim to be a *reductive* account of reasons. We can appeal to reactive attitudes to allow us to identify the valency of a normative reason, without thereby implying that normative reasons are reducible to the appropriateness of holding those attitudes.¹⁵

3 A better buck-passing view

In this section I formulate my version of the deontic buck passing view. To prepare the ground for the analysis, we need first to be sensitive to two ways in which reasons can *combine*—two *verdictive* claims: first, reasons can be weighed for and against the same response; second, the net verdict of reasons against can be weighed against the net verdict of reasons for an alternative response. I call the former verdictive claim, having *sufficient reason*:

Sufficient reason A has sufficient reason to Φ if, and only if, the reasons *for* A's Φ -ing outweigh the reasons *against* A's Φ -ing

Where, conversely, to have sufficient reason not to Φ is for A's reasons against Φ -ing to outweigh A's reasons for Φ -ing. Second, call the latter verdictive claim having *greater reason*:

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 an alternative, or set of alternatives to Φ -ing)

Again, where, conversely, having greater reason against Φ -ing is for there to be more net reason against Φ than to Ψ .

In formulating my version of the deontic buck-passing view I make use of both of these verdictive claims, in addition to the distinction between reasons for and reasons against. Let me begin by first stating my analyses of requirement and ought, and I shall then comment on them at length. To anticipate, in an intuitive sense, a requirement functions to rule out courses of action. On the other hand, where an agent ought to do something, this (very roughly) expresses the idea that it would be

¹⁵ What about epistemic reasons? Is having evidence against p and nevertheless believing p somehow epistemically worse/more blameworthy than having evidence for q and not believing that q ? In general I think there may be a parallel here: if you have evidence *against* p then you are at fault for believing p anyway, but you may have weak evidence *for* a huge number of propositions, and yet not at all be epistemically at fault for failing to believe them all. This is similar to the practical case in which having reason against renders you blameworthy for doing so anyway, while having for reason for does not imply that you are blameworthy for failing to do so, because (for one thing) you may have other perfectly acceptable options open to you.

a good (perhaps the best) course of action. I analyse these deontic properties as follows:

- Required to Φ A is required to Φ if, and only if, some consideration(s), F, give(s) A reason *for* Φ -ing and F also gives A sufficient reason *against* Ψ -ing¹⁶ (where Ψ -ing is an alternative, or set of alternatives to Φ -ing)
- Ought to Φ A ought to Φ if, and only if, there is a set of considerations that gives A greater reason to Φ

In what follows I explain these analyses further, argue that they are independently plausible, and that they allow us to meet the two objections to the standard view. In Sect. 4 I discuss my view in relation to other buck passing views.

3.1 Required to Φ

Intuitively, a requirement expresses the idea either of ruling out any alternative course of action, or of ruling out a single action. This is not true for ‘ought’ claims. In order to respect this difference, a key move that I make in distinguishing these two deontic verdicts in my analyses, builds on the claim that a requirement is typically concerned with, or takes as its object, *both an act and its alternatives*. ‘Ought’, on the other hand, is typically concerned only with, or takes as its object, a *single act*, and bears on alternatives only *indirectly*.¹⁷

Given this way of understanding a requirement, I need to defend the claim that a consideration that favours Φ -ing in the analysis of a requirement itself *also* counts against those actions that conflict with Φ -ing. For example, according to me, if A is required to keep her promise to B, then that A promised B that A would return B’s £10 this afternoon counts against A’s going to the cinema if, and only if, going to the cinema is a way of reneging on that promise.

By contrast, I claim, if A merely ought to visit the restaurant on the basis of its tasty fare, then that the restaurant’s fare is tasty does *not* count against, say, A’s going to the cinema, *even if going to the cinema is a way of not going to the restaurant*. In cases like this, the explanation why A should go to the restaurant is given by the fact that A has a good reason to go. On the other hand, promise-giving reasons count, in addition, against performing alternative actions, not just because there is something better to do, but because promise-giving reasons *bear directly* on those alternatives.

The intuition I’m trying to elicit is that the tastiness of the fare at the restaurant doesn’t give A reason *against* going to the cinema, but the fact that A promised to return the £10 this afternoon *does* give A reason against going to the cinema this afternoon. These latter sorts of considerations do more than just make some option,

¹⁶ In the sense of sufficient reason defined above.

¹⁷ As may be obvious already, this 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.

so to speak, attractive: they themselves count against other competing options. Thus, there appears to be a structural difference between pleasure-based and promise-based reasons. The object of a pleasure-based reason is the consideration that reason favours or disfavors.¹⁸ The object of a promise-based reason is wider, ranging over both the act favoured/disfavoured, and the alternatives. Or, so at least, appears intuitive given the examples.

Note that the fact that A promised to return the £10 gives A reason against going to the cinema only if going to the cinema conflicts with A's being able, within a reasonable amount of time, to return the £10 to B. But isn't that true of *any* reason? For example, isn't 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 say that the reason to go to the restaurant is stronger than the reason to go to the cinema. 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 A should go to the cinema rather than the 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.

But there is a subtle difference in the way we express a requirement. It is not *merely* that there are reasons favouring Φ -ing and reasons favouring Ψ -ing and what A is required to do is given by the fact that A has a stronger reason to Φ . The fact that A promised to return the £10 (now) *is also* a reason not to do whatever conflicts with giving back the £10. Here the *same* consideration both favours Φ -ing *and* disfavors Ψ -ing. Thus it seems that requirements express the claim that you should do something both because the balance of reasons falls on one side, and that those reasons that fall on one side themselves also count against (bear on) acting otherwise.

We can make this point clearer by applying the distinction between reasons for and reasons against to these cases. As I claimed above, we should accept that reasons display a kind of asymmetry. Where A has reason *against* acting in some way, all else being equal she is appropriately subject to negative personal reactive attitudes, perhaps in light of her recognition of this reason. On the other hand, where A has reason *for* acting in some way, and A decides not to act in that way, A may be appropriately subject only to negative self-reactive attitudes.

There is, however, a problem for my 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 the three options is the degree of harm that they will cause to B, with X causing the least amount of harm, Y more, and Z still more. There is nothing to be said *for* doing

¹⁸ Again, given the assumption that Snedegar's "opportunity cost" principle is false, which I have argued above that we have good reason to believe.

any of these things. However, A has no further choices—these are the only options open to him.

We can get around this objection by dropping the condition that if you are required to do something, there must be something to be said for doing that thing:

Required to Φ * A is required to Φ if, and only if, there is some consideration(s), F, that gives A sufficient reason against Ψ -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. *Required** reflects this.

Consider now the following, second, objection to *Required*. Suppose that the fact that it is sunny outside gives A sufficient reason to go to the plaza, sufficient reason against going to the cinema, and sufficient reason against staying at home. And suppose that these are the only options available to A now. Since there is a consideration that gives A reason to Φ , which also itself gives A reason not to act in any of the ways that conflict with A's Φ -ing, where the set of options is reduced to these three, *Required* tells us that A is required to go to the plaza. But, in an intuitive sense, A is not subject to a requirement to go to the plaza. So *Required* is false because its extension is incorrect.

First, the case is potentially misleading. That it is sunny outside is a reason *for* going outside, not a reason against staying in, at least according to my schema of reasons. To count against acting in some way, a consideration must make appropriate negative reactive attitudes for acting in that way, highlighting 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.

Second, we need to specify more precisely what the relevant grounds are in this case. Suppose 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.

Third, I don't find cases that severely limit the alternatives open to an agent to be particularly illuminating. In reality, we have a very large number of alternatives open to us at any one time, and given this very large number, having sufficient reason *against* each of these relevant alternatives looks relevant to what one is required to do. The most obvious cases that *do* make available only a very small set of alternatives are those that involve restricted agency or coercion. But in these sorts of cases, the external factor that limits an agent's alternatives to a small set will likely itself imply that the situated agent is no longer in realm of "mere" choice.

3.2 Required not to Φ

My analysis of a requirement, however, is incomplete. Clearly requirements *not* to act in some way cannot be modelled on my analysis of requirements *for* acting in some way, because they rule out acting in *one* way, not *all alternatives* to acting in that way. The most obvious view here is the following:

Required not to Φ A is required not to Φ if, and only if, (i) some consideration(s), F, give A sufficient reason against Φ -ing, and that reason has a significant degree of strength

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 sufficient reason against Φ -ing, 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 subject to negative reactive attitudes for Φ -ing but the strength of those attitudes is significant enough to identify this as a case in which A is required not to Φ .¹⁹

However, my definition of a requirement not to do something cannot be right as it stands. Consider again a scenario in which you *must* perform your least bad option. Assuming that you have sufficient reason against performing each of these options, and that there are no reasons favouring any of the options, according to my definition of a requirement not to Φ , as it stands, you are required not to perform *any* of these options. Furthermore, according to *Required 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:

Required not to Φ^* A is required not to Φ if, and only if, there is some consideration(s), R, which give A sufficient 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 sufficient 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 .

The questions ‘What am I required to do now?’ and ‘What am I required not to do now?’ raise different issues. In the first case, an answer places you under a

¹⁹ Note that this view is not inconsistent with absolute prohibitions. To say that Φ -ing is absolutely prohibited is to say that there sufficient reason for A not to Φ and there is no consideration that *could* justify A’s Φ -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 for A not to Φ .

demand to perform a single action, and only a single action.²⁰ In the second case, an answer may place you under a demand not to perform many actions, or many demands not to perform many different actions. It is a defensible view that you cannot, at t_1 , be required to do more than one thing. It is a platitude that 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* status, is that when considering whether to perform some act, comparison classes are often important. Often you decide 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 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 that A ought to Φ , involves weighing competing reasons for competing actions (this may be why there is conceptual pressure to accept that there is an ‘ought, all-things-considered’). On the other hand, having sufficient reason *not* to act in some way involves (a) only competing reasons for and against that omission and (b) that there can be multiple non-competing ‘most reason’ verdicts supporting different omissions.

I claim that to be required not to act in some way is for there to be sufficient reason not to act in that way of significant strength. To have sufficient reason not to act in some way is for there to be more reason against Φ -ing than for Φ -ing, and to have reason against Φ -ing is to explain why A, in Φ -ing, is appropriately subject to negative reactive attitudes. Such negative attitudes, 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.²¹ For example, consider the difference between prudential and moral reasons: I have prudential reason to avoid acting in a way that causes harm to me, if I do it. I have moral reason to avoid acting in a way that causes harm to 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 harmed because of the badness of harm. I am subject to criticism on the basis of my failure to care about the badness of the harm. In the second case, it is less plausible that this can be entirely captured in teleological terms. It is not merely my failure to recognise the badness of harm that forms the basis of appropriate criticism against me. It is also 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

²⁰ 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.

²¹ I say “can determine” here as I am not attempting to offer the final word on this, and am quite open to alternative interpretations of how we should understand the “significant degree of strength”.

against promoting harm, and another is associated with a respect-based reason against harming agents.

We can also note the difference between cases in which you have sufficient reason against, for which there are *no* reasons for, and cases in which you have sufficient 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 waive 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 waivable by the agent.

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

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 the 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 in 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.²²

On one plausible interpretation, 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 reason against (and thus makes appropriate negative reactive attitudes) and (ii) because an agent cannot waive another agent's reason. That is at least one plausible line that we could take in distinguishing moral from prudential requirements, in addition to a difference in the strengths and kinds of the grounding reasons.

If we appeal to intuition about what would count as appropriate criticism it seems plausible that you are not appropriately subject to negative reactive attitudes for

²² See Wallace (2013: 150–8) for further discussion on this point.

failing to do what you had less-than-most reason to do (as long as you do what you had most reason to do).

4 Rival views and *ought*

Although the literature on deontic buck-passing is slight, with only one theory explicitly advanced in the recent literature,²³ Justin Snedegar has helpfully shown how a number of different theories of reasons can be utilised to provide different buck passing views.²⁴ All of the views that he considers are built on the basis of a common claim—that normative reasons can be distinguished into two kinds. Given that a desideratum of a good buck passing view is to be able to distinguish cases like *Lunch* and *Theatre* a distinction among reasons can, the thought goes, be mapped structurally onto the distinction among deontic verdicts—between *ought* and requirement.

There are a number of such putative distinctions among reasons on the market. Matt Bedke distinguishes between reasons for acting, and reasons to perform speech acts (reasons for *others* to respond to the agent in certain ways).²⁵ Joshua Gert distinguishes between reasons which *require* acting, and reasons which *justify* acting.²⁶ Joseph Raz distinguishes between first and second order reasons.²⁷

We can use these distinctions to build buck passing views that respect the difference between the deontic verdict appropriate in *Lunch* from the verdict appropriate in *Theatre*. However, the fundamental problem, identified by Snedegar, shared by the resultant views, is that because each distinction among reasons is conceptually independent, that implies that the deontic verdicts that they reduce do not bear structural relations as a matter of conceptual necessity to one another.²⁸ Moreover, some of the theorists who defend these distinctions—notably Gert and Bedke—*insist* that the kinds of reasons that they identify are conceptually independent of each other.

This would not matter if *ought* and requirement did not bear structural relations to one another; however it is very plausible that they do. If you are required to act in some way, this entails that you *ought* to do it. But if you *ought* to act in some way, this does *not* entail that you are required to act in that way.

Snedegar's response to this problem is to build in both kinds of reason into the stronger deontic verdict (requirement). For example, an interpretation of Bedke's analysis of requirement is:

²³ Which is due to Matt Bedke. See Bedke (2011).

²⁴ See Snedegar (2016).

²⁵ Bedke (2011).

²⁶ Gert (2003). Again, his distinction maps on to Greenspan's negative/positive reasons distinction. Douglas Portmore makes use of a similar distinction in Portmore (2011).

²⁷ See, for example, Raz (1975).

²⁸ Snedegar (2016: 175–7).

$\text{Bedke}_{\text{required}}$ A is required to Φ if, and only if, (i) A has most reason to Φ , and
(ii) there is most reason to require A to Φ ²⁹

While ‘ought’ is defined simply in terms of what A has most reason to do. The immediate problem with this move is that condition (i) appears to be unmotivated—it is just added into get the right result.

One way to respond here is to follow Stephen Darwall in arguing that it is just *incoherent* to require somebody to act in a way that fails to be supported by what they have reason to do, because we cannot *blame* somebody for this.³⁰ Snedegar points out, however, that a new problem emerges—it seems possible to suppose both that existence-internalism about reasons is true and that we can require an agent to act in a way that does not serve their desires. But this combination of views is ruled out by $\text{Bedke}_{\text{required}}$. Thus, again, what is needed is an explanation of the conceptual relations between the two kinds of reasons identified—why *must* it be the case that we have reason to require you to Φ only if you have most reason to Φ ?³¹

My view, however, avoids the need to answer this question, since according to my analyses, what it is to be required to do something is that you have most reason against acting on the alternatives. The conceptual relationship between requiring and having most reason is built into the view, without the need to appeal to requiring reasons.

Second, by appealing to the reasons for/against distinction, we can make use of a distinction among reasons, while at the same time retaining a clear conceptual relation between the two reasons. That is because reasons for and reasons against are not *fully* conceptually independent due to the role that reasons for and against must play in *reasoning*.

Judgements about what an agent has most reason to do or believe involve the weighing and balancing of reasons for and against. That is, the weighing of reasons, for and against, is *partly constitutive* of reasoning. This is true anyway, independently of the truth of the asymmetry claim about reasons. Now, clearly, given the truth of the latter claim, weighing and balancing reasons is more complicated than that suggested by the metaphor of a mechanical scale, with weights on either side. But many theorists have come to recognise that the notion of weighing must be more complex than that anyway.³²

Note that the putative kinds of reasons identified above are not constitutive of reasoning in the same way. One can have justifying reasons *for and against* acting in some way, and requiring reason *for and against* acting in some way. Or, one can have requiring reasons *for and against* requiring somebody to respond in some way. Each of the distinctions broached above cuts across the for/against distinction.

²⁹ Snedegar (2016: 175).

³⁰ See Darwall (2006, ch.6).

³¹ Snedegar (2016: 176–7). .

³² See, for example Dancy (2004a, b), Lord and Maguire (2016).

In order to demonstrate how my view guarantees the desired entailments, let us now consider my analysis of ‘ought’:

Ought to Φ A ought to Φ if, and only if, there is a set of considerations that gives A greater reason to Φ

Again, just as requirements *not* to act in some way cannot be modelled on the analysis of requirements *for* acting in some way, that you *ought not* cannot be modelled on its being the case that you *ought*:

Ought not to Φ A ought not to Φ if, and only if, A has greatest reason to Ψ (where Ψ -ing is an alternative, or set of alternatives to Φ -ing)

What the claim that A ought to Φ is intended to express is that some course of action is either a good, or the best option from some set of alternatives (ought to Φ) or that, of some set of actions, A should do otherwise than Φ -ing (ought not to Φ). Note that one way in which ‘ought not to Φ ’ differs from having a *requirement* not to perform some action is that, if A ought not to Φ , there may be some reasons in favour of Φ -ing, but there is *still more* to be said for every alternative. On the other hand, where A ought to Φ , she is not required to Φ because the considerations that give her reason to Φ do not *also* rule out those considerations that conflict with her Φ -ing.

It is important to note that my analysis of the claim that A ought not to do something does not include reasons against. That is because, as I have tried to make clear, the presence of reasons against implies that it is appropriate to hold negative reactive attitudes either from the first person, or third person perspective (however weak and potentially excusable). Thus, I have avoided analysing ‘A ought not to Φ ’ in terms of reasons against in order to avoid collapsing this analysis into *Required not to Φ* . So, on my analysis is it true that if A is required to Φ , this entails that A ought to Φ , but *not* vice versa?

First, it should be clear that *Required to Φ* entails *Ought to Φ* , since if A has sufficient reason not to do *otherwise* than Φ -ing, then A has more net reason to Φ than to Ψ , as her reasons against defeat any reasons for any conflicting action. Conversely, *Ought to Φ* clearly does not entail *Required to Φ* , since being required to Φ entails all of A’s alternatives to Φ -ing being ruled out, which *Ought to Φ* does not imply. So the desired entailments are preserved here.

What about the entailment from *Required not to Φ* to *Ought not to Φ* ? I defined ‘ought 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 *Ought 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 (*Required 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 are preserved here too.

5 Resolving the supererogation problem

The second difficulty with standard view is that it appeared to conceptually rule out supererogation. Suppose that A has sufficient reason to sacrifice her life to save five people in a burning building. It seems plausible to suppose that such a reason could outweigh opposing reasons A has, i.e. it seems possible that A can have most reason to perform a supererogatory act, in this sense. According to the standard view, A therefore *ought* to sacrifice herself to save the five (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, it cannot be the case that you ought to perform a supererogatory action. Thus, the standard view rules out the conceptual possibility of supererogation.

When we use the term ‘ought’ we can sometimes mean ‘required’. According to my analysis of ‘ought to Φ ’, for it to be the case that A ought to Φ , 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 that you ought to, 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 F gives A sufficient reason to Φ and F gives A sufficient 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 five people only by risking their own life by entering the burning building. If the fact that it would save five lives counts decisively in favour of risking their own life by entering the building, *and* that fact also decisively counts against any alternative, then according to my analysis A is required to enter the building. 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 five lives *does* plausibly count against A’s refraining from entering the building; 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 five lives counts decisively against their risking their own life.

However, we can allow that A has *most* reason to save the five in that he has *sufficient* reason to save the five (A’s reasons for saving the five outweigh the reasons against saving the five) and *greater* reason to save the five (A has more reason to save the five than to do anything else). Still, according to my analysis, A is *not required* to save the five. A is only required to save the five *if* (and *only if*) the consideration that gives him reason to save the five *also* counts decisively against the 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 the possibility: supererogatory acts which represent your best option can make it the case that you ought to perform them, in the sense above, but they cannot be required in my sense.

6 Conclusion

My version of the deontic buck-passing view makes use of a distinction among normative reasons to reduce the deontic properties of ought and requirement to reasons. This view allows us to capture the different deontic verdicts in *Lunch* and *Theatre* above, and does so in a way that captures the structural relations between these verdicts. This gives the view an advantage over the most prominent rival view, due to Bedke. Finally, the view avoids the supererogation problem by making room for supererogatory acts, by clearly distinguishing what an agent can be required to do, from what an agent ought to do.

Acknowledgements I would like to express my gratitude especially to Luke Elson, Sarah Fisher, Brad Hooker, Justin Snedegar, Philip Stratton-Lake, and Jonathan Way for their comments and feedback on this paper.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Bedke, M. (2011). Passing the deontic buck. *Oxford Studies in Metaethics*, 6, 128.
- Dancy, J. (2004a). Enticing reasons. In W. R. Jay (Ed.), *Reason and value: Themes from the moral philosophy of Joseph Raz* (pp. 91–118). Oxford: Oxford University Press.
- Dancy, J. (2004b). *Ethics without principles*. Oxford: Oxford University Press.
- Darwall, S. (2006). *The second-person standpoint: Morality, respect, and accountability*. Cambridge: Harvard University Press.
- Gert, J. (2003). Requiring and justifying: Two dimensions of normative strength. *Erkenntnis*, 59(1), 5–36.
- Greenspan, P. (2004). Asymmetrical practical reasons. In Reicher, M. E., & Marek, J. C. (Eds.), *Experience and analysis: Papers of the 27th international wittgenstein symposium* (pp. 387–394).
- Greenspan, P. (2007). Practical reasons and moral ‘Ought’. In R. Schafer-Landau (Ed.), *Oxford studies in metaethics* (Vol. II, pp. 172–194). Oxford: Clarendon Press.

- Greenspan, P. (2011). Craving the right: Emotions and moral reasons. In C. Bagnoli (Ed.), *Morality and the emotions* (pp. 39–61). Oxford: Oxford University Press.
- Lord, E., & Maguire, B. (2016). An opinionated guide to the weight of reasons. In E. Lord & B. Maguire (Eds.), *Weighing reasons*. Oxford: Oxford University Press.
- Portmore, D. (2011). *Commonsense consequentialism: Wherein morality meets rationality*. Oxford: Oxford University Press.
- Raz, J. (1975). *Practical reason and norms*. London: Hutchinson.
- Raz, J. (2011). *From normativity to responsibility*. Oxford: Oxford University Press.
- Snedegar, J. (2016). Reasons, oughts, and requirements. *Oxford Studies in Metaethics*, 11, 155–181.
- Snedegar, J. (2017). *Contrastive reasons*. Oxford: Oxford University Press.
- Snedegar, J. (2018). Reasons for and reasons against. *Philosophical Studies*, 175(3), 725–743.
- Strawson, P. F. (1974). *Freedom and resentment and other essays*. Abingdon: Routledge.
- Wallace, R. (2013). The deontic structure of morality. In D. Bakhurst, M. O. Little, & B. Hooker (Eds.), *Thinking about reasons: Themes from the philosophy of Jonathan Dancy* (pp. 137–167). Oxford: Oxford University Press.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.