

What We May Expect of Each Other

Benjamin Kiesewetter, Bielefeld University

Accepted for publication in *Ethics* (expected January 2026)

This is the Author Submitted Manuscript (ASM), the version originally submitted to the journal prior to peer review. The peer-reviewed Accepted Author Manuscript (AAM) is available upon request by emailing the author.

Abstract: What is it for an agent to have a moral obligation, or, equivalently, to be overall morally required, to perform an act? In this article, I present an informative account of moral obligations in terms of reasons and show how it can be put to use to solve some pressing problems in moral philosophy. The core idea is that morality is the realm of what we can legitimately expect of one another: for agents to be morally required to ϕ is for them to have reasons to ϕ that are at the same time reasons for everyone to *expect* them to ϕ . This account is superior to standard rationalist and sentimentalist approaches in capturing both the first-personal deliberative as well as the socio-emotional dimension of moral obligations. Over and above, it provides illuminating explanations of supererogation and the imperatival character of moral requirements.

What is it for an agent to have a moral obligation, or, equivalently, to be overall morally required, to perform an act?¹ In this article, I present an analysis of moral obligations in terms of reasons, argue that it provides illuminating explanations of some of their potentially puzzling features, and suggest how it can be put to use to solve some pressing problems in moral philosophy. I will start by highlighting some of the general attractions of the project – sometimes called *deontic buckpassing* – of explaining deontic properties or concepts in terms of reasons (§1). I will then turn to two prominent approaches in the literature – a rationalist and a sentimentalist one – and discuss some of their benefits but also some of their specific challenges and limitations (§§2–3). Subsequently, I will present and explain my own proposal, according to which, roughly, agents are morally required

¹ I am here concerned only with *overall* moral obligations, and I will use the terms ‘moral obligation’, ‘moral requirement’ and ‘moral duty’ interchangeably to refer to them. See, however, §6 for a brief outlook on how the ideas developed here can be applied to *pro tanto* duties.

when their reasons to act are also reasons for everyone to *expect* them to act (§4). After discussing some potential objections to this *expectation view* (§5), I highlight the ways in which it combines the virtues of the rationalist and the sentimentalist approach while avoiding their problems. Finally, I show that the expectation view has virtues that neither of the existing views can claim for itself. In particular, I argue that it provides illuminating accounts of supererogation and the imperatival character of moral obligation (§6).

§1 Deontic Buckpassing

The Deontic Buckpassing Approach aims to analyse deontic properties or facts (or, alternatively, deontic concepts)² in terms of normative reasons. Deontic properties are properties such as rightness, wrongness, obligatoriness, permissibility, oughtness, being required, or having a claim right. By normative reasons, I mean the factors that speak in favour of or against actions or attitudes – the ‘pros’ and ‘cons’ that count when making up our mind. To illustrate this notion by way of an example, imagine yourself wondering about whether to spend the winter holidays in Hawai’i. There are factors that speak in favour of doing so: the weather will be nice, your kids love to surf, and you will get to spend time with your friends in Honolulu. But there are also factors speaking against doing so: Hawai’i is expensive, it’s far away, and your neighbours have asked you to help renovate the shared space. These factors are *reasons* for and against spending the winter holidays in Hawai’i, these reasons come with certain weights or strengths and there will be facts of the matter about how they compete with each other. This notion of a reason is, I submit, a pre-theoretical one. Every person who has ever made a reflective decision involving the weighing of considerations is already familiar with the idea of a reason.

To get a sense of the Deontic Buckpassing Approach, consider the proposal that ought facts can be analysed in terms of facts about the overall balance of reasons:

Ought: For it to be the case that A ought to ϕ is for it to be the case that A’s reasons for ϕ -ing are on balance stronger than competing reasons.³

² In what follows, I will focus on properties and facts, and I will understand the proposed analytical claims as metaphysical rather than conceptual or semantical claims. Much of what I go on to say could be employed on behalf of an analysis of deontic concepts as well, however.

³ See e.g. Schroeder (2007, 130); Parfit (2011, 33); Schmidt (2024).

This is an attractive proposal. It constitutes an informative, reductive account of ought facts in terms of facts that we are independently familiar with: facts about reasons and how they weigh against each other. If it is correct, we can explain what an ought fact is in terms of facts we are all familiar with just in virtue of being capable of making reflective choices.

It is natural to think of the Deontic Buckpassing Approach as part of a bigger approach to normativity that takes reasons to be fundamental. According to the so-called Reasons-First Approach, *all* normative properties – i.e. not only deontic but also, for example, evaluative and epistemically normative properties – can be analysed in terms of reasons.⁴ One reason to be attracted by this picture is that it accounts for the pervasive assumption that the distinction between the normative and the non-normative is an important general and unified distinction. And it is clear that the Deontic Buckpassing Approach sorts very well with the Reasons-First Approach to normativity. It is worth noting, however, that passing the deontic buck to reasons can be attractive to opponents of the Reasons-First Approach as well. For example, proponents of the so-called Fittingness-First Approach can and often do accept that deontic concepts are to be analysed in terms of reasons.⁵ They just hold that the buck doesn't stop there, because reasons are to be analysed in terms of fittingness. More generally, everyone who thinks that reasons are analysable in terms of a non-deontic notion N should find deontic buckpassing attractive, insofar as it indirectly delivers an analysis of the deontic in terms of N.⁶

One of the attractions of the Deontic Buckpassing Approach thus lies in its potential to contribute to a global unified account of the normative, for example in terms of reasons or fittingness. But deontic buckpassing can also be attractive to those who are sceptical with regard to this global project.⁷ This is because every *local* analysis of a deontic concept or phenomenon in terms of reasons improves our understanding, as it is an informative analysis in terms of a notion we are independently familiar with. Ideally such an analysis provides illuminating explanations of deontic phenomena that seem puzzling or even paradoxical. Consider, for example, the phenomenon of supererogation, i.e. the (presumed) fact that it can be morally permissible to act in a way that is morally suboptimal. Though it has been disputed that there is a real paradox here, it is widely agreed that the possibility of supererogation at least calls for explanation. One of the aims

⁴ See esp. Skorupski (2010); Schroeder (2007; 2021).

⁵ See esp. McHugh and Way (2016; 2022).

⁶ Even those who propose to analyse reasons in terms of ought facts (see e.g. Broome 2004; Kearns and Star 2009; Nebel 2019) might find it attractive to analyse other deontic properties, such as the property of being morally required, in terms of reasons.

⁷ See e.g. Wodak (2020) and Reisner (2023).

of this essay is to demonstrate that an analysis of moral requirements in terms of reasons can provide such an explanation.

One thing that we need to keep in mind when passing the deontic buck is that in order for the analysis to be informative, we should shy away from referring to deontic properties in the analysans. Consider, for example:

Bedke's View: "A is required to ϕ iff there is most reason (of the right kind) to require A to ϕ ."⁸

This proposal makes essential use of the verb-phrase 'to require', which arguably cannot be understood independently of the idea of a requirement. On its most natural reading, to require A to ϕ is to perform some kind of authoritative speech act with what Austin called "illocutionary force", i.e. an act that *makes it the case* that A is required. On this reading, the suggested analysis is both false and circular: false because there are requirements that are not authority-related and circular because it is difficult to see how the power to make it the case that A is required could be understood without reference to the property of being required. Now, Bedke is explicit that he doesn't intend to make use of this authority-related notion of the verb-phrase 'to require'. Instead, he suggests that "requiring A to ϕ is to answer 'no' to the following hypothetical question from A: 'May I not ϕ ?'"⁹, and he maintains that one need not have any kind of authority or standing to do this. I think it is doubtful that this characterization picks out an ordinary rather than stipulated notion of 'to require'. If you ask me "May I break my promise to ϕ ?", and I sincerely answer "no" because I think that you may not break it, I do not thereby seem to require you to ϕ in an ordinary sense of that term. There is a difference between asserting that an act is required and requiring that act (in the ordinary sense). But more importantly in our context, Bedke's explication of the verb-phrase 'to require' appeals to the ideas that one may or may not do something, and these *just are* the ideas of a permission and a requirement. Bedke's account of what it is to require thus refers to the property of being required: it says that to require something is to give a certain kind of answer to the question of whether actions have that property. This renders his analysis circular.¹⁰

⁸ Bedke (2011, 137). See also Snedegar (2016) and Lösche (2021), who both follow Bedke.

⁹ Bedke (2011, 138).

¹⁰ Note that this objection differs from the horse/cart objection that Bedke discusses. I don't object that the only right kind of reason to require is that the action is required. My objection is that we cannot understand the act of requiring without appealing to the property of being required. Some have argued that there can be analyses that are non-viciously circular and that tell us something essential about the relation between concepts or properties. It's not my purpose here

§2 Rationalism

Let's take a closer look at two different approaches to deontic buckpassing. What I will call the "rationalist" approach analyses the moral requirement an agent *A* is subject to in terms of *A*'s reasons to act. An immediate implication of this view is that there is necessarily a reason – depending on the specific proposal perhaps even *most* reason – to be moral (hence the label "rationalism"). In my view, this is a welcome feature of the rationalist approach. Firstly, it explains why moral requirements have first-personal normative authority, i.e. why they are relevant in an agent's deliberation what to do. Secondly, it does this in a way that provides a particularly attractive response to the so-called why-be-moral challenge. If moral requirements are to be analysed in terms of reasons for action, then there is no need to come up with an 'extra' reason to be moral. Moral requirements are automatically supported by reasons, and these reasons are not external to morality. Moreover, the rationalist gets reasons to be moral without making the substantive normative assumption that moral requirements are reasons – an assumption that has been criticized as fetishistic (Smith 1994), and which may be argued to beg the question against the moral skeptic (cf. Prichard 1912). Finally, and relatedly, rationalists are in a good position to explain a moderate and plausible form of judgement internalism, according to which rational agents are motivated to act in accordance with their judgements about what they are morally required to do. They can do this because it is independently plausible to think that rational agents are motivated to act in accordance with their judgements about what they have reason to do, and according to rationalist buckpassing, judgements about moral requirements are judgements about reasons (or at least judgements about facts that consist in reasons facts).¹¹

So there is much to like about the rationalist approach. But there are also problems. To begin with, consider:

Simple Rationalism: For *A* to be morally required to ϕ is for *A* to have most reason to ϕ .¹²

to rule out this general possibility. But as I understand the Deontic Buckpassing Approach, it takes reasons to be not only interestingly related but also *more fundamental* than the deontic.

¹¹ This is a point where the difference between a conceptual and a metaphysical analysis can become relevant: the latter allows that an agent who is mistaken about the nature of moral obligations judges an act to be morally required without judging himself to have reason to perform it. The variant of judgement internalism that those rejecting the conceptual claim can explain is thus limited to rational agents that are not mistaken about the nature of what their moral judgements refer to.

¹² Cf. Smith (1994, 182); Stratton-Lake (2002, 15).

The problem for simple rationalism is that there are situations in which we have most reason to ϕ although we are morally permitted to not- ϕ . If your reasons to eat chocolate ice cream are stronger than all your competing reasons in a given situation, it doesn't seem to follow that you are morally required to eat chocolate ice cream.

To fix this problem, rationalists might appeal to the notion of a *moral* reason in their analysis:

Moralized Rationalism: For A to be morally required to ϕ is for A to have most *moral* reason to ϕ .¹³

But this poses new difficulties. The first of these difficulties is that moralizing rationalists owe us an account of moral reasons that does not presuppose deontic notions and is sufficiently neutral with respect to first-order disagreements about morality. With regard to the first of these points, moralizing rationalist cannot characterize moral reasons as those that are related to moral requirements if they are to avoid circularity. With regard to the second point, they also have to avoid characterizing moral reasons in a normatively contentious way, for example as impersonal welfare-based reasons or reasons arising from hypothetical contracts (to mention just two possibilities). At least they have to avoid doing this as long as they want to allow that proponents of different ethical theories can share the same concept of a moral requirement (or the same metaphysical account of the property of being morally required) without contradicting themselves.

The second, perhaps even more severe problem for Moralized Rationalism stems from the possibility of supererogation. For cases of supererogation seem to be those in which agents have most moral reason for an action and are nevertheless morally permitted to do something else. Thus, for example, in a case in which you can save two people by risking your life and running into a burning building, it seems plausible to think that the balance of your moral reasons alone favours saving the two, even though it is not the case that you are morally required to do so.¹⁴ This case (henceforth *Burning Building*) is a direct counterexample to Moralized Rationalism.

¹³ Smith instead suggests to restrict the relevant act-types to those that are picked out by the platitudes on morality. This proposal faces a number of difficulties that the presented one avoids. For one, not all moral requirements are related to specific act-types that can be identified by appeal to platitudes about morality. You may be morally required to eat chocolate ice cream if doing so would cheer up your sick child, for example. For another, the very same act-type can be favoured by moral and non-moral reasons. But if you have most reason to perform such an act-type because of non-moral reasons, it does not seem to follow that you are morally required.

¹⁴ This verdict can be avoided by assuming that the reason against risking one's own life is a moral reason that counterbalances the reasons to save the two, but this doesn't strike me as plausible. It also misses the sense in which saving the two in such a situation would be morally superior or heroic: how could that be so if the moral reasons do not favour saving the two?

Philosophers with consequentialist leanings often deny that supererogation is possible, because they believe that such cases are ruled out by the correct moral theory. But even those philosophers should concede that there can be substantive moral disagreement about the existence of such cases, and it seems reasonable to suppose that such first-order moral disagreement need not necessarily involve mistaken views about the concept or metaphysical nature of moral obligation. There is thus reason to reject a supererogation-excluding analysis of moral obligations even for those who hold moral views that deny that supererogation is possible.

A *prima facie* promising way for the rationalist to accommodate supererogation is to include a further condition into the analysis:¹⁵

Dual Ranking Rationalism: For A to be morally required to ϕ is for A to have most moral reason *and most reason overall* to ϕ .

The idea behind this proposal is that cases of supererogation should be seen as cases in which moral reasons are counterbalanced by non-moral reasons. Accordingly, in *Burning Building* there is a strong non-moral reason against saving the two, which is provided by the risks of doing so, and counterbalances the moral reasons for saving the two. By doing so, it ensures that the condition that there is most reason overall to save the two is not satisfied and that therefore, even though the act is favoured by the balance of *moral* reason, there is no moral requirement to save the two.

This story might work in what we might call heroic cases of supererogation – cases like *Burning Building* – but the problem is that it does not work in non-heroic cases. Consider the following example of supererogation, presented by Horgan and Timmons (2010, 47):

Widowed Neighbour. “Olivia and her husband Stan have recently moved to St. Louis... During their first week in their new home, Olivia attends a block party organized by one of their new neighbors where she meets a recently widowed woman, Mary, a neighbor who lives a few doors down from Olivia and Stan. In conversation, Olivia learns that Mary lost her husband to cancer after forty-eight years of marriage. She also learns that Mary is an avid baseball fan and that she and her husband used to regularly attend Cardinals games. But

¹⁵ This option is overlooked by some of those who charge rationalist buckpassers for being unable to account for supererogation, such as Darwall (2022, 113). In Darwall’s case this is especially noteworthy because his account of moral reasons as “moral obligation-making considerations” (*ibid.*, 115) seems to leave no room for the kind of objection to the *Dual Ranking Rationalism* that I go on to discuss, which relies on the possibility of moral reasons that are not obligation-makers. It thus seems that Darwall has no resources to argue against this version of the rationalist approach.

without anyone to go with, she doesn't go anymore. The next day, it occurs to Olivia that it would be a nice gesture to offer to go to a Cardinals game with Mary, although she herself has no particular interest in the game... She calls Mary, who is delighted by the invitation, and they end up going to a game."

Again, this is a case in which it seems plausible to say that Olivia is not morally required to invite Mary even though her moral reasons favour that action. The fact that it would cheer her up is, plausibly, a reason to invite Mary, and it seems to belong to the category of moral reasons.¹⁶ Further, we can stipulate that no competing moral reasons are present, which means that the balance of moral reasons favours inviting Mary. And yet it seems false to think of Olivia as being morally required to invite Mary; she clearly goes beyond the call of duty when inviting her. But is that because she has a non-moral reason against inviting Mary that counterbalances her moral reason to do so? This does not strike me as a plausible diagnosis. Suppose that Olivia also has a non-moral reason to invite Mary (she will enjoy Mary's company and even find interest in the game), and her competing non-moral reasons are weaker (if she doesn't go to the game, she will have an average evening at home). Does it now follow that she is morally required to invite Mary to go to the game? Would it be morally wrong for Olivia to spend her evening at home instead of inviting Mary to the game? These conclusions strike me as false. It is, morally speaking, up to Olivia whether she wants to invite Mary to the game or spend the evening at home, and this is not changed by the fact that she would enjoy the former more. This means, however, that *Dual Ranking Rationalism* needs to be rejected.¹⁷

Here is another kind of case in favour of this conclusion: Suppose you have to decide what to do with your sabbatical. You can pursue your plan on starting to write a novel, or you volunteer for an NGO that supports an orphanage in India. The moral reasons favour volunteering for the

¹⁶ There may be conceptions of moral reasons that rule out, per stipulation, that 'non-requiring' reasons are moral reasons, but recall the point that rationalists cannot make use of such a stipulation, on pain of circularity. Apart from stipulative purposes, it is difficult to see how reasons like Olivia's could fail to be a moral reason, as acting on them seems to be morally praiseworthy, and the disposition to act on them seems to be a moral virtue. Defending an account of moral reasons is beyond the scope of this essay, but I am sympathetic to the hypothesis that a reason R for A to ϕ is a moral reason iff A's ϕ -ing for the reason that R has a certain kind of non-instrumental value (sometimes referred to as moral worth), which in turn might be spelled out in terms of reasons for pro-attitudes towards A's ϕ -ing for the reason that R. This conception might also help rationalists with the first problem mentioned above, which is why I think that the second problem is more severe.

¹⁷ De Kenessey's "decisive moral reason view", according to which "an action is morally required just in case the moral reasons in favor of that action are enough on their own to outweigh all of the reasons, moral and nonmoral, to perform any alternative" (2023, 1) is subject to the same objection: The moral reason for inviting Mary may be stronger on its own to outweigh all reasons against it without thereby rendering the invitation obligatory. The problem is that reasons for doing others a favor are not in the business of making actions morally required.

NGO. There are important non-moral reasons for and against both options, but on balance they favour spending your sabbatical with the NGO. You decide to use your sabbatical for writing your novel nevertheless. As you act against the balance of reasons, you are less than fully rational, but is your decision morally impermissible? I don't think so.

The problem with non-heroic supererogation seems to be avoided by Portmore's version of the Dual Ranking View, which qualifies the relevant moral reasons:

Portmore's Dual Ranking View: A is morally required to ϕ iff A has most morally *requiring* reasons and most reason overall to ϕ .¹⁸

This view seems to accommodate cases like *Widowed Neighbour* on the assumption that the reason to cheer up Mary by inviting her is a morally *enticing* rather than a morally *requiring* reason, i.e., a reason that can make actions morally praiseworthy but not morally required. The problem, however, is that it does this at the cost of giving up on the project of an informative analysis of moral requirements, as the notion of a morally requiring reason just is the notion of a reason that tends to make actions morally required.¹⁹

We might summarize the difficulties for rationalist buckpassers that I have hitherto discussed as follows. Those who want to analyse moral requirements in terms of reasons for action face two wrong-kind-of-reasons problems: one with respect to non-moral reasons and another one with respect to morally enticing or 'non-requiring' reasons. The challenge is to exclude such reasons from the analysis without presupposing deontic notions and without introducing substantive moral assumptions into the analysis of moral concepts or properties, and this challenge has not yet been dispelled.

Let me finally turn to another problem for rationalism, one that would persist even if these wrong-kind-of-reason problems could be solved. Morality has a special social and emotional dimension. If we act in ways that are morally impermissible, this is a matter that concerns other people in a way that other violations of practical reason don't. Whether I act in accordance with

¹⁸ As he states it: "S's performing ϕ is morally permissible if and only if there is no available alternative, ψ , that S has both more requiring reason and more reason, all things considered, to perform, where a requiring reason is just a reason that has some moral requiring strength" (Portmore 2011, 137).

¹⁹ The same is true of the account defended by Schmidt (2023), who works with a definition of moral reasons that stipulates that moral reasons make actions required in the absence of competitors (cf. Schmidt 2023, 99, fn. 3). This shouldn't be seen as an objection, as neither Portmore nor Schmidt purport to provide an informative analysis of requirements but are rather concerned with an account of the nexus between requirements and reasons. The point is that such a view is of no use to the Deontic Buckpassing Approach.

my non-moral reason is, at least other things equal, ‘my own business’ and I am at liberty to deny other people’s involvement with it. This is different with moral requirements. If I am criticized for acting wrongly, it’s not appropriate to reply “Mind your own business” or “I haven’t asked you for advice”. Notably, this is so even if the criticism is raised by a person who is not directly affected by my wrongdoing. In a good sense, whether one’s actions are morally permissible concerns everyone. This is related to the fact that performance of moral duty can be *demanded*. This is not to say that it is advisable to demand such performances in each and every case, but that such performances are in principle fitting objects of demands. Relatedly, it seems appropriate (again, other things being equal) to respond to violations of moral requirements with certain emotions that are often associated with the attitude of blame: emotions like resentment, indignation, or guilt.

The problem is that it is difficult to see how rationalism can fully account for this socio-emotional dimension of moral requirements. If moral requirements are just norms of practical reason, why do they differ so markedly from other norms of practical reason? Suppose I suffer from chronic neck tension, and because of that, I have prudential reasons to exercise daily, which are on balance decisive. It doesn’t follow that whether I exercise is anyone else’s concern, it doesn’t follow that exercising can be demanded of me, and it also doesn’t follow that it’s appropriate to react in emotional ways to my not exercising. It would follow, however, if I were morally required to exercise daily. That moral obligations have these implications calls for explanation, and it is far from obvious that rationalists can provide one.

§3 Sentimentalism

This is a good point to turn to the sentimentalist version of deontic buckpassing. According to the sentimentalist, the moral requirements that an agent is subject to are not to be analysed in terms of the reasons of this agent, but rather in terms of the reasons of *other* (or *all*) persons; and not in terms of reasons for compliance, but in terms of reasons to *react* to *non-compliance* by way of having certain emotions – namely, the *reactive* emotions that Strawson (1962) famously drew attention to and which (at least according to a widely shared conception) constitute attitudinal blame.²⁰ Thus, on a simple construal, sentimentalist buckpassers claim:

²⁰ Throughout this essay, I am concerned with the *attitude* of blame, which should be distinguished from the act of blame that communicates the attitude. See Wallace (1994) for an elaborate emotional theory of blame and Menges (2017) for a recent defence.

Simple Sentimentalism: For A to be morally required to ϕ is for it to be the case that there would be sufficient reason to blame A if A were not to ϕ .²¹

It is evident that the sentimentalist has great advantages over the rationalist when it comes to explaining the socio-emotional dimension of moral requirements. The social dimension is written into the analysis by the assumption that moral requirements are constituted by reasons for *others*, the emotional dimension by the assumption that they are constituted by reasons for *emotions*. This means that sentimentalists are also in a good position to account for what is plausible about non-cognitivism about moral judgements, which I take to be that moral judgements *often* come with certain emotions, and that this does not seem to be a coincidence. Emotivist theories of moral judgement go too far in identifying moral judgements with emotions, but if sentimentalism is correct, then we can be cognitivists and still account for a strong connection between moral judgements and emotions, for according to the sentimentalist, moral judgements are judgements about reasons for emotions (or at least judgements about facts that consist in reasons for emotion facts).

While these are great advantages, sentimentalism also faces important challenges. One set of challenges arises from conceptions that take moral blame to involve the belief, or at least some kind of representation or thought, that the agent violated a moral obligation.²² This threatens to make the sentimentalist analysis circular. The objection echoes the one that Ross put forward against Ewing's analysis of goodness in terms of fitting admiration and other pro-attitudes: "admiration is not a mere emotion; it is an emotion accompanied by the thought that that which is admired is good."²³ Ewing's reply was that admiration is a response to good-making features rather than goodness and therefore need not be seen as involving thoughts about goodness.²⁴ Sentimentalist deontic buckpassers can give an analogous response: the reactive emotions that constitute blame are responses to wrong-making features and need not involve thoughts about wrongness.

This seems to me a tenable position about blame. It also seems in a better position than its rival to explain *recalcitrant* blame, which is held despite the subject's judgement that the agent did

²¹ Cf. Gibbard (1990, 42); Darwall (2006, 91–99); Skorupski (2010, 291–95). A predecessor of sentimentalism is Mill's view of moral obligations, which connects wrongness to what ought to be punished, where the punishment in question may be understood in terms of negative reactive attitudes (cf. Mill 1861, 72–73).

²² For example, Wallace maintains that "moral reactive attitudes are explained by the belief that some moral obligation has been violated" (1994, 37) and Rosen holds that "resentment of X for A involves the deontic thought that ... A was wrong or impermissible" (2015, 76).

²³ Ross (1939, 278).

²⁴ Cf. Ewing (1948, 172); see also Rabinowicz and Rønnow-Rasmussen (2004, 396).

not actually do anything wrong.²⁵ But it is worth mentioning that sentimentalist buckpassers can also accept conceptions of blame that take blame to involve normative content. For one, they can agree with views on which that normative content is not deontic. This is true, for example, of views that take blame emotions, like resentment or indignation, to be a combination of anger directed at A for A's ϕ -ing and a judgement or thought that that very anger is justified or deserved.²⁶ For another, sentimentalist buckpassers could even accept the view that blame attitudes involve deontic thoughts but then analyse moral obligations in terms of reasons for the emotional components of blame that do not presuppose deontic thoughts. For example, if resentment is anger caused and accompanied by the thought that the action was wrong, sentimentalists may still define wrongness in terms of reasons for anger. Although these matters need more careful treatment than I can give them here, it seems premature to regard the circularity objection as decisive, as there are various options that the sentimentalist can pursue in response.

Another set of challenges for sentimentalists arises because the *extension* of wrongness and blameworthiness seems to come apart in certain cases. It is widely (although not universally)²⁷ acknowledged that wrongness is not sufficient for blameworthiness, and some authors argue that it is not even necessary.²⁸ Some of these issues may be solvable by introducing further conditions into the analysis. For example, sentimentalists might say that for an action to be morally required is for it to be the case that it would be blameworthy to omit *without excuse*.²⁹ Other issues will depend on substantive disagreements: it is difficult to see how sentimentalists can avoid the implication that wrongness is necessary for blameworthiness, but the arguments against this implication can also be questioned.³⁰ I will leave open whether sentimentalists can resolve the extensional worries with their view. Instead, I will focus on two deeper problems with sentimentalism that would persist even if these worries can be resolved.

The first of these problems has to do with the normativity of moral obligations. According to sentimentalism, the normative force of moral obligations does not primarily apply to the agent who

²⁵ For example, it seems that one can irrationally feel guilty without believing that one acted wrongly (cf. Gibbard 1990, 148–49), which means that guilt cannot involve the belief that one acted wrongly. This problem of recalcitrance has led some cognitivists to focus on *representations* or *thoughts* rather than judgments. On these views, irrational guilt represents one's action as wrong even though the agent believes it to be permissible. But it is contentious whether this "quasi-judgmentalist" view solves the problem (cf. D'Arms and Jacobson 2003; 2023, Ch. 5).

²⁶ See Pereboom (2014, 128). Deserved or justified anger is in turn naturally understood in terms of the reasons that A's ϕ -ing provided for being angry with A.

²⁷ Compare Rivera-Lopéz (2006) and Bruno (2023).

²⁸ See e.g. Driver (1992); Capes (2012); Graham (2014).

²⁹ See e.g. Darwall (2010, 149).

³⁰ For defences of this implication, see e.g. Rosen (2015, 76) and Kiesewetter (2017, 28–31).

is subject to the obligation, but to others, or at least equally well to others. In my view, this gets things wrong: moral obligations are in the first instance normative for the agent. The point here is not that sentimentalists in principle cannot accommodate the assumption that morality is normative for the agent. It is open to sentimentalists to argue that facts about what kind of behaviour gives rise to reasons for blame themselves constitute (or entail in other ways) reasons to avoid the relevant behaviour.³¹ But this would mean that moral requirements are normative for the agent only because they are normative for everyone, i. e. only because they are constituted by reasons for everyone to react to the agent in certain ways and such reasons entail or presuppose reasons for the agent to comply. In my view, this explanation is implausibly indirect. The point of moral obligations is first and foremost to guide agents towards moral action. Even if it also provides a guide to emotional reactions to other people's immoral behaviour, it does not seem correct to think of this function as the primary one that explains the former.

The second problem for sentimentalism is that the attitudes it focuses on are backwards-looking, while moral obligations have (again: primarily) a forward-looking character. This is of course related to the point that moral obligations have an important role to play in first-personal deliberation about what to do. But even if we are concerned with the relevance that moral obligations have for others to relate to the subject of the obligation, it strikes me as misleading to put the focus on attitudes that make sense only after the deed rather than attitudes that may be able to exert an influence on whether the agent complies. Moral obligations play an important forward-looking role in deliberation, advice and in exerting pressure on others that isn't well-captured in terms of reasons for backwards-looking attitudes.³²

§4 The expectation view

With all of these potential benefits and difficulties for deontic buckpassing views on the table, let me now turn to a new proposal of how moral obligations can be analysed in terms of reasons. I

³¹ See Darwall (2006, 97–99) and Skorupski (2010, 295–301).

³² Nye's sentimentalist analysis, according to which "to judge that it is morally wrong for X to ψ is to judge that it is, on balance, fitting for X to feel obligated not to ψ " (2022, 131) captures this prospective character of moral obligations better than other sentimentalist views, which focus on blame emotions. However, Nye's proposal still faces the first problem, and it also faces additional worries about circularity. By contrast, Rowland's "amends-based" analysis, according to which "what it is for A's ϕ -ing to be morally ... required is just for A to have sufficient reason to ϕ and for there to be sufficient reason non-role dependent reason for A to have (non-instrumental) pro-attitudes towards her making amends for not ϕ -ing if she does not ϕ " (2019, 174) incorporates a first-personal action-guiding dimension, but it still entails that what distinguishes moral requirements from other norms of practical reason are reasons for backwards-looking attitudes.

want to suggest that moral obligations can be understood in terms of both reasons for the agent to act and reasons for everyone to *expect* the agent to act.

The expectation view: For A to be morally required to ϕ is (i) for A to have decisive reasons to ϕ that are (ii) also sufficient agent-neutral reasons to expect A to ϕ .

Before turning to the virtues of this account in §6, §5 will discuss two worries one might have with the expectation view, and this section will focus on explaining its content. I will start with condition (i). By A's having 'decisive reason' to ϕ , I mean that there is a non-empty set of reasons for A to ϕ , such that this set of reasons is stronger than any competing set (i.e. stronger than any set of reasons *against* ϕ -ing or set of reasons for relevant alternatives to ϕ -ing).³³ This much is familiar from rationalist proposals. But condition (ii) qualifies the relevant reasons, by requiring that A's reasons for compliance are at the same time sufficient agent-neutral reasons to expect A to ϕ .

By an 'expectation', I mean an expectation not in the *predictive sense*, in which expecting essentially involves having a positive credence that something will happen, but in what I will call, for lack of a better term, the *directive sense*.³⁴ While expecting in the predictive sense is belief-like (perhaps even identical to belief), expecting in the directive sense is more like a desire. These notions are independent. I certainly expect Donald Trump to lie in the predictive sense, while I do not expect him to lie but rather expect him *not* to lie in the directive sense.³⁵

Although a directive expectation plausibly involves the desire that a person acts in a certain way, it is much more complex than a simple desire with that content. Two characteristics make the notion more precise. The first is that directive expectations involve dispositions to reactive attitudes. These involve the angry emotions that we have already encountered in the discussion of sentimentalism, but in other cases might simply amount to a sort of personal disappointment rather than resentment or indignation. Expectations (I will henceforth omit the qualification 'directive')

³³ There is some controversy about what competing reasons are (cf. Schroeder 2015, 163–64), but this controversy can be bracketed for the purposes of this essay. It is worth noting that talk of sets of reasons being stronger or weightier is non-committal with respect to how and which reasons aggregate and therefore does not imply a naïve 'kitchen scale model' of weighing reasons. It is also compatible with widespread incommensurability between reasons.

³⁴ Directive expectations are sometimes called "normative expectations". I don't object to this terminology, but I want to avoid its (potential) connotation that the attitude in question involves a belief or representation with normative content. This may be so (see §5), but it is controversial and I don't want to presuppose it.

³⁵ As this example demonstrates, I use 'directive expectation' in a purely directive sense that involves no predictive elements at all (directive expectations might involve beliefs about what the agent *can* do, but this is a different matter). I do not deny that there are also notions of 'expectation' that involve both predictive *and* directive aspects, but these are not notions I focus on.

thus differ from other desires that someone will act in a certain way by way of disposing us to be disappointed in that person if she does not perform the action. Crucially, the kind of disappointment at issue is “reactive” or “agent-directed”: it is essentially disappointment *in* a person, not merely disappointment *that* a person has acted in a certain way (the latter being a kind of disappointment that all sorts of desires or hopes can give rise to, while only the former is characteristic for the frustration of a directive expectation).³⁶

Secondly, it seems to me plausible that there is an important connection between the attitude of an expectation and the speech act of a *demand*. One of the main ideas of Austin’s speech act theory was that speech acts can be identified by appeal to an attitude that they purport to express. For example, it seems plausible to think that purporting to express belief is part of what it is for an act to be an *assertion*. That does not mean, of course, that one cannot assert what one does not believe, but it means that such an assertion is insincere. I want to suggest that the same is true for the speech act of a demand and the attitude of an expectation: it is in the nature of a demand to purport to express an expectation, which means that to demand what one does not expect the other to do involves some kind of insincerity or at least a lack of genuine engagement. Expectations can thus be understood as the kind of desires concerning other people’s behaviour that we purport to express when performing the speech act of a demand.

The expectation view appeals only to *agent-neutral* reasons for expectations. By an agent-neutral reason, I mean a reason that is necessarily shared by all agents, unless there is some particular explanation of why an agent lacks them – for example because of incapacity or lack of standing. Practical reasons need not be agent-neutral in order to give rise to moral obligations, as many obligations do depend on contingent properties of the obligated agents. Promissory obligations are a case in point. I maintain, however, that reasons for expectations must be agent-neutral in order to give rise to moral obligations. Most importantly, role-dependent reasons for expectations – the reasons we have for expecting certain kinds of behaviour because we are, for example, parents, partners, supervisors or superiors do not count. More exactly, they count only insofar as they are *also* agent-neutral reasons, and so they are covered by the present formulation of the expectation view.

An example may help to illustrate this point. Johnny is a clerk in a local organic grocery store. His schedule today involves stocking up the shelves in various sections. As Johnny likes his job and

³⁶ For the difference between “reactive” and “propositional disappointment”, see also Telech and Dahan Katz (2022, 861–63), although I disagree with them that reactive disappointment is essentially related to comembership in what they call “thick evaluative relations” (a topic that is beyond the scope of this essay).

finds pleasure in being a reliable employee, the fact that his schedule involves stocking up the shelves is a reason for him to do so, and supposing that no other important considerations are in play, we can assume that it is stronger than competing reasons. The same fact also gives Johnny's boss, Dana, sufficient reason to expect Johnny to stock up the shelves. Intuitively, however, Johnny is not morally required to stock up the shelves. The suggested analysis avoids this result because Dana's reason to expect Johnny to stock up the shelf is not agent-neutral, but depends on her role as Johnny's boss.

Contrast this with the case of Francesco, the captain of a huge cruise ship that collided with a rock. The responsibilities of a captain involve staying on the ship until all passengers have left it. Suppose this is a decisive reason for him to stay on the ship. Plausibly, the same fact also gives the passengers and the members of his crew a role-dependent reason to expect him to stay on the ship. But in this case, we might well want to say that Francesco does not merely have a role obligation, but also a *moral* obligation to stay on the boat. According to the expectation view, this will be so if there are not only role-dependent reasons to expect him to stay on the ship, but if, in addition, facts about his responsibilities as a captain give *everyone*, independently of their role (e.g. as a passenger or crew member), reasons to expect him to stay on the ship. This strikes me as a plausible verdict in this case. The expectation view not only suggests a natural way to distinguish between moral and role obligations but also gives a plausible account of the cases where the two seem to coincide.

You may disagree with me that Johnny lacks a moral obligation to stock up the shelves or that Francesco has a moral obligation to stay on the ship. This is not a problem for the expectation view as long as you are then also willing to accept that there are sufficient agent-neutral reasons to expect Johnny to stock up the shelves or no sufficient agent-neutral reasons to expect Francesco to stay on the ship. These strike me as the right conclusion to draw if you disagree with my first-order verdicts on the cases. The examples were only meant to illustrate the general point that cases in which role obligations go along with moral obligations are plausibly cases in which agent-relative reasons for expectations go along with agent-neutral ones.

By the phrase 'sufficient agent-neutral reasons', I mean agent-neutral reasons that are not weaker than competing *agent-neutral* reasons. That is, the sufficiency relates only to the balance of agent-neutral reasons, while agent-relative reasons are bracketed. (I realize that this use of 'sufficient agent-neutral reasons' is somewhat artificial, but it will facilitate the following discussion.) This poses two questions: Firstly, why sufficient and not decisive reasons to expect? Secondly, why are the reasons sufficient only with respect to the agent-neutral reasons and not all things considered?

The answer to the first question is that obligation-corresponding reasons for expectations should be seen as *justifying* expectations without necessarily making their absence unjustified. It can be fully rational to refrain from expecting people to comply with their moral obligations, even if such expectations would be justified. Just think about how cognitively demanding it would be for you to have all the expectations that you could justifiably have! This is why these reasons should be seen as (*ceteris paribus*) sufficient (i.e., not weaker than competing reasons) rather than decisive (i.e., stronger than competing reasons).

The answer to the second question is that obligation-corresponding reasons for expectations can be defeated by agent-relative reasons in special cases. To see this, consider an example. Ivan is fighting in an unjust war, and he is about to execute a group of innocent civilians. Tanya, an enemy combatant, can prevent this, but only by killing Ivan. Let us assume that it is Tanya's moral obligation to prevent the execution and kill Ivan. The expectation view entails that there is an agent-neutral reason to expect Tanya to kill Ivan, a reason that, in principle, everyone has. But now consider Ivan's mother Olga. Would she be justified in expecting Tanya to kill Ivan? We should not assume that Ivan's actions are malicious: he may be a very young soldier, misled by manipulation and misinformation, acting under pressure or even force. Surely, Olga has strong agent-relative reasons to desire that Tanya doesn't kill Ivan. Intuitively, these reasons compete with her reasons to expect Tanya to kill Ivan (which, to recall, is itself a desire-like state), and they may be strong enough to make it inappropriate for Olga to have such expectations. In this case, Olga's agent-neutral reasons to expect Tanya to kill are plausibly outweighed by agent-relative ones, and yet this does not seem to affect Tanya's obligation to kill. Again, you may not agree with the substantive moral assumptions driving this particular example, but I hope that the example still illustrates the possibility of a case that has this general structure. And since such cases seem possible, we should not assume that the agent-neutral reasons for expectations that correlate with moral obligations must be all things considered sufficient for everyone in each particular situation.

To see how the expectation view works, let us consider three possibilities. The first is that both conditions (i) and (ii) hold. Suppose that, in a variant of *Widowed Neighbour*, Olivia has promised Mary to take her to the game, and suppose that her promissory reason is stronger than competing reasons in the situation. In such a case, I would hold that the promise not only provides a decisive reason to take Mary to the game but also a sufficient reason to expect of Olivia that she takes Mary to the game. The expectation view entails that this action is morally required. Next, suppose that, in another variant of *Widowed Neighbour*, Olivia has promised Mary to take her to the game, but her promissory reason is outweighed by a more important consideration (assume that Olivia has to

take care of her sick child, for example). In this case, the first condition is not satisfied and the promise does not generate a moral requirement.³⁷ Finally, consider the original example, in which Olivia has not promised to take Mary to the game, but just considers it as a nice gesture. In this case, I argued that the fact that Mary would enjoy going to the game might be a decisive reason to invite her even though it does not plausibly generate a moral requirement. This is because condition (ii) fails to apply, i.e. Olivia's reasons to invite Mary are not sufficient agent-neutral reasons to expect her to do that. This is because the fact that Mary would enjoy going to the game is in fact *no reason at all* to expect her to do it.

A natural question is whether there could be cases in which there are sufficient agent-neutral reasons to expect an agent to ϕ without the agent having decisive reason to ϕ , and why, if there are no such cases, the reference to decisive reasons to ϕ is not redundant for the analysis. My answer is that it is indeed difficult to see how there could be sufficient agent-neutral reasons for expecting agents to perform actions they have sufficient reason not to perform. This means that if our aim was merely to state necessary and sufficient conditions for moral obligations, we could indeed focus on reasons for expectations alone. However, my aim in this essay is not merely to state necessary and sufficient conditions, but to give an account of what moral obligations are. And as I made clear in the discussion of sentimentalism, I take it to be essential for obligations that they involve reasons for the subject to act as they are obligated to. In my view, it is these reasons that play the primary role in constituting obligations and the function of condition (ii) is mainly to qualify the relevant subset. While the biconditional 'A is morally required to ϕ iff there are sufficient agent-neutral reasons to expect A to ϕ ' may well be true, it is in my view a mistake to think that this biconditional could capture the nature of moral obligation (neither could it capture the meaning of the concept).

Finally, how does the expectation view relate to Kant's (1785) famous claim that moral requirements are categorical, i.e. independent of the subject's contingent desires? Like other rationalist views, the expectation view entails that *if* there are no categorical reasons for action (e.g. Williams 1979), then there are no categorical moral requirements – i.e., then either moral requirements are not categorical (e.g. Foot 1972) or there are no moral requirements (e.g. Mackie 1977). But if one allows for categorical reasons for action, as many contemporary authors do, the expectation view allows for existing moral requirements that are categorical. The expectation view is thus neutral on the categoricity of moral requirements. However, those who think that

³⁷ It may still generate a *pro tanto* moral obligation, but I am here concerned with the notion of an overall moral obligation. See Kiesewetter (2023) on *pro tanto* moral obligations.

categoricity is built into the nature of moral obligations are free to qualify the reasons in the expectation view's first condition as categorical.

§5 Two worries

A natural worry about the expectation view is that expectations have representational content, and that this introduces circularity problems, analogous to those that have been discussed in the context of sentimental analyses of deontic or evaluative properties (see §3). As in these other cases, it is important to note that proponents of the expectation view can take on board much of what moderately cognitivist views about the relevant attitude might claim. While they have to deny that expectations involve representations with contents involving the concept of moral obligation, they can agree that expectations involve representations of other kinds. In particular, they can accept that expectations represent their objects as having certain non-normative properties (properties that constitute or provide reasons for expectations) and they can accept that expectations represent their objects as having certain normative properties that do not presuppose deontic properties (or at least not the property of being morally required). For example, proponents of the expectation view can accept that expectations not only involve dispositions to react with the kind of emotions characteristic for blame, but also the thought that these emotions would be (*ceteris paribus*) justified or deserved.

Insofar as it is possible to have recalcitrant expectations, i.e. to expect actions while judging that these actions cannot justifiably be expected, this puts pressure on the view that expectations involve normative representations.³⁸ But most importantly for my purpose, there are strong independent reasons for thinking that expectations need not involve representations with moral content.³⁹ For one, it is implausible to think that expectations that we have for role-dependent reasons represent agents as morally required. By expecting Johnny to stock up the shelves, for example, Dana does not plausibly represent him as being morally required to do that. For another, openly immoral expectations seem possible without creating much mental tension. We have no difficulties imagining the openly immoral mafia boss to expecting strict loyalty from his clan members, hence performing acts that he considers morally impermissible and blasphemous. But if expectations represented agents as morally required, such expectations would be very hard to make sense of.

³⁸ See also note 25.

³⁹ See also Wallace (1994, 36–37).

It might be replied that while these cases show that expectations do not represent their objects as *morally* required, expectations could still represent their objects as required in a more generic sense. I must admit that I find the supposed content of such a representation elusive – it’s not clear to me what it means to say that an action is required in a generic sense if not that the action is subject to legitimate expectations.⁴⁰ In any case, even if all cases of expectations could unproblematically be seen as cases of representations with deontic content, I don’t see the force of why we *should* see expectations in this way. As indicated, there are alternative conceptions of expectations that are compatible with the expectation view, and the attractions of the expectation view might be reason enough to reject a moralized or otherwise ‘deonticized’ conception of an expectation.

A related worry might be that the particular idea of an *agent-neutral reason* for an expectation presupposes the notion of a moral obligation. If the facts that constitute the relevant agent-neutral reasons are facts about moral obligations, we would again run into a circle. But as before, I don’t see the force of thinking that agent-neutral reasons must be seen as facts about moral obligations. On an at least equally plausible view, agent-neutral reasons for expectations are provided by the non-normative facts that make an action morally required (i. e., the right-makers), such as the fact that one has promised to do something, or the fact that an action would prevent harm, etc. Indeed, on reflection this view has significant advantages over its rival.

Firstly, it is a widely accepted point that morally good persons care non-derivatively about right-making features and are not primarily motivated by thoughts about rightness.⁴¹ This, in turn, is widely taken to support the view that moral reasons are not (at least not primarily) provided by rightness facts, but by the features of actions that make them right. This argument completely carries over to reasons for expectations. If morally good persons are primarily concerned with right-makers rather than rightness facts, then not only their actions, but their expectations, too, will be based in right-making rather than rightness facts, and this supports the view that agent-neutral reasons for expectations are primarily constituted by right-makers rather than rightness facts.⁴²

⁴⁰ There may be a generic “standard-relative” sense of being required in which all sorts of conventions generate requirements, no matter what reasons there are to comply (cf. Broome 2007). But since holding an act to be required in this sense does not involve any kind of normative endorsement, it does not seem a plausible sense in which expectations could be said to represent their objects as required.

⁴¹ Famously, some have argued that there is even something objectionably fetishistic about motivation by rightness, see e.g. Smith (1994, 74–76). For recent criticism of this view, see Johnson King (2022).

⁴² Note that the expectation view is consistent with the assumption that facts about rightness or moral requirements are *derivative* reasons for expectations with no independent normative force (cf. Parfit 2011, 39, on the buckpassing account of value). In the words of Johnson King (2019), “we can have our buck and pass it, too”.

Secondly, there are *pro tanto* agent-neutral reasons for expectations that correspond with outweighed moral reasons and therefore cannot be provided by the fact that an action was morally required. For example, if you are morally permitted or required to break a promise, there is still an outweighed agent-neutral reason to expect you to keep your promise, but since you are not morally required to keep your promise, the reason cannot be the fact that you are morally required.

§6 Some virtues of the expectation view

The expectation view combines many of the virtues of the rationalist and the sentimentalist analysis, while at the same time avoiding their problems. Like the rationalists, proponents of the expectation view can give an elegant account of the first-personal normative authority of morality in terms of reasons for compliance, which avoids the typical pitfalls of the why-be-moral challenge. And like the rationalist, they can account for the plausible assumption that rational agents act in accordance with their moral judgements. But the focus on reasons to act that are also reasons to expect seems to solve the kind of wrong-kind-of-reasons problems that plagues the rationalist account. At the same time, it seems to account for the social and emotional dimension of moral obligations, as it entails that moral obligations come with reasons not only for the agent but also for all of us, and reasons for attitudes that involve dispositions to certain emotions. For this reason, proponents of the expectation view are in as good a position as the sentimentalists to explain why there is a non-contingent relation between moral judgements and emotions, without giving up on a cognitivist account of moral judgement. However, they are in a much better position to explain the deliberative and forward-looking character of moral obligation – not only because of their essential appeal to reasons for action, but also because they focus on the forward-looking attitude of an expectation rather than the backward-looking reactive attitudes.

The expectation view does not merely combine the benefits of rationalism and sentimentalism, however; it has genuinely new and distinct virtues. In what follows, I will argue that it provides illuminating answers to two longstanding challenges of moral philosophy: the challenge to explain the possibility supererogation, and the challenge to explain the imperatival character of morality.

Let me start with the so-called “paradox” of supererogation: If it’s possible to go beyond one’s duty, then it must be morally permissible to refrain from actions that are morally speaking better. But how could it be permissible to act in morally suboptimal ways? Not everyone agrees, of course, that there is a real paradox here. But almost everyone agrees that the phenomenon calls for explanation. An analysis of the property of being morally required is a very natural place to look for such an explanation. But neither the rationalist nor the sentimentalist analysis is helpful here.

The rationalist, as we have seen, has serious difficulties accommodating supererogation without referring to the very property that it aims to analyse. The sentimentalist does not face this extensional problem: although there are reasons to act in supererogatory ways, plausibly there are no reasons to blame people for refraining from supererogatory actions that would have to be excluded from the analysis of wrongness.⁴³ But we cannot say that the sentimentalist analysis does much to illuminate the phenomenon of supererogation.

The expectation view, I think, provides a distinct and plausible explanation of supererogation. According to the expectation view, morally suboptimal actions can be permissible if there are no sufficient agent-neutral reasons to expect the agent to perform the action. This just follows from the suggested analysis of moral obligation: in order for an obligation to be in place, we need not only reasons to act, but also sufficient agent-neutral reasons to expect the agent to act. Cases of supererogation are thus cases in which the moral reasons of an agent favour an action, but we lack sufficient agent-neutral reasons to expect the agent to perform the act. This can happen in two ways. Firstly, there may be agent-neutral reasons *against* expecting the agent to perform the morally favoured act, which are strong enough to undermine the sufficiency of the agent-neutral reasons *for* expecting the act. This seems to me a plausible analysis of the heroic cases of supererogation. In *Burning Building*, the moral reasons the agent has to save the lives of others are plausibly also reasons to expect the agent to act. But the risks that she would take in order to save the lives of others are also weighty reasons against expecting her to do it, and plausibly they are strong enough to make expectations in such cases unjustified. Secondly, there can be cases in which the agent's moral reasons for action are *no reason at all* to expect the agent to act, and this is why the second condition for moral obligations does not apply. This is what seems to be going on in the non-heroic cases of supererogation, such as *Widowed Neighbour*. My claim is that the fact that it would be nice for a person if one performed a certain act, while being a moral reason to perform this act, is *no reason at all* to expect one to perform it.

The expectation view thus offers an answer to the question how it can be possible that a morally suboptimal act is morally permissible. In addition, it yields the intuitively correct results when

⁴³ More exactly, it is plausible that there are no object-given reasons for blame in such cases. If there are state-given reasons for blame, these would still pose a wrong-kind-of-reasons problem for sentimentalists. However, I agree with those who think that putative state-given reasons for attitudes are better understood as object-given reasons for *wanting* to have these attitudes rather than reasons for these attitudes themselves (see e.g. Skorupski 2010, 87–89; Parfit 2011, App. A; Way 2012). So I assume throughout this essay that there are no state-given reasons for attitudes (note that this is consistent with accepting the existence of the kind of state-given reasons *against* attitudes that Schroeder (2012) has argued are reasons of the right kind). If there were such reasons, they would have to be excluded from both the sentimentalist as well as the expectationalist analysis, just as they would have to be excluded from the buckpassing analysis of value in terms of reasons for attitudes (see Gertken and Kiesewetter [2017] for an overview over this debate).

combined with pre-theoretically plausible judgements about reasons for expectations and their strength. These are, of course, substantive normative judgements, but this is as it should be. Those who reject supererogation are free to reject the first-order assumptions about reasons for expectations that I have made. An analysis of moral obligation can only locate the disagreement over whether there are cases of supererogation, and in my view the expectation view not only locates it correctly, but also in a way that sheds a new light on it and may be instructive for the philosophical dispute.

Another virtue of the expectation view that I wish to highlight is its ability to explain the imperatival character of morality. Recall G.E.M. Anscombe's thesis, put forward in her famous essay "Modern Moral Philosophy" (1958), that the concept of a moral obligation is unintelligible – and should be jettisoned – outside the context of a divine command conception of ethics. While many would reject her radical conclusion, it is widely agreed that Anscombe raises an important challenge for non-theistic modern moral philosophy. As I see it, this challenge is best described as a challenge about how to make sense of the imperatival character of moral obligation. For one, how can there be demands without demandants? For another, why do moral norms amount to demands in the first place? It's clear why they do if moral norms are divine commands, but far from obvious if they are not. Moreover, not all practical norms are demands. If I have most reason to eat chocolate ice cream, it does not seem to follow that I am under a demand to eat it. So there is definitely something to explain why moral obligations, in particular, have imperatival character.

There is a group of views in moral philosophy, according to which moral obligations arise out of some sort of authoritative act, and such *voluntarist* views are, in principle, in a good position to explain the imperatival nature of moral obligations. Famously, Kant (1785) claims that moral obligations originate in a law that we give ourselves. Darwall (2006) holds that moral duties are based in demands of other persons and Wolf (2009) maintains that they are "social commands". Others again attempt to revive a divine command conception of ethics (Adams 1999). This is not the place to discuss such views, but it seems clear that they all either face significant challenges with capturing pretheoretically plausible judgements about moral obligation, or have highly contentious theoretical commitments, or both.

Wallace's relational approach to morality purports to account for the imperatival character of morality as well (Wallace 2019). According to this approach, all moral requirements arise from *directed obligations* – obligations that are owed to claimholders. Assuming that "it is characteristic of the paradigm cases of directed obligation that we register them in practical thought as

demands”⁴⁴, Wallace argues that his view is in an exceptionally good position to explain the imperatival character of morality. One problem with this explanation is that it faces challenges with cases of duties that do not seem to be owed to claimholders – such as certain kinds of environmental or aggregative duties. While Wallace avoids voluntarist assumptions, his explanation, too, involves highly contentious theoretical commitments. Another and even more significant problem is that Wallace’s explanation presupposes that *directed* obligations are “correctly registered”⁴⁵ as demands, and this seems no less in need of an explanation than the assumption that moral obligations are correctly registered as demands. We still want to know *why* they are registered as demands if, firstly, there is no demandant and, secondly, other practical norms are not registered as demands.

The expectation view, I think, provides an illuminating explanation of the sense in which moral obligations are demands. Consider the following argument:

1. Moral obligations entail sufficient reasons for expectations.
2. It is constitutive for demands that they purport to express expectations.
3. For every speech act S that constitutively purports to express attitude α : S is fitting (appropriate, correct) if, and only if, there are sufficient reasons for α .
4. Therefore, if there are sufficient reasons for expecting A to ϕ , then it is fitting to demand of A that s/he ϕ -s (from 2&3).
5. Therefore, moral obligations entail the fittingness of a demand (from 1&4).

Premise (1) is not exactly accurate. As we have seen, the expectation view provides a constitutive explanation only for the weaker claim that moral obligations entail sufficient reasons for expectations *unless* these reasons are defeated by competing agent-relative reasons. The conclusion of this argument must therefore be restricted in a similar way: moral obligations entail that a demand is fitting only for speakers who lack stronger agent-relative reasons against expectations. This does not limit the explanatory power of the expectation view in any relevant respect. We do not want to explain why it is fitting *for Olga* to demand that Tanya complies with her moral obligations – because plausibly, it isn’t. What we want to explain is rather why compliance with moral obligations can be fittingly demanded by speakers who are not in such extraordinary circumstances. It is therefore unproblematic, for the present purposes, to abstract away from agent-

⁴⁴ Wallace (2019, 55)

⁴⁵ Wallace (2019, 53)

relative reasons for and against expectations, and for the sake of simplicity I will disregard them in what follows.

Premise (2) is an assumption I have appealed to above when elucidating the notion of a prescriptive expectation. I don't know how to argue against it other than by counterexample, and I cannot think of one. (3) expresses a plausible general connection between reasons for attitudes and fittingness norms for speech acts. (4) and (5) follow logically from the premises.

It is important to stress that the fittingness of a speech act, in the sense of 'fittingness' relevant for this argument, does not entail a reason (let alone sufficient reason) to *perform* that speech act. Just as the correctness of an assertion does not, by itself, entail a reason for making the assertion, the fittingness of a demand does not entail a reason to demand.⁴⁶ And note that this is also not what we want to explain: in my view, it's not a desideratum for a theory that moral obligations entail *reasons* to demand. We often have strong reasons not to demand compliance with a moral obligation, and it's difficult to see why moral obligations should necessarily come with *pro tanto* reasons to perform certain speech acts. Rather, what we want to explain is that morally required actions are fitting objects of demands. This is what the expectation view explains if we assume the kind of connection between fitting speech acts and reasons for attitudes expressed in (3).⁴⁷

A potential counterexample to (3) is the speech act of assertion. Isn't assertion fitting if, and only if, it is correct and correct if, and only if, it is true? If so, (3) could be true only if 'p if, and only if, there are sufficient reasons to believe p' holds true. But this assumption seems false.

In defence of (3), the following two responses seem to me tenable. Firstly, the fittingness of a speech act should plausibly be understood in relation to *objective* reasons, i.e. reasons that need not be part of anyone's body of evidence. Once this is made explicit, the assumption that 'p if, and only if, there are sufficient reasons to believe p' is no longer implausible. Arguably, p is necessarily a sufficient *objective* reason to believe p, and not-p is necessarily an *objective* reason *against* believing p that is stronger than any objective reason *for* believing p.⁴⁸ It follows from these assumptions that p is true if, and only if, there are sufficient objective reasons to believe p.

Alternatively, proponents of (3) might also reject the assumption that fittingness and correctness are co-extensional and argue that assertions are fitting only if the speaker has sufficient

⁴⁶ Or so it seems plausible to me. Some authors aim to analyse fittingness and correctness in terms of reasons, and they thereby commit themselves to the entailment (cf. Schroeder 2010, 40). In my view this constitutes a problem for the analysis.

⁴⁷ A possibility worth thinking about is that the fittingness of speech acts can be *analysed* in terms of the reasons for the *attitude* that the speech act purports to express.

⁴⁸ See also Schroeder (2021, 226).

evidence for believing the content of the assertion.⁴⁹ For example, defenders of the so-called knowledge norm of assertion might argue that it is fitting to assert (and one has sufficient reason to believe) only what one knows or is in a position to know.⁵⁰ Others again might hold that some sort of evidential support is sufficient for the fittingness of assertion.⁵¹ Discussing these options is beyond the scope of this essay, but the important point is that once we allow for a distinction between correctness and fittingness, it's unclear why the case of assertion should be a counterexample to (3).

If an argument along the suggested lines can be defended, it would provide a fundamental explanation of the imperatival character of moral obligations without any voluntarist or relational assumptions. This is a powerful reason to accept the expectation view.

§7 Conclusion and outlook

In this essay, I have argued that moral requirements can be understood in terms of reasons. Such an analysis is generally attractive because it is informative and contributes to a unified understanding of normativity. The expectation view illuminates several features of moral obligations, including their first-personal authority, as well as their social, emotional, and motivating dimension, and it avoids the major difficulties of the main approaches to a reasons-based analysis of moral obligations. In addition, I hope to have shown that the expectation view promises to solve some of the most pressing problems of moral philosophy, including the why-be-moral challenge, the paradox of supererogation and Anscombe's challenge to modern moral philosophy.

The focus of this essay was on the notion of an overall moral obligation or requirement. It is natural to ask what the expectation view implies for related notions, in particular the notion of a directed duty that is owed to a claim holder. Directedness is a particular challenge for deontic buckpassers, as reasons do not seem to have directed structure: in contrast to directed duties, they do not seem to be "owed" to others, and it is not obvious how this owing relation can be analysed in terms of reasons.

I believe, however, that the expectation view can be used to reformulate classic theories of rights in terms of reasons, thereby yielding a reasons-based analysis of directed duties. The general idea of the interest theory of rights is that reasons or obligations correspond to claims if they are

⁴⁹ See e.g. Howard and Leary (2022) for opposition to the identification of fittingness with correctness.

⁵⁰ See e.g. Williamson (2000, Ch. 11) and Littlejohn (2012, Ch. 5) for claims along these lines.

⁵¹ Compare Lackey's *Reasonable to Believe Norm of Assertion* (Lackey 2007, 608).

grounded in an interest (e.g. Raz 1986). The general idea of the will theory is that they do so if they are subject to a person's will (e.g. Hart 1982). These ideas give us the following expectational versions of the interest and the will theory of rights:

The expectational interest theory of rights

For A to have a right against B (and B to have a corresponding duty owed to A) is for it to be the case that (i) R is a reason for B to ϕ and an agent-neutral reason to expect B to ϕ , and (ii) an aspect of A's interest explains (i).

The expectational will theory of rights

For A to have a right against B (and B to have a corresponding duty owed to A) is for it to be the case that (i) R is a reason for B to ϕ and an agent-neutral reason to expect B to ϕ , and (ii) A has the ability to cancel (i) by declaration (e.g. by consent).

As reformulations of the classical theories, these views are likely to be subject to similar problems as their original statements, and I would be surprised if one these formulas survived critical reflection in its present form. Nevertheless, my hope is that analysing rights in terms of reasons can shed new light on the debate between interest and will theorists. This, however, is a topic for another day.⁵²

⁵² [Acknowledgments.] Work on this article was funded by the European Union (ERC Grant 101040439, REASONS FIRST). Views and opinions expressed are however those of the author only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

References

- Adams, Robert Merrihew. 1999. *Finite and Infinite Goods: A Framework for Ethics*. New York: Oxford University Press.
- Anscombe, G E M. 1958. 'Modern Moral Philosophy'. *Philosophy* 33 (124): 1–19.
- Bedke, Matthew S. 2011. 'Passing the Deontic Buck'. *Oxford Studies in Metaethics* 6: 128–52.
- Broome, John. 2004. 'Reasons'. In *Reason and Value. Themes from the Moral Philosophy of Joseph Raz*, edited by R. Jay Wallace, Philip Pettit, Samuel Scheffler, and Michael Smith, 28–55. Oxford: Oxford University Press.
- . 2007. 'Requirements'. In *Hommage à Wlodek. Philosophical Papers Dedicated to Wlodek Rabinowicz*, edited by Toni Rønnow-Rasmussen, Björn Petersson, Jonas Josefsson, and Dan Egonsson, 1–41. URL = <http://www.fil.lu.se/hommageawlodek/>.
- Bruno, Daniele. 2023. 'Being Fully Excused for Wrongdoing'. *Pacific Philosophical Quarterly* 104 (2): 324–47. <https://doi.org/10.1111/papq.12425>.
- Capes, Justin A. 2012. 'Blameworthiness without Wrongdoing'. *Pacific Philosophical Quarterly* 93 (3): 417–37.
- D'Arms, Justin, and Daniel Jacobson. 2003. 'The Significance of Recalcitrant Emotion (or, Anti-Quasijudgmentalism)'. *Royal Institute of Philosophy Supplements* 52: 127–45.
- . 2023. *Rational Sentimentalism*. Oxford University Press.
- Darwall, Stephen. 2006. *The Second-Person Standpoint. Morality, Respect and Accountability*. Cambridge, MA: Harvard University Press (paperback ed. 2009).
- . 2010. 'But It Would Be Wrong'. *Social Philosophy and Policy* 27 (2): 135–57.
- . 2022. 'A Gibbardian Account of (Narrow) Moral Concepts'. In *Meaning, Decision, and Norms: Themes from the Work of Allan Gibbard*, edited by Billy Dunaway and David Plunkett, 109–24. University of Michigan: Maize Books.
- de Kenessey, Brendan. 2023. 'The Relation Between Moral Reasons and Moral Requirement'. *Erkenntnis*. <https://doi.org/10.1007/s10670-023-00755-7>.
- Driver, Julia. 1992. 'The Suberogatory'. *Australasian Journal of Philosophy* 70 (3): 286–95.
- Ewing, A.C. 1948. *The Definition of Good*. New York: Macmillan.
- Foot, Philippa. 1972. 'Morality as a System of Hypothetical Imperatives'. *The Philosophical Review* 81 (3): 305–16.
- Gertken, Jan, und Benjamin Kiesewetter. 2017. 'The Right and the Wrong Kind of Reasons'. *Philosophy Compass* 12 (5): 1–14.

- Gibbard, Allan. 1990. *Wise Choices, Apt Feelings. A Theory of Normative Judgment*. Oxford: Clarendon Paperbacks (repr. 2002).
- Graham, Peter A. 2014. 'A Sketch of a Theory of Moral Blameworthiness'. *Philosophy and Phenomenological Research* 88 (2): 388–409.
- Hart, H.L.A. 1982. *Essays on Bentham: Jurisprudence and Political Philosophy*. Oxford: Oxford University Press.
- Horgan, Terry, and Mark Timmons. 2010. 'Untying a Knot From the Inside Out: Reflections on the "Paradox" of Supererogation'. *Social Philosophy and Policy* 27 (2): 29–63.
- Howard, Christopher, and Stephanie Leary. 2022. 'In Defence of the Right Kind of Reason'. In *Fittingness*, edited by Christopher Howard and R. A. Rowland, 221–42. Oxford: Oxford University Press.
- Johnson King, Zoë. 2019. 'We Can Have Our Buck and Pass It, Too'. *Oxford Studies in Metaethics* 14: 167–88.
- . 2022. 'Deliberation and Moral Motivation'. *Oxford Studies in Metaethics* 17: 254–79.
- Kant, Immanuel. 1785. *Groundwork of the Metaphysics of Morals*. Edited by Jens Timmermann. Translated by Mary Gregor. Cambridge: Cambridge University Press (2011).
- Kearns, Stephen, and Daniel Star. 2009. 'Reasons as Evidence'. *Oxford Studies in Metaethics* 4: 215–42.
- Kiesewetter, Benjamin. 2017. *The Normativity of Rationality*. Oxford: Oxford University Press.
- . 2023. 'Pro Tanto Rights and the Duty to Save the Greater Number'. *Oxford Studies in Normative Ethics* 13: 190–214.
- Lackey, Jennifer. 2007. 'Norms of Assertion'. *Noûs* 41 (4): 594–626.
- Littlejohn, Clayton. 2012. *Justification and the Truth-Connection*. Cambridge University Press.
- Löschke, Jörg. 2021. 'Reasons to Act, Reasons to Require, and the Two-Level Theory of Moral Explanation'. *Philosophical Studies* 178 (1): 169–85. <https://doi.org/10.1007/s11098-020-01426-x>.
- Mackie, J.L. 1977. *Ethics. Inventing Right and Wrong*. London: Penguin (repr. 1990).
- McHugh, Conor, and Jonathan Way. 2016. 'Fittingness First'. *Ethics* 126 (3): 575–606.
- . 2022. *Getting Things Right: Fittingness, Reasons, and Value*. Oxford University Press.
- Menges, Leonhard. 2017. 'The Emotion Account of Blame'. *Philosophical Studies* 174 (1): 257–73. <https://doi.org/10.1007/s11098-016-0680-9>.
- Mill, John Stuart. 1861. *Utilitarianism*. Cambridge: Cambridge University Press (2nd ed. 1864/2014).

- Nebel, Jacob M. 2019. 'Normative Reasons as Reasons Why We Ought'. *Mind* 128 (510): 459–84. <https://doi.org/10.1093/mind/fzy013>.
- Nye, Howard. 2022. 'Morality and the Bearing of Apt Feelings on Wise Choices'. In *Meaning, Decision and Norms: Themes from the Work of Allan Gibbard*, edited by Billy Dunaway and David Plunkett, 125–44. Maize Books.
- Parfit, Derek. 2011. *On What Matters*. Vol. 1. Oxford: Oxford University Press.
- Pereboom, Derk. 2014. *Free Will, Agency, and Meaning in Life*. Oxford: Oxford University Press.
- Portmore, Douglas W. 2011. *Commonsense Consequentialism: Wherein Morality Meets Rationality*. New York: Oxford University Press.
- Prichard, H.A. 1912. 'Does Moral Philosophy Rest on a Mistake?' In *Moral Writings*, 7–20. Oxford: Clarendon Press (2002).
- Rabinowicz, Wlodek, and Toni Rønnow-Rasmussen. 2004. 'The Strike of the Demon: On Fitting Pro-Attitudes and Value'. *Ethics* 114 (3): 391–423.
- Raz, Joseph. 1986. *The Morality of Freedom*. Oxford: Oxford University Press (repr. 1988).
- Reisner, Andrew. 2023. 'Against the "First" Views: Why None of Reasons, Fittingness, or Values Are First'. In *Value, Morality & Social Reality*, edited by Andrés G. Garcia, Mattias Gunnemyr, and Jakob Werkmäster. Department of Philosophy, Lund University. <https://doi.org/10.37852/oblu.189.c534>.
- Rivera-López, Eduardo. 2006. 'Can There Be Full Excuses for Morally Wrong Actions?' *Philosophy and Phenomenological Research* 73 (1): 124–42.
- Rosen, Gideon. 2015. 'The Alethic Conception of Moral Responsibility'. In *The Nature of Moral Responsibility: New Essays*, edited by Randolph K. Clarke, Michael McKenna, and Angela M. Smith, 65–87. New York: Oxford University Press.
- Ross, W. David. 1939. *Foundations of Ethics*. Oxford: Clarendon Press (repr. 1968).
- Rowland, Richard. 2019. *The Normative and the Evaluative: The Buck-Passing Account of Value*. Oxford: Oxford University Press.
- Schmidt, Thomas. 2023. 'How Reasons Determine Moral Requirements'. *Oxford Studies in Metaethics* 18: 97–115.
- . 2024. 'The Balancing View of Ought'. *Ethics* 134 (2): 246–67.
- Schroeder, Mark. 2007. *Slaves of the Passions*. Oxford: Oxford University Press.
- . 2010. 'Value and the Right Kind of Reason'. *Oxford Studies in Metaethics* 5: 25–55.
- . 2012. 'The Ubiquity of State-Given Reasons'. *Ethics* 122 (3): 457–88.

- . 2015. ‘What Makes Reasons Sufficient?’ *American Philosophical Quarterly* 52 (2): 159–70.
- . 2021. *Reasons First*. Oxford: Oxford University Press.
- Skorupski, John. 2010. *The Domain of Reasons*. Oxford: Oxford University Press.
- Smith, Michael. 1994. *The Moral Problem*. Oxford: Blackwell.
- Snedegar, Justin. 2016. ‘Reasons, Oughts, and Requirements’. *Oxford Studies in Metaethics* 11: 155–81. <https://doi.org/10.1093/acprof:oso/9780198784647.003.0007>.
- Stratton-Lake, Philip, ed. 2002. *Ethical Intuitionism: Re-Evaluations*. Oxford: Clarendon Press.
- Strawson, P. F. 1962. ‘Freedom and Resentment’. In *Freedom and Resentment and Other Essays*, 1–28. London: Routledge (2008).
- Telech, Daniel, and Leora Dahan Katz. 2022. ‘Condemnatory Disappointment’. *Ethics* 132 (4): 851–80. <https://doi.org/10.1086/719512>.
- Wallace, R. Jay. 1994. *Responsibility and the Moral Sentiments*. Cambridge, MA: Harvard University Press.
- . 2019. *The Moral Nexus*. Princeton: Princeton University Press.
- Way, Jonathan. 2012. ‘Transmission and the Wrong Kind of Reason’. *Ethics* 122 (3): 489–515.
- Williams, Bernard. 1979. ‘Internal and External Reasons’. In *Moral Luck*, 101–13. Cambridge: Cambridge University Press (1981).
- Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.
- Wodak, Daniel. 2020. ‘Who’s on First’. *Oxford Studies in Metaethics* 15: 49–71.
- Wolf, Susan. 2009. ‘Moral Obligations and Social Commands’. In *Metaphysics and the Good: Themes From the Philosophy of Robert Merrihew Adams*, edited by Samuel Newlands and Larry M. Jorgensen, 343–67. Oxford: Oxford University Press.