
Too Far Beyond the Call of Duty: Moral Rationalism and Weighing Reasons

Chris Tucker 10/25/21
William & Mary cstucker@wm.edu

Forthcoming in *Philosophical Studies*

Abstract: The standard account of supererogation holds that Liv is not morally required to jump on a grenade, thereby sacrificing her life, to save the lives of five soldiers. Many proponents defend the standard account by appealing to moral rationalism about requirement. These same proponents hold that Bernie is morally permitted to jump on a grenade, thereby sacrificing his life, to spare someone a mild burn. I argue that this position is unstable, at least as moral rationalism is ordinarily defended. The proponent of the standard account of supererogation must either reject moral rationalism or endorse that Bernie is morally required to remain in safety. Along the way, this paper brings together three neglected topics: going **too far** beyond the call of duty, moral rationalism about **permission**, and how to weigh reasons when some reasons have a different proportion of justifying and requiring weight than others.

0. Introduction

Supererogation involves going beyond the call of (moral) duty in a morally good way. Such acts are *morally optional*, i.e. permissible but not required. Yet not just any morally optional act is supererogatory. It has to be in some sense *morally better* than some other permissible option. For example:

Five Lives: Liv realizes that an enemy grenade will kill five of her fellow soldiers if she remains in safety. She can save the five soldiers by jumping on the grenade, thereby sacrificing her own life. Or she can remain in safety and allow the fellow soldiers to die.

Liv's jumping on the grenade would constitute a paradigmatic instance of supererogation. It is morally optional for Liv to jump on the grenade, and it is morally better if she does.

Otherist accounts of morality hold that there is an asymmetry between an agent's interests and the interests of the other, i.e., the members of the moral community who aren't the agent: while altruistic reasons are moral reasons, self-interested reasons are not.¹ This picture explains how Liv's jumping on the grenade is morally better, for Liv's self-interested reason to remain in safety is a non-moral reason. This is a genuine achievement. Otherism's most prominent rival is impartialism, the view that the interests of all members of the moral community—including the agent—are equal moral reasons. Even if Liv had to choose between her life and that of a single fellow soldier, it would be morally better for her to sacrifice her life. Yet if Liv's self-interest is a moral reason in exactly the same way as that of a fellow soldier, then how it could be *morally better* to sacrifice? Impartialists must make some concession to otherism if they wish to capture

¹ Portmore endorses what I call 'otherism' in his 2011: 96 (including nt 39), 128-9.

the idea that, other things being equal, altruism is morally better than self-interested action.

On the other hand, otherists must make concessions of their own. If Liv's self-interest isn't a moral reason, then how can it make it morally permissible to remain in safety? How can the otherist make sense of the sacrifice's being a moral option rather than a moral requirement?

The standard otherist account of optionality holds that Liv's self-interested reason is a non-moral reason which is nonetheless morally relevant. It is relevant insofar as it has moral justifying weight that outweighs the pro tanto requirement to save the five lives. If the self-interested reason did have moral requiring weight, then the standard otherist view is that it would be a moral reason after all (e.g., Portmore 2011: 128). Hence, the account further contends that the self-interested reason cannot have moral requiring weight. I challenge this further contention.

While we can go beyond the call of duty in a morally good way, we can also go *too far* beyond the call of duty. Consider:

Mild Burn: Bernie realizes that an enemy grenade will mildly burn one of his fellow soldiers if he remains in safety. He can spare the soldier the mild burn by jumping on the grenade, thereby sacrificing his own life. Or he can remain in safety and allow the fellow soldier to be mildly burned.

My view is that Bernie is *morally* required to remain in safety. To account for the moral requirement against pursuing this mild altruistic benefit, I hold that the (allegedly non-moral) self-interested reason has moral requiring weight.

The otherist may protest. "If Bernie jumps on the grenade, he makes a mistake from the perspective of rationality but not morality. He is, in fact, morally permitted to jump on the grenade."² As an impartialist sympathizer, I find this protest hard to swallow. Bernie is one person among many; his wellbeing matters to morality just like that of any other person; his humanity deserves just as much respect as everyone else's; and so forth. Yet these tropes won't convince the otherist. Indeed, the otherist can deny that Bernie is morally required to remain in safety...but only at cost.

Moral rationalism about requirement is the thesis that moral requirements are rational requirements. Many otherists endorse such moral rationalism and use it to motivate their standard account of optionality.³ My goal is to force these otherists into choosing between: (a) giving Bernie's wellbeing *moral* requiring weight and thereby granting that Bernie is *morally* required to remain in safety or (b) reject the moral rationalism that many otherists use to motivate their account of optionality.

In §1, I briefly review Gert's justifying/requiring distinction. In §2, I present the standard otherist account of supererogation, as it is developed by Portmore (2008; 2011: ch 5).⁴ Portmore's version is influential,⁵ and it is the clearest about how to weigh the altruistic moral reasons against the self-interested non-moral reasons. In §3, I present the parallel reasoning that commits Portmore to holding that Bernie is morally required to

² Those who would make such an objection include Curtis (1981: 314), Gert (2014), and maybe Portmore (2011: ch 5, especially pg 139).

³ Many, not all. Gert (2014) is an otherist who rejects moral rationalism. While I reject moral rationalism, I am sympathetic with Portmore's (2014) reply to Gert's brand of anti-rationalism.

⁴ Similar accounts are given by Clark (1978), Curtis (1981), and Muñoz (forthcoming).

⁵ For example, it is assumed by both Massoud (2016) and Archer (2016).

remain in safety. This parallel reasoning depends on two moves. The first is extending moral rationalism about *requirement* to moral rationalism about *permission*, i.e., moral permissions are rational permissions. I defend this extension in §4. The second move concerns how to weigh reasons to determine an overall deontic verdict. I defend the second move in §5. In §6, I deny Portmore’s last chance to break the parallel between his argument and the extension.

This paper brings together three neglected topics. First, there is little sustained reflection on how to understand going *too far* beyond the call of duty or how it relates to existing normative theories. Second, while there is a rich debate concerning the merits of moral rationalism about *requirement*, there is little discussion of moral rationalism about *permission* or whether one can coherently endorse one kind of rationalism and reject the other.

The third neglected topic is how to weigh the reasons needed to account for supererogation. A natural model of weighing reasons is a single scale on which the reasons for ϕ go in one pan and the reasons for $\sim\phi$ go on the other. It assumes that as a reason has justifying weight (as it pushes ϕ down toward permissibility), it also has requiring weight (it pushes $\sim\phi$ up toward impermissibility). This assumption is incompatible with the standard otherist account of supererogation, which holds that self-interested reasons are *merely justifying* (have justifying but not requiring weight). Merely justifying reasons push ϕ toward permissibility *without* pushing $\sim\phi$ toward impermissibility. What sort of model do we use to weigh merely justifying reasons?

1. Justifying vs Requiring, Requirements, and Moral Options

A consideration has **justifying weight for ϕ (JW_ϕ)** iff the consideration makes ϕ permissible in the absence of sufficiently weighty countervailing considerations. I can accept or decline some surgery for my child. *That the surgery will cause my child pain for weeks* has justifying weight for declining: in the absence of sufficiently weighty countervailing considerations, it is permissible to decline the surgery.

A consideration has **requiring weight for ϕ (RW_ϕ)** iff it makes the alternative, $\sim\phi$, impermissible in the absence of sufficiently weighty countervailing considerations. *That the surgery will cause my child pain for weeks* also has requiring weight for declining: in the absence of sufficiently weighty countervailing considerations, it is impermissible to *accept* the surgery.

A requirement to ϕ is a compound deontic verdict: ϕ is a **requirement** iff ϕ is permissible and $\sim\phi$ is impermissible.⁶ Justifying and requiring weight work together to make one required to ϕ . When the pain’s justifying weight makes it permissible to decline the surgery *and* its requiring weight makes it impermissible to accept it (so there are no sufficiently weighty countervailing considerations), the pain *requires* me to decline the surgery. On the other hand, if the benefits of the surgery outweigh the pain’s justifying and requiring weight, then I’m not (all-in) required to decline the surgery. I am only *pro tanto* required to decline the surgery, i.e., required in the absence of sufficiently weighty countervailing considerations.

⁶ Here I follow Portmore (2011: 88-9) and Dorsey (2016: 166). I also assume that we always have exactly two options. For how to think about weighing reasons in the context of any number of options, see my manuscript.

When reasons have more justifying than requiring weight, they have some tendency to make actions permissible without making them required. Such reasons are useful when, like Portmore, you are trying to explain **moral options**, cases in which both ϕ and $\sim\phi$ are permissible.

2. Making Sense of Going Beyond the Call of Duty: Optionality

In this section, I present Portmore's otherist account of how it is morally optional for Liv to jump on the grenade. The upshot is that Liv's self-interested non-moral reason must have moral justifying weight.

In a nutshell, the account is this. The lives of the five other people pro tanto—but not all-in—require Liv to jump on the grenade. Consequently, there must be some reason to remain in safety with at least as much justifying weight for remaining in safety as there is requiring weight to jump. The only candidate for having justifying weight to remain in safety is Liv's self-interested reason, so self-interested reasons have moral justifying weight. More formally:

1. The prevention of five deaths is a reason to jump on the grenade that has moral requiring (and justifying) weight, and thus Liv is morally required to jump on the grenade in the absence of countervailing reasons.⁷
2. Liv is not (all-in) morally required to jump on the grenade and is permitted to remain in safety.
3. If an agent has some moral requiring reason, MRR, for ϕ but is permitted to $\sim\phi$ (and so is not all-in required to ϕ), then there is a reason with at least as much justifying weight for $\sim\phi$ as MRR has moral requiring weight for ϕ .⁸
4. In Five Lives, Liv's self-interested reason is the only candidate for having justifying weight for remaining in safety.

5. Therefore, the self-interested reason has moral justifying weight for not jumping on the grenade that is at least as weighty as the moral requiring weight of the five lives.

1 is very plausible. There is general agreement that we are pro tanto morally required to prevent the death of others and, per §1, both justifying and requiring weight work together to explain (pro tanto) moral requirements.

2 is debatable. For example, Dorsey (2016: chs 3, 4) argues that Liv is (all-in) morally required to save the five lives. Dorsey and Portmore do agree, however, with the near undeniable claim that (2') Liv is not (all-in) *rationaly* required to jump on the grenade, where **rationality** is the unique practical perspective that has final authority over what to do.⁹

Portmore and Dorsey part ways precisely because they disagree about whether morality's requirements are "rationally authoritative" (Portmore 2011: 4). In other words, they disagree over **moral rationalism about requirements** (MR_{REQ}), i.e., the claim that ϕ is morally required only if it is rationally required. Dorsey denies MR_{REQ} and Portmore

⁷ Portmore defends 1 in his 2008: 373-5.

⁸ For Portmore's defense of 3, see his 2008: 375, 2011: 127.

⁹ Sobel (2007b: 14-16) likely joins Dorsey in endorsing 2' but rejecting 2. And so would anyone else who takes the demandingness objection to consequentialism to reveal, not that consequentialism is a false moral theory, but that morality lacks final authority over how we are to live.

endorses it. Since Liv is not rationally required to jump on the grenade, given MR_{REQ}, it follows that she isn't morally required either (2008: 377, nt 15; 2011:130, nt 15).

3 is very plausible, and the model of weighing reasons in §5 will confirm it.

Portmore assumes that 4 is true. As Five Lives was described, Liv has *at least* two reasons, namely the altruistic reason (the lives of the five soldiers) and Liv's self-interested reason (her life). Portmore tells us to "Assume that these are the only morally relevant facts" (2011: 126). We'll revisit the assumption in §6, but for now let's play along and screen off any other morally relevant considerations. Once we do, Liv's self-interested reason is the only candidate for having the requisite justifying weight to prevent Liv's pro tanto requirement from being an all-in requirement.

The picture that emerges from Portmore's 1-5 provides an explanation of how jumping on the grenade can be morally optional. Our self-interested reasons have enough justifying weight to prevent jumping on the grenade from being morally required, despite its being morally better to do so.¹⁰ If Liv jumps on the grenade, she really does go *beyond* the call of duty.

3. Making Sense of Going *Too Far Beyond the Call of Duty*

This section argues that, if Portmore's strategy shows that self-interested reasons have moral justifying weight, then it can be extended to show that self-interested reasons have moral *requiring* weight. One can't endorse Portmore's account and then deny that Bernie is *morally* required to remain in safety.

In a nutshell, the account is this. We are pro tanto justified in taking the necessary means of preventing others from suffering mild harms. Hence, the mild burn pro tanto justifies Bernie to jump on the grenade. But it doesn't all-in justify it. Perhaps one can permissibly sacrifice their life to secure some altruistic benefits that fall short of saving a life. For example, perhaps Bernie could sacrifice his life to prevent someone's paralysis. Yet one can't permissibly sacrifice their worthwhile life to prevent a mild burn.¹¹ Consequently, there must be some reason to remain in safety with at least as much requiring weight for remaining in safety as there is justifying weight to jump. The only candidate is Bernie's self-interested reason, so self-interested reasons have moral *requiring* weight. More formally:

1*. The prevention of the mild burn is a reason to jump on the grenade that has moral justifying weight, and thus Bernie is morally *permitted* to jump on the grenade in the absence of countervailing reasons.

2*. Bernie is not (all-in) morally *permitted* to jump on the grenade.¹²

¹⁰ On Portmore's view, moral betterness involves a third dimension of reasons, favoring. The altruistic act is morally better than the self-interested act, because requiring reasons (like the altruistic prevention of death) have favoring strength and merely justifying reasons (like self-interested reasons are on his view) do not have favoring strength (2019: 188, nt 4; cf. 2011: 122-4, including nts 6-7, and 128). My view of moral betterness also involves a third role of reasons (see my manuscript).

¹¹ Where do you draw the line for how much self-sacrifice is permissible? The model in §5 answers: you are permitted to self-sacrifice until the requiring weight for the self-interested action > the justifying weight for the self-sacrificial action. Portmore can't complain about this answer because we can always ask him, where do we draw the line for how much self-interested action is permitted? His answer, also given by the model in §5, is that you are permitted to act self-interestedly until the justifying weight for the self-interested action is < the requiring weight for the self-sacrificial alternative.

¹² Stangl (2016, especially 355) endorses 2*, and Clark (1978: 32) is at least sympathetic.

3*. If an agent has a moral justifying reason, MJR, for ϕ but is not all-in morally permitted to ϕ , then there is a reason with more moral *requiring* weight for $\sim\phi$ than MJR has justifying weight for ϕ .¹³

4*. In Mild Burn, Bernie's self-interested reason is the only candidate for having requiring weight for remaining in safety.

5*. Therefore, the self-interested reason has moral *requiring* weight that is weightier than the moral justifying weight of the mild burn.

1* is just as plausible as 1. Just as we are pro tanto required to prevent the deaths of five people, we are pro tanto justified in preventing small harms to others.

2* will be a no brainer for those who, like Dorsey (2016, ch 3), endorse impartialism about morality. For impartialism holds that the agent's wellbeing has the same moral weight as everyone else's.

Portmore endorses the otherist conception of morality, so he cannot accept the impartialist rationale for 2*. Yet consider (2*'): Bernie is not (all-in) *rationally* permitted to jump on the grenade. This claim seems near undeniable. Jumping on the grenade in such circumstances would display "recklessness" (Massoud 2016: 706) or "foolishness" (Stangl 2016: 355). Bernie would have "decisive reason" not to jump on the grenade (Portmore 2011: 4).¹⁴

2*' gives us 2 as long as we accept **moral rationalism about permission** (MR_{PERM}), i.e., ϕ is morally permitted only if it is rationally permitted. MR_{REQ} justified the inference from 2' (Liv is not *rationally* required to sacrifice) to 2 (Liv is not *morally* required to sacrifice). Likewise, MR_{PERM} justifies the inference from 2*' (Bernie is not *rationally permitted* to sacrifice) to 2* (Bernie is not *morally permitted* to sacrifice). In the next section, I argue that Portmore's defense of moral rationalism about requirement can be extended to an equally plausible defense of moral rationalism about permission. So 2 and 2* stand or fall together.

In §5, I defend the move from 3 to 3*, as well as defend a general model of weighing reasons.

For now, we are just assuming that 4* is true. As Mild Burn was described, Bernie has at least two reasons, namely the altruistic reason (preventing the mild burn) and Bernie's self-interested reason (his life). Portmore tells us to "Assume that these are the only morally relevant facts" (2011: 126). We'll revisit this assumption in §6, but for now screen off any other morally relevant considerations. Once we do, Bernie's self-interested reason is the only candidate for having requiring weight for remaining in safety.

The picture that emerges from 1*-5* explains how one can go *too far* beyond the call of moral duty. Bernie's self-interested reasons have enough moral requiring weight to prevent jumping on the grenade from being morally permitted, despite the altruistic benefit. If Bernie jumps on the grenade, he makes a moral mistake, not just a rational one. He goes *too far* beyond the call of duty.

Premise 1* of my extension seems straightforward. The rest of the paper defends 2* 3*, and 4*. The ultimate payoff is that the otherist must either (i) grant that Bernie's self-

¹³ Massoud (2016: 706) rejects 3*. She mistakenly assumes that if the justifying weight of remaining in safety is great enough it will defeat the pro tanto permission to jump on the grenade. Yet justifying weight is not in the business of making anything impermissible no matter how much of it you have. Only requiring weight does that.

¹⁴ Also see Curtis 1981: 314 and Clark 1978: 32.

interest has *moral* requiring weight and thus Bernie is morally required to remain in safety, or (ii) reject the moral rationalism that motivates the standard otherist account of supererogation's optionality.

4. Defending the Extension to 2*: Moral Blameworthiness and Rationality

Portmore's second premise, 2, is that Liv is morally permitted to remain in safety. To defend 2, he appeals to moral rationalism about requirement (MR_{REQ}). The extension's second premise, 2*, is that Bernie is morally required to remain in safety. To defend 2*, the extension appeals to moral rationalism about *permission* (MR_{PERM}). I argue that the best argument for moral rationalism about requirement can be extended into an equally good (or, as I prefer to think of it, equally bad) argument for moral rationalism about permission. Thus, Portmore's defense of Liv's moral permission commits him to Bernie's moral requirement.

4.1. Portmore's Argument for Moral Rationalism about Requirement

Portmore (2011; cf. Harman 2016) endorses moral rationalism about requirement but denies moral rationalism about permission. This position is unstable, at least when one defends the former moral rationalism by appealing to a connection with blameworthiness. This is what Portmore and many others do.¹⁵ Here is the main argument:

The Blameworthiness and Requirement Argument

- A. If *S* is morally required to $\sim\phi$, then *S* would be morally blameworthy for freely and knowledgeably ϕ -ing.¹⁶
- B. *S* would be morally blameworthy for freely and knowledgeably ϕ -ing only if *S* is *rationally* required to $\sim\phi$.¹⁷

- C. Therefore, if *S* is morally required to $\sim\phi$, then *S* is rationally required to $\sim\phi$.¹⁸

A links moral requirements with moral blameworthiness. B links moral blameworthiness with rational requirements. Together A and B give us C, which is just moral rationalism about requirement, the idea that moral requirements are also rational requirements. Before I present the parallel argument for moral rationalism about permission, it will be useful to first explain why this argument is so unsuccessful.

Moral rationalists, including Portmore, recognize that there is a *conceptual* distinction between morality and rationality, where only the latter is defined as having final authority over what to do. Of course, distinct concepts need not entail distinct referents. 'Biden' is conceptually distinct from 'the President of the US' even though Biden is (identical to) the President of the US. Likewise, perhaps some or all of morality's deontic verdicts (e.g., its verdicts of required and permissible) have final authority even though they do

¹⁵ See, for example, Portmore (2008, 2011: 43-4), Darwall (2006: ch. 5; 2016), and Murphy (2017: 26). I focus on Portmore's 2011 presentation of this argument, but he provides the same argument in his 2021: 54-56.

¹⁶ This premise is especially popular. J. S. Mill provides the classic statement of it. In addition to those who explicitly endorse the whole argument, contemporary support for this premise includes Gibbard (1990: 40-45) and Skorupski (1999: ch 7).

¹⁷ Darwall (2016: 269) takes this premise to be conceptually true. Portmore (2011: 44) initially claimed that it was *not* a conceptual truth, but he seems to have changed his mind (2014: 241). Sobel (2007b) rejects moral rationalism, but even he admits that B is plausible (2007a: 155-6).

¹⁸ The standard presentations of this argument replace 'is rationally required to $\sim\phi$ ' with 'does not have sufficient reason to ϕ '. This amounts to the same thing, as Portmore (2011: 42-3) makes clear.

not have such authority by definition. The conceptual distinction is nonetheless important. It shows us that we need an argument to insist that certain verdicts of morality have final authority (e.g., moral requirements are also rational requirements).

Yet once we make the conceptual distinction between *morality* and *rationality*, we must distinguish between *moral* blameworthiness and *rational* blameworthiness. The former is blameworthiness from the moral perspective and the latter is blameworthiness from the rational perspective. Again, distinct concepts need not entail distinct referents. It may be that there is just one kind of blameworthiness and that moral and rational blameworthiness are identical. The conceptual distinction is nonetheless important. It shows us that we need an argument to insist that blameworthiness from the moral perspective is or entails blameworthiness from the rational perspective.

The literature on moral rationalism tends to assume that there is just one kind of blameworthiness and it is *moral* blameworthiness.¹⁹ You'll notice that the argument for moral rationalism about requirement is stated in terms of *moral* blameworthiness.²⁰ This is why A is arguably a conceptual truth, as Portmore (2011: 44) and Darwall (2016: 269) contend. You might think that a coherent practical perspective is going to blame you just in case you freely and knowledgably make a mistake from that perspective. If you don't do what you are morally required to do, then you are making a moral mistake. When you freely and knowingly make such a mistake, it is your fault that you made such a mistake and thus morality blames you for it. So, A is plausibly (conceptually) true.

Yet B is not a conceptual truth. B holds that blameworthiness from the *moral* perspective entails a *rational* mistake (violating a rational requirement). We need some argument to justify linking *moral* blameworthiness to *rationality* in this way. To avoid begging the question, this argument should not itself assume that moral blameworthiness just is rational blameworthiness.

Portmore does provide an argument for B, but the argument itself assumes that moral blameworthiness is rational blameworthiness. It begins by claiming that:

B1. S is morally blameworthy for some action only if S has the capacity to respond to both moral and non-moral reasons.²¹

B1 says that responsiveness to both moral and non-moral reasons "opens the door" to assessments of moral blameworthiness (2011: 48). Such responsiveness is, in other words, what makes an agent's actions eligible to be assessed for moral blameworthiness (and moral praiseworthiness).

Portmore remarks, "Surely, it cannot be that the very capacity that opens the door to an agent's being blameworthy is the one that leads her to perform blameworthy acts" (48). In other words:

B2. If B1, then S can't be morally blameworthy for flawlessly responding to both moral and non-moral reasons. [cf. Portmore's 2.19, 2011: 48]

The conjunction of B1 and B2 entails:

¹⁹ It isn't just the literature on moral rationalism. See, for example, Tognazzini and Coates' SEP entry on blame.

²⁰ Portmore explicitly says he has in mind *moral* blameworthiness (2011: 43). Dorsey's (2016: 56-7) objection to the Blameworthiness and Requirement Argument seems to assume that there is just one kind of blameworthiness, but he insists that it is *rational* blameworthiness.

²¹ B1 is a simplified version of Portmore's 2.17 (2011: 48). 2.17 doesn't specify that both moral and non-moral reasons are needed, but 2.16 does.

B3. Therefore, S can't be morally blameworthy for flawlessly responding to both moral and non-moral reasons.²²

If B3 is true, then B arguably follows.²³ Nonetheless, this argument for B is a complete failure.

An initial problem is the ambiguity that runs through B1-B3. Start with B1. Does it claim that moral blameworthiness requires the capacity to respond to *some* or *all* non-moral reasons? It is plausible that moral blameworthiness requires the capacity to respond to *some* moral reasons, such as epistemic reasons. If you aren't capable of responding to epistemic reasons about what your moral reasons are, it isn't clear that you can be capable of responding to your moral reasons at all. Yet moral blameworthiness doesn't require the capacity to respond to *every* non-moral reason. Why would moral blameworthiness about any arbitrary ϕ require the capacity to respond to, say, aesthetic reasons?

The mere fact that moral blameworthiness requires the capacity to respond to *some non-moral reason or another* is useless in this context. Portmore needs to block the claim that *self-interested* reasons can make it rationally permissible to act against your moral requirements. Hence, Portmore needs B1 to entail that moral blameworthiness requires, more specifically, the capacity to respond to *self-interested* moral reasons. We should understand all of B1-B3 in like manner. For example, B3 should be interpreted to entail that S can't be morally blameworthy for flawlessly responding to both moral and *self-interested* reasons.

Now that we've clarified the ambiguity, we can see that the premises aren't plausible unless we assume that moral blameworthiness is rational blameworthiness. Regarding B1, we have no reason to suppose that moral blameworthiness requires the capacity to respond specifically to *self-interested* reasons. Suppose that there were *moral weirdos*, creatures capable of responding to moral reasons but *incapable* of responding to self-interested reasons.²⁴ I submit that such a creature could be morally blameworthy if it freely and knowingly acted against his decisive moral reasons. For example, the creature might be morally blameworthy for murdering his rich neighbor to benefit his relatively poor neighbor. If we are really concerned with *moral* blameworthiness and are not assuming that moral blameworthiness just is rational blameworthiness, it is hard to see why the capacity to respond to *self-interested* reasons is required to be morally blameworthy. And so B1 is in doubt.

B2 suffers from a similar problem. Suppose that morality cares only about some—not all—reasons that rationality cares about. That is, suppose that morality's deontic verdicts are a function solely of moral reasons and it ignores non-moral (or self-interested) ones. Then it would be no surprise that morality would blame you for

²² I focus on Portmore's 2011 presentation, but his 2021: 57-61 provides the same argument that I present as B1-B3.

²³ Suppose B is false: suppose that you are morally blameworthy for ϕ -ing despite being rationally permitted to ϕ . Since you are rationally permitted to ϕ , it seems that your ϕ -ing can be a flawless response (from the rational perspective) to your moral and non-moral reasons. And yet you are still blameworthy for that flawless response (from the moral perspective). This violates B3. It seems, then, that B3 entails B.

²⁴ These creatures can do the act that their self-interested reasons require of them; however, they can't do such acts in response to their self-interested reasons. For example, they might survive another day, not because it is in their best interest, but because they can't benefit anyone tomorrow if they don't. Also, thanks to Daniel Muñoz for the name 'moral weirdos'.

something that you were rationally permitted to do. That is, it would be no surprise that you could be morally blameworthy for flawlessly—from the point of view of rationality—responding to moral and non-moral (or self-interested) reasons. Once we take the distinction between moral and rational blameworthiness seriously, it isn't clear why we should endorse B1 or B2.

There is a grain of truth (grain of soundness?) in Portmore's argument for B. If we switch to *rational* blameworthiness, something in the neighborhood of B1-B3 *will* be sound. Perhaps a certain psychopath has no capacity to respond to moral reasons. He nonetheless can be rationally blameworthy for failing to respond appropriately to his self-interested reasons. For example, he might freely and knowledgeably choose a trivial immediate benefit and thereby sacrifice his long-run wellbeing. But he can't be rationally responsible for failing to respond to moral reasons, for he lacks the capacity to do so (cf. Portmore 2011: 48, nt 54). So:

B1'. S is rationally blameworthy for failing to respond correctly to moral reasons only if S has the capacity to respond correctly to moral reasons.

“Surely, it cannot be that the very capacity that opens the door to an agent's being blameworthy [for failing to respond correctly to moral reasons] is the one that leads her to perform blameworthy [failures to correctly respond to moral reasons]” (48).

B2'. If B1', then S can't be rationally blameworthy for flawlessly (from the rational perspective) responding to moral reasons.

Put the two premises together and you get:

B3'. Therefore, S can't be rationally blameworthy for flawlessly (from the rational perspective) responding to moral reasons.

(I do *not* claim that B1'-B3' is an equally plausible extension of Portmore's B1-B3. I claim that B1-B3 fails, whereas B1'-B3' is plausibly sound and seems to capture the grain of truth in B1-B3.)

Yet B1'-B3' is of no use to the moral rationalist. Morality and rationality presumably agree that one can flawlessly respond to a moral reason even if you act against it. They might disagree, however, about the conditions under which acting against a reason is flawless. Morality might demand an opposing *moral* reason of at least equal weight. Rationality might make a less onerous demand, namely that there be some moral *or non-moral* reason of at least equal weight. Since morality and rationality might disagree about what counts as a flawless response to moral reasons, it shouldn't be a surprise if one can be morally blameworthy without being rationally blameworthy.

Of course, Portmore's account of morality contends that morality takes into account non-moral reasons by giving them merely justifying weight. He will, therefore, reject the idea that morality blames you even when rationality doesn't. Yet this contention can't help him here without vicious circularity. His argument for that contention relies on moral rationalism as a (sub-)premise (§2). Hence, he can't use his contention that non-moral reasons have merely justifying moral weight as a (sub-)premise in his argument for moral rationalism.

The second premise in Portmore's argument for moral rationalism is B. His argument for B fails. Just as arguments for moral rationalism must be sensitive to the conceptual distinction between *morality* and *rationality*, such arguments must also be sensitive to the distinction between *moral* and *rational blameworthiness*. Portmore's argument for B fails to have the latter sensitivity, so it can't be used to show that B is plausible.

4.2. An Equally Plausible Argument for Moral Rationalism about Permission

The best (at least from Portmore's perspective) argument for moral rationalism about requirement is the Blameworthiness and Requirement Argument. Now that we've seen why Portmore's argument for the controversial second premise fails, it will be easy to see why there is an equally plausible argument for the claim that moral rationalism about permission is true.

To extend the Blameworthiness and Requirement Argument, let's use 'morally blameless' as equivalent to 'not morally blameworthy'. That gives us:

The Blameworthiness and Permission Argument

A*. If S is morally *permitted* to ϕ , then S would be morally *blameless* for freely and knowledgeably ϕ -ing.

B*. S would be morally *blameless* for freely and knowledgeably ϕ -ing only if S is rationally permitted to ϕ .

C*. Therefore, if S is morally permitted to ϕ , then S is rationally permitted to ϕ .

A* links moral permissions with moral blamelessness. B* links moral blamelessness with rational permissions. Together A* and B* give us C*, which is just moral rationalism about permission, the idea that moral permissions are also rational permissions.

If A is a conceptual truth, then A* is too. If you do what you are morally permitted to do, then you've had a moral success. And morality is not going to blame you for succeeding (cf. Skorupski 1999: 146, 150). Consider an analogy. Dad says to Kid, "I prefer that you wash my car, but it is permissible to clean up your room instead." Kid cleans his room instead. Dad can't sensibly blame Kid for freely and knowledgeably doing what he was expressly permitted to do.

Furthermore, Portmore's B seems committed to A*. A* says that *moral* permissions guarantee moral blamelessness. Portmore's B entails that *rational* permissions guarantee moral blamelessness.²⁵ If *moral* permissions don't guarantee *moral* blamelessness, it is hard to see why *rational* permissions would do so. In other words, if Portmore rejects A*, it makes it even harder for him to make B plausible.²⁶

²⁵ B explicitly says that, if you are morally blameworthy for ϕ -ing, then you are rationally required to $\sim\phi$. After contraposition, B says that if you are not rationally required to $\sim\phi$, then you are morally blameless for ϕ -ing. Now consider the antecedent 'you are not rationally required to $\sim\phi$ '. If prohibition dilemmas are impossible—i.e., if it is *impossible* that both ϕ and $\sim\phi$ are *impermissible*—then this antecedent entails that ϕ is permissible. And Portmore's account of reasons entails that such dilemmas are impossible. The only way to get prohibition dilemmas is for there to be reasons that have more requiring weight than justifying weight [my forthcoming: §§4.1, 7.2]. Yet Portmore denies that such reasons are possible (2011: 137-143). So, in his hands, B entails that, if you are rationally permitted to ϕ , then you are morally blameless for ϕ -ing.

²⁶ A referee wonders whether that A* rules out quasi-supererogation (an action that you are praiseworthy for performing but blameworthy for omitting). Maybe, but an appeal to quasi-supererogation would be at odds with Portmore's approach to supererogation. Quasi-supererogation is supposed to contrast with 'supererogation' understood as an action that is praiseworthy for performing but *not* blameworthy for omitting. Portmore explicitly disavows such approaches to supererogation: "the idea that an act is supererogatory is, for me, solely a deontic notion and does not depend on whether its performance is, given the agent's motives and intentions, praiseworthy" (2011: 5, nt 4; cf. 97, nt. 41). And, again, Portmore's B seems committed to A* whatever A* might entail.

I admit that the second premise of the argument for MR_{PERM} , B^* , is not particularly plausible; however, this admission doesn't break the parity. The best argument for B was a complete failure. That's why I tend to think that the Blameworthiness Arguments for moral rationalism about requirement and moral rationalism about permission are equally *bad* arguments. With that said, there is a case to be made that they really are *good* arguments.

Darwall's actual arguments for moral rationalism fail for the same reason that Portmore's do.²⁷ Yet he has the resources to accept both Blameworthiness arguments, both kinds of moral rationalism, and allow that Bernie is morally required to remain in safety. On his view, morality is equal accountability to all members of the moral community, including oneself (2006: 100-4, especially 102). Since an agent is accountable to oneself just as much as he is accountable to others, it is no mystery why Bernie's self-interested reasons would have moral requiring weight and would make him morally required to remain in safety. Of course, that's no good for Portmore's purposes. He needs a picture that vindicates moral rationalism about requirement without vindicating the claim that self-interested reasons have (moral) requiring weight.

I have provided an initial case, then, for the claim that the best argument for moral rationalism about requirement can be extended to an equally good (bad) argument for moral rationalism about permission. Further reflection confirms this initial case.

Two obvious worries about the argument for MR_{PERM} are also worries about the argument for MR_{REQ} . First, one might worry that A^* is false because someone, Arthur, might perform a morally permitted act for morally illicit motives. Portmore's argument for B faces the same worry, and he responds that "although it is clear that Arthur is blameworthy, it far from clear that Arthur is blameworthy for [performing the morally permitted act]" (2011: 45).²⁸

Second, there are plausible counterexamples to MR_{PERM} (Harman 2016). But there are also plausible counterexamples to MR_{REQ} . Many philosophers hold that MR_{REQ} should be rejected because the correct moral theory makes unreasonable demands, i.e., some moral requirements are not rational requirements (e.g., Wolf 1982; Sobel 2007; Dorsey 2016: ch 3). Given the conceptual nature of Portmore's argument for MR_{REQ} , he uses MR_{REQ} as a constraint on moral theorizing.²⁹ Putative counterexamples to MR_{REQ}

²⁷ Like Portmore, Darwall simply assumes that there is only one kind of blame/blameworthiness: "When we blame someone, we presuppose that the person we are blaming cannot adequately answer for what he has done [from the rational perspective]" (2016: 268; cf. 2016: 270, 2006: 98). X 's blaming Y from a perspective P may presuppose both that X and Y share P and that Y violated a requirement of P (cf. Pufendorf's point in Darwall 2006: 112). Yet it does not follow that P must be the rational perspective. If an eccentric billionaire pays you a \$50k to stop being my friend, I might fittingly blame you for being a bad friend and fittingly take a range of negative reactive attitudes toward you. I might nonetheless hold that you had adequate reason to do what you did (from the rational perspective). A police officer might fittingly blame you and hold you legally responsible for the minor traffic violation by writing you a ticket. The officer might nonetheless concede that you had adequate reason to commit the minor traffic violation from the rational perspective. Fitting blame is not necessarily fitting blame *from the rational perspective*.

²⁸ An alternative response would be to revise the two arguments so that they concern *underived* blameworthiness. Suppose that, contra Portmore, blameworthiness for one's morally illicit motives or past moral failures can transmit to one's now performing a morally permissible act. The blameworthiness for now acting rightly is *derived* from some moral failure or another. The key intuitions behind A and A^* is that wrong action can—and right action can't—be an original, or underived, source of blameworthiness.

²⁹ Dorsey (2016: ch 2) also understands Portmore this way.

must be interpreted instead as counterexamples to the putatively correct moral theory. Yet the parallel argument for MR_{PERM} is equally conceptual. Thus, if MR_{REQ} should be used as a constraint on moral theorizing, then MR_{PERM} should too.

At the end of the day, Portmore will likely note that MR_{REQ} seems more plausible to him than MR_{PERM} . So, here's a challenge for him to think about as he drifts off to sleep. The conjunction of A^* and B^* entails MR_{PERM} . If the latter is false, then so is the former. I predict that however he explains why that conjunction is false, he will find an equally plausible explanation for why the conjunction of A and B (the premises of Portmore's argument for MR_{REQ}) is false. If I'm right, his denial of MR_{PERM} would commit him to rejecting his own strategy for defending MR_{REQ} .³⁰ Since I reject both versions of moral rationalism and both blameworthiness arguments, I would consider that a win.^{31,32}

The goal of this section was to defend the transition from Portmore's 2 (Liv is morally permitted to remain in safety) to my 2* (Bernie is morally *required* to remain in safety). The only vulnerable part of Portmore's defense of 2 is moral rationalism *about requirement*. The only vulnerable part of my extension to 2* is moral rationalism *about permission*. I argued that the best argument for the former moral rationalism can be extended into an equally good (or equally bad) argument for the latter moral rationalism. In short, I have shown that Portmore's 2 and my 2* stand or fall together.

5. Defending the Extension to 3*: Weighing Reasons


5.1. Dual Scale

³⁰ Presumably, there is also some explanation of why MR_{REQ} would be true even though MR_{PERM} is false. Portmore may think he has such an explanation. His principle, META, entails both MR_{REQ} and $\sim MR_{\text{PERM}}$ (2011: 137). Yet Portmore's argument for META simply assumes that MR_{PERM} is false (2011: 139), so we shouldn't take META for granted in a context in which MR_{PERM} is up for debate.

³¹ A referee worries that, since requirement is a logically stronger deontic status than permission, moral rationalism about requirement is correspondingly more plausible than moral rationalism about permission. In reply, there is nothing special about requirements or the reasons that generate requirements. A requirement to ϕ is just the combined deontic verdict that ϕ is permissible and $\sim\phi$ is impermissible. To anticipate the model in the next section, permissions to ϕ are the result of a single competition between the justifying weight for ϕ (JW_{ϕ}) and requiring weight for $\sim\phi$ ($RW_{\sim\phi}$). Requirements to ϕ are the result of this same competition and an additional competition which determines the (im)permissibility of $\sim\phi$, namely $JW_{\sim\phi}$ vs RW_{ϕ} . I don't see why this difference gives us much, if any, reason to expect that moral requirements have a greater connection to rational statuses than moral permissions do.

³² A referee wonders: if moral rationalism about requirement and permission stand or fall together, does it follow that moral rationalism about requirement also stands or falls with any other normatively significant status (e.g., every morally best action is rationally best)? No. Portmore's argument for moral rationalism is driven by intuitions about *when actions are blameworthy*. Whether an action is required or permissible seems to have a tighter connection to when actions are blameworthy than whether an action is morally best. You aren't morally blameworthy for freely and knowledgeably failing to do the *morally best* action when it is merely supererogatory; however, you may be morally blameworthy for freely and knowledgeably performing an *impermissible* action or an action that you were *required* not to do. Furthermore, permissions and requirements are functions of the competition between justifying and requiring weight. Some other normatively significant statuses are not (merely) functions of justifying and requiring weight. For example, Portmore's account of morally best appeals to a third feature of reasons (recall note 10 above). Consequently, it is less surprising if moral rationalism about requirement and permission stand or fall together than that moral rationalism about requirement and bestness stand or fall together.

In this section, I argue that 3 entails 3*. 3 and 3* are principles concerning how to weigh reasons against one another in order to determine a deontic verdict. It is tempting to think that reasons must be weighed on a “single scale” (cf. Curtis 1981: 31). The reasons for ϕ , R_ϕ , go in one pan and the reasons against ϕ , $R_{\sim\phi}$, go in the other. ϕ is permissible iff the reasons against ϕ are not weightier than the reasons for ϕ . Such a view is simple and popular, but it cannot be Portmore’s.

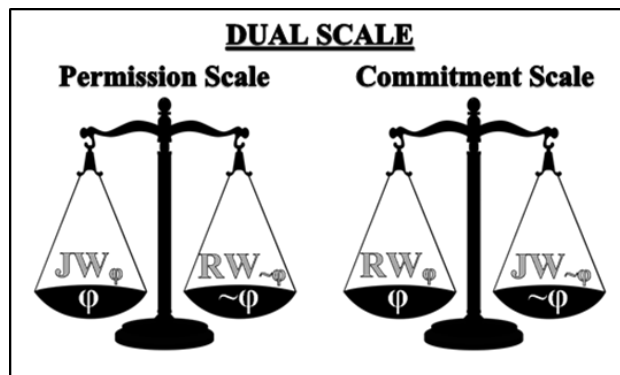


Portmore’s account of supererogation depends on the assumption that different kinds of reasons (altruistic vs self-interested) have different proportions of justifying and requiring weight. Such an assumption is incompatible with the single scale model of weighing reasons (Gert 2004: ch 5; 2007; Tucker forthcoming: §3-5). The problem isn’t with the image of the scale, but the image of a *single* scale.

Like the single scale model, **Dual Scale** holds that the deontic status of an action is determined by the relative weights of the reasons for and against it. But it holds that, to fully capture the two different weights (justifying vs requiring), we must appeal to two scales. **Permission Scale** determines whether an act is permissible: ϕ is permissible iff the justifying weight for ϕ (JW_ϕ) is at least as weighty as the requiring weight against ϕ ($RW_{\sim\phi}$).

We should introduce one more term, so that we can understand my name for the second scale. If your only goal in life is to eat every rock that you find, I might hold up a rock and remark that your aim *commits* you to eating this rock. In this sense of commitment, ϕ is a **commitment** iff $\sim\phi$ is impermissible. If you can be committed to eating this rock without it being permissible to do so, then you are in what is often called a prohibition dilemma (both ϕ and $\sim\phi$ are *impermissible*). But set such things aside. More relevant to this paper is that ϕ is required iff ϕ is both permissible and a commitment (i.e., $\sim\phi$ is impermissible).

Commitment Scale determines whether the act is a commitment (whether the alternative is *impermissible*): ϕ is a commitment iff the requiring weight for ϕ (RW_ϕ) is weightier than the justifying weight against it ($JW_{\sim\phi}$). The two scales work together to determine whether ϕ is required. An act is required just when Permission Scale says that ϕ is permissible *and* Commitment Scale says that $\sim\phi$ is impermissible.



Dual Scale is a model for how reasons are to be *weighed*. It does not itself come with a view about the *weights* of various kinds of reasons.

Elsewhere I develop and defend this model at length (forthcoming, 2022, manuscript). Here I work with a simplified version. In the rest of this sub-section, I use the Liv and Bernie cases to illustrate how the model works. In the next sub-section, I show that, given

three assumptions, Portmore’s 3 commits Portmore to both 3* of the extension and Dual Scale.

Liv must choose whether to jump on the grenade to save five soldiers (Sacrifice) or remain in safety (Safety). The standard otherist account supposes that the lives of the five soldiers have both justifying and requiring weight (say, 500 $JW_{Sacrifice}$ and 500 $RW_{Sacrifice}$). The account treats self-interested reasons, such as the value of Liv’s life, as weighty merely justifying reasons (say, 1,000 JW_{Safety} and 0 RW_{Safety}). Permission Scale entails that Sacrifice is permissible, because the justifying weight of Liv’s altruistic reason is weightier than the non-existent requiring weight of her self-interested reason (500 $JW_{Sacrifice} > 0 RW_{Safety}$). Commitment Scale entails that the Sacrifice isn’t a commitment, because the requiring weight of the altruistic reason is outweighed by the justifying weight of her self-interested reason (500 $RW_{Sacrifice} < 1,000 JW_{Safety}$).

(deontic status of Liv’s Sacrifice, assuming standard account of supererogation)

Permission Scale		Commitment Scale	
500 $JW_{Sacrifice}$	0 RW_{Safety}	500 $RW_{Sacrifice}$	1,000 JW_{Safety}
Verdict: Sacrifice is permissible		Verdict: Sacrifice is not a commitment	
Combined verdict: Sacrifice is permissible but not required			

Since the altruistic act is permissible but not a commitment (and so not required), one goes beyond the call of duty in performing it. At least, one goes beyond the call on the plausible assumption that it is morally better to save the five soldiers than to save one’s own life.

Bernie must choose whether to jump on the grenade to spare a soldier from a mild burn (Sacrifice) or remain in (Safety). Let’s suppose that the prevention of a mild burn has 1 unit of justifying and requiring weight (1 $JW_{Sacrifice}$ and 1 $RW_{Sacrifice}$). A minor change to the standard otherist account will allow us to make sense of Bernie’s being *morally* required to remain in safety. Let’s say that self-interested reasons have 100 times less requiring than justifying weight. In the Liv case, we assumed that the agent’s life had 1,000 justifying weight for Safety. So that means Bernie’s life has a measly 10 requiring weight for Safety; however, that measly amount is enough to make Bernie required to remain in safety. Permission Scale holds that Sacrifice is *impermissible* (1 $JW_{Sacrifice} < 10 RW_{Safety}$). Commitment Scale says that Sacrifice is *not* a commitment (1 $RW_{Sacrifice} < 1,000 JW_{Safety}$). Thus, Sacrifice is neither permissible nor a commitment. Bernie is, in other words, morally required to remain in safety.³³

(deontic status of Bernie’s Sacrifice, assuming a slight modification to standard account of supererogation)

Permission Scale		Commitment Scale	
1 $JW_{Sacrifice}$	10 RW_{Safety}	1 $RW_{Sacrifice}$	1,000 JW_{Safety}
Verdict: Sacrifice is <i>impermissible</i>		Verdict: Sacrifice is not a commitment	
Combined verdict: required to <i>not</i> Sacrifice			

³³ Recall that ϕ is a commitment iff $\sim\phi$ is impermissible. Since Sacrifice is *not* a commitment, then Safety is permissible. Since Sacrifice is impermissible, then Safety is a commitment. So Safety is a permissible commitment, which is just for Safety to be required.

5.2. From 3 to Both 3* and Dual Scale

Now that we understand Dual Scale, I will rely on three assumptions and Portmore's 3 to establish both 3* and Dual Scale. The details are tedious. If you are willing to assume that Dual Scale is true, then defending the transition from 3 to 3* is easy: the Permission Scale vindicates Portmore's 3 (Portmore's 3 is essentially the "only if" part of the Permission Scale) and the Commitment Scale vindicates the extension's 3* (3* is essentially the "only if" part of Commitment Scale). If you are satisfied, then feel free to skip to §6. If you don't want to assume that Dual Scale is true, then you'll have to wade through the tedium (sorry!).

Recall:

3. If an agent has some moral requiring reason, MRR, for ϕ but is permitted to $\sim\phi$ (and so is not all-in required to ϕ), then there is a reason with at least as much justifying weight for $\sim\phi$ as MRR has moral requiring weight for ϕ .

3 entails that $\sim\phi$ is permissible *only if* $JW_{\sim\phi} \geq RW_{\phi}$. My first assumption is one that Portmore seems to make, namely (**Sufficiency**) $\sim\phi$ is permissible *if* $JW_{\sim\phi} \geq RW_{\phi}$.³⁴ The conjunction of Portmore's 3 (the "only if" part) and Sufficiency (the "if" part) gives us the full:

Permission Scale: ϕ is permissible iff $JW_{\phi} \geq RW_{\sim\phi}$.

Permission Scale entails Commitment Scale. The picture of Dual Scale from §5.1 illustrates this. Notice that Commitment Scale is just the mirror image of Permission Scale after swapping ϕ and $\sim\phi$. For example, after you swap ϕ and $\sim\phi$, the left side of Permission Scale is the right side of Commitment Scale. This shouldn't be surprising. Permission Scale models whether ϕ is permissible. Commitment Scale models whether the *alternative* is impermissible (aka, whether ϕ is a commitment). Put the two models together, and you model whether ϕ is required. To verify that Permission Scale entails Commitment Scale, we just need to clarify the relationship between assignments of permissibility of ϕ and impermissibility of $\sim\phi$.

My second assumption is:

No Overlap: no act is both (morally) permissible and impermissible.³⁵

Given (the "if" part of) Permission Scale and No Overlap, it follows that ϕ is impermissible *only if* $JW_{\phi} < RW_{\sim\phi}$. My third and final assumption is:

No Gaps: every act is (morally) permissible or impermissible.

³⁴ Quibbling about whether " \geq " is the right place to draw the line for permissibility won't be productive. As long as there is some amount of net justifying weight ($JW_{\phi} - RW_{\sim\phi}$) that is sufficient for the permissibility of ϕ , 3 will entail a version of 3* that works for my purposes. It seems hard to deny that ϕ is permissible when the justifying weight for it outweighs the requiring weight against it *by a wide margin*. Given that a mild burn has very little requiring weight, I can weaken (ii*) to: the requiring weight for $\sim\phi$ is not, by a wide margin, outweighed by MJR's justifying weight. I would then just need to make the corresponding changes to the assignment rules for Permission and Commitment Scales.

³⁵ Do not confuse No Overlap with **No (Prohibition) Dilemma**: there is never a case in which both it is impermissible to ϕ and impermissible to $\sim\phi$. I allow for the conceptual possibility of dilemmas, which partly explains why I distinguish between commitment and requirement. Yet I do not allow for the conceptual possibility that the same act is all-in both permissible and impermissible (see my manuscript for a partial defense of this claim).

Given (the “only if” part of) Permission Scale and No Gaps, it follows that that φ is impermissible *if* $JW_{\varphi} < RW_{\sim\varphi}$. Put these new “if” and “only if” parts together, and we get the intermediate **Impermissibility Rule**: φ is impermissible iff $JW_{\varphi} < RW_{\sim\varphi}$. To get Commitment Scale, just apply the Impermissibility Rule to $\sim\varphi$: $\sim\varphi$ is impermissible (aka φ is a commitment) iff $JW_{\sim\varphi} < RW_{\varphi}$ (equivalently: $RW_{\varphi} > JW_{\sim\varphi}$). And then we have:

Commitment Scale: φ is a commitment (the alternative is impermissible) iff the $RW_{\varphi} > JW_{\sim\varphi}$.

So far, I’ve shown that once Portmore endorses 3 (and my three assumptions), he is stuck with Dual Scale whether he likes it or not. And once he’s stuck with Dual Scale, he is stuck with 3* too. 3* is essentially the “only if” part of the Commitment Scale. Recall:

3*. If an agent has a moral justifying reason, MJR, for φ but is not all-in morally *permitted* to φ , then there is a reason with more moral *requiring* weight for $\sim\varphi$ than MJR has justifying weight for φ .

3* says that φ is impermissible ($\sim\varphi$ is a commitment) only if $RW_{\sim\varphi} > JW_{\varphi}$. When we are evaluating the deontic status of $\sim\varphi$, Commitment Scale tells us the same thing, that φ is impermissible (aka: $\sim\varphi$ is a commitment) *only if* $RW_{\sim\varphi} > JW_{\varphi}$. In short, I have argued that endorsing 3 commits one to endorsing 3*, as well as the Dual Scale model of weighing reasons.

6. Portmore’s Last Chance

Portmore assumed 4 (Liv’s self-interested reason is the only candidate to explain Liv’s moral permission to remain in safety). My extension assumed 4* (Bernie’s self-interested reason is the only candidate to explain his moral requirement to remain in safety). Portmore’s last chance to break the parallel is for his assumption to be more plausible than mine. But it isn’t.

6.1. An Objection to Both 4 and 4*

Standard otherists deny that self-interest has moral requiring weight, but they don’t necessarily deny that agents have moral duties to themselves. Thus, they might concede that Bernie is morally required to remain in safety but then find some reason besides self-interest to explain this moral requirement. The following reasons are distinct from Bernie’s self-interest and yet are candidates to explain his moral requirement to remain in safety:

Autonomy: Bernie has moral requiring weight to prevent the destruction of his rational autonomy, which would result from his death (Schofield 2019: 228).

Talents: Bernie has moral requiring weight to develop his worthwhile talents (Portmore 129). Chances are that his death would prevent at least one worthwhile talent from being sufficiently developed.

Promises: Promises to oneself have moral requiring weight, at least until we let ourselves off the hook (Rosati 2011). Given that certain goals or career aspirations may involve making promises to oneself, it is plausible that Bernie’s death would break some promise to himself.

Hence, 4* is false. There are other candidates besides Bernie’s self-interested reasons to explain why he is morally required to remain in safety.

Nonetheless, the parallel between Portmore's argument and the extension remains intact. 4 is also false. There are other candidates to explain Liv's permission to remain in safety, such as altruistic but partial reasons. Liv's death will harm those beloved who love her, especially any dependents, and it is widely held that we have very weighty reasons to promote the interests of our beloved.

The Liv case is mine, but 4 is also false for Portmore's Fiona case (125). Fiona can use her savings self-interestedly (as the down payment that finalizes her home purchase) or altruistically (as a donation to Oxfam). Fiona is morally permitted to forgo the vastly more beneficial Oxfam donation, but her self-interest is not the only candidate to explain her moral permission. If Fiona doesn't complete the home purchase, she presumably breaks significant promises and harms people, such as the owners of the house she promised to buy and the realtors who don't get paid. These harms to others are especially weighty because she would be causing rather than merely allowing them.

The above objections to 4 and 4* win the battle but lose the war. They win the battle, because they show that *in the single Liv (Fiona) and Bernie cases at hand* there is another candidate explanation. They lose the war, because 4 and 4* are probably true for some Liv and Bernie case or another. In the rest of the section, I defend a general strategy for defending 4* and explain why you can't reject that general strategy without rejecting Portmore's 4.

6.2. *Why the Extension Wins the War*

Consider what happens if we gradually decrease the cost of jumping on the grenade. We notice a corresponding decrease in how much moral requiring weight there is to remain in safety. There is less moral requiring weight when Bernie is just paralyzed from the neck down and less still when he just loses his non-dominant hand. Eventually, we reach a point at which he is no longer required to remain in safety, holding fixed that the only benefit to jumping on the grenade is preventing the mild burn.³⁶ In other words:

Requiring Weight Correlation: the weight of Bernie's self-interested reasons is systematically correlated with the moral requiring weight that he has to remain in safety.

The extension's conclusion—that self-interested reasons have moral requiring weight—if true, provides a simple, straightforward explanation of Requiring Weight Correlation. Unless a potential alternative explains the Requiring Weight Correlation or at least covaries with Bernie's self-interested reasons, it will fail to explain Bernie's moral permission in some version of the Bernie case.

For example, consider reasons to protect one's autonomy, reasons to develop one's talents, and reasons to keep promises to oneself. They neither explain Requiring Weight Correlation, nor covary with self-interested reasons. There is, therefore, some version of the Bernie case in which these alternatives are not candidate explanations. Simply revise the Bernie case as follows: jumping on the grenade will not kill Bernie, but will paralyze him from the waist down; Bernie has a unique but undeveloped talent for musical composition, one worth developing at the expense of any other talent he might have (cf. Portmore 129); and he has made no promises to himself about the future, except perhaps that he will try to be more altruistic.

³⁶ I also leave it open that, after a certain point, self-interested reasons are sufficiently weak that they no longer have moral requiring weight at all.

In this revised case, Bernie is still *rationally* required to remain in safety. It is foolish to endure lifelong paralysis for the sake of preventing someone else's mild burn. Thanks to moral rationalism about permission and Dual Scale (or 3*), it follows that Bernie is *morally* required to remain in safety and these self-interested reasons suffice for (some reason that has) moral requiring weight. Now, however, the only candidate to explain his moral requirement to remain in safety are those very self-interested reasons. In the revised case, jumping on the grenade is no obstacle to preserving his rational autonomy, developing his talents, or keeping promises to himself.

Perhaps I've missed some other candidate explanation. But don't get tunnel vision. Don't forget that an alternative needs to do more than explain Bernie's moral requirement in the above two versions of the Bernie case. My revisions to the Bernie case were an instance of:

*General Strategy**: modify the Bernie case so that (i) the alleged alternative is not a candidate to explain Bernie's moral requirement to remain in safety, and (ii) the greater self-interested benefits make it *rationally impermissible* to secure the smaller altruistic benefits.

The resulting Bernie case, when combined with moral rationalism about permission and Dual Scale (or 3*), underwrites the extension. To block the extension's conclusion (while endorsing 1*-3*), you'll need an alternative that blocks *General Strategy**. To block *General Strategy**, you'll need an alternative that explains *Requiring Weight Correlation* or covaries with self-interested reasons. (Portmore's 4 is false in the original Liv case, but he can defend 4 with a parallel strategy.³⁷)

If *General Strategy** fails, then Portmore's 4 is false. Bernie's moral requiring weight to remain in safety systematically covaries with the weight of his self-interested reasons. So, to block *General Strategy**, you need an alternative that has both moral requiring weight and systematically covaries with the weight of his self-interested reasons. Suppose there is one. Call it **Alt**. Since Alt co-varies with self-interest, it will be present not only in all Bernie cases, but also in all Liv cases. Portmore does not allow requiring weight to outstrip justifying weight (2011: 137-143). Hence, Alt would have justifying weight too, and therefore would be an alternative to self-interest's having justifying weight. In short, any alternative that makes the extension's 4* false for every Bernie case will make Portmore's 4 false for every Liv case. 4 and 4* stand or fall together.

Conclusion

In this paper, I brought together three neglected topics: going *too far* beyond the call of duty, moral rationalism about *permission*, and how to weigh reasons when some reasons have a different proportion of justifying and requiring weight than others. I brought them together to challenge the standard otherist account of supererogation, as championed by Portmore.

If self-interest isn't a moral reason, as the otherist contends, it is initially unclear how the otherist can make sense of supererogation's optionality. By relying on moral

³⁷ Since self-interested reasons are also systematically correlated with justifying weight (cf. Portmore 2011: 126), Portmore can defend 4 using:

General Strategy: modify the Liv case so that (i) the alleged alternative is not a candidate to explain Liv's moral permission to remain in safety, and yet (ii) the smaller self-interested benefits make it *rationally* permissible to forgo the larger potential altruistic benefits.

rationalism about requirements, Portmore argues that self-interested reasons have moral justifying weight, even though they are non-moral reasons. He assumes that *if* self-interested reasons had moral requiring weight, they would be moral reasons, which would be tantamount to giving up on otherism. And thus Portmore will balk at the idea that Bernie's self-interested reasons have *moral* requiring weight.

I have argued, however, that Portmore's position is unstable. Portmore's argument that self-interested reasons have moral justifying weight can be extended to show that they also have moral requiring weight. Consequently, Portmore and like-minded otherists must give up the moral rationalism that motivates their account of supererogation's optionality *or* they must concede both that Bernie's self-interested reason has moral requiring weight and that Bernie is morally required to remain in safety. In short, I have shown that Portmore must choose between two things near and dear to his heart: his moral rationalism and his otherism.³⁸

References

- Archer, Alfred. 2016. "Moral Obligation, Self-Interest, and the Transitivity Problem." *Utilitas* 28: 441-464.
- Clark, Michael. 1978. "The Meritorious and the Mandatory". *Proceedings of the Aristotelian Society* 79: 22-33.
- Curtis, Barry. 1981. "The Supererogatory, the Foolish, and the Morally Required." *Journal of Value Inquiry* 15: 311-318.
- Darwall, Stephen. 2016. "Making the 'Hard Problem' of Moral Normativity Easier." *Weighing Reasons*. Errol Lord and Barry Maguire (eds.). Oxford University Press, 257-278.
- Dorsey, Dale. 2016. *The Limits of Moral Authority*. Oxford University Press.
- Gert, Josh. 2014. "Moral Rationalism and Commonsense Consequentialism." *Philosophy and Phenomenological Research* 88: 217-224.
- _____. 2007. "Normative Weight and the Balance of Reasons." *Philosophical Review* 116: 533-562.
- _____. 2004. *Brute Rationality: Normativity and Human Action*. Oxford University Press.
- Harman, Elizabeth. 2016. "Morally Permissible Moral Mistakes." *Ethics* 126: 366-393.
- Massoud, Amy. 2016. "Moral Worth and Supererogation." *Ethics* 126: 690-710.
- Mill, J. S. (1991). [1861]. *Utilitarianism*. In J. M. Robson, ed., *Collected Works of John Stuart Mill, Vol. 10*. London: Routledge, pp. 203–259.
- Muñoz, Daniel. Forthcoming. "Three Paradoxes of Supererogation." *Nous*. Early View version available at DOI: <https://doi.org/10.1111/nous.12326>.
- Murphy, Mark. 2017. *God's Own Ethics: Norms of Divine Agency and the Argument from Evil*. Oxford University Press.

³⁸ As usual, I have been extremely fortunate to have the support of many people and institutions. William & Mary provided an invaluable sabbatical and research grant. Helpful comments on earlier drafts were provided by Julia Driver, Josh Gert, Jonah Goldwater, Aaron Griffith, Andrew Moon, Daniel Muñoz, Doug Portmore, Sue Su, Philip Swenson, Ethan Terrill, Chris Tweedt, anonymous referees, and the audiences at the 2020 *Obligations to Oneself Workshop* and the 80th Meeting of the Virginia Philosophical Association. I greatly appreciate their support, and I'm especially grateful to my fabulous former student, Sue Su. I formulated the original draft of this paper in my head, thanks to Sue Su's many excellent questions relating to her own term paper. Without her questions, this paper probably would not exist.

- Portmore, Douglas. 2021. *Morality and Practical Reasons* Cambridge.
- _____. 2014. "Replies to Gert, Hurley, and Tenenbaum." *Philosophy and Phenomenological Research* 88: 241-55.
- _____. 2011. *Commonsense Consequentialism: Wherein Morality Meets Rationality*. Oxford University Press.
- _____. 2008. "Are Moral Reasons Morally Overriding?" *Ethical Theory and Moral Practice* 11: 369-388.
- Rosati, Connie. 2011. The Importance of Self-Promises. *Promises and Agreements: Philosophical Essays*. Hanoch Sheinman (ed.). Oxford University Press: 124–55.
- Schofield, Paul. 2019. "Practical Identity and Duties to the Self." *American Philosophical Quarterly* 56: 219-232.
- Sobel, David. 2007a. "Subjectivism and Blame." *Canadian Journal of Philosophy* 33 (Supplement): 149-170.
- _____. 2007b. "The Impotence of the Demandingness Objection." *Philosopher's Imprint* 7: 1-17.
- Stangl, Rebecca. 2016. "Neo-Aristotelian Supererogation." *Ethics* 126: 339-365.
- Tognazzini, Neal and D. Justin Coates. 2018. "Blame." *Stanford Encyclopedia of Philosophy*. Stable URL: <https://plato.stanford.edu/entries/blame/>.
- Tucker, Chris. Manuscript. *The Weight of Reasons*.
- _____. Forthcoming. "The Dual Scale Model of Weighing Reasons." *Nous*. Early View Version available at doi: <https://doi.org/10.1111/nous.12361>.
- _____. 2022. "Weighing Reasons Against." 2022. "Weighing Reasons Against." *Oxford Studies in Metaethics* 17: XX-XX.
- Wolf, Susan. 1982. "Moral Saints." *The Journal of Philosophy* 79: 419-439.