

The Relation between Moral Reasons and Moral Requirement

Brendan de Kenessey

Forthcoming at *Erkenntnis* – please use published version for citation

What is the relation between moral reasons and moral requirement? Specifically: what relation does an action have to bear to one's moral reasons in order to count as morally required? This paper defends the following answer to this question: 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. I argue that this *decisive moral reason view* satisfies three key desiderata: it is compatible with either affirming or denying the existence of moral options; it vindicates moral rationalism, the thesis that an action can be morally required only if one ought to do it all things considered; and most distinctively, it explains why unexcused moral wrongdoing necessarily shows disregard for moral reasons.

Philosophers talk about what we have *moral reason* to do, and about what we are *morally required* to do. This paper investigates the relation between these two notions. Specifically, it asks: what relation does an action have to bear to my moral reasons for it to be the case that I am morally required to do it?

I will defend the following answer:

THE DECISIVE MORAL REASON VIEW: For any person X and action ϕ , X is morally required to ϕ just in case, for any alternative option ψ , X 's moral reasons to ϕ rather than ψ are enough on their own to outweigh all of X 's reasons, moral and nonmoral, to ψ rather than ϕ .

Put more intuitively, the decisive moral reason view says that an action is morally required when one ought to do it *because of* the moral reasons in its favor. Moral requirements arise when moral reasons settle what one ought to do all things considered.¹

In support of this view, I argue that it satisfies three important desiderata for a theory of the relation between moral reasons and moral requirements.

Section 1 argues that our theory should be compatible with either affirming or denying the existence of *moral options*: i.e., that it is sometimes morally permissible to perform actions other than the morally best action available. This leads me to reject the *most moral reason view*, according to which an action is morally required just in case there is most moral reason to do it.

Section 2 argues that our theory should vindicate *moral rationalism*, the thesis that an action can be morally required only if one ought to do it, all things considered. Affirming moral rationalism makes possible an attractive explanation of moral options: moral options arise when *nonmoral* reasons give us sufficient reason to do something other than the morally best action. This leads me to reject the *only moral reasons view*, according to which what actions are morally required is fully determined by moral reasons alone.

Section 3 argues that our theory should capture the observation that unexcused moral wrongdoing necessarily shows disregard for moral reasons. That is, if a person responsibly chooses a morally wrong action, they must not have given adequate weight to moral reasons. This leads me to reject the *dual-ranking view* defended by Douglas Portmore (2011) and Thomas Schmidt (forthcoming), according to which an action is morally required just in case it is most supported by moral reasons and is what one ought to do all things considered.² Though the dual-ranking view

¹ Muñoz (2021) proposes a similar view; see §4 for discussion of the differences.

² Both Portmore and Schmidt's formulations of the view are somewhat more complex, but this is the main idea.

satisfies our first two desiderata, I argue that it incorrectly implies that one can responsibly choose a morally wrong action without showing moral disregard.

The decisive moral reason view satisfies all three of these desiderata. The view is neutral on whether there are moral options, while illuminating what is at stake in the debate over their existence. It vindicates moral rationalism: for if the moral reasons to ϕ outweigh all of the reasons to do otherwise, it follows that one ought to ϕ all things considered. Most distinctively, the decisive moral reasons view explains why unexcused moral wrongdoing necessarily shows moral disregard. On this view, if an action is morally wrong, then one's moral reasons not to do it are enough on their own to outweigh all of the reasons in its favor. So, if a person responsibly chooses to perform a morally wrong action, she cannot have given her moral reasons adequate weight. These advantages lead me to conclude that we should accept the decisive moral reason view.

0. Preliminaries

This section introduces the three main normative concepts we will be operating with, and the terminology I will use for each.

(1) *Moral requirement*. I use 'morally permissible' and 'morally wrong' as logical correlates of 'morally required': ϕ ing is morally wrong just in case not- ϕ ing is morally required; ϕ ing is morally permissible just in case it is not morally wrong.

One helpful way to get a grip on the concept of moral requirement is by means of its connection with moral accountability. Following Stephen Darwall (2006) and others, I will assume that a characteristic feature of moral requirement is that other people can usually fittingly hold me accountable to my moral requirements by *demanding* that I abide by them and *blaming* me if I flout them without excuse.

(2) *The all-things-considered 'ought'*. I assume that there is a meaningful distinction between the claim that one is morally required to perform an action and the claim that one *ought*, all things considered, to do it. Though some philosophers treat the phrases 'ought' and 'morally required' as synonyms, many others distinguish between these statuses.³ Some motivating examples: I ought to stretch after I go on a run, wash the dirty dishes in my sink, and read new philosophy articles in my area, but it would be strange to say that I am *morally required* to do any of these things. For another contrast, not everything I ought to do is something that others can fittingly demand of me or blame me for omitting. While I ought to wash my dishes, it intuitively seems that no one else can fittingly demand that I wash my dishes or blame me if I leave them festering in the sink.

The sense of 'ought' I have in mind is the all-things-considered, practical 'ought'. This is the 'ought' that answers the question of practical deliberation, "What to do?" I ought to do an action just in case it is the correct answer to this question – the thing to be done.

(3) *Normative reasons*. On a now-standard view, whether one ought to do an action is a function of the *normative reasons* in favor of that action and its alternatives. I will assume that one ought to ϕ just in case one has most reason to ϕ , i.e. more reason to ϕ than to perform any alternative ψ . As I will use the phrase, one *has decisive reason* to ϕ just in case one ought to ϕ all things considered. And one *has sufficient reason* to ϕ just in case it is not the case that one ought not ϕ .

The distinction between *moral reasons* and *nonmoral reasons* will play an important role in this paper. I take moral reasons to be normative reasons that are appropriately related to morality – reasons that speak *morally* in favor of an action. (I won't try to say what it takes for a reason to

³ For example, Darwall (2006), Portmore (2011; 2019), Dorsey (2016), and Wallace (2019).

be related to morality in this way). Paradigm cases include our reasons to keep our promises, to respect others' rights, to prevent harm, and so on. On some views, these are all the normative reasons for action there are. But I wish to allow for the possibility that there are normative reasons that count in favor of actions without counting morally in their favor. I will call these *nonmoral reasons*. Intuitive cases include reasons of self-interest, reasons based in personal projects, and (perhaps) reasons arising from relationships such as friendship.

I will call an action ϕ *morally best* just in case one has more moral reason to ϕ than to perform any alternative. I will call an action *morally better* if it is more favored by moral reasons. And I will call actions that are not morally best *morally disfavored*.

1. The most moral reason view and the options desideratum

We begin with a simple hypothesis:

THE MOST MORAL REASON VIEW: For any person X and action ϕ , X is morally required to ϕ just in case X has more moral reason to ϕ than to perform any alternative action ψ .

One way to motivate the most moral reason view is by analogy with 'ought'. Just as I ought to do whatever I have most reason to do, perhaps I am morally required to do whatever I have most moral reason to do.

Despite its *prima facie* plausibility, the most moral reason view conflicts with our intuitive judgments. As both Portmore (2009) and Darwall (2016: 266) have observed, the most moral reason view is incompatible with the commonsense belief that we have *moral options*: that we are sometimes morally permitted to do actions other than the morally best action available. Suppose it is Saturday morning, and I am choosing between two ways of spending my day: I could volunteer

for Oxfam, or I could play video games. Intuitively, it is morally permissible for me to play video games, even though volunteering would be morally better. The most moral reason view contradicts this judgment. On this view, if I have more moral reason to volunteer, it follows that I am morally required to do so.

Another way of putting the same point is that the most moral reason view does not make room for *supererogation*. A supererogatory action is an action that is morally better than the minimum one is morally required to do. For an action to be supererogatory, there must be some morally worse alternative to it that is morally permissible. Since the most moral reason view says that the morally best action is always morally required, it excludes the possibility of supererogation.

For those of us who believe in options and supererogation, this is enough to reject the most moral reason view. But of course, there are venerable theories in normative ethics – most prominently act consequentialism – that deny these beliefs. Should skeptics about options accept the most moral reason view?

It depends on their theoretical aims. If they are looking for a substantive normative thesis that tells us which actions are morally required, then options skeptics have every reason to accept the most moral reason view. But I have a different theoretical aim in mind. I am looking for an account that captures the *structural* relations between the properties of moral requirement and moral reasons. And even options skeptics have reason to doubt that the most moral reason view is well-suited to this purpose.

In addition to the substantive questions that are the traditional fare of normative ethics (e.g., should you push the man in front of the trolley?), we can also ask questions about the structural relations between the properties these debates are about. What I mean by ‘structural relations’ is,

roughly, metaphysically necessary relations between the properties that hold independently of, and constrain, the properties' extensions. For instance, it is a substantive, extensional fact that there is water in my glass, but it is a structural fact that water contains hydrogen.

I am interested in the metaphysical relationship between the *properties* of having moral reason and being morally required, rather than between our *concepts* of these properties.⁴ However, people competent with the concepts of moral reasons and moral requirement will typically have an implicit grasp of their necessary relations. Since our target concepts are employed by those with opposing normative views, it is a desideratum for accounts of this structural kind to be neutral on substantive disagreements wherever possible. The default presumption should be that any widely-held ethical theory is not metaphysically confused.

So, though I believe there are moral options, I want our account to allow for the coherence of denying their existence. Similarly, I think a skeptic about moral options should want to allow for the coherence of affirming their existence. If this is right, then even options skeptics have reason to reject the most moral reason view *as a structural account*. The intelligibility of believing that there are moral options suggests that the most moral reason account cannot be the right account of the relation between the properties of moral reason and moral requirement.

This motivates a first desideratum for our account:

THE OPTIONS DESIDERATUM: Our account of the relation between moral reasons and moral requirement should be compatible with both affirming and denying the claim that it is sometimes morally permissible to do morally disfavored actions.

⁴ Thanks to a referee for suggesting I clarify this.

2. The only moral reasons view and the rationalism desideratum

It is natural to feel somewhat baffled by the failure of the most moral reason view. If we line up possible ways I could spend my Saturday on a spectrum from morally worst (going on a murder spree) to morally best (volunteering for Oxfam), the intuitive line between moral wrongness and moral permissibility seems to be plopped arbitrarily in the middle: somewhere between throwing eggs at passing cars (mildly wrong) and playing video games (permissible but not great). What could explain why the threshold for moral permissibility falls *there* rather than somewhere else on the spectrum? (Compare Muñoz 2021: 3).

The source of this bafflement, I suspect, is the assumption that our moral requirements are determined *solely* by our moral reasons. Making this explicit, we get:

THE ONLY MORAL REASONS VIEW: For any person X and action ϕ , whether X is morally required to ϕ is fully determined by facts about the moral reasons in favor of ϕ ing and its alternatives.

This view is compatible with the most moral reason view, but it is also compatible with more options-friendly theories like the *enough moral reason view*, which says that an action ϕ is morally required just in case there is *strong enough* moral reason to ϕ . The enough moral reason view can accommodate options by saying that the moral reason to do the morally best action is not always strong enough to ground a moral requirement.

Of course, the challenge for the enough moral reason view is to say what counts as ‘strong enough’ moral reason, and one might worry that there’s no non-arbitrary way to set this threshold. But there is a deeper problem: our commonsense judgments suggest that whether an action is morally required is not a monotonic function of the strength of the moral reason to do it. Consider this example from Justin Snedegar:

Suppose that someone has left \$1,000 in a bank account and set up the following mechanism. If you press button A, all \$1,000 will be transferred to Oxfam; if you press button B, \$999 will be transferred to Oxfam, and the remaining \$1 will be burned (Snedegar 2016: 163).

Snedegar judges, and I agree, that you are morally required to press button A. But the moral reason to press button A rather than B is quite weak: all that is at stake is a single dollar in Oxfam's budget. Compare our earlier example, supposing that, if I volunteered for Oxfam on Saturday, I would raise \$100 in donations. My moral reason to volunteer (\$100 more for Oxfam) is clearly stronger than the reason that favors pressing button A over B (\$1 more for Oxfam). Yet I am not morally required to volunteer. It is hard to see how the enough moral reason view could explain this pattern of judgments. How could the moral reason to add \$1 to Oxfam's budget be strong enough to yield a moral requirement, while the moral reason to add \$100 to Oxfam's budget is too weak?

The problem is that we are looking only at moral reasons. If we take *nonmoral* reasons into account, this pattern of judgments looks less puzzling. You are morally required to press button A in Snedegar's example because there is no reason whatsoever to choose button B over button A. The Saturday example is different because there *is* a reason for me to play video games: the fact that I will enjoy it. And importantly, this reason's weight *for me* exceeds its weight from the moral point of view. This explains why I am morally permitted not to volunteer: because my moral reasons to do so are counterbalanced by my nonmoral reasons of self-interest.

Most accounts of moral options take this line. On the accounts offered by Scheffler (1982), Slote (1991), Parfit (2011), Portmore (2011), and Schmidt (forthcoming) we are morally permitted

to do morally disfavored actions when and because these actions are favored by sufficiently strong nonmoral reasons, such as reasons of self-interest.

If reasons of self-interest are not moral reasons, however, how can they influence what we are morally required to do? My answer: because nonmoral reasons influence what one ought to do all things considered, and an action cannot be morally required unless it is also what one ought to do. Behind this answer lies a thesis widely known as

MORAL RATIONALISM: For any person X and action ϕ , X is morally required to ϕ *only if* X ought to ϕ all things considered.

Moral rationalism is the key to understanding how nonmoral reasons can generate moral options. If we suppose that my nonmoral reasons of self-interest give me sufficient reason, all things considered, to spend Saturday playing video games, then this explains why I am not morally required to volunteer. Given moral rationalism, I cannot be morally required to volunteer unless I ought to do so all things considered.⁵

I believe that we should accept moral rationalism, in part because it yields this attractive explanation of moral options. However, moral rationalism is a controversial thesis. Accordingly, the rest of this section will defend it.

Portmore (2011: 43-44) and Darwall (2016: 268-269) both defend moral rationalism by appeal to the connection between moral requirement and blame. They argue that it is fitting to

⁵ If one denies moral rationalism, one could hold that nonmoral reasons sometimes give us sufficient reason to do a morally disfavored action, and yet we are still morally required to do the morally best action. (This is Dorsey (2013)'s view). So, the claim that we sometimes have the rational option to do morally disfavored actions does not entail that we have the moral option to do so. But for moral rationalists like me, these views go hand in hand. Thanks to a referee for prompting me to clarify this.

blame someone for failing to ϕ only if she ought to have ϕ ed all things considered. Darwall puts the point nicely: “it is incoherent to blame someone for wrongdoing while accepting that he had sufficient reason to act as he did” (2013: 34).⁶ If we accept the popular view that it is typically fitting to blame people for violating moral requirements without excuse,⁷ then the fact that fitting blame requires ‘ought’ seems to show that moral requirement does too.

Though I find this argument compelling, it is not inescapable. Someone who is skeptical about moral rationalism may be equally skeptical of rationalism about blame. As Dale Dorsey has argued (2016: 56-59), it is hard to give an argument for moral rationalism that does not assume some quasi-rationalist commitment as a premise. This leads me to think that the best way to defend moral rationalism is not to offer further arguments in its favor, but instead to remove the obstacles that have led others to doubt it.

The obstacle I aim to clear is the common assumption that moral rationalism implies that morality *overrides*, *trumps*, or is generally *more important than* other considerations in determining what one ought to do. Sarah Stroud takes moral rationalism to imply “that morality takes precedence over other concerns” (1998: 171), while Dorsey glosses it as claiming that “in normative competition with prudence and various other domains, moral requirements have greater weight” (48). Let’s call this the *moral overridingness interpretation* of moral rationalism.

This interpretation takes moral rationalism to be a substantive claim about the relative weights of normative reasons. But, as I argued in §1, we should want a thesis about the structural relations between normative reasons and moral requirement – as moral rationalism seems to be –

⁶ Similarly, Benjamin Kiesewetter argues that “a person is criticizable only if she violates a decisive reason” (2017: 24). If blame is a form of criticism, this also supports moral rationalism. Thanks to a referee for this reference.

⁷ See, e.g., Mill (1861/2001: 48-49), Darwall (2006), and Wallace (2019: Ch. 3).

to be neutral on such substantive debates. Dorsey argues that we should reject moral rationalism for this reason: “To understand whether ... moral requirements have greater weight [than other reasons], it is not enough *simply* to understand the notion of a moral requirement” (2016: 48). I agree. But I think this gives us reason to reject not moral rationalism, but the moral overridingness interpretation of it.

The moral overridingness interpretation of moral rationalism is optional. We are forced to adopt it only if we accept the only moral reason view. If the only moral reason view were true, then moral rationalism would imply that moral reasons rationally override all others. On this view, moral reasons fully determine what I am morally required to do; and given moral rationalism, this fully determines what I ought to do all things considered.

If we allow that moral requirements depend on nonmoral reasons as well as moral reasons, however, a different interpretation becomes available. On this interpretation, both ‘ought’ and moral requirement depend on the balance of *all* of one’s normative reasons for action, moral and nonmoral. This allows us to flip the order of dependence. We can say that moral rationalism is true not because moral requirements determine what we ought to do, but because whether an action is morally required depends in part upon whether one ought to do it. In other words, moral rationalism says not that morality dictates ‘ought’, but that ‘ought’ constrains moral requirement. Call this the *rational constraint interpretation* of moral rationalism.

On this interpretation, moral rationalism makes *no claim whatsoever* about the relative weight of moral and nonmoral reasons. It is compatible with egoism, consequentialism, or any other view of what reasons we have. Moral rationalism says that moral requirement must align with ‘ought’, but leaves open how both of these properties depend on moral vs. nonmoral reasons.

The rational constraint interpretation of moral rationalism thus avoids Dorsey's objection, because it is not a substantive normative thesis.

With the rational constraint interpretation of moral rationalism in hand, we can explain moral options by appeal to nonmoral reasons. If my nonmoral reasons give me sufficient reason to play video games, then this action cannot be morally wrong, as it is not the case that I ought not do it. However, this explanation requires the substantive premise that nonmoral reasons are sometimes weighty enough to give us sufficient reason to perform morally disfavored actions. This means, importantly, that accepting moral rationalism is also compatible with denying the existence of moral options. For example, if the only normative reasons for action are moral reasons, then the morally best action will also always be what one ought to do. An account that vindicates the rational constraint interpretation of moral rationalism will thus satisfy the options desideratum from §1. Even better, moral rationalism illuminates what is at stake in this debate: whether we have moral options hinges on whether nonmoral reasons can ever give us sufficient reason to do morally disfavored actions.

This motivates a second desideratum for our account:

THE RATIONALISM DESIDERATUM: Our account of the relation between moral reasons and moral requirement should vindicate (the rational constraint interpretation of) *moral rationalism*, the claim that one can be morally required to ϕ only if one ought to ϕ all things considered.

3. Two views: the dual-ranking view and the decisive moral reason view

In this section I introduce the two views that I take to be the most promising contenders for a theory of the relation between moral reasons and moral requirement: the dual-ranking view,

defended by Portmore (2009, 2011) and Schmidt (forthcoming), and my own proposal, the decisive moral reason view. Both views satisfy our two desiderata so far. Once both views are on the table, I will argue that the decisive moral reason view is superior to the dual-ranking view.

3.1. The dual-ranking view

Douglas Portmore has defended versions of our first two desiderata. Portmore rejects the most moral reason view because it excludes moral options; he defends moral rationalism; and he explains moral options by appeal to nonmoral reasons (Portmore 2009, 2011). Here's a simplified version of Portmore's positive view:

THE DUAL-RANKING VIEW: For any person X and action ϕ , X is morally required to ϕ just in case, for any alternative ψ , (i) X has more reason to ϕ than to ψ all things considered and (ii) X has more moral reason to ϕ than to ψ .

The dual-ranking view combines the two main ideas we have seen so far: the idea that moral requirement entails 'ought' (moral rationalism), and the idea that moral requirements must be most supported by moral reasons (the most moral reason view). The dual-ranking view says that an action is morally required just in case it meets both of these conditions at once: that is, just in case it is both morally best *and* best all things considered.

Portmore's official view is somewhat more complex:

META: "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)

The difference between META and my formulation of the dual-ranking view is that META talks about ‘requiring reason’ rather than ‘moral reason’, incorporating Joshua Gert’s distinction between moral requiring and justifying strength (Gert 2007). In §3.5 I consider whether this distinction helps Portmore reply to my objection.

Thomas Schmidt (forthcoming) has also advocated a version of the dual-ranking view. His view is

MR₄: “An action ϕ is morally required if, and only if, and because the set of moral reasons M for ϕ is non-empty and it is not defused, i.e., for each incompatible alternative ψ to ϕ , either (i) the set S of (moral and/or non-moral reasons) for ψ is unable to defuse M (i.e., M is weightier than S), or else (ii) S is able to defuse M , but S can be prevented from defusing (i.e., M is weightier than the set of moral reasons in S), and the set N of nonmoral reasons for ϕ prevents S from defusing M (i.e., the union of M and N is weightier than S .)” (Schmidt forthcoming: 22)

Despite its complexity, this principle is extensionally equivalent to the dual-ranking view, as Schmidt acknowledges (20). A set of moral reasons M to ϕ is not defused just in case either (i) M is weightier than all the reasons against ϕ ing, or (ii) M is weightier than the moral reasons against ϕ ing, and all of the reasons to ϕ (moral and nonmoral) outweigh all of the reasons against ϕ ing. These conditions are fulfilled just in case the dual-ranking view’s two conditions hold.

The dual-ranking view handily satisfies our first two desiderata. Most obviously, it satisfies the rationalism desideratum, as it builds moral rationalism directly into condition (i). And as I argued in §2, any theory that accepts (the rational constraint interpretation of) moral rationalism will satisfy the options desideratum as well. If nonmoral reasons can sometimes give us sufficient

reason to do a morally disfavored action, then the dual-ranking view will imply that in such cases it is permissible to do less than the moral best.

3.2. The decisive moral reason view

What I find lacking in the dual-ranking view is an explanatory connection between its two conditions. If I ought to ϕ all things considered and I have most moral reason to ϕ , then the dual-ranking view says that I am morally required to ϕ , even if the explanation of *why* I ought to ϕ has little to do with moral reasons. I will argue shortly that in such cases the dual-ranking view delivers counterintuitive verdicts. This motivates my proposal:

THE DECISIVE MORAL REASON VIEW: For any person X and action ϕ , X is morally required to ϕ just in case, for any alternative option ψ , X 's moral reasons to ϕ rather than ψ are enough on their own to outweigh all of X 's reasons, moral and nonmoral, to ψ rather than ϕ .

The intuitive idea is that I am morally required to ϕ when I ought to ϕ *because of* my moral reasons. To see whether I am morally required to ϕ , the decisive moral reason view asks: if my moral reasons were the only considerations in favor of ϕ ing, but all of my reasons not to ϕ remained, would it still be the case that I ought to ϕ ? I am morally required to ϕ just in case the answer is yes.

Note that the decisive moral reason view entails that an action can be morally obligatory *only if* it meets the dual-ranking view's two conditions. If the moral reasons to ϕ are enough to outweigh all of the reasons not to ϕ , that entails both that the moral reasons in favor of ϕ ing are stronger than the moral reasons for any alternative, and that one ought to ϕ all things considered. Where I differ from the dual-ranking view is that I don't take meeting its two conditions to be *sufficient* for moral requirement.

The decisive moral reason view meets our first two desiderata, for much the same reasons as the dual-ranking view. First, it is compatible with both affirming and denying the existence of moral options. If we hold that my nonmoral reason to play video games on Saturday is at least as strong as my moral reason to volunteer for Oxfam, then the decisive moral reason view implies that playing video games is morally permissible, because my moral reason to volunteer does not outweigh my nonmoral reason to play video games. But if we hold that my moral reason to volunteer is strong enough, on its own, to outweigh my reason to play video games, then the decisive moral reason view will imply that I am morally required to volunteer.

Second, the decisive moral reason view entails (the rational constraint interpretation of) moral rationalism. On the decisive moral reason view, if X is morally required to ϕ , then X 's moral reasons to ϕ outweigh all of the reasons to do otherwise, and so X ought to ϕ all things considered. Crucially, the decisive moral reason view does not imply anything about the relative weight of moral reasons. It says that moral requirement entails 'ought' not because moral reasons always outweigh nonmoral reasons, but because moral requirements only arise in cases where moral reasons *do in fact* outweigh nonmoral reasons.

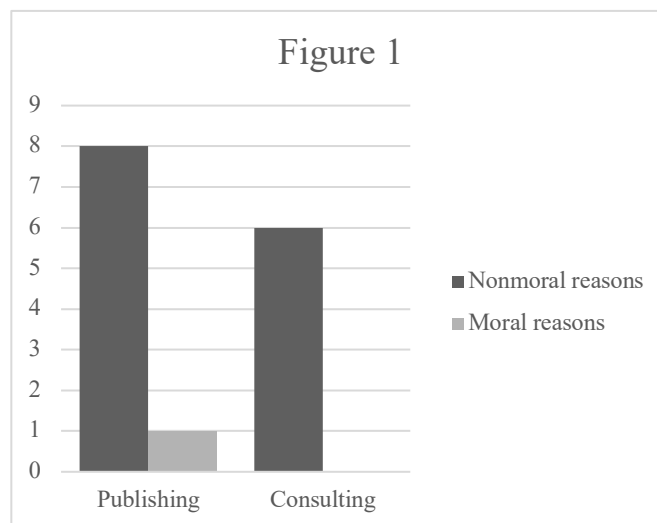
3.3. Comparing the two views

The dual ranking view and decisive moral reason view come apart in cases with a specific structure. These are cases where (i) an person's moral reasons most favor ϕ ing, and (ii) she ought to ϕ all things considered, but (iii) her moral reasons to ϕ are not strong enough on their own to outweigh all of her reasons to perform some alternative. In such cases, the dual-ranking view implies that the person is morally required to ϕ , because of (i) and (ii), while the decisive moral reason view implies that she is not morally required to ϕ , because of (iii).

To see which view is right, let's consider a case with this structure:

TWO JOBS: Joan, a recent graduate, is choosing between two job offers. One is a consulting job based in Dallas, the other is at a publishing company in New York. On the one hand, Joan's parents live in Dallas, and she prefers the warm weather there. On the other hand, Joan has always dreamed of a career in publishing. Moreover, moral reasons slightly favor the publishing job. Knowing that Joan is passionate about environmental issues, the publishing company has offered to sponsor a volunteer day in which the company's 30 employees will spend a day picking up trash at a local beach – but only if Joan accepts their offer. Suppose that other than this one moral benefit, the two jobs are morally equivalent. All things considered, Joan ought to choose the publishing job: following her passion for publishing is more important to Joan's well-being than living closer to family and in a warmer climate.

Figure 1 represents the weight of Joan's reasons. (To simplify the figure, I represent all reasons as positive considerations in favor of actions).



In this case, Joan has most moral reason to choose the publishing job, and she ought to choose it all things considered. The dual-ranking view thus implies that Joan is morally required to take the publishing job. But this is counterintuitive. Suppose that Joan, against her reasons, accepts the consulting job. The dual-ranking view implies that Joan thereby does something morally wrong, and thus that her action warrants others' blame. This strikes me as clearly false. The moral reason in favor of Joan's choosing the publishing job is simply too weak, relative to the life-altering nonmoral reasons at stake, to ground a moral requirement for her to do so.⁸ To see this, hold fixed Joan's nonmoral reasons in favor of choosing the consulting job – the warm weather and proximity to family – and imagine that the *only* consideration in favor of choosing the publishing job over the consulting job is that doing so will slightly reduce trash on one beach. Would that be enough to make it the case that Joan ought to choose the publishing job? No: in such a case, Joan would have sufficient reason to choose the consulting job. Joan's moral reasons in favor of the publishing job are not enough, on their own, to outweigh her reasons to take the consulting job. And so the decisive moral reason view says that she is not morally required to choose the publishing job – which, in my view, is the right result.

However, Schmidt (forthcoming) considers structurally similar cases and reaches the opposite conclusion.⁹ His rationale is that “whether an action is morally required is, as far as the non-moral side of things is concerned, a matter of its overall non-moral costs, with these being a function not only of the non-moral reasons against the action, but also of the non-moral reasons for it” (Schmidt forthcoming, 7). In other words, to figure out whether ϕ ing is morally required,

⁸ This description of the case assumes that nonmoral reasons exist and can generate moral options; I am asking what predictions our two theories make conditional on this assumption.

⁹ Thanks to the editor and referees at *Erkenntnis* for encouraging me to engage with Schmidt's arguments.

we should look at whether the moral reasons to ϕ are outweighed by the *net nonmoral cost* of ϕ ing. And in TWO JOBS, the net nonmoral cost of accepting the publishing job for Joan is zero, or negative: accepting the publishing job is *better* for Joan! And so Joan is in a position to provide a moral benefit (less trash on the beach) at *no net cost* to herself. Why wouldn't she be morally required to do so?¹⁰

Our disagreement concerns whether the moral reasons need to outweigh the *net* nonmoral cost or the *pro tanto* nonmoral cost of performing an action to make that action morally required. The dual-ranking view implies the net cost view. The decisive moral reason view implies the *pro tanto* cost view: in order to make ϕ ing morally required, the moral reasons to ϕ must be able to outweigh the nonmoral reasons in favor of any alternative on their own, without the help of the nonmoral reasons in favor of ϕ ing. In other words, the moral reasons to ϕ must outweigh the nonmoral costs of ϕ ing considered independently of its nonfungible nonmoral benefits: such as, for Joan, giving up the higher salary and warmer climate that the consulting job would afford.

'Nonfungible' is key here. The net cost / dual-ranking view gives the right verdicts in cases where the nonmoral costs and benefits of an action are fungible.¹¹ Consider a variant of the Snedegar case we considered in §2 (2016: 163):

¹⁰ Thanks to a referee for also suggesting this reply.

¹¹ Schmidt's case is plausibly of this kind: "if A performed ϕ , then this would save B from great pain now and save A from great pain later. But it would cause A great pain now" (forthcoming, 6). Assuming that pain now is fungible with pain later, the costs to A of ϕ ing are compensated for by fully fungible benefits. If this is right, then the decisive moral reason view agrees with Schmidt that A is morally required to ϕ in this case (see below).

TWO GOOD BUTTONS: I can push only one of two buttons. If I push button A, I will receive \$1,000,000 and \$100 will be donated to Oxfam.. If I push button B, I will receive \$1,000,000 and no money will be donated to Oxfam.

It seems clear that I am morally required to push button A, because by doing so I will provide \$100 to Oxfam at no cost to myself. The net cost view delivers this verdict: the benefit of pushing B, receiving a million dollars, is fully canceled out by the benefit of pushing A, and so there is no net nonmoral cost of pushing A. The decisive moral reason view appears to stumble here: if we understand the pro tanto cost of pushing A as forgoing the \$1,000,000 I would receive from pushing B, then plausibly my moral reason to donate \$100 to Oxfam does not outweigh my nonmoral reason to take the million, and so the view appears to imply, incorrectly, that I am not morally required to push A.

Not so fast. If we move back from talk of costs to talk of reasons, it becomes clear that the decisive moral reason view gets the right verdict. The fact that button B will give me a million dollars is *not* a reason to push B rather than A, because it is not a *difference* between B and A: A will also give me a million dollars. ‘Rather than’ is key here: the decisive moral reason view is formulated in terms of contrastive reasons. It says that for ϕ to be morally required, it must be that for any alternative ψ , my moral reasons to ϕ *rather than* ψ outweigh all of my reasons to ψ *rather than* ϕ .¹² While the fact that B will give me \$1,000,000 is a reason to push B rather than do nothing, it is not a reason to push B rather than A. In general, when the costs of an action are compensated for by *fungible* benefits, then it is the net cost of the action we should consider, because that is

¹² We can adopt this formulation without taking a stand on whether *all* normative reasons for action are contrastive, as argued by Snedegar (2017). All we need is that the reasons that determine moral requirement are contrastive. See also Muñoz (2021: 9).

what gives reason to choose it over its alternatives. Since your moral reason to push A rather than B clearly outweighs your non-existent reason to push B rather than A, the decisive moral reason view can agree with the dual-ranking view that you are morally required to push A.

Where the views come apart is when the costs of not ψ ing are not fungible with the benefits of ϕ ing. This is the case in TWO JOBS. By taking the publishing job, Joan is giving up a warmer climate and proximity to her parents. Neither of these costs are compensated for by fungible benefits of taking the publishing job. It is not as if New York also has a warm climate and is also near her parents. Instead, the primary benefit of taking the publishing job – following her passion – is of a different kind from the costs of turning down the consulting job. This means that the facts that Dallas is warmer and closer to Joan’s family are reasons to take the consulting job rather than the publishing job. These reasons are outweighed, not canceled out. And so the decisive moral reason view says that, for Joan to be morally required to take the publishing job, her moral reasons to do so must be enough on their own to outweigh her reasons to live in a warmer climate and closer to family – in other words, to outweigh the *pro tanto* cost of taking the publishing job. Since this is plausibly not the case, Joan is not morally required to take the publishing job. In contrast, given that she has more moral reason to take the publishing job, all the dual-ranking view requires is that Joan’s *total* reasons to take the publishing job outweigh her reasons to take the consulting job. Since this is the case, the dual-ranking view predicts – counter to my intuitions – that Joan is morally required to take the publishing job.

An independent argument for considering *pro tanto* rather than net cost comes from Alfred Archer (2016). Archer considers cases of *beneficial altruism*, in which “people who perform acts that seem like paradigmatic examples of supererogation report that it was in their self-interest to act as they did” (2016: 343). One example is a man dubbed “Free Help Guy” who publicly offers

help to anyone who contacts him, performing numerous acts of kindness as a result, but says “I wouldn’t call this altruism because I think I’ve got more out of this than anyone else” (343). Suppose Free Help Guy does benefit, overall, from his altruistic acts. Does this mean he is morally required to give out free help, since doing so is morally best and comes at no net cost to him? Archer thinks this is implausible. Free Help Guy’s altruism is intuitively supererogatory. If he were to stop giving free help, we would not hold him to be doing wrong.

The decisive moral reason view agrees with Archer’s verdicts, while the dual-ranking view cannot. The dual-ranking view implies that, since offering free help is both morally best and best all things considered for Free Help Guy, it is morally required. The decisive moral reason view says that, if Free Help Guy’s moral reasons to offer free help are not strong enough on their own to outweigh his (moral and nonmoral) reasons to do otherwise, then he is not morally required to do so. Free Help Guy’s situation might be like my Saturday: to give out free help, he has to give up an enjoyable afternoon of playing video games. This cost is, by his own testimony, outweighed by the fulfillment he gets from helping. But fulfillment and enjoying video games are not fungible, and so the pro tanto cost of giving up video games is not canceled out by the benefits of fulfillment, even though it is outweighed. If the moral reasons to offer help are not enough on their own to outweigh the cost of giving up video games, but require the ‘help’ of Free Help Guy’s prudential reasons to do so, then the decisive moral reason view will vindicate the intuitive verdict that Free Help Guy’s actions are supererogatory.

3.4. The disregard desideratum

So far my arguments for preferring the decisive moral reason view to the dual-ranking view have appealed to intuitions about cases: TWO JOBS and Free Help Guy. Such arguments are limited,

however, since case intuitions can differ. They also lack explanatory depth: sure, Joan isn't intuitively required to take the publishing job, but why? So let me now offer an argument for the decisive moral reason view that is more explanatory and less dependent on case judgments.

To introduce the argument it will help to return to TWO JOBS. I have supposed that Joan's moral reason to accept the publishing job (which, recall, is that doing so will prompt a trash-pickup day) is not weighty enough on its own to outweigh her nonmoral reasons to accept the consulting job. This implies that Joan could choose the consulting job while giving fully adequate weight to her moral reasons for choosing the publishing job. Suppose that Joan chooses the consulting job solely because she underestimates the prudential importance of passion and overestimates the importance of wealth, and so does not give enough weight to her nonmoral reason to take the publishing job. She takes the fact that taking the publishing job will lead to a slight reduction in trash on the beach to be a moral reason in favor of doing so, and gives this reason the right amount of weight relative to her nonmoral reasons. But she, correctly, takes this moral reason to be outweighed by her nonmoral reason to take the consulting job. Her mistake is only to misjudge the relative weights of her *nonmoral* reasons.

The dual-ranking view holds that it would be morally wrong for Joan to choose the consulting job, since doing so is both morally worse and worse all things considered. But Joan can coherently choose the consulting job while giving the correct weight to her moral reasons. Thus the dual-ranking view implies that Joan can intentionally and responsibly choose a morally wrong action *while giving her moral reasons fully adequate weight in her reasoning*.

This means that the dual-ranking view must deny a plausible thesis about moral requirement: that unexcused moral wrongdoing necessarily shows *moral disregard*. As I will use the term, a person shows disregard for a body of reasons just in case the weight she gives to those

reasons in her deliberation or motivation is less than the normative weight those reasons have as a matter of fact. Spelled out, then, my claim is that if a person performs a morally wrong action without excuse, then she must have given moral reasons less weight in her deliberation or motivation than they in fact have.

An argument for this claim comes from the connection between moral requirement and blame. Unexcused violations of moral requirement are fitting objects of blame. On a widely held view, blame is best understood as a response not to the wrongful action per se, but rather to the objectionable moral *attitude* expressed by that action (Strawson 1967/2008: 6-7; Hieronymi 2001: 546; Scanlon 2008: 128; Wallace 2019: 10-11). This is how excuses render blame unfitting: by showing that the person's wrongful action was not due to an objectionable attitude. What is the attitude to which blame is a fitting response? I suggest that it is moral disregard. A person is blameworthy when and because she does not show adequate care for moral reasons. If a person gives adequate weight to moral reasons in her deliberation, but makes some other, nonmoral rational mistake, it seems unfitting to morally blame her for her action. Moral disregard thus connects the dots between moral wrongdoing and blame: unexcused moral wrongdoing necessarily shows moral disregard, and moral disregard is a fitting object of blame.

The dual-ranking view is incompatible with this plausible account of why unexcused moral wrongdoing warrants blame. For as TWO JOBS shows, the dual-ranking view implies that one can intentionally, responsibly do wrong without showing any moral disregard. If Joan chooses consulting because she underestimates the strength of her nonmoral reasons to choose publishing, while giving fully adequate weight to her weak moral reasons to do so, then she shows no moral disregard. But on the dual-ranking view, she is doing moral wrong, because choosing the publishing job is both morally best and best all-things-considered.

This means that advocates of the dual-ranking view must allow that some unexcused moral wrongdoing is not blameworthy. Joan clearly does not warrant moral blame for her choice, as she gave her moral reasons fully adequate weight. And she is fully responsible for her choice: no excusing conditions obtain. So if Joan's choice is morally wrong, as the dual-ranking view implies, it is a case of unexcused moral wrongdoing that warrants no moral blame.

Making this explicit, here is the third desideratum for our account:

THE DISREGARD DESIDERATUM: Our account of the relation between moral reasons and moral requirement should vindicate the claim that unexcused moral wrongdoing necessarily shows moral disregard: that is, if a responsible agent performs a morally wrong action without excuse, she must have not given adequate weight to her moral reasons in her deliberation or motivation.

The dual-ranking view's failure to satisfy the disregard desideratum not only leads it to endorse a counterintuitive verdict about TWO JOBS, but also renders it unable to capture the necessary connection between unexcused moral wrongdoing and blameworthiness. For these reasons, I think we should reject the dual-ranking view.

In contrast, the decisive moral reason view implies that unexcused moral wrongdoing necessarily shows moral disregard. Here's how. On the decisive moral reason view, if I am morally required to ϕ , then my moral reasons alone are enough to make it the case that I ought to ϕ all things considered. So, if I gave moral reasons their proper weight, I would necessarily be led to the conclusion that I ought to ϕ . Therefore, if I voluntarily choose not to ϕ (and no excusing conditions obtain), then I must either be giving moral reasons insufficient weight in my deliberation, or acting akratically. If I act akratically, judging that I ought to ϕ and then not- ϕ ing

anyway, then I am not giving moral reasons adequate weight in my *motivation*, even if I properly weigh them in my normative judgment. Either way, my action shows moral disregard.

As a result, the decisive moral reason view neatly explains the necessary connection between unexcused moral wrongdoing and blameworthiness. On the decisive moral reason view, one cannot responsibly choose a morally wrong action while giving adequate weight to moral reasons. So, we can directly infer from the fact that someone chose a morally wrong action without excuse that they have the objectionable attitude of moral disregard. This moral disregard, in turn, is what makes them blameworthy.

The disregard desideratum also helps us explain why the difference between fungible and nonfungible costs matters, as I argued above. The trouble for the dual-ranking view in TWO JOBS is that Joan could choose the ‘morally wrong’ action of taking the consulting job solely because she misjudges how her nonmoral reasons weigh against each other, while correctly weighing her moral reasons against her nonmoral reasons and thus showing no moral disregard. But Joan can only coherently misjudge the relative weights of her nonmoral reasons because they are not fungible. Since the benefit of pursuing her passion is not fungible with the cost of living in colder weather, Joan could coherently make the mistake of thinking the latter outweighs the former. Compare TWO GOOD BUTTONS, where the pro tanto cost of pushing button A (forgoing the \$1,000,000 provided by B) is compensated for by a fungible benefit of pushing A (that A also provides \$1,000,000), while A also donates \$100 to Oxfam. Here, in order to choose A, Joan would have to make the much stranger mistake of thinking her reason to take the million from A is somehow stronger than her reason to take the million from B. She could only coherently make this mistake if she thought these were *not* fungible – that the money from A was special in some way. So, if Joan is coherent and recognizes that the nonmoral benefits of A and B are fungible, she will

see that she has no nonmoral reason to choose A rather than B, and so that her moral reason to choose B is decisive. Thus pushing button A would necessarily show moral disregard. The disregard desideratum thus gives us a principled rationale for treating fungible cases like TWO GOOD BUTTONS differently from nonfungible cases like TWO JOBS.

While both the dual-ranking view and the decisive moral reason view meet the options desideratum and the rationalism desideratum, only the decisive moral reason view meets the disregard desideratum. I take this to be the strongest consideration in its favor. The fact that unexcused moral wrongdoing is blameworthy is widely taken to be a signature, even defining, feature of moral requirement.¹³ The decisive moral reason view offers an elegant explanation of this fact, while the dual-ranking view appears to imply that some unexcused moral wrongdoing shows no moral disregard and is thus not blameworthy. At minimum, the dual-ranking theorist owes us an account of why unexcused moral wrongdoing is blameworthy, if it need not show moral disregard. While if we adopt the decisive moral reason view, the link between moral wrongdoing and blame is baked in to the relation between moral reasons and moral requirement.

3.5. Replies

I will now consider two replies on the dual-ranking view's behalf. The first reply appeals to Joshua Gert's distinction between justifying and requiring strength (Gert 2007). Gert argues that we should distinguish between a reason's *requiring* strength, the extent to which it counts towards an action's being (rationally or morally) required, and its *justifying* strength, the extent to which the reason counts towards an action's being (rationally or morally) permissible. Portmore's official thesis, META, incorporates Gert's distinction. It says that an action is morally required just in case

¹³ See the references in footnote 7.

it is both what one has most reason to do all things considered and what one has most moral *requiring* reason to do (Portmore 2011: 137).

This suggests a reply: perhaps the moral reasons in favor of Joan's taking the publishing job have only morally justifying strength. And so if we adopt META, on which morally required actions must be most favored by moral requiring reasons, we can avoid the prediction that Joan is morally required to take the publishing job.

The problem with this reply is that Joan's reason to reduce trash on the beach plausibly *does* have morally requiring strength. For instance, if Joan were to throw a candy wrapper on the beach rather than taking 30 seconds to walk to a trashcan, that would be morally wrong, albeit mildly. When there are no significant countervailing reasons, weak environmental considerations can ground moral requirements.¹⁴ Joan has more moral requiring reason to take the publishing job, and so META predicts that she is morally required to do so.

The second reply appeals to the idea that Joan's reasons to reduce trash on the beach are *imperfect*, in that she has discretion regarding when and how to do so. If she takes the consulting job, for example, she might volunteer to pick up trash to compensate. What Joan has reason to do is adopt a plan over the course of her life that gives adequate weight to her environmental reasons, which does not require her to take up any particular opportunity to reduce trash.¹⁵

But I don't think this helps. First, granting that Joan's environmental reasons are imperfect, this does not change the fact that she has more moral reason to take the publishing job than the consulting job. From the moral perspective, the more Joan reduces trash on the beach, the better.

¹⁴ Portmore uses a similar case to argue that reasons to donate to charity have moral requiring strength (2009: 374). This is also part of the point of Snedegar's two buttons case (2016: 163).

¹⁵ Portmore argues that imperfect reasons help to explain moral options in general (2011: Ch. 6; 2019: Ch. 4.4-4.5)

The fact that taking the publishing job will reduce trash on the beach is surely a moral reason in its favor.

Second, my intuitions about the case remain the same even if we modify it to encompass Joan's whole life plan – say, a choice between spending her whole career in publishing vs. consulting. Suppose that Joan anticipates that, if she takes the career in consulting, she will drive more often, and thus have a slightly larger carbon footprint, over the whole course of her life. Thus Joan's imperfect moral reasons favor the publishing career over the consulting career. However, it is still implausible to claim that Joan is morally required to choose the publishing career.

One might object that the last paragraph assumed *actualism*, as it assumed that the fact that Joan will not make the lifestyle changes to compensate for her larger carbon footprint is relevant to her reasons to take the consulting career.^{16,17} On the alternative, *possibilist* view, the fact that Joan could possibly change her lifestyle enough to compensate for her extra driving might mean that she does not have more moral reason to take the publishing job. But this does not help. For it is plausible that the smallest carbon footprint Joan could *possibly* have if she chooses the consulting career is bigger than the smallest carbon footprint she could have if she chose the less-driving-heavy publishing career. So even on possibilism, Joan has more moral reason to choose the publishing career. But if the difference in lowest-possible carbon footprints is small enough, it is still implausible to claim, as the dual-ranking view must, that Joan is morally required to take the publishing job.

¹⁶ Thanks to a referee for raising this objection.

¹⁷ For an overview of the actualism/possibilism debate, see Portmore (2019: Ch. 5).

4. Basic permission views

Before we close, I would be remiss if I didn't consider one final class of views, which I will call *basic permission views*.¹⁸ On basic permission views, our moral requirements are determined by the balance between our moral reasons and a different kind of normative factor that counts towards actions being morally permissible without counting in favor of their being morally required. We have already encountered one version of this view: Gert (2007) argues that there are *justifying reasons*, which count towards actions being permissible without counting towards their being required. Two other similar views are offered by Hurka and Shubert (2011), who call the permitting factors '*prima facie* permissions', and Muñoz (2021), who calls them 'prerogatives' (following Scheffler 1982).

These views are pretty close to the decisive moral reason view. Consider, for example, Muñoz's *prerogatives principle*: "An option x is obligatory just if, for any alternative y, there is more reason to choose x than there is combined reason and prerogative to choose y" (4). If you substitute 'moral reason' for 'reason' and 'nonmoral reason' for 'prerogative', you get the decisive moral reason view.¹⁹ So what's the difference?

It's subtle. The decisive moral reason view allows us to say that there is only one fundamental kind of input to what one ought to do: normative reasons. As far as 'ought' goes, we can say that every reason's justifying strength is the same as its requiring strength. (The decisive moral reason view does not entail this, but allows for it). However, moral and nonmoral reasons

¹⁸ Thanks to two referees for prompting me to add this section.

¹⁹ In fact, Muñoz considers just this view: "An option x is obligatory just if, for any alternative y, there is more moral reason to choose x than there is combined moral and non-moral reason to choose y" (2021: 4). He thus noticed the decisive moral reason view before I did, even though he ultimately prefers the prerogatives principle.

contribute differently towards an action's being *morally* required. In fact, the decisive moral reason view entails that nonmoral reasons have only morally justifying strength, while moral reasons have morally justifying and requiring strength. For it implies that nonmoral reasons in favor of an action can make that action morally permissible, but not morally required. So the decisive moral reason view allows us to explain the distinction between morally justifying and requiring strength without having to posit this as a basic feature of normative reasons – the distinction instead falls out from the relation between reasons and moral requirements. In contrast, advocates of the basic permission view tend to hold that the difference between moral reasons and the permitting factor is irreducible and fundamental.

I take it to be an advantage of the decisive moral reason view that it explains the possible divergence between morally justifying and requiring strength without needing to posit a primitive difference between justifying and requiring reasons. However, it does require us to distinguish between moral and non-moral reasons. So which theory one prefers may come down to which distinction one thinks is more plausibly regarded as basic.

Ultimately, the common ground between the decisive moral reason view and the basic permission view is more important than this subtle difference in conceptual furniture. Both views hold that an action is morally required when the moral reasons in its favor are enough, on their own, to outweigh all of the factors that count towards its alternatives being rationally permissible. Thus the contribution of this paper has not been to discover a completely novel view. It has instead been to develop the decisive moral reason view and present new arguments for it – most importantly, the argument that this view is uniquely well-placed to explain the connection between moral wrongdoing and moral disregard.

5. Conclusion

This paper has defended an account of the relation between moral reasons and moral requirement I call *the decisive moral reason view*. This view says that 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, in favor of any alternatives. I have argued that this view satisfies three plausible desiderata. It allows for the coherence of either affirming or denying the existence of moral options (the options desideratum). It entails moral rationalism, the thesis that an action can be morally required only if one ought to do it all things considered, and specifically vindicates the more plausible rational constraint interpretation of this thesis (the rationalism desideratum). And, in contrast with the dual-ranking view, the decisive moral reason view explains why unexcused moral wrongdoing necessarily shows moral disregard and is thus blameworthy (the disregard desideratum).

I'll conclude by pointing to one further attraction of the decisive moral reason view. Recall that a central motivation for the most moral reason view was that it yielded an attractive symmetry between moral requirement and the all-things-considered 'ought' (§1). Just as you ought to do what you have most reason to do, you are morally required to do what you have most moral reason to do. However, this analogy ran into immediate trouble, as it ruled out moral options.

The decisive moral reason view draws a different analogy between moral requirement and 'ought'. One ought to ϕ , all things considered, just in case one's normative reasons to ϕ are enough, on their own, to outweigh all of the normative reasons not to ϕ . In parallel, one is morally required to ϕ just in case one's *moral* reasons to ϕ are enough, on their own, to outweigh all of the normative reasons not to ϕ . Just as we ought to do what all of our reasons decisively favor, we are morally required to do what our moral reasons decisively favor.

References

- Archer, Alfred. 2016. Supererogation, Sacrifice, and the Limits of Duty. *The Southern Journal of Philosophy* 54: 333-354.
- Darwall, Stephen. 2006. *The Second-Person Standpoint: Morality, Respect, and Accountability*. Cambridge: Harvard University Press.
- Darwall, Stephen. 2013. Bipolar Obligation. In *Morality, Authority, and Law: Essays on Second-Personal Ethics I*, 20-39. Oxford: Oxford University Press.
- Darwall, Stephen. 2016. Making the 'Hard' Problem of Moral Normativity Easier. In *Weighing Reasons*, edited by Errol Lord and Barry Maguire, 258-277. Oxford: Oxford University Press.
- Dorsey, Dale. 2013. The Supererogatory, and How to Accommodate It. *Utilitas* 25: 355-382.
- Dorsey, Dale. 2016. *The Limits of Moral Authority*. Oxford: Oxford University Press.
- Gert, Joshua. 2007. Normative Strength and the Balance of Reasons. *Philosophical Review* 116: 533-562.
- Hieronymi, Pamela. 2001. Articulating an Uncompromising Forgiveness. *Philosophy & Phenomenological Research* 62: 529-555.
- Hurka, Thomas, and Esther Shubert. 2011. Permissions to Do Less Than the Best: A Moving Band. In *Oxford Studies in Normative Ethics*, Vol. 2, edited by Mark Timmons, 1-27. Oxford: Oxford University Press.
- Kiesewetter, Benjamin. 2017. *The Normativity of Rationality*. Oxford: Oxford University Press.
- Mill, John Stuart. 1861/2001. *Utilitarianism* (Second Edition), edited by George Sher. Indianapolis: Hackett Publishing Company.
- Muñoz, Daniel. 2021. Three Paradoxes of Supererogation. *Noûs* 55: 699-716.

- Parfit, Derek. 2011. *On What Matters* (Vol. 1). Oxford: Oxford University Press.
- Portmore, Douglas. 2009. Are Moral Reasons Morally Overriding? *Ethical Theory & Moral Practice* 11: 369-388.
- Portmore, Douglas. 2011. *Commonsense Consequentialism: Wherein Morality Meets Rationality*. Oxford: Oxford University Press.
- Portmore, Douglas. 2019. *Opting for the Best: Oughts and Options*. Oxford: Oxford University Press.
- Scanlon, T. M. 2008. *Moral Dimensions*. Cambridge: Belknap Press.
- Scheffler, Samuel. 1982. *The Rejection of Consequentialism*. Oxford: Clarendon Press.
- Schmidt, Thomas. Forthcoming. How Reasons Determine Moral Requirements. In Russ Shafer-Landau (Ed.), *Oxford Studies in Metaethics* (Vol. 18). Oxford: Oxford University Press. Available at <https://philpapers.org/archive/SCHHRD-8.pdf>.
- Slote, Michael. 1991. Review: Shelly Kagan's *The Limits of Morality*. *Philosophy and Phenomenological Research* 51: 915-917.
- Snedegar, Justin. 2016. Reasons, Oughts, and Requirements. In Russ Shafer-Landau (Ed.), *Oxford Studies in Metaethics* (Vol. 11), 155-179. Oxford: Oxford University Press.
- Snedegar, Justin. 2017. *Contrastive Reasons*. Oxford: Oxford University Press.
- Strawson, P. F. 1967/2008. Freedom and Resentment. In *Freedom and Resentment and Other Essays*, 1-28. New York: Routledge.
- Stroud, Sarah. 1998. Moral Overridingness and Moral Theory. *Pacific Philosophical Quarterly* 79: 170-189.
- Wallace, R. Jay. 2019. *The Moral Nexus*. Princeton: Princeton University Press.