



Conditional Obligation, Permissibility, and the All or Nothing Problem

Xueshi Wang¹

Received: 13 December 2023 / Revised: 6 May 2024 / Accepted: 11 August 2024
© The Author(s), under exclusive licence to Springer Nature B.V. 2024

Abstract

In Horton's All or Nothing Problem, the agent has three options: a permissible act that saves no one, a wrongful act that saves only one child, or a supererogatory act that saves two children. Some may argue that if the agent is not going to save two children, she should save none rather than just one. However, this conclusion is counterintuitive. Although there are many proposed solutions to this problem, none of them provides a fully satisfactory answer. In this article, I argue that it is plausible to accept the Individualist Obligation Principle according to which, the agent is morally permitted to fulfill a *prima facie* obligation to save an individual in grave peril, provided that doing so demonstrates a substantial engagement with the individualized obligations to each individual involved.

Keywords Conditional obligation · Moral obligation · Supererogation · The All or Nothing Problem

1 Introduction

The All or Nothing Problem is a paradox raised by Joe Horton (2017) in developing his argument for the conditional obligation of effective altruism. The paradox asks how a wrong act could be better than a permissible act in a three-way choice: Suppose two children are to be crushed by a collapsing building, and you can either (i) do nothing, (ii) save one child by sacrificing your arms, or (iii) save two children by the same cost (the Rescue Case). Intuitively, there are two plausible claims:

- (1) It is *permissible* for you to save no one.
- (2) It is *wrong* for you to save only one child.

✉ Xueshi Wang
wangxueshi@cuhk.edu.cn

¹ School of Humanities and Social Science, The Chinese University of Hong Kong, Shenzhen 518172, China

It is believed that both claims are plausible, since morality does not ask you to sacrifice your own arms to save another, and saving two requires no more of you than saving one.¹ However, if you are not going to save two, (1) and (2) would imply a counterintuitive conclusion:

(3) You *ought to* save neither child rather than save only one.

It seems implausible to say that you should save zero rather than save one. The best moral view, as Horton argues, would not discourage you from doing some good if you refrain from doing nothing. In short, (1) and (2) are plausible claims, and yet these two plausible claims would lead to an implausible conclusion (3). This is so-called the ‘All or Nothing Problem’.²

In literature there are roughly five responses to the problem. Horton (2017) rejects (1) by suggesting that if you are willing to make a sacrifice to save at least one child, you should save two. Pummer (2019) and Muñoz (2021) attempt to avoid (3) by rejecting a bridge principle, which states that you should choose a permissible act with a less favorable outcome over a wrong act with a significantly better outcome. Portmore (2019) rejects the bridge principle for a different reason and proposes an alternative solution: It is not true that there is more reason to save one child than there is to save none. Sinclair (2018) agrees with the view that you are conditionally obligated to do more good by doing some good, at least when it comes to claims. However, he maintains that the conditional obligation of effective altruism does not apply to benefits beyond claims. McMahan (2018) rejects (2), suggesting that it is wrong yet permissible to incur great personal cost to do the lesser good. Jin (2021) even contends that saving only one child is not morally wrong.

In Section 2, I consider Horton’s Optimific Altruism, and argue that willingness-to-save cannot simply determine the deontic status of an act that brings about only a suboptimal outcome. In Section 3, I examine Pummer’s Conditional-Permission Claim and argue that this claim is not sufficient to resolve the All or Nothing Problem. In Section 4, I argue that Muñoz’s Justifying and Requiring Principle may not provide a satisfactory solution to certain cases falling under the problem’s scope. In Section 5, I consider Sinclair and McMahan’s accounts and argue that it is plausible to accept the Individualist Obligation Principle according to which, each individual in grave peril is owed a *prima facie* individualist obligation due to their moral value, and fulfilling the obligation at great personal cost may be morally permissible, even if it is morally wrong to do so.

¹ In this article, I simply follow Horton’s suggestion that there is no obligation to maximize impartial good or minimize harm in cases of the form picked out by the All or Nothing Problem. For a defense of this view, see Scheffler (1994), Kamm (1996).

² Generally, the problem is identified as a paradox of supererogation in which a wrong option is morally better than a permissible option protected by a kind of moral reasons not to supererogate (Muñoz, 2021).

2 Optimific Altruism

In order to avoid the counterintuitive implication (3), Horton suggests that we can reject (1) by accepting the following claim:

(1*) If you were not willing to save either child, it would be permissible for you to do nothing, but because you are willing to save one child at great personal cost, you should save both children.³

(1*) and (2) imply (4), namely, if you can save one, you should ideally save two, but if you do not save both, you should still save one as the next best action. In general terms, for any agent x and any two distinct individuals y and z , if x is willing to save at least one individual, then the following holds: If x ought to save both y and z but does not do so, then x ought to save y or z . Replacing (1) with (1*) would avoid (3), even if (2) is accepted. This response is generalized by Horton (2017: 99) as Optimific Altruism according to which, if you are not willing to bear great personal cost to do some good, it is permissible for you to do nothing, but if you are willing to bear the cost and there is no adequate agent-relative reason to bring about a sub-optimal outcome, you are required to do the most good at the same cost.

There are two problems with Horton's solution to the problem. First, there is a gap between what you will and what you actually do. Rejecting (1) in favor of (1*) is implausible in cases in which you are willing to perform supererogatory acts (Pummer, 2019: 280). This point can be illustrated by the presence of the multifaceted nature and inherent ambivalence of willingness. You may simultaneously have two genuine forms of willingness—a willingness to make a rescue effort at great personal cost and a willingness to let both children die to keep your arms. It may happen that you decide to refrain from making a sacrifice without giving up willingness-to-save. In this case, Horton's suggestion would imply that it is wrong for you not to make a sacrifice to save both children. However, this implication seems overly demanding.

Second, willingness-to-save cannot simply determine the deontic status of an act that brings about only a suboptimal outcome, since it can be easily affected by psychological factors such as regret, especially in demanding rescue situations. As Horton's argument relies on the notion of willingness-to-save, it is crucial to explore the potential scenarios where one might withdraw their willingness to save others. I argue that the *feeling of regret* may make it harder for you to continue with the rescue effort, even if the welfare costs of saving one and saving two are the same.

³ Horton's (2017: 97–98) notion of willingness is understood in terms of justification. An act is morally wrong if it cannot be justified to those affected by it. If you are willing to save one person, you cannot reasonably appeal to the amount of sacrifice needed for the rescue to justify saving only one, because saving two requires no greater sacrifice. Saving one person when you could have saved two at no additional cost lacks justification to the second person.

The feeling of regret is a complex psychological state that may arise when you are engaged in a highly demanding rescue effort. It is a growing sense of unease and doubt, a nagging feeling that the situation is becoming too much to bear.⁴ As you immerse yourself in the rescue effort, the reality of the situation begins to sink in. You may experience a profound internal conflict. On the one hand, you may feel a strong moral pull to help the children in need. This might be driven by empathy, a sense of beneficence, or a general commitment to doing what is right. On the other hand, as the costs of the rescue become clearer, you may start to question whether you are truly prepared to make such an extreme sacrifice. This ambivalence can be deeply unsettling, as your moral convictions clash with growing reservations. As the rescue effort unfolds, you may develop an intense desire to remove yourself from the stressful and emotionally charged situation. This aversion to the immediate circumstances of the rescue can overwhelm your original noble intentions.

The feeling of regret can be especially acute when the act of saving the first child is likely to be emotionally exhausting. Having made the difficult decision to sacrifice their arms and having followed through with it, you may feel a deep sense of psychological depletion. This depletion can make saving the second child feel like an insurmountable challenge, even if it does not require additional physical costs. Also, you may experience cognitive dissonance as your initial belief in the value of the sacrifice clashes with growing doubts and begin to discount the future value of your sacrifice as the psychological costs become more salient. This cognitive dissonance and temporal discounting can make it harder for you to maintain your commitment to the rescue effort over time.

The feeling of regret can arise even when the tangible costs of continuing the rescue effort remain the same. To illustrate this point more clearly, let us consider a variant of the Rescue Case. Imagine you rush into a burning building to save a child. In the process of rescuing the child, you sustain severe burns over a portion of your body. Now imagine that after exiting the building with the first child, you learn that there is another child still trapped inside. Let us assume that rescuing the second child would not expose you to any additional physical risk, and the physical pain you would experience while rescuing the second child is relatively negligible.

This example parallels the Rescue Case in key ways. The physical costs are the same whether you save one or two individuals, but the psychological costs are cumulative and can reach a breaking point after the first rescue where we can reasonably expect you not to continue. In demanding situations like these, there may exist a gap between what morality requires and what we can reasonably expect of you given human limitations and psychological realities in general. The feeling of regret is a way of conceptualizing this gap—it is the complex psychological state that can make continuing a rescue effort feel subjectively harder, even if it remains objectively required.

⁴ The feeling of regret I depict here is different from the regret one has towards one's own past behaviors (Williams & Bernard, 1981), from the regret that can be felt about states of affairs constituting the source of guilt (Miller, 2020), and also from the regret that emerges from the pernicious consequences of believing an evidence-based proposition (Kelly, 2002).

This conceptualization aligns with Brian McElwee's (2023) notion of psychological difficulty in a certain sense. McElwee argues that demandingness has two aspects: cost and psychological difficulty. The cost refers to the sacrifices the agent must make, while the psychological difficulty refers to how hard it is for the agent to bring themselves to perform the act. Just as an act can be permissible to omit due to significant physical costs, an act might also be rendered permissible to omit due to significant psychological difficulty. The feeling of regret could potentially make saving two so psychologically demanding that saving one becomes permissible.

However, unlike McElwee's view, the permissibility of saving one child does not alter the moral wrongness of this act. The rightness or wrongness of an act is determined by weighing the moral reasons supporting and opposing it. These reasons, grounded in requiring and justifying moral considerations (see Section 4), provide a framework for evaluating the moral status of an act. The primary factors that determine an act's rightness or wrongness are moral reasons, such as the value of saving a life or the prerogative to avoid self-harm. These moral considerations hold true regardless of the specific psychological costs to the agent. Psychological difficulties, however, operate on a different level. They do not alter the fundamental moral calculus but instead introduce a separate consideration of what can be reasonably expected of you, given the general limitations that humans face. When an act becomes psychologically overwhelming and reaches a breaking point, it can become permissible to omit that act, even if it remains the morally optimal choice.

The feeling of regret is not a moral reason that directly determines the rightness or wrongness of an act. Instead, it is a psychological condition that we can reasonably anticipate you might experience while engaging in the rescue effort.⁵ Reasons provide basis for the rightness or wrongness of an action, while psychological difficulties create space for its permissibility. This permissibility is a separate consideration that tempers our practical expectations of you without altering our fundamental moral assessment. This framework provides a nuanced understanding of the relationship between moral requirements and psychological difficulties. It upholds the primacy of moral reasons in determining rightness and wrongness of an action while still recognizing the ethical significance of psychological realities. There is a set of actions that can be categorized as morally wrong yet permissible. A detailed argument supporting this category of actions will be presented in Section 6.

Noteworthy, it is important to distinguish between two types of situations involving supererogatory acts (McMahan, 2018). Cases where an agent first makes a sacrifice that renders some act supererogatory and then faces a choice between doing that supererogatory act or doing an even better supererogatory act. In these cases, once the initial sacrifice is made, doing the best act becomes morally required because it no longer involves any or enough additional cost to the agent. Cases where an agent faces a choice between multiple supererogatory acts, before making

⁵ Even though the feeling of regret is a factor that can be weighed against the moral reasons, its strength is not sufficiently strong enough to outweigh the reason to save the second child. While we should not be entirely insensitive to the psychological difficulty the agent faces, we should be very cautious about letting those burdens override moral obligations, especially when lives are at stake.

the sacrifice that renders each of them supererogatory. In these cases, there is no point at which the agent has already absorbed the cost, so doing the best act never transitions from supererogatory to obligatory. Cases like the Rescue Case fall into the first type, while most instances of charitable giving belong to the second type.

In this article, I am focused mainly on the first type. Imagine a scenario where you have three options: doing nothing, pressing Button 1 to save one child at the cost of your arms, or pressing Button 2 to save the same child and another one at the same cost. In the Rescue Case, saving the second child remains possible after you have incurred the cost necessary to save the first child. However, in this case, if you have incurred the cost associated with saving only one child, it is not possible for you to save two children without incurring additional cost. My discussion may not apply to cases where the *choice* among supererogatory acts precedes the *sacrifice* that makes each act supererogatory.

3 The Conditional-Permission Claim

Pummer's (2019) solution to the problem lies in rejecting a bridge principle implicitly employed to support the counterintuitive implication (3). He argues that (3) cannot be derived from the mere conjunction of (1) and (2) unless a bridge principle between non-conditional and conditional obligations is assumed:

The Bridge Principle. For any agent S and any three mutually exclusive options A, B, and C, if (1) A and C are morally permissible, (2) B is morally wrong, and (3) S is not willing to do C, then S ought to do A.

The Bridge Principle would say that if you are not willing to save both children, you are obligated to do nothing over saving one child. In some cases, it seems counterintuitive to think that people should choose a permissible option over a wrong one. These cases often share a common structure: A person, who faces a choice between three options, can either refrain from action, perform a supererogatory act that involves substantial personal sacrifice but leads to a highly desirable outcome, or choose a similar supererogatory act with the same cost that not only achieves the same result but also offers an additional benefit. Pummer (2019: 284) finds it plausible to reject the Bridge Principle and instead accept the Conditional-Permission Claim.

The Conditional-Permission Claim. For any agent S and any set of options {A, B, C, ...} ranked in order of their goodness, such that A is the best option, B is the second-best option, and C is worse than both A and B, if (1) C is morally permissible for S unconditionally, (2) A is morally better than B, and B is morally better than C, and (3) S is not going to do A, then S is morally permitted to do B.

The Conditional-Permission Claim offers a framework for thinking about the permissibility of second-best options in contexts where the best option is not going to be chosen. The problem with Pummer's argument is that the Conditional-Permission Claim is not sufficient to explain the data in some cases with similar structural characteristics to the Rescue Case. Within these cases, there exist three options: refrain from taking any action, perform a supererogatory act that moderately improves the

well-being of many people at great personal cost, or carry out the aforementioned supererogatory act and an additional one that significantly benefits another individual at the same cost. Consider the following case:

Finger Pain: You can either (i*) do nothing, (ii*) relieve 100 children from severe finger pain by sacrificing your fingers, (iii*) relieve the same 100 children from severe finger pain and save a separate child's life, at the same cost.

In Finger Pain, it is plausible to think that both (i*) and (iii*) are permissible, and (ii*) is wrong. The Bridge Principle implies that if you are not going to do (iii*), you should do (i*). If the Conditional-Permission Claim is plausible, it is believed to have explanatory power to account for the absurdity of the implication of the Bridge Principle when applied to Finger Pain. (iii*) is morally better than (ii*) in a way that makes (ii*) wrong, since gratuitously leaving a child to die is not acceptable and there is no other justification for doing that. In Pummer's term, (iii*) is seriously morally better than (ii*). Is (ii*) morally better than (i*) in a way that makes (i*) wrong? To answer this question, it is necessary to evaluate states of affairs according to their goodness and see whether (ii*) is better than (i*) at the first place.

For the sake of argument, let us assume that the children's suffering caused by finger pain does not really endanger their health, and it is also assumed that the aggregation of individuals' suffering is intelligible and its application to the case under discussion is not questioned.⁶ Even though the finger pain suffered by each child can be aggregated in one way or another, it is not quite obvious that the benefits derived from sacrificing your fingers could outweigh the great personal cost you incur in doing the supererogation. I believe that relieving 100 children from finger pain is not as important as keeping your fingers, even though the prerogative not to self-harm is simply not taken into consideration. So (ii*) is not morally better than (i*).

If (ii*) is not morally better than (i*), it does not make sense to ask whether (ii*) is morally better than (i*) in a way that makes (i*) wrong. According to the Conditional-Permission Claim, if you are not going to do the best thing, you are permitted to do the next best thing. In Finger Pain, given (i*) is better than (ii*), it implies that if you are not going to choose (iii*), it is permissible for you to do nothing.⁷ However, this implication does not offer any new information but rather tells us what we already know to be the case. It says nothing about the permissibility of (ii*) in Finger Pain. While it may not be feasible to maximize the value of goodness by choosing (ii*), a plausible moral view would not discourage you from helping others, even if it entails a significant personal sacrifice, provided that the benefits resulting from such an act are substantial and potentially meaningful in aggregate. An altruistic individual may choose to put their own life at risk in order to prevent others from

⁶ Noteworthy, not all philosophers acknowledge the legitimacy of aggregation of individuals' claims in moral deliberation. For the relevant discussion, see Hirose (2014).

⁷ It is permissible for you not to sacrifice your fingers, and the Conditional-Permission Claim implies that if you are not going to do the best thing, it is the next best thing you are permitted to do. As a result, it seems that you are morally required to do so.

experiencing severe harm or to enable them to lead a more fulfilling life, even if the benefits of such a sacrifice are not as considerable as those achieved by saving multiple lives.⁸ If she places a higher value on others' well-being than her own life, there appears to be no compelling reason to stop her from doing supererogation.

The Conditional-Permission Claim fails to predict the right answer in Finger Pain in which the best act is seriously better than the wrong act, while the wrong act is not substantially better than the least best yet permissible act. However, it is intuitively appealing to think that you are still permitted to do the wrong act, conditional on your not doing what is best.⁹ The Conditional-Permission Claim is not sufficient to account for the absurdity of the implication of the Bridge Principle when applied to certain cases falling under the purview of the All or Nothing Problem.

4 The Justifying and Requiring Principle

Muñoz (2021) argues against the view that the understanding of right and wrong depends solely on one moral dimension. Instead, he suggests that permissibility should be understood in terms of two independent dimensions—requiring reason and justifying reason. Requiring reasons tend to make an act morally required, while justifying reasons tend to prevent an act from being made impermissible by requiring reasons.¹⁰

The Justifying and Requiring Principle. For any agent *S* and any two distinct acts *x* and *y*, *x* is morally obligatory for *S* if (1) there are requiring reasons R_x or justifying reasons J_x for *S* to do *x*, (2) there are requiring reasons R_y or justifying reasons J_y for *S* to do *y*, and (3) the combined strength of R_x and J_x is greater than the combined strength of R_y and J_y .

In the Rescue Case, what makes saving one morally wrong is that it is unjustifiably worse than saving two. Although saving zero is worse than saving one, there is an adequate justifying reason (compensating the deficit in requiring reason) not to sacrifice your arms to do supererogation. By contrast, there is more requiring reason to save two than there is to save one, but there is no justification to make up for the deficit in requiring reason. The interaction between requiring reason and justifying reason is constructed by Muñoz (2021: 704) exactly as follows:

⁸ Some may argue that it is wrong to make silly sacrifices (Raz, 2018). However, I believe that sacrificing your fingers to relieve 100 children from finger pain is not a silly thing. It improves their situation in a substantial way, and this is true particularly for those who do not reject moral aggregation.

⁹ Consider the following options: (1) saving group A consisting of 999 people at no cost; (2) saving group B consisting of 999 other people by sacrificing your arms; (3) saving group C consisting of all individuals in group B plus a separate person by sacrificing your arms. Intuitively, (3) is seriously better than (2), but (2) is not substantially better than (1). Pummer's Conditional-Permission Claim would mistakenly imply that if you are not going to save group C, you are not permitted to save B. This is a counterintuitive implication.

¹⁰ In principle, these two dimensions can be considered conceptually distinct from one another. An act may rank higher on one dimension but not on the other, resulting in a moral mismatch between requiring reasons and justifying reasons (Muñoz & Pummer, 2022).

	Save Zero	Save One	Save Two
Requiring Reason	1	5	10
Justifying Reason	10	0	0

By comparing these two dimensions of reasons, it is possible that a wrong act is better than a permissible act in a specific choice set. You have the most requiring reason to save two, the second most requiring reason to save one, and the least requiring reason to save zero. You have a sufficient justifying reason not to harm yourself but only have a negligible justifying reason to make a rescue effort. The Justifying and Requiring Principle implies that saving one is worse than saving two but still better than saving zero. It seems that this two-dimensional model can explain people’s judgements across various scenarios, including the Rescue Case.

I argue that a crucial challenge for Muñoz’s Justifying and Requiring Principle is the question of how to assign appropriate weights to justifying and requiring reasons. This challenge is underscored by the risks of setting the value of justifying reasons either too low or too high and the implications it has for decision-making. If the weight assigned to justifying reasons is minimal compared to the weight of requiring reasons, this would render justifying reasons inconsequential, as they would rarely, if ever, be strong enough to tip the balance in favor of an action that goes against the requiring reasons. For instance, if the justifying reason to protect one’s bodily integrity is given a very low weight compared to the requiring reason to save others, it could lead to the counterintuitive conclusion that individuals are morally obligated to sacrifice their own well-being for the sake of others in nearly all circumstances.

Imagine you have three options: do nothing, save two children but lose your arms, or save three children at the same cost. Based on Muñoz’s suggestion mentioned above, the interaction between requiring reason and justifying reason would be constructed by as follows:

	Save Zero	Save Two	Save Three
Requiring Reason	1	10	15
Justifying Reason	10	0	0

In this case, the Justifying and Requiring Principle implies that you are required to save three lives instead of protecting your own arms. It seems to me that this requirement is too demanding because it fails to recognize the moral significance of an individual’s right to bodily integrity.¹¹ Importantly, setting the value of justifying reasons too low could undermine the very purpose of introducing justifying reasons into the moral calculus. The concept of justifying reasons is meant to provide a countervailing consideration that can sometimes outweigh requiring reasons, allowing for the

¹¹ Non-consequentialists and some of consequentialists are reluctant to accept the view that, in the absence of special obligations, you are morally required to sacrifice your arms to save three persons’ lives. For discussion of the demandingness of the requirement, see Williams (1973), Scheffler (1994), and Mulgan (2005). Muñoz’s view makes it hard to recognize this non-consequentialist element in his account of moral reasons.

permissibility of acts that would otherwise be considered wrong. If justifying reasons can be consistently overridden by requiring reasons, they would fail to serve the intended function and would become largely irrelevant in decision-making.

On the other hand, setting the value of justifying reasons extremely high also poses significant challenges. If the weight assigned to justifying reasons is so high that it can rarely, if ever, be outweighed by requiring reasons, it could lead to a moral framework that is excessively permissive. If the justifying reason to protect one's bodily integrity is given such a high weight that it consistently overrides the requiring reason to save others, it could lead to the counterintuitive conclusion that individuals are not required to save others even in cases of large-scale life-or-death scenarios.

Imagine you have three options: do nothing, save 100 children but lose your arms, or save 101 children at the same cost. The interaction between requiring reason and justifying reason may be constructed by as follows:

	Save Zero	Save 100	Save 101
Requiring Reason	1	100	105
Justifying Reason	> 105	0	0

In this case, the Justifying and Requiring Principle implies that it is permissible for you to keep your arms instead of saving lives. Some might find this implication counterintuitive, because it fails to capture the moral intuition that there are some circumstances in which the moral requirement to save others is so strong that it overrides personal interests. Setting the value of justifying reasons extremely high could effectively nullify the role of requiring reasons in most moral deliberations. If requiring reasons are consistently outweighed by justifying reasons, they would cease to serve as genuine moral requirements.

It might be suggested that we can assign a value to the justifying reason of saving zero in a way that works for both sets of choices {saving zero; saving two; saving three} and {saving zero; saving 100; saving 101}. If we say that the justifying reason for saving zero has a value of 50, it means that keeping one's arms is less important than saving more than ten children, but more important than saving just three children. However, the specific value of 50 assigned to saving zero seems deeply arbitrary and lacking in principled justification. It is not clear why this particular value would be chosen, or how it could be defended against alternative values that might yield different implications in different choice sets. Attempting to reduce moral considerations to numerical values may seem to offer a kind of clarity, but it often does so at the cost of oversimplification. It can lead us to arrive at conclusions that are more driven by the arbitrary assumptions of value assignments than by a deep engagement with the underlying moral issues.

The Justifying and Requiring Principle is based on a particular view of values. On this approach, values can have both requiring and justifying dimensions. By evaluating each option in a choice set along these two dimensions and aggregating them additively, the Justifying and Requiring Principle provides a way of comparing options.

The Multidimensionality of Values

$$(X_{d1}, X_{d2}) > (Y_{d1}, Y_{d2}) > (Z_{d1}, Z_{d2}) \text{ iff } (X_{d1} + X_{d2}) > (Y_{d1} + Y_{d2}) > (Z_{d1} + Z_{d2})$$

The principle states that for any three options X, Y, and Z in a choice set, X is morally better than Y if and only if the sum of X's scores on the requiring and justifying dimensions is greater than the sum of Y's scores on those dimensions, and the same with Y and Z. However, as discussed above, this quantitative approach either results in counterintuitive implications or relies on arbitrary value assignments. It reveals the difficulties involved in trying to reduce complex moral considerations to numerical scores and aggregating them in a purely additive way.

These challenges indicate that to adequately assess the relative importance of moral options, we may need to move beyond just quantitatively balancing reasons and also include qualitative factors in the moral framework. The qualitative model can provide a more plausible explanation for the different intuitive judgments in the choice sets: {saving zero; saving two; saving three} and {saving zero; saving 100; saving 101}. It considers the special moral status of the right to bodily autonomy while also recognizing the increasing urgency of saving lives as the numbers grow larger. The central idea is that the right to bodily autonomy, while carrying special weight, is not an absolute or lexically prior constraint that always trumps other moral considerations. Rather, it is a presumptive right that can be overridden in circumstances where adhering to it would lead to catastrophic outcomes.

In the first choice set, the number of lives at stake is relatively small, and the consequences of respecting the right to bodily autonomy are not so severe as to justify overriding it. The loss of two lives, while tragic, is not a catastrophic outcome that would warrant violating a fundamental human right. The justifying reason to preserve bodily integrity thus outweighs the requiring reason to save two lives. In the second choice set, however, the moral stakes are much higher, and the consequences of respecting the right to bodily autonomy are much more severe. The loss of 100 or 101 lives is a catastrophic outcome that would have devastating consequences for the individuals involved. In this extreme case, the moral urgency of saving so many lives is so overwhelming that it justifies overriding the right to bodily autonomy, even though this right continues to carry special weight.

The qualitative model can thus explain the shift in intuitive judgments between these two choice sets by appealing to the idea of a threshold of moral urgency at which the presumptive right to bodily autonomy can be justifiably overridden. This threshold is not reached in the first choice set, but it is reached in the second choice set. Importantly, this explanation does not rely on a purely quantitative aggregation of lives saved, but rather on a qualitative judgment about the severity and disproportionality of the consequences in each case. The right to bodily autonomy is not simply outweighed by a larger number of lives saved, but rather overridden in an extreme case where the moral stakes are exceptionally high and the consequences of inaction are profoundly dire.

This qualitative approach may somehow avoid the ‘small-benefits-can-tip-the-scales’ objection. Comparing the following two choice sets: {doing nothing; saving 50 children by sacrificing your arms; saving 100 children at the same cost} and {doing nothing; saving-49-children-from-death-and-one-child-from-paralysis by sacrificing your arms; saving 100 children at the same cost}.¹² For the sake of argument, let us suppose that saving 50 is a bright line separating the permissible from the impermissible, since some may think that it is deemed substantial enough to be an adequate part of the ideal supererogatory act of saving 100.¹³ In this case, saving 50 is permissible, but saving-49-from-death-and-one-from-paralysis is impermissible, even though the difference in outcomes is small. It seems implausible that such a small difference could make the difference between an act being permissible and impermissible.

This concern raises questions about how to specify the threshold in the qualitative framework. A sharp line can lead to counterintuitive results in cases just above and below the cutoff point. However, the qualitative view could be blurred or smoothed out to avoid the problematic implication. One way is to introduce a range of cases near the threshold where the permissibility of choosing the lesser good gradually tapers off, rather than suddenly flipping from a full permission to a full prohibition. Another related approach is to say that in borderline cases extremely close to the threshold, choosing either the lesser or greater good is permissible. If saving 55 lives is squarely above the threshold and saving 45 is squarely below it, cases like saving 49 lives plus preventing paralysis could be treated as a sort of moral tie with saving 50. The agent would be permitted to choose either.¹⁴

5 The Individualist Obligation Principle

To solve the problem, Sinclair (2018: 53–55) distinguishes between the realm of claims, where a failure to satisfy claims requires justification based on agent-relative sacrifice, and the realm of benefits, where a failure to provide greater benefits to potential recipients at no extra cost does not trigger a demand for justification. According to his Thoroughgoing Non-Consequentialist View, one has a *prima facie* obligation to satisfy claims when encountering emergency cases. If they are willing

¹² I thank one of the reviewers for giving me this example. It made me think more about the ‘small-benefits-can-tip-the-scales’ objection.

¹³ Notably, my argument here does not hinge on whether saving 50 lives is considered substantial enough to be an adequate part of the ideal supererogatory act of saving 100 lives. If you find it implausible, it may be helpful to draw the separating line according to your own assessment.

¹⁴ These approaches capture the intuition that small differences around the threshold should not completely flip an act from permissible to impermissible. However, it should be admitted that any view which sometimes permits choosing a marginally lesser good over a marginally greater one will be vulnerable to some version of the ‘small-benefits-can-tip-the-scales’ objection. While blurring or smoothing the threshold may mitigate the specific problem, it does not completely eliminate the difficulty of weighing small differences in large-scale outcomes. Nonetheless, I believe some form of blurred or smoothed threshold is likely the most promising approach for the qualitative approach to take in response to this concern.

to make a sacrifice to satisfy any claim, they would be obligated to satisfy all claims through this sacrifice.

First, Sinclair does not really give a satisfactory answer to the problem. The distinction between the realm of claims and the realm of benefits is helpful in the sense that you are not required to be an effective altruist described as doing as much good as possible with whatever good you have. However, the All or Nothing Problem still remains in the realm of claims, since the counterintuitive implication, namely that you act wrongly by saving one rather than saving zero, does not simply disappear.

Second, Sinclair (2018: 52–53) claims that in the Rescue Case you have a *prima facie* obligation to save both children. However, the Thoroughgoing Non-Consequentialist View is not the only plausible interpretation of the conditional obligation in emergency situations. I argue that it may be misleading to say that there is a *prima facie* obligation to save both children. An alternative framework of obligations can be employed to analyze the Rescue Case, which can be understood as follows: After having saved the first child, it is then found that an additional child can be rescued without incurring any further cost. Specifically, before entering the building, you are obligated to save one of the two children, even though the failure to fulfill the obligation can be justified by great personal cost. As you enter the building and have saved the first child, you subsequently incur an additional obligation to save the second child, since the child can be saved without any effort. The obligation should be fulfilled, and to do otherwise is morally wrong. This structure of obligations in this case can be understood as follows: You have a *prima facie* obligation to save the first child, and then incur an additional obligation to save the second child conditional on your fulfilling the obligation of saving the first child. Strictly speaking, you only have a *prima facie* obligation to save each child, but not to both.¹⁵

While Sinclair places significant value on the notion of the separateness of persons, this two-step structure of obligations is more sensitive to this idea. Sinclair (2018: 55–58) argues that Pummer (2016) and Horton's half-hearted non-consequentialist account of conditional obligation fails to acknowledge the distinction between the cases in which claims conflict with each other and the cases in which claims do not. Sinclair believes that their account is insensitive to the notion of the separateness of persons in the sense that there is no impersonal evaluative perspective from which to compare the well-being of different individuals. Although Sinclair's account performs much better than its rivals do, as it does not appeal to the default duty to promote impartial goodness, it does not give enough weight to this notion. The Individualist Obligation Principle appears to be a better fit for the notion of the separateness of persons compared to the Thoroughgoing Non-Consequentialist View.

¹⁵ If you save one and leave another to die, your act only treats the second child wrongly. The first child cannot reasonably complain that her claim is not given due moral weight. However, Sinclair's account seems to imply that, as long as you are under a *prima facie* obligation to satisfy as many claims as possible, the first child can also make a reasonable complaint from the point of view of abstract reasoning, as failing to save either would normally constitute the violation of the obligation, regardless of whether her claim is satisfied or not.

Another solution to the problem is to suggest that it is wrong yet permissible to incur great personal cost to do the lesser good. McMahan (2018) believes that it would be helpful to make a distinction between the wrongness of an act and the impermissibility of it. If you have acted to save the first child, you should save the second, and you act wrongly if you save only one. But it does not mean that it is not conditionally permissible for you to save one if you refrain from saving both. This suggests a special category of acts that are wrong yet permissible. McMahan is quite right in pointing out that an act might be wrong yet permissible, but he does not provide a detailed argument for it. In the following, I will argue that the permissibility of saving one is inferred from what I call ‘the individualist obligation’.¹⁶ The core of solving the problem is to reject (2). Accurately speaking, (2) should be revised so that it can accommodate a category of acts that are morally wrong yet permissible.

What does ‘wrong’ mean, if it means something other than impermissible? The rightness or wrongness of an act is determined by the balance of requiring reasons and justifying reasons. A requiring reason is an aspect of an act that inclines to make this act morally required. A justifying reason is an aspect of an act that contributes to it being the case that the act is not morally wrong (Pummer, 2021). An act is *wrong* when it unjustifiably falls short of the alternative there is most requiring reason to do. But it may still be *permissible* if it represents a substantial partial fulfillment of the *prima facie* obligations generated by the intrinsic moral value of each individual involved, even though it remains suboptimal in comparison to the ideal act.

An act can be conditionally permissible in a qualified sense if it crosses a threshold of adequate engagement with the overall set of *prima facie* obligations binding the agent, where those obligations are grounded in the separate moral value of each individual rather than in the impersonal aggregation of value. This permissibility judgment is based on a holistic assessment of how well the act appreciates the ethical significance of each person’s situation and respects the obligations to them qua individuals. (Numbers can be relevant to this judgment insofar as they bear on the pattern of the obligations fulfilled or unfulfilled, but not in a way that treats the individuals’ claims as simply additive.) The core insight is that there can be permissibility even among wrong acts, depending on how substantially they reflect proper regard for the intrinsic value of individuals and the moral significance of the obligations to them.

Let us write $WR(X/\neg M)$ to mean that X is wrong conditional on not doing an ideal supererogatory act M , and write $WRPE(X/\neg M)$ to mean that X is wrong yet permissible conditional on not doing an ideal supererogatory act M .

Conditionally Wrong. $WR(X/\neg M)$ iff (1) X falls short of the ideal supererogatory act M there is more requiring reason to do; (2) this shortfall cannot be justified by

¹⁶ The term I am using here is informed by Timmermann (2004), who suggests the individualist lottery procedure as a way to represent claims in conflict cases where they are owed to individual persons rather than to a collective group. Saunders (2009) supports this approach, which is seen as an attractive alternative to the perspectives put forth by Taurek (1977), Parfit (1978), Broome (1990), Kumar (2001), Hirose (2004), and Munoz-Dardé (2005) on the normative role of numbers.

any countervailing moral reasons; and (3) the alternative that is compatible with $\neg X$ and $\neg M$ is unconditionally permissible.

Conditionally Wrong yet Permissible. WRPE($X/\neg M$) iff (1) X unjustifiably falls short of the ideal supererogatory act M there is more requiring reason to do; (2) X still represents a substantial partial fulfillment of the *prima facie* obligations generated by the intrinsic moral value of each individual involved, such that it crosses a threshold of adequate engagement with the overall set of obligations binding the agent in light of their supererogatory decision; (3) X does not infringe on the rights of others.

In the Rescue Case, although you have a prerogative not to self-harm, it does not change the normative fact that you have a *prima facie* obligation to each individual child to save their life. This obligation should be understood in terms of individualist obligation according to which, the obligation to provide aid is owed to each child separately, rather than to the group of children as a whole. The key idea is to analyze the conditional obligation on an individualized basis, taking seriously the intrinsic moral value of each person qua individual. On this view, your obligation to rescue is grounded in the distinct value of each child's life, not in the impersonal maximization of overall welfare.

Imagine there is only one child facing a life-threatening situation. You have a *prima facie* obligation to save the child's life, but this obligation can be outweighed by your prerogative not to harm yourself. Saving the child at great personal cost is praiseworthy, but not morally required. Now, imagine that after saving this child, you find it possible to save another child without any difficulty. At this point, saving only one child becomes morally wrong. The obligation to save the second child is a separate consideration that cannot be outweighed by the prerogative not to self-sacrifice, since you have already incurred the cost. Failing to save the second child without a justifiable reason is morally wrong. Nevertheless, even though it is wrong to let the second child die pointlessly, it is still permissible to save only one child, given that you are not saving both. This scenario involves two separate moral obligations, each with different moral bases. The first obligation involves a conflict between self-interest and altruism, while the second obligation is simply about what we owe to each other.¹⁷

In general, morality does not require you to be a perfect altruist, always ready to save as many people as possible at great personal harm. The significant cost to yourself can exempt you from an obligation you would otherwise have, rendering it 'inactive'. But if you choose to provide aid despite the cost, the now 'reactivated' obligation to save the first child is considered permissible to fulfill.¹⁸ By doing so, you demonstrate respect for the unique value of an individual, irrespective of whether you also fulfill a distinct obligation to save another individual. Failing to save the second child does not negate the moral value of saving the first, given the profound impact it has on that child's life.

¹⁷ For the claim that what we owe to each other takes priority over other reasons for acts, see Scanlon (1998).

¹⁸ By saying 'reactivated obligation', I mean that a *prima facie* obligation, which is initially deactivated by a prerogative not to self-harm, can be reactivated by your voluntarily giving up the prerogative.

Indeed, the second child can reasonably complain that she is gratuitously left to die and that the moral value of her life is not being properly respected. This is precisely what makes saving one morally wrong. However, it seems not appropriate to suggest that if you do not save both children, you should save neither. The first child could also reasonably object to the idea that you are not allowed to save her life, even though you are willing to perform a supererogatory act. Her complaint is not that you are protected by a prerogative not to self-harm, which prevents the requiring reasons from making it wrong to do nothing. Rather, her complaint is that you are not permitted to save her life, given that you have chosen to relinquish the prerogative for the sake of doing something supererogatory.

The first child may be persuaded to think that if you are to save her life at great personal cost, you ought to save both at the same cost, and failing to do so is wrong. However, she might also question why the obligation to save the second child can overwhelmingly affect the obligation to save her own life. It appears that the moral significance of the first child's life is not being adequately considered, even though saving only one child is deemed morally wrong. Morality is concerned not only with the complaint of the second child but also with that of the first child. According to the second child's complaint, saving one is deemed morally wrong, while the first child's complaint makes it conditionally permissible to do so.

Each person's intrinsic moral worth generates a distinct *prima facie* obligation to aid them if they are in grave peril. If you choose to bear the relevant cost to do something good, a plausible moral view should not actively dissuade you from fulfilling the *prima facie* obligation to do that good. Each claim to aid exerts moral force irrespective of anyone else you could help in addition. It is permissible for you to save one child, not because saving one is better than saving zero from impartial concern, but because it reflects proper regard for the intrinsic value of her life and the moral significance of your obligations to her.

The Individualist Obligation Principle. For any alternative act X compatible with not doing an ideal supererogatory act M, X is conditionally permissible if: (1) There exists an alternative act L with a sufficiently strong justification, such that for every individual *i* involved, X better fulfills the *prima facie* obligation arising from *i*'s intrinsic moral value compared to L; (2) The degree to which X partially fulfills the overall set of *prima facie* obligations is greater than some threshold *t*, representing adequate engagement with those obligations; (3) For every individual *i*, X does not violate *i*'s fundamental rights.

This principle states that X is conditionally permissible if X constitutes a partial fulfillment of the *prima facie* obligations, crossing a threshold of adequate engagement with the overall set of these obligations. In the Rescue Case, saving one is wrong since it is unjustifiably worse than saving two, but it is still permissible in the sense that it represents a substantial fulfillment of the *prima facie* obligations to each child, such that it crosses a threshold of adequate engagement with the moral significance of each individual's life. I think the key here is to focus on the notion of what constitutes an adequate part of an ideal supererogatory act, and how that relates to properly appreciating the significance of the *prima facie* obligations at stake.

What counts as substantial enough to reach a given threshold depends on a variety of contextual factors, like the overall magnitude of the obligations at stake, the demandingness of the sacrifice involved, and the relative difference between the good achieved and the good foregone. There may be no precise formula for working this out. The threshold for what counts as substantial may vary and scale up as the number of individual claims increases. Saving one out of two might be a satisfactory response to the overall moral math, whereas saving one out of 100 is clearly unsatisfactory.

Comparing the following choice sets: {saving zero, saving one, saving 100} and {saving zero, saving 99, saving 100}. Saving 99 people at great personal cost, while not ideal, does enough to engage with and respond to the stringent moral reasons in play to cross the threshold of permissibility. It demonstrates a commitment to make a great sacrifice to engagement with the significance of each individual's life. In contrast, saving only one person when you could save 100 at no extra cost seems to fall too far short of appreciating the moral weight of the unfulfilled *prima facie* obligations to be even permissible. The gap between the good you have done and the good you have left undone is much wider, to the extent that your act clearly does not exhibit adequate respect for the moral significance of the additional individuals you might have saved.

Importantly, this is not a matter of aggregation of claims for an act to be (wrongly) permissible, namely the aggregation of 99 individual claims is greater than that of one claim. Rather, it flows from recognizing the intrinsic moral worth of each person and the individualized obligations their peril generates. If you decline to save 99 people when you could easily do so, you are not merely failing to bring about a better outcome. You are failing to properly appreciate the moral significance of individuals' claims. The more obligations left unsatisfied, the less positive attitudes and regard you express for the significance of individuals' claims. You are expressing a much weaker commitment to the significance of those individuals, such that these unsatisfied claims make your act of saving one impermissible. Thus, it is about the attitudes and values expressed by your response to the overall moral situation. Saving more people when you easily can reflects a more substantial commitment to the moral worth of each individual and the stringency of your obligations to them, over and above the commitment evinced by saving just one.

There is another way to put this point. In the situations like the Rescue Case, where you can save more lives at no extra cost, the number of people you choose to save reflects something important about your moral commitments and attitudes. Saving one when you could easily save 100 shows far less respect for the importance of each individual compared to saving 99 out of 100. The good you have achieved is 'swamped' by the good you have neglected, reflecting an inadequate regard for the value of the satisfied claims or even the claim that has been satisfied. Crucially, this way of thinking does not treat the claims of the 99 as aggregating or outweighing the claim of the one. It focuses on the *overall* ethical quality of your act, not a quantitative summing of value units. Each individual's claim retains its moral weight, but leaving more of those claims unmet may alter the weight of the one claim you do meet.

An analogy to promises might help clarify the idea. If you have made promises to ten different people to benefit them at great personal cost, you have an obligation to each of them individually. There is clearly a difference between keeping nine of those promises and keeping only one. Keeping nine promises demonstrates a certain degree of respect for the moral weight of your commitments, even if it falls short of perfection. But keeping only one promise when you could easily have kept nine more shows a severe lack of regard for the seriousness and stringency of the responsibilities you undertake. The nine unmet obligations ‘swamp’ the significance of the one fulfilled obligation in a way that renders your overall conduct ethically inadequate. Similarly, saving nine when I could save ten might be a sufficient partial fulfillment of the individualist obligations that reaches the threshold of permissibility, while saving one is an inadequate response to the obligations that falls below the threshold of minimal decency. The nine unmet obligations ‘swamp’ the significance of the one met obligation without aggregating the distinct individuals’ claims.

6 Conclusion

Julia Driver (1992) explores the concept of suberogatory acts, which are characterized as those acts that are bad to do but permissible. This is meant to mirror the concept of supererogatory acts which are good to do but not required. The suberogatory includes cases such as taking the last adjoining seat on a train, knowing that a couple behind you wants to sit together, or sometimes feeling that others owe you a favor based on their past good deeds, even though favors are not obligatory.

Driver wants to carve out conceptual space for acts that are permitted but still deserving of disapproval. Deontic status (supererogatory, permissible, obligatory, forbidden) and goodness/badness are considered distinct dimensions that should not be conflated. Just as the supererogatory may not be better than the obligatory, the permissible may not be better than the forbidden. Suberogatory acts are permissible but still bad in some way. They are permissible because they do not violate any moral obligations. However, they are still open to moral disapproval because they fall short of moral ideals.

There is a range of permissible acts, some of which are morally good (supererogatory or obligatory), some of which are morally neutral, and some of which are morally bad (suberogatory). The suberogatory helps capture the common intuition that some acts, while not forbidden, are still subject to moral blameworthiness. However, I think we should consider another subcategory of permissibility: supererogatory but blameworthy acts. First, unlike Driver’s ‘train seating’ and ‘owed favors’ examples, the Rescue Case can be presented as a supererogatory scenario where there is a clear point at which an obligation can be exceeded for an act to be supererogatory.

Second, the failure to do the ideal act in the scenarios characterized by All or Nothing Problem is fundamentally different from that in scenarios involving suberogatory acts. When you do not offer the last adjoining seat to the couple or do not do a favor for the friend you often help, you are completely avoiding doing something supererogatory. This failure may be considered worse than a merely permissible act.

However, when you save one child instead of two, it still constitutes a significant part of meeting the requiring reason to do the ideal. This is much better than just doing something permissible. Although this act is better than a merely permissible act, it is still morally wrong. This implies that we need to go beyond supererogatory acts, which are permissible but merely bad, to fully capture this category of acts that are permissible.

Right acts	<ul style="list-style-type: none"> • Supererogatory and praiseworthy acts • Obligatory and praiseworthy acts
Wrong acts	<ul style="list-style-type: none"> • Permissible and not blameworthy acts • Permissible but blameworthy acts (Supererogatory but blameworthy acts; Suberogatory acts) • Forbidden acts

It is right to point out that permissibility and moral goodness/badness are separate dimensions that can come apart, and that we should make room for acts that are permitted but still open to negative evaluation. On the one hand, some acts are morally worse than merely permissible (or even impermissible) acts, but these worse acts are still permissible (supererogatory acts). On the other hand, some acts can be better than what you are allowed or required to do, but these better acts are still subject to criticism (supererogatory-but-blameworthy-acts). It is this category of acts that give rise to the All or Nothing Problem. They are supererogatory because they go beyond the call of duty by being notably beneficial. They are blameworthy because as supererogatory acts, they are bad to do, even though they do enough to engage with and respond to the stringent moral reasons in play to cross the threshold of permissibility.

I reject the notion that permissibility entails moral neutrality. I also deny that permissibility entails moral rightness. The rightness or wrongness of an act operates on a distinct level and is determined by whether it is unjustifiably worse than the alternative there is more requiring reason to do, while the permissibility of an act is determined by whether it crosses a threshold of adequate engagement with the overall set of moral reasons binding the agent, where those reasons are grounded in the separate moral status of each individual rather than in the impersonal aggregation of value. This would be helpful to distinguish permissible-but-blameworthy-acts from permissible-and-not-blameworthy-acts, including merely-permissible-acts and obligatory-but-not-praiseworthy-acts. Making this distinction helps to better understand the nuances in moral judgments.

Funding This study was funded by the National Social Science Fund of China (grant number: 22FZXB098) and the MOE Project of Humanities and Social Sciences (grant number: 22YJC720018).

Declarations

Conflict of Interest The author declares no conflict of interest.

References

- Broome, J. (1990). Fairness. *Proceedings of the Aristotelian Society*, 91(1), 87–101.
- Driver, J. (1992). The suberogatory. *Australasian Journal of Philosophy*, 70(3), 286–295.
- Hirose, I. (2004). Aggregation and numbers. *Utilitas*, 16(1), 62–79.
- Hirose, I. (2014). *Moral aggregation*. Oxford University Press.
- Horton, J. (2017). The all or nothing problem. *The Journal of Philosophy*, 114(2), 94–104.
- Jin, H. (2021). A holist deontological solution to the All or Nothing Problem. *Philosophia*, 49(5), 2067–2079.
- Kamm, F. (1996). *Morality, mortality, volume II: Rights, duties, and status*. Oxford University Press.
- Kelly, T. (2002). The rationality of belief and some other propositional attitudes. *Philosophical Studies*, 110(2), 163–196.
- Kumar, R. (2001). Contractualism on saving the many. *Analysis*, 61(2), 165–170.
- McElwee, B. (2023). Cost and psychological difficulty: Two aspects of demandingness. *Australasian Journal of Philosophy*, 101(4), 920–935.
- McMahan, J. (2018). Doing good and doing the best. In *The ethics of giving: Philosophers' perspectives on philanthropy* (pp. 78–102). Oxford University Press.
- Miller, R. B. (2020). Augustine, moral luck, and the ethics of regret and shame. *The Journal of Religion*, 100(3), 361–385.
- Mulgan, T. (2005). *The demands of consequentialism*. Oxford University Press.
- Muñoz, D. (2021). Three paradoxes of supererogation. *Noûs*, 55(3), 699–716.
- Muñoz, D., & Pummer, T. (2022). Supererogation and conditional obligation. *Philosophical Studies*, 179(5), 1429–1443.
- Munoz-Dardé, V. (2005). The distribution of numbers and the comprehensiveness of reasons. *Proceedings of the Aristotelian Society*, 105(2), 207–233.
- Parfit, D. (1978). Innumerate ethics. *Philosophy & Public Affairs*, 7(4), 285–301.
- Portmore, D. W. (2019). *Opting for the best: Oughts and options*. Oxford University Press.
- Pummer, T. (2016). Whether and where to give. *Philosophy & Public Affairs*, 44(1), 77–95.
- Pummer, T. (2019). All or nothing, but if not all, next best or nothing. *The Journal of Philosophy*, 116(5), 278–291.
- Pummer, T. (2021). Impermissible yet praiseworthy. *Ethics*, 131(4), 697–726.
- Raz, J. (2018). On the moral significance of sacrifice. *International Journal of Philosophical Studies*, 26(3), 308–314.
- Saunders, B. (2009). A defence of weighted lotteries in life saving cases. *Ethical Theory and Moral Practice*, 12(3), 279–290.
- Scanlon, T. (1998). *What we owe to each other*. Harvard University Press.
- Scheffler, S. (1994). *The rejection of consequentialism: A philosophical investigation of the considerations underlying rival moral conceptions*. Oxford University Press.
- Sinclair, T. (2018). Are we conditionally obligated to be effective altruists? *Philosophy and Public Affairs*, 46(1), 36–59.
- Taurek, J. M. (1977). Should the numbers count? *Philosophy & Public Affairs*, 6(4), 293–316.
- Timmermann, J. (2004). The individualist lottery: How people count, but not their numbers. *Analysis*, 64(2), 106–112.
- Williams, B. (1981). *Moral luck: Philosophical papers 1973–1980*. Cambridge University Press.
- Williams, B. (1973). A critique of utilitarianism. In *Utilitarianism: For and Against* (pp. 75–150). Cambridge University Press.

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

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.