Philosophical Studies (2023) 180:2399–2413
https://doi.org/10.1007/s11098-023-01991-x

Can a single account of supererogation handle both finite and infinite cases?

Holly M. Smith[1,2]

Accepted: 21 May 2023 / Published online: 19 June 2023

Abstract
Discussions of supererogation usually focus on cases in which the agent can choose among a finite number of options. However, Daniel Muñoz has recently shown that cases in which the agent faces an infinite chain of increasingly less good options make trouble for existing definitions of supererogation. Muñoz proposes a promising new definition as a solution to the problem of infinite cases. I argue that any acceptable account of supererogation must (1) enable us to accurately identify supererogatory acts in both finite and infinite option cases. It must also (2) include a suitably related account of what makes one act more supererogatory than another for finite, infinite, single-choice (one agent choosing among several supererogatory options) and inter-choice (two different agents, each choosing a supererogatory option) cases. I further argue that the best current account of supererogation for finite cases works well for finite cases, but cannot handle infinite cases. However, Muñoz's proposal cannot handle inter-choice cases in either finite or infinite cases. I conclude we still need an account for infinite cases, and may have to settle for separate definitions of supererogation for finite and infinite cases.

Keywords: Commendatory value · Comparative supererogation · Finite options · Infinite options · Moral minimum · Supererogation

It has long been recognized that people sometimes perform actions that are morally good, indeed morally better than other permissible options, but that are not morally required. Philosophers call these acts "supererogatory." Someone who happens on a burning house may permissibly do nothing more than call the fire department. But

✉ Holly M. Smith
hsmith@philosophy.rutgers.edu

[1] Philosophy Department, University of California, Berkeley, USA

[2] Philosophy Department, Rutgers, the State University of New Jersey, New Brunswick, USA

if she risks her life to save a child from the fire, she performs a morally better act – an act that is supererogatory although not obligatory. It is less well recognized, but equally obvious on reflection, that some supererogatory acts are better than others. A soldier who loses his life in saving a comrade does something more supererogatory than another soldier who loses his arm in saving a comrade. This is an inter-choice case in which we compare the acts of two different agents. Indeed, the United States military recognizes different grades of supererogation, and accords them differential material benefits, in honoring soldiers with three levels of medals of valor: the Medal of Honor, the Distinguished Service Cross, and the Silver Cross (U. S. Department of Defense (Undated); Wallace (Undated a and b)). There are also single-choice cases in which an agent may have several supererogatory acts to choose among on a given occasion. A person who loses her life in saving two children does something more supererogatory than she would have done in burning her hands to save only one child.

Most cases of supererogation are ones in which the agent can choose among a finite number of options. Frank is considering whether to donate money to an unhoused person he encounters. Because he has only five ten-dollar bills in his wallet, he has six options, all permissible: donate nothing, donate a single ten-dollar bill, donate two ten-dollar bills, etc., up to donating five ten-dollar bills. But, as we'll see later, there are cases in which the agent faces an *infinite* number of options, all of which are supererogatory.

These considerations suggest that an adequate account of supererogation must meet (at least) the following criteria:

1. The account must enable us to accurately identify supererogatory acts in both finite and infinite option cases, and explain what features make these acts supererogatory.
2. The account must include a suitably related account of what makes one act more supererogatory than another for finite, infinite, single-choice and inter-choice cases.

I shall argue that no existing account of supererogation can meet both of these criteria, and that the best current account of supererogation for infinite cases cannot meet Criterion 2.

1 Defining "supererogatory"

The concept of a supererogatory act has been filled out in a variety of ways. For purposes of this paper, I shall examine the following definition.[1]

Simple Definition: An act X is supererogatory if and only if, and because:

1. X is all things considered (ATC) permissible but not ATC required.

[1] I argue for a more complex version of the Simple Definition in Smith (2023). For the purposes of this paper, the complexities of that definition are not needed.

2. X is morally good, that is, it has greater excess moral value than the permissible alternative with the least excess moral value.
3. X would impose a personal cost on the agent, that is, it has greater opportunity cost for its agent than the permissible alternative with the least opportunity cost.

Some of this terminology requires explanation. I use "moral value" here to refer to the kind of value that makes a supererogatory act good. It is an important feature of this kind of value that it does not affect whether an act is permissible, obligatory, or wrong. For example, on many moral theories, repaying a debt to a friend *earlier* than promised is morally good, but not obligatory. Different accounts of this kind of value have been offered by different authors, but there is no need here to litigate among them.[2]

Like most accounts, mine assumes that any supererogatory act must be costly to the agent. If Tina saves Tom's life by merely pushing a button, at no cost to herself, her act is morally good but hardly supererogatory.

To determine whether a given act is supererogatory, we cannot just examine how much positive moral value it would produce and how burdensome it would be for the agent. For example, in some situations an agent must choose among a set of bad options. All her options have negative moral value since all would lead to another person suffering some injury. Nonetheless some of these acts are morally superior because they would lead to less suffering than other permissible options. Similarly, an agent can perform a supererogatory act even though all her options would result in a positive welfare state for her, so long as the act in question leaves her personally worse off than other permissible alternatives would have done. The relevant comparisons are to the least morally valuable *permissible* option and to the least costly *permissible* option available to the agent.

The *excess (moral) value* of an act is the difference between its intrinsic moral value and the intrinsic moral value of its least valuable permissible alternative. For example, the excess moral value of Frank's donating fifty dollars is the difference in moral value between his donating nothing (leaving the unhoused person in his current unhappy state) and his donating fifty dollars (enabling the unhoused person to secure a meal and shelter for the night).

The *opportunity cost* of an act is the difference between its intrinsic cost to the agent and the intrinsic cost of its least costly permissible alternative. The opportunity cost of Frank's donating fifty dollars is the difference in cost to him between his donating fifty dollars and his donating nothing.

[2] In Smith (2023), I argue for a particular account of this kind of value, which I label "erogatory value." Its closest analog may be the "commendatory value" discussed in Little and MacNamara (2017).

2 Defining "even more supererogatory"

Criterion 2 requires an adequate account of supererogation to include a suitably related account of what makes one act more supererogatory than another, for finite, infinite, single-choice and inter-choice cases. In Smith (2023) I develop and argue for the merits of such an account to supplement the Simple Definition of supererogation: the Comparative Formula I account. Inspection of cases indicates that both the excess moral goodness of the act and its opportunity cost to the agent contribute positively to its comparative supererogatoriness. A soldier who loses both legs in saving a comrade's life does something more supererogatory than a soldier who loses only one leg in similarly saving a comrade's life, while a soldier who loses a hand in saving a comrade's leg does something more supererogatory than a soldier who loses a hand in merely saving a comrade's foot.[3] Thus someone who sustains Opportunity Cost X to bring about a given excess moral good does something more supererogatory than someone who sustains a *lesser* opportunity cost to bring about an equivalent excess moral good, while someone who sustains Opportunity Cost X to bring about a substantial excess moral good does something more supererogatory than someone who sustains Opportunity Cost X to bring about a *lesser* excess moral good. An action's comparative supererogatoriness seems to be a positive function of its opportunity cost and excess moral value: as its opportunity cost and excess value rise, so does its comparative supererogatoriness. This conclusion suggests the approach to making comparisons between supererogatory acts articulated in Comparative Formula I.[4]

Comparative Formula I:

A. Supererogatory act X is more supererogatory than supererogatory act Y if and only if, and because, the sum of the opportunity cost plus the excess moral value of X is greater than the sum of the opportunity cost plus the excess moral value of Y.
B. Supererogatory act X and supererogatory act Y are equally supererogatory if and only if, and because, the sum of the opportunity cost plus the excess moral value of X is equal to the sum of the opportunity cost plus the excess moral value of Y [BLINDED Sect. 2.2].

Although Formula I requires adding the opportunity cost and excess value of an act, it is not possible to add two kinds of units which are not comparable. An analogy would be adding a policy's ecological value and economic effects in order to compare it to a rival policy. To surmount this difficulty, in Smith (2023) I propose a version

[3] A rival account of comparative supererogatoriness offered by Hurka and Tsagarakis (2021) rejects the second claim about the soldier cases in this paragraph.

[4] There are many positive functions. Comparative Formula I characterizes the relevant function as an additive one; one referee for this paper suggested that it might be multiplicative. A multiplicative account would directly capture the fact that an action with 0 opportunity cost could not be supererogatory, although this assumption would require independent argument. However, these two accounts, applied as they would be to positive figures for excess value and opportunity cost, would yield the same rank-ordering of any two supererogatory acts. I shall use the simpler additive account.

of Formula I which instead requires adding the z-scores of the act's opportunity cost and excess value to provide an aggregate score for the act that can be compared with the aggregate scores for other acts. The z-score of a data point measures how many standard deviations below or above a population mean the data point is.[5] This is a standard statistical normalization technique used in social science; its appropriateness for making inter-theoretic comparisons, interpersonal comparisons of utility, and possibly aggregating incommensurable goods has recently been argued for by MacAskill and colleagues (MacAskill et al., 2020). However, to simplify presentation, in this paper I will use Formula I, and simply apply it to values representing the opportunity costs and excess moral values of the various acts. No distortion of the comparisons is introduced thereby.

As examination of the cases we've discussed reveals, the Simple Definition and Comparative Formula I appear to work well for both single-choice cases and inter-choice cases in which the agent(s) confront a finite set of options.

3 Problems raised by cases with an infinite number of options

The Simple Definition of supererogation is an example of what Daniel Muñoz calls *Incompatible with the Moral Minimum (IMM)* definitions (Muñoz, 2021a: 2066). For him the "moral minimum" can be understood as "doing the least good that one can permissibly do" (Muñoz 2021a: 2063-4). Although several theorists[6] have proposed *IMM* definitions, Muñoz points out that while *IMM* accounts of supererogation may work well with ordinary cases in which the agent faces a finite number of alternatives, they fail in cases in which the agent faces a set of alternatives forming an infinitely descending chain of moral goodness. His *Tradeoff Dial* case provides an example of such a case.

> A *Tradeoff Dial* … can be set to any real number between 0 and 1—except 0. If you set it to x, you will be subjected to x days of quite serious pain, and you will also relieve Ingrid, an innocent stranger, of $2x$ days of even worse pain. The bigger your sacrifice, the greater Ingrid's benefit. Clearly, since you have a right not to incur pain, you are not obligated to ratchet up the dial. But it sure would be nice of you (Muñoz 2021a: 2065.).

As Muñoz points out, in the *Tradeoff Dial*, you have infinitely many permissible options—setting the dial to 0.1, to 0.01, to 0.0001, etc., and he argues persuasively that each of these options is supererogatory:

> Imagine that you have tentatively chosen to turn the dial to 0.5, and you are considering whether to move to something higher, like 0.6. Either option is permissible, but it is clearly better, in every relevant sense of 'better,' to turn the dial up. We are talking about a big benefit for Ingrid at a smaller cost to

[5] For details, see [BLINDED § 2.2].

[6] For example, McNamara (1996).

> you. We would be right to praise you if you spared her pain as an expression of altruism. You would be doing something you have more moral reason to do. You would also be producing better consequences, impartially considered. Indeed, the more you turn up the dial, the better things get, in all of these ways. It is therefore supererogatory to set the dial to 0.6 instead of 0.5. But the same reasoning shows that *any* setting is supererogatory, since any such setting will be better, in every relevant sense, than any of the infinitely many (permissible) settings below (Muñoz 2021a: 2065-6; 2069).

It is clear that no permissible option available to the agent in *Tradeoff Dial* is incompatible with the permissible option that is the least good the agent can do, since there is no permissible option that is the least good the agent can do. However slight the amount of good an option would do for Ingrid, the agent can always do slightly less good for her by setting the dial lower and still be doing something permissible. Since, according to *IMM* accounts, an act is supererogatory only if it is incompatible with the least good one can permissibly do, none of the options in *Tradeoff Dial* meets that requirement, so none is supererogatory. This is highly counterintuitive, since it seems irrefutable that *every* option in this case is supererogatory.

This argument applies to my Simple Definition of supererogation, since it requires that a supererogatory act have greater moral value than the permissible alternative with the least moral value. In *Tradeoff Dial* there is no permissible alternative with the least moral value, so the Simple Definition implies that no option in this case counts as supererogatory—even though it seems clear that they should all count.[7]

Moreover, *Tradeoff Dial* shows that the Simple Definition's associated Comparative Formula I delivers the wrong answers when used to evaluate whether one supererogatory act is better than another in cases involving infinitely descending chains of moral goodness or cost to the agent. Formula I requires us to compare acts' excess values and opportunity costs. Since *Tradeoff Dial* has no least valuable permissible act, or least costly permissible act, none of its options has either excess value or opportunity cost, and none will be evaluated as better than any others (or even equally good). The Simple Definition, and its associated Formula I for comparing supererogatoriness between acts, cannot handle cases involving infinitely descending chains of moral goodness or cost to the agent.

4 Muñoz' solution

Muñoz offers a solution to the problem of defining supererogatory for infinite cases by proposing a new definition for supererogation:[8]

[7] If the *IMM* definition also requires a supererogatory act to be more costly than the least costly permissible act (as the Simple Definition does), *Tradeoff Dial* shows doubly that the definition does not count any acts in this scenario as supererogatory.

[8] Muñoz (2021a: 2070). He characterizes *Big Picture* as a definition he "tentatively endorses" (2021a: 2070), since he foresees that even trickier cycle cases than those he considers may be devised (2021a: 2073).

Big Picture Better Than a Permissible Alternative (for short, Big Picture):[9]

> An option A is supererogatory just if (i) A is permissible, (ii) A is better than some permissible alternative B, and (iii) for any alternative C, if B is better than C, then A is better than C, and if C is not better than B, then C is not better than A.

Applied to *Tradeoff Dial*, *Big Picture* implies, correctly, that every permissible option available to you as the agent is supererogatory. Moreover, *Big Picture* appears to deliver correct identifications of supererogatory acts in finite-option cases such as the case of Frank's donation to the unhoused person. In this case, the more ten-dollar bills Frank donates, the cost to him will rise but the benefit to the unhoused person will rise even more, making the larger donation a better act. Hence, for example, Frank's donating thirty dollars (A) is permissible, and it is better than his permissibly donating twenty dollars (B). Donating twenty dollars (B) is better than donating only ten dollars, and donating thirty dollars is also better than donating only ten dollars. *Big Picture* correctly classifies any of Frank's donations as supererogatory.

Muñoz' arguments in favor of *Big Picture* as the best definition for supererogation cases involving infinitely descending chains are highly persuasive. It is tempting to hope (as Muñoz implicitly seems to do) that *Big Picture* provides a single account of supererogation that works for both finite and infinite cases.

5 An important tweak to *Tradeoff Dial* as an example

Tradeoff Dial establishes Muñoz' point more convincingly if we tweak it slightly. His presentation focuses on *Tradeoff Dial* as an example in which the agent has an infinitely descending chain of permissible options: setting the dial to 0.6 is worse than setting it to 0.7, setting it to 0.5 is worse than setting it to 0.6, and so forth. An important part of what makes each of these settings worse than a higher setting is that the higher setting relieves Ingrid of more pain. Muñoz doesn't tell us what level of pain Ingrid would be "relieved of" if you do set the dial to some number. However, since he talks about Ingrid being "relieved of" pain, it must be the case that in setting the dial you are benefitting her by *reducing* her pain compared to pain she is already undergoing, or that she would undergo without your intervention. But compared to what is Ingrid being relieved of pain by your setting the dial to, say. 0.5? The natural thought is that your setting the dial relieves Ingrid of the pain she would experience

[9] Muñoz introduces *Big Picture Better Than a Permissible Alternative* to handle both the problem of infinite alternatives and a second problem, that of cases in which there are intransitive "better than" cycles among the agent's options. Because of complex questions involving cycle cases, I shall not discuss them here. Muñoz also considers a definition called *Incompatible with Descending Levels (IDL)*, as advocated by McNamara (2011, 169): An option A is supererogatory just if A is permissible and, for some permissible option B, doing A is incompatible with all permissible options worse than, or as good as, B (Muñoz 2021a: 2071). Arguing that IDL can handle infinite descending levels cases but not cycles, he rejects it in favor of *Big Picture* (Muñoz 2021a: 2071-2).

if you didn't set the dial at all. For example, the case might be one in which if you don't set the dial, Ingrid will experience five days of pain worth −20 for each day.

The best way to insert this assumption into Muñoz' example and still retain the features it needs is to stipulate that the agent has the option of not setting the dial at all, but that not setting the dial would be *wrong*.[10] For the sake of argument let us accept that doing nothing, at no cost to yourself, but at the price to Ingrid of five days of -20 units of pain, is morally prohibited—while also accepting that any slight setting of the dial, however small, which would relieve Ingrid of some pain (as compared to her pain if you don't set the dial), is permissible.

Assume, then, that the agent has the wrongful option of not setting the dial, and that any actual setting is permissible. Then not all the agent's options are permissible. However, if the dial can be set at any amount greater than zero, it is still true that your permissible options involve an infinitely descending chain of worse and worse—but still permissible—options. Since each of these options is better than some other permissible alternative, according to *Big Picture* each of the dial-settings is supererogatory. Let us call this revised case, incorporating a prohibited option of not setting the dial, *Tradeoff Dial II*.

Big Picture correctly classifies each dial setting in *Tradeoff Dial II* as supererogatory: any dial setting A has some permissible alternative B at a lower setting that is worse than A, insofar as B's cost to the agent is smaller but its benefit to Ingrid is even smaller still.

6 Flaws in Muñoz' solution

While *Big Picture* has its virtues in dealing with infinite cases, examination shows that it cannot always correctly assess which of two acts in inter-choice cases is more supererogatory, either in finite-option or in infinite-option cases. Recall that our criteria for an acceptable account of supererogation included the following:

2. The account must include a suitably related account of what makes one act more supererogatory than another for finite, infinite, single-choice and inter-choice cases.

Such an account should be able to tell us, for example, that a person who risks serious burns to her hands to save a child in a burning house does something less supererogatory than the act of another person who risks her life to save the child from a burning house. Although Muñoz does not provide any such comparative account as part of *Big Picture*, he clearly holds that some supererogatory acts are better than others. As we've seen, when introducing *Tradeoff Dial*, he says "Imagine that you have tentatively chosen to turn the dial to 0.5, and you are considering whether to move to something higher, like 0.6….it is clearly better, in every relevant sense of

[10] Not setting the dial could not be either permissible or supererogatory, since it would then provide a moral minimum (not relieving Ingrid of any pain), since no other option is worse than this. But a moral minimum option can't be allowed for Muñoz' purposes.

'better', to turn the dial up". *Big Picture* itself refers to an act A as better than some permissible alternative B. Any account of comparative supererogatoriness associated with *Big Picture* must be constructed from the factors that *Big Picture* identifies as making an act supererogatory at all, just as any account of what it is for one person to be larger than another must be constructed from a comparison between the height and weight of the two people, since largeness itself is a function of a person's height and weight.[11]

Proposing an account of what makes one act more supererogatory than another that is appropriate to Muñoz 's *Big Picture* takes speculation, since Muñoz himself offers no such account. However, it seems clear from his discussion of "betterness" that the key factors are the beneficiary's benefit and the cost to the agent, even though he does not explicitly mention cost as a component of supererogation.[12] ("We are talking about a big benefit for Ingrid at a smaller cost to you" (Muñoz, 2021a: 2065)). Thus both cost to the agent and value to the beneficiary are available as components for a *Big Picturesque* account of comparative supererogatoriness.

This might tempt us to think that the Formula I comparative account of supererogatoriness, appropriate to the Simple Definition of supererogatoriness, is also appropriate for *Big Picture*.

Comparative Formula I:

A. Supererogatory act X is more supererogatory than supererogatory act Y if and only if, and because, the sum of the opportunity cost plus the excess moral value of X is greater than the sum of the opportunity cost plus the excess moral value of Y.
B. Supererogatory act X and supererogatory act Y are equally supererogatory if and only if, and because, the sum of the opportunity cost plus the excess moral value of X is equal to the sum of the opportunity cost plus the excess moral value of Y (Smith, 2023 Section 2.2).

However, as we've seen, the approach embodied in Formula I for comparing which of two acts is more supererogatory cannot work in cases such as *Tradeoff Dial II*. In

[11] Wilson et al. (2017: 60). There's an unavoidable cart-before-the-horse problem with *Big Picture* in this connection, since it assumes an account of what makes an option "better" (which here must mean "more supererogatory") in defining what makes an act supererogatory at all.

[12] In explaining "betterness" Muñoz also cites the fact that we would praise you more if you set the dial to 0.6, that you would be doing something you had more moral reason to do, and that you would be producing better consequences, impartially considered. But these are not variables independent of cost and value. We would only praise you more if setting the dial at 0.6 is morally better than setting at 0.5; you would only have more moral reason to set it at 0.6 if doing so were better (for example, either because of its benefit or both its benefit and its cost); and the impartial consequences include both the cost to the agent and the benefit to Ingrid. In other work on supererogation Muñoz (2021b: 702) explicitly introduces an agent's prerogative (which in his examples involves cost to the agent) as an important feature of supererogatory actions. Any account of "betterness" that included only the moral value of acts being compared would have the unwonted implication that an agent who merely pricks her finger with a pin to extend a stranger's life by twenty years does something more supererogatory than someone who sacrifices both her arms to extend a stranger's life by nineteen years.

Tradeoff Dial II, the agent faces a set of permissible alternatives whose moral values and costs to the agent both form infinitely descending chains. These acts can't be accurately compared using Formula I, since this formula requires that the agent's set of alternatives includes a permissible alternative with the least excess value and a permissible alternative with the least opportunity cost.

However, a cousin of Formula I, which we can call Formula II, appears to work well for comparing alternatives in infinite cases like *Tradeoff Dial II*.

Comparative Formula II:

A. Supererogatory act X is more supererogatory than supererogatory act Y if and only if, and because, the sum of X's intrinsic cost plus its intrinsic moral value is greater than the sum of the Y's intrinsic cost plus its intrinsic moral value.
B. Supererogatory act X and supererogatory act Y are equally supererogatory if and only if, and because, the sum of X's intrinsic cost and its intrinsic moral value is equal to the sum of Y's intrinsic cost plus its intrinsic moral value.[13]

The difference between these two formulas is that Formula I utilizes differences in the sums of the acts' *excess* moral values and *opportunity* costs, whereas Formula II utilizes differences in sums of the acts' *intrinsic* moral values and their *intrinsic* costs. Formula II must stick to intrinsic value and cost because there is no least valuable or least costly permissible act in an infinite series of the kind exemplified in *Tradeoff Dial II*, making it impossible to identify an excess value or opportunity cost for any supererogatory act.

When used to compare the supererogatory acts in single-choice *Tradeoff Dial II*, Formula II provides the correct evaluations. To calculate the comparative supererogatoriness of your acts according to Comparative Formula II, we add, for each act, its intrinsic cost (to you) plus its intrinsic value (given Ingrid's pain). For example, suppose setting the dial at 1.0 costs you a single day of pain (each day being worth −10), while it relieves Ingrid of two days of greater pain from the five she would otherwise experience. So she experiences three painful days (each day being worth −20). Then the sum of intrinsic cost plus intrinsic value for setting the dial at 1.0 equals −70 [(-10+3 (-20)]. Setting the dial at 0.5 costs you half a day of pain [0.5 (-10) = -5], while doing so relieves Ingrid of a single day of pain, leaving her to experience four days' pain worth −80 ([4 (-20) = -80]. This act's sum of intrinsic cost plus intrinsic value equals −85. Since −85 is worse than −70, Formula II correctly evaluates setting the dial at 1.0 as better—more supererogatory—than setting it at 0.5. As the dial setting goes up, the agent endures more pain in order to relieve Ingrid of even greater pain. Because Ingrid's pain drops faster than the agent's pain rises, as each higher setting is utilized, moving the dial to the higher setting is more supererogatory.

[13] The official version of Formula II would incorporate suitable references to the z-scores of the acts' costs and moral values.

Despite Comparative Formula II's apparent success in infinite cases such as *Trade-off Dial II*, it does not always deliver the correct comparison in finite cases involving the comparison of supererogatory acts that figure in *different* choice situations and so are not alternatives to each other. Case I, shown in Table 1 (in which OC = opportunity cost and EV = excess moral value) provides this kind of counterexample. In this case, if the first agent (faced with Choice i) performs Act A, she will experience −50 and produce a moral value of +100. If instead she performs Act B, she will experience a positive 25 and produce a moral value of 0. If the second agent (faced with Choice ii) performs Act C, she will experience −50 and produce a moral value of +100. If instead she performs Act D, she will experience 0 and produce a moral value of 0.

All four acts are morally permissible. Both Acts A and C, but not Acts B and D, qualify as supererogatory according to *Big Picture*.

The intrinsic costs (-50) and moral values (+100) of Act A and Act C are the same. Thus, if we use Formula II to compare Acts A and C, we must conclude that these two acts are equally supererogatory, since each of them has a composite value of +50. But intuitively, Act A is more supererogatory than Act C. The two acts result in the same experience for the agent (-50), but Act A degrades the agent's outcome, as compared to what she would experience if she performed Act B, by 75 points (its OC or opportunity cost). By contrast, Act C only degrades the agent's outcome, as compared to what she would experience if she performed Act D, by 50 points (its OC). In this crucial sense Act A costs its agent more than Act C costs its agent.

To support our intuition that the agent in Act A performs a more supererogatory act, we can't merely compare the *intrinsic* costs of Acts A and C, but rather must compare the *opportunity costs* of these two acts—their costs as compared with their least costly permissible alternatives. The first agent suffers a greater opportunity cost to achieve the same good. Using intrinsic cost and value, as Formula II requires, rather than using opportunity cost and excess value, renders our comparison insensitive to the true cost of the supererogatory act in question. Thus using Formula II to compare supererogatory acts in inter-choice finite choice situations can yield unacceptable assessments.[14] Using Comparative Formula I, however, yields the correct comparison between Acts A and C: it implies that Act A, having the higher OC + EV, is more supererogatory than Act C.

Not only does Formula II fail to deliver the correct comparative evaluations in some finite-option inter-choice cases, but it also fails to deliver the correct compara-

Table 1 Case I

Act	Act	Cost to agent	OC	Moral value	EV	Cost to agent + Moral value	OC + EV
Choice i	A	-50	75	+100	100	+50	175
	B	+25	0	0	0		
Choice ii	C	-50	50	+100	100	+50	150
	D	0	0	0	0		

[14] A parallel case could be devised to show equally counterintuitive results when the two agents' acts would result in similar opportunity costs, but one act would produce a higher excess value than the other, even though their intrinsic values are the same.

Table 2 Tradeoff Dial II

Dial setting	Deontic status	Agent days of pain	Agent cost	Ingrid days of pain	Ingrid pain	Agent cost + Ingrid pain
1.0	Permitted	1.0	1 x -10 = -10	5 − 2 = 3 days	3 x -20 = -60	-70
0.9	Permitted	0.9	0.9 x -10 = -9	5–1.8 = 3.2 days	3.2 x -20 = -64	-73
0.6	Permitted	0.6	0.6 x -10 = -6	5–1.2 = 3.8 days	3.8 x -20 = -76	-82
0.5	Permitted	0.5	0.5 x -10 = -5	5 − 1 = 4 days	4 x -20 = -80	-85
Act X						
0.4	Permitted	0.4	0.4 x -10 = -4	5 − 0.8 = 4.2 days	4.2 x -20 = -84	-88
0.3	Permitted	0.3	0.3 x -10 = -3	5 − 0.6 = 4.4 days	4.4 x -20 = -88	-91
0.2	Permitted	0.2	0.2 x -10 = -2	5 − 0.4 = 4.6 days	4.6 x -20 = -92	-94
0.1	Permitted	0.1	0.1 x -10 = -1	5 − 0.2 = 4.8 days	4.8 x -20 = -96	-96
0	Not available					
None	*Wrong*	0	0	5 days	5 x -20 = -100	-100

Table 3 Tradeoff Dial III

Dial setting	Deontic status	Agent days of pain	Agent cost	Ingrid days of pain	Ingrid pain	Agent cost + Ingrid pain
1.0	Permitted	1.0	1 x -10 = -10	5 − 2 = 3 days	3 x -20 = -60	-70
0.9	Permitted	0.9	0.9 x -10 = -9	5–1.8 = 3.2 days	3.2 x -20 = -64	-73
0.6	Permitted	0.6	0.6 x -10 = -6	5–1.2 = 3.8 days	3.8 x − 20 = -76	-82
0.5	Permitted	0.5	0.5 x -10 = -5	5 − 1 = 4 days	4 x -20 = -80	-85
Act Y						
0.4	Permitted	0.4	0.4 x -10 = -4	5 − 0.8 = 4.2 days	4.2 x -20 = -84	-88
0.3	Not available					
None	*Wrong*	0	0	5 days	5 x -20 = -100	-100

tive evaluations in some *infinite*-option cases.[15] To see the problem, we need to present *Tradeoff Dial II* in a little more detail.

Again assume that Ingrid would experience − 20 units of pain each day for five days if you do nothing, for a total of -100. Assume further that you would experience no pain if you do nothing, but if you do set the dial you would experience pain worth − 10 units on each full day you are in pain. Ingrid's days of pain are reduced by twice the days of pain that you would experience at a given setting.

Table 2 displays some of your infinite set of alternatives and their upshots

Now consider *Tradeoff Dial III*, a case just like *Tradeoff Dial II* except that the available dial settings descend from 1.0 towards 0.3, which itself is unavailable, just as setting it at 0 is unavailable in *Tradeoff Dial II.*

Compare setting the dial to 0.5 in *Tradeoff Dial II (*call this Act X) with setting it to 0.5 in *Tradeoff Dial III* (Act Y). The sum of X's intrinsic cost plus its intrinsic moral value is -85, while the similar sum for Y is also − 85. According to Comparative Prin-

[15] Thanks to an anonymous referee for pressing me on whether such counterexamples exist.

ciple II, then, Act X and Act Y are equally supererogatory. However, reflecting on the two cases, it seems clear that Act X (in *Dial II*) is *more* supererogatory than Act Y (in *Dial III*). *Dial III*, like *Dial II*, is an example of the kind of cases about which Muñoz is concerned: ones in which the agent faces a set of alternatives forming an infinitely descending chain of moral goodness (and agent cost) in which there is no worst permissible option. Nonetheless each of Act X (setting the dial to 0.5) in *Dial II* and Act Y (setting the dial to 0.5) in *Dial III* occur in a series of settings that has what we can call a "lower practical limit," namely an infinitely approachable but practically unavailable setting.[16] For Act X this setting is 0, while for Act Y this setting is 0.3. Crucially, Act X in *Dial II* is *further above* its lower practical limit (0) than Act Y is above its lower practical limit (0.3). Similarly, in the real number line, 4 is arithmetically further than 3 beyond 2, even though between each of these pairs of numbers there are infinitely many more. In setting their dials to 0.5 the agent in *Dial II* does better relative to her lower practical limit than the agent in *Dial III* does relative to her lower practical limit. This means that the agent in *Dial II* has the option, which she rejects, of doing worse than the agent in *Dial III.* For example, the *Dial II* agent could permissibly set her dial at 0.1, whereas setting the dial that low is not available to the agent in *Dial III*. In setting the dial at 0.5, the *Dial II* agent rises higher above an inferior option (0.1) than the *Dial III* agent does when she sets the dial at 0.5, since the *Dial III* agent has no equally inferior option. The contrast between the two agents' respective accomplishments is analogous to the contrast between the accomplishments of the Case I agents in Choice i and Choice ii. Thus to do adequate justice to the relative supererogatoriness of setting the dial to 0.5 in *Tradeoff Dial II* and in *Tradeoff Dial III*, we need to assess these two acts in the context of their respective lower practical limits.

This might suggest that the *Big Picture* account of the supererogatory needs to be supplemented, not with the Principle II account of comparative supererogatoriness, but rather with an alternative account of comparative supererogatoriness that depends on comparing how much better each of two acts is than its lower practical limit. Unfortunately, such an account of comparative supererogation has no foundation in *Big Picture*, since *Big Picture* does not utilize the concept of a lower practical limit in defining supererogation. Indeed, its whole rationale is to reject this kind of comparison to a unique lower index in order to determine whether or not an act is supererogatory.

Lacking supplementation by a satisfactory principle of comparative supererogation, *Big Picture* must be rejected as an inadequate account of supererogation in finite and infinite cases. However, perhaps a variant of the (IMM) Simple Definition of supererogation, as well as the concepts of excess value and opportunity cost, could be stated using the concepts of lower and upper practical limits. An appropriate account of comparative supererogation, using the same new concepts, could be adjoined. This variant should satisfactorily handle both finite and infinite cases of supereroga-

[16] If setting the dial to 1.0 were unavailable, as the 0 and 0.3 settings are unavailable, then 1.0 would be the upper practical limit in each case. Patrick Wu (personal communication) has pointed out that my notions of the "lower practical limit" and the "upper practical limit" are recognized in mathematics as the "infimum" (i.e., greatest lower bound) and the "supremum" (i.e., least upper bound) of a set of possibly infinite numbers.

tion. Possibly the account could incorporate *Big Picture's* Clause (iii), which enables *Big Picture* to deal with cycles. We would then have a hybrid account that handles cycles as well as finite and infinite cases. This proposal is well worth pursuing.

7 Conclusion

I started with the claim that any acceptable account of supererogation must satisfy two criteria:

1. The account must enable us to accurately identify supererogatory acts in both finite and infinite option cases, and explain what features make these acts supererogatory.
2. The account must include a suitably related account of what makes one act more supererogatory than another for finite, infinite, single-choice and inter-choice cases.

As we've seen, the Simple Definition of supererogation meets both these criteria for finite option cases, but fails both of them for infinite option cases. Although limited, it can be accepted as an adequate account for finite cases. On the other hand, *Big Picture*, while it meets Criterion 1 for both finite and infinite option cases, nonetheless fails Criterion 2 in both finite and infinite option inter-choice cases. The Simple Definition is successful within the sphere of finite cases, but *Big Picture* is unsuccessful within the spheres of both finite and infinite cases. However, neither provides a complete picture of supererogation and comparative supererogation for the full range of possible cases. We may need to concede that different accounts are needed for different kinds of cases.[17] We've learned along the way that no account of supererogation, however otherwise attractive, is adequate unless it provides the foundation for an account of comparative supererogation that is equally satisfactory for both single-choice and inter-choice cases.

Acknowledgements I am grateful to helpful comments from James Goodrich, D. Black, Patrick Wu, audiences at the University of California, Davis and the April 2023 Pacific Division meeting, and especially the anonymous referees of this journal, who provided unusually thoughtful comments on earlier versions of this paper.

[17] If so, we then face the problem of how to compare two supererogatory acts when one is a member of a finite number of alternatives and the other is a member of a chain of infinitely descending alternatives.

References

Hurka, T., & Evangeline, T. (2021). More supererogatory. In J. McMahan, T. Campbell, J. Goodrich, & K. Ramakrishnan (Eds.), *Principles and persons: The legacy of derek parfit* (pp. 463–77). Oxford University Press.

Little, M., & MacNamara, C. (2017). For better or worse: Commendatory reasons and latitude. In M. Timmons (Ed.), *Oxford studies in normative ethics* (7 vol., pp. 138–160). Oxford University Press.

MacAskill, W., Cotton-Barratt, O., & Ord, T. (2020). Statistical normalization methods in interpersonal and intertheoretic comparisons. *Journal of Philosophy*, *117*, 61–95.

McNamara, P. (1996). Making room for going beyond the call. *Mind*, *105*, 415–450.

McNamara, P. (2011). Praise, Blame, Obligation, and DWE: Toward a framework for classical supererogation and kin. *Journal of Applied Logic*, *9*, 153–170.

Muñoz, D. (2021a). Infinite options, intransitive value, and supererogation. *Philosophical Studies*, *178*, 2063–2075.

Muñoz, D. (2021b). Three paradoxes of supererogation. *NOÛS*, *55*, 699–716.

Smith, H. M. (2023). Even more supererogatory. *Australasian Journal of Philosophy.* https://doi.org/10.1080/00048402.2022.2153369

U. S. Department of Defense (Undated). Military Awards for Valor - Top 3. https://valor.defense.gov/

Wallace, J. (Undated a). Distinguished Service Cross. Veteran.com. https://veteran.com/distinguished-service-cross/

Wallace, J. (Undated b). Medal of Honor: Benefits, History and Facts. Veteran.com. https://veteran.com/medal-of-honor-benefits/

Wilson, J. P., Hugenberg, K., & Rule, N. O. (2017). Racial bias in judgments of physical size and formidability; From size to threat. *Journal of Personality and Social Psychology*, *113*(1), 59–80.