

# Ethics: An International Journal of Social, Political, and Legal Philosophy

## How Should We Aggregate Competing Claims?

--Manuscript Draft--

<b>Manuscript Number:</b>	12546R3
<b>Full Title:</b>	How Should We Aggregate Competing Claims?
<b>Article Type:</b>	Original Article
<b>Corresponding Author:</b>	Alex Voorhoeve  UNITED KINGDOM
<b>Corresponding Author Secondary Information:</b>	
<b>Corresponding Author's Institution:</b>	
<b>Corresponding Author's Secondary Institution:</b>	
<b>First Author:</b>	Alex Voorhoeve
<b>First Author Secondary Information:</b>	
<b>Order of Authors:</b>	Alex Voorhoeve
<b>Order of Authors Secondary Information:</b>	
<b>Additional Information:</b>	
<b>Question</b>	<b>Response</b>
Word count of your submission (including notes and references)	10694
<b>Abstract:</b>	<p>How Should We Aggregate Competing Claims?</p> <p>Many believe that we ought to save a large number from being permanently bedridden rather than save one from death. Many also believe that we ought to save one from death rather than a multitude from a very minor harm, no matter how large this multitude. I argue that a principle I call "Aggregate Relevant Claims" satisfactorily explains these judgments. I offer a rationale for this principle and defend it against objections.</p>

## How Should We Aggregate Competing Claims?\*

*Alex Voorhoeve*

---

Many believe that we ought to save a large number from being permanently bedridden rather than save one from death. Many also believe that we ought to save one from death rather than a multitude from a very minor harm, no matter how large this multitude. I argue that a principle I call “Aggregate Relevant Claims” satisfactorily explains these judgments. I offer a rationale for this principle and defend it against objections.

---

**I.** Suppose that you are a morally-motivated stranger who, using your own resources, can help others at negligible cost to yourself. Consider the following cases:

**Case 1.** You can either fully cure one young man’s terminal illness (thereby restoring him to good health for a normal lifespan) or fully cure a number of other young people of an illness which will not shorten their lives, but which will leave them completely disabled (they will be mentally unimpaired, but permanently bedridden, which for them will entail a life just somewhat better than an early death).<sup>1</sup>

Is there a number of people you can save from being completely disabled for which you ought to save them from this disability instead of the one from death?

**Case 2.** You can either cure one young man's terminal illness or cure some number of other young people of an illness which will cause a very minor harm (they will be bedridden for a day). Those facing this minor harm otherwise have good prospects.

Is there a number of people you can spare this very minor harm for which you ought to spare them this harm instead of saving the one from death?

Many philosophers answer these questions with “yes” and “no,” respectively.<sup>2</sup> That is, they believe in both a limited form of aggregation (viz., that you ought to save a large number of people from complete disability rather than one from premature death) and in a limited form of non-aggregation (viz., that there is a harm small enough such that *no* number of such very minor harms to people who will in any case have good lives can outweigh curing one young person's terminal illness). Survey results suggest that a substantial share of subjects agrees.<sup>3</sup>

If these moral judgments are both correct, then this spells trouble for familiar views in distributive ethics. On standard forms of utilitarianism and prioritarianism, curing one young person from a terminal illness has a fixed, positive, non-infinite moral value. So does sparing one well-off person the harm of being bedridden for a day—albeit that the latter value is much lower. Moreover, sparing  $N$  people from the latter harm generates  $N$  times as much value as sparing one person this harm, so that, on these theories, for some  $N$ , the moral value of saving  $N$  well-off people from being bedridden for a day will exceed the moral value of saving one young person's life. These theories therefore cannot accommodate the judgment that there is no number of very minor harms to those who will in any case have good lives that one ought to prevent instead of saving one person's life.<sup>4</sup> Moreover, as Marc Fleurbaey, Bertil Tungodden, and Peter Vallentyne

have shown, standard forms of egalitarianism cannot accommodate these judgments in combination.<sup>5</sup>

I am interested in the question whether there is an appealing moral view that explains the aforementioned pair of judgments. In this paper, I shall examine the following candidate, which I shall call Aggregate Relevant Claims (or ARC, for short).

*ARC:* 1. Each individual whose well-being is at stake has a claim on you to be helped.

(An individual for whom nothing is at stake does not have a claim.)

2. Individuals' claims *compete* just in case they cannot be jointly satisfied.

3. An individual's claim is *stronger*:

a. the more her well-being would be increased by being aided; and

b. the lower the level of well-being from which this increase would take place.

4. A claim is *relevant* if and only if it is sufficiently strong relative to the strongest competing claim.

5. You should choose an alternative that satisfies the greatest sum of strength-weighted, relevant claims.

ARC takes this relatively simple form in the limited class of cases I shall consider. In these cases, you must decide whom to benefit by improving their well-being over what would otherwise happen due to natural causes. (No alternative involves making anyone worse off than they would have been but for your intervention.) To make it plausible that the well-being at stake is of concern, I focus on cases in which you can improve people's health-related well-being. I

also assume that no one has claims based on desert; that no one is responsible for her level of well-being; and that no one has prior entitlements. In addition, my cases involve neither risk nor changes in the population. Furthermore, the alternatives under consideration either give an individual the best or the worst feasible outcome for him. (I therefore shall not consider cases in which some alternatives are of intermediate value for an individual.) Finally, claims to an alternative are either all relevant or all irrelevant.

In Case 1, ARC yields the following conclusions. The young man who can be cured of a terminal illness has a very strong claim. This is because so much is at stake for him and because he would be very badly off if he were not aided. A person who can be cured of the complete disability also has a strong claim, because there is a great deal at stake for him and if he were not aided, he would have a life only just worth living. Since both competing claims are so strong, they are relevant to each other. On ARC, you therefore ought to save a large number of people from becoming completely disabled rather than save one from death.

In Case 2, the claim of the young man is very strong for the aforementioned reasons. By contrast, a person who can be cured of the very minor ailment has only a very weak claim, because so little is at stake for her and she would be well-off even if unaided. The latter's claim is therefore irrelevant to the former. No number of claims to be rid of the very minor harm can therefore jointly outcompete the young man's claim to be saved from death.

In sum, unlike standard theories of distributive ethics, ARC accommodates both an attractive form of aggregation and a plausible limit on aggregation.<sup>6</sup> Some have therefore endorsed it or something akin to it.<sup>7</sup> However, we require of a moral principle not merely that it accommodates firmly held case judgments, but also that it offers a plausible rationale for these

judgments. In what follows, I shall offer such a rationale. I shall then appeal to this rationale to defend ARC against several objections.

**II.** The following is a thumbnail sketch of the proposed rationale for ARC.<sup>8</sup> There are two conflicting ways of responding to the equal objective moral value of each individual's well-being. One of these is aggregative and the other non-aggregative. Both of these ways rightfully have a hold on us. ARC partially accommodates and arbitrates between these ways of responding to people's claims in a plausible way.

Let me elaborate. On the aggregative approach, the satisfaction of an additional person's claim of a given strength is always just as important, no matter how many other people's claims of that strength are at issue. It is, for example, just as important to save 1001 over 1000 people from premature death as it is to save 2 over 1 from this fate. On this approach, then, satisfying  $N$  claims of a given strength is  $N$  times as important as satisfying one claim of a given strength. This naturally yields the judgment that what is most important is that one satisfies the greatest possible sum of strength-weighted claims. This approach is attractive because it asserts the unvarying and equal marginal importance of every additional person's claim of a given strength.

By contrast, the non-aggregative approach involves imaginatively placing oneself, one person at a time, in the position of each person who has a claim and viewing the situation through her eyes. On this approach, for each individual taken separately, one takes in what she would have to give up if another person's competing claim were satisfied. After one has performed this imaginative exercise for all individuals, one at a time, one does not aggregate their claims. Instead, one takes an objective perspective on the importance of each claim taken separately. When one does so, it will appear most important to satisfy the strongest claim.<sup>9</sup>

Moreover, if one instead were to satisfy a different claim, then the larger the gap in strength between the strongest claim and this satisfied claim, the more one would depart from what is most important, on this approach.

This approach appears to be rooted in one element of our moral psychology, namely sympathetic identification. Adam Smith famously argued that the process of placing oneself in another's shoes is the central moral impulse.<sup>10</sup> I am making the more modest proposal that it lies at the heart of a non-aggregative approach to distributive ethics. It seems to me a crucial fact that this form of sympathy is essentially non-aggregative. I can sympathetically imagine myself in one person's position, but I cannot imagine myself in the position of a collection of individuals, except by imagining myself in the position of a single (perhaps representative) member of the collective, or by imagining myself in each of their positions *one at a time*. Psychological research suggests that, as a consequence, moral reasoning that relies solely on this process of sympathetic identification will not account for numbers.<sup>11</sup> Of course, the mere fact that this approach draws on a psychologically real form of moral concern does not show that the approach is right. Its justification is that this form of concern for each person taken alone is a natural expression of our appreciation of the separateness of persons. The supreme importance of satisfying the largest claim can therefore be seen as an expression of respect for this separateness.<sup>12</sup>

ARC embodies one way of partially accommodating and arbitrating between these aggregative and non-aggregative approaches. It tells you to do what is best on the aggregative approach under the constraint that this does not depart too far from what is most important on the non-aggregative approach. It thus takes the form of maximisation under a constraint: it maximises the sum of strength-weighted claims that are satisfied under the constraint that these

claims are relevant. The demand for maximisation stems from the aggregative approach. The constraint stems from the non-aggregative approach.

**III.** The foregoing is a first step towards a rationale for ARC. We should now ask: Why should one accommodate the non-aggregative approach by adopting the “relevance” constraint? This question can also be put as follows. Imagine a situation in which unrestricted aggregation would favour saving a number of people with weaker claims rather than one person with a competing stronger claim. In this situation, the “gap” between, on one side, the strongest claim and, on the other side, the strength of each of the weaker claims that would be satisfied if we followed the aggregative approach is a measure of the departure from what is most important on the non-aggregative approach. Why should there be a point at which this measure crosses a morally crucial boundary, giving rise to a “veto” of the non-aggregative approach over the aggregative approach? And what determines this point? As Gabriel Wollner puts it, before the view can be accepted, we must have an “independently plausible and principled reason why [and where] the veto should apply.”<sup>13</sup> I shall now propose an answer to these questions that draws on a fuller characterisation of what is involved in taking up the perspective of each person within the terms set by the non-aggregative approach.

From the first-person perspective, we typically are much more concerned for our own well-being than for the well-being of strangers. As Frances Kamm writes of a case in which two individuals (who are strangers to each other) have competing claims to be saved from imminent death: “*for each, from his subjective view, his not surviving will be almost as if [neither] survives.*”<sup>14</sup> Nonetheless, for most of us, even from our personal perspective, the well-being of strangers is of concern; the value of their well-being is visible and can motivate us from the



personal point of view, albeit that it does not have the importance to us that it has from an objective point of view. From our personal perspective, the death of a stranger may matter little, but it matters some.

There are, of course, individuals who care little about their own well-being. There are also those whose dedication to the well-being of strangers is so complete that, even from their personal point of view, their own well-being has no special significance. Furthermore, there are people who care nothing for the well-being of strangers. But the pattern of relative concern outlined applies to most well-socialised and psychologically normal individuals.

Up to a point, common-sense morality regards as permissible such greater self-concern. In circumstances in which no other moral considerations besides one's own well-being and the well-being of a stranger are at issue (e.g., one has no special obligations and would not violate anyone's rights by acting on one's greater concern for oneself), it also regards it as morally permissible to act on it. This is clear in cases of the following kind. You and a stranger each face a distinct threat of harm. These threats are not due to human agency (they are, say, posed by naturally occurring diseases). The stranger is powerless to avert either threat, but using only your own resources, you can either save yourself from harm or save the stranger from the harm he faces (you cannot save both yourself and him; moreover, no one else can save you or the stranger). In such cases, common-sense morality judges it permissible to save yourself rather than a stranger from an equally large harm. It also regards it as permissible to save yourself from a lesser harm rather than a stranger from a somewhat greater harm. For example, if you can either save yourself from the aforementioned complete disablement or save a stranger from death, then you are permitted to prevent your disablement. But there are limits to the extra concern for your dear self that everyday morality permits. If you face a very minor harm (such as

an illness that will leave you bedridden for a day) and can either prevent this harm to yourself or prevent the death of a stranger, then it holds that you must save the stranger.

Of course, it is a matter of dispute whether common-sense morality's judgments in such cases are defensible or whether, instead, morality always requires perfectly equal, impartial concern.<sup>15</sup> In this paper, I shall simply assume that, up to a limit, one is indeed morally permitted to be more concerned for oneself than for a stranger and to act on this pattern of concern when no other moral considerations (such as rights or special ties) stand in the way.

I shall refer to the point of view that results from taking on a person's maximally permissible degree of self-concern and minimally required other-concern as "the permissible personal perspective."<sup>16</sup> We can draw on this notion to add to our characterisation of the non-aggregative, sympathetic approach. On this augmented sympathetic approach, when one imaginatively places oneself in the shoes of a person whose well-being is at stake, one takes up his permissible personal perspective, from which his claim looms larger than it does from a purely objective perspective. According to the scale of values that one adopts when one takes on his permissible personal perspective, in a one-to-one comparison of this person's claim vis-à-vis the competing claim of a stranger, the satisfaction of the former takes on extra significance. The satisfaction of his claim will therefore take priority over the satisfaction of an equally strong competing claim of a stranger. The same will be true in some cases in which his claim is weaker than a stranger's claim. Nonetheless, it is also true that if he has a very weak claim and a stranger has a very strong competing claim, then from his permissible personal perspective, he will give priority to the stranger's claim.

We can draw on this idea to formulate the following criterion for when a claim is relevant. Person's A's claim is relevant to B's competing claim just in case, in a one-to-one

comparison of their claims from A's permissible personal perspective, A's claim is at least as important as B's. Or, to put this test in terms of permissible action: A's claim is relevant to B's competing claim just in case, in a situation in which A must choose whether to satisfy his claim or B's claim and in which no morally relevant factors besides the minimally required concern for B stand in the way of A acting on his self-interest, he would be permitted to satisfy his own claim.

By way of illustration, consider again Case 1, in which there is one young man with a claim to be saved from premature death and there are a number of people with competing claims to be saved from complete disablement. The proposed test for whether the latter's claims are relevant is this. Imagine that you could either save yourself from complete disablement or save a stranger from death, and that in making this decision, you need consider only your own well-being and that of the stranger. Would it be permissible to save yourself? Plausibly, the answer is "yes." If so, then the claim to be saved from complete disablement is relevant to the claim to be saved from a terminal illness.

Next consider Case 2, in which there is one young man with a claim to be saved from death and there are a number of people with competing claims to be spared being bedridden for a day. Imagine that you could either spare yourself from the latter very minor harm or save a stranger from death, and that in making this decision, you need consider only your own well-being and that of the stranger. Would it be permissible to spare yourself? Plausibly, the answer is "no." If this answer is correct, then the claim to be saved from this very minor harm is not relevant to the young man's claim.

Why should this test determine the (ir)relevance of a claim? The answer has several parts. First, this test embodies the key elements of the non-aggregative, sympathetic approach. The

non-aggregative element is captured by considering claims one-to-one. The sympathetic element is captured by taking on each person's maximally permissible degree of self-concern (and minimally required degree of other-concern). Of course, as noted above, this pattern of concern may differ from the person's actual degree of self-concern. But this idealisation is warranted. Suppose, first, that A- assigns greater importance to his well-being compared to the well-being of a stranger than is morally permissible—say, he would save himself from a very minor harm rather than save a stranger from an early death. It would be improper for this outsized self-concern to lead to A-'s claim being judged relevant.<sup>17</sup> Suppose, next, that A+ has more than the minimally required degree of altruism. This fact alone should not lead us to judge a claim of his irrelevant when he could permissibly give his claim priority over a stranger's competing claim. Imagine that both A and A+ each have claims of equal strength (they both stand to lose the same amount of well-being from the same level) and that each of their claims competes with B's stronger claim. Suppose further that A has precisely the maximally permissible degree of self-concern and the minimally required degree of other-concern. Finally, imagine that from A's personal perspective, his claim has priority over B's in a one-to-one comparison, so that his claim is relevant, but that from A+'s more altruistic personal perspective, B's claim has priority. It would be implausible to maintain that A's claim is relevant and A+'s claim is irrelevant. The claims of people who are more altruistic than morality requires do not thereby alone diminish in moral significance.<sup>18</sup>

The second reason for employing this criterion for (ir)relevance is that it tracks an important result of the process of imaginatively taking up each person's permissible personal perspective within the terms set by the non-aggregative approach. When one does so in Case 2, *a form of unanimity emerges*. From the permissible personal perspective of someone with the very

weak claim, the satisfaction of the competing very strong claim will take priority over the satisfaction of her own claim (when these claims are compared one-to-one). The same will be true from the perspective of every other person with the very weak claim. Naturally, the person with the very strong claim will agree. From *his* personal perspective, *his* claim to be saved from death takes priority over a stranger's claim to be saved from a very minor harm. The process of imaginative sympathy—the process of judging the situation from each person's permissible subjective perspective—will therefore yield *agreement* between all concerned. From every person's perspective, in a one-to-one comparison of the two competing claims, the strongest claim ought to take priority.

In Case 1, by contrast, there is no such unanimity. The process of imaginative sympathy yields *conflict* between, on one side, every person with a claim to be saved from the complete disability, and, on the other side, the single person with a claim to be saved from death. Now, it remains true that when the observer has finished his “round” of identifying with the perspective of each, the non-aggregative approach will judge it most important to help the one facing death. Nonetheless, the difference between Case 1 and Case 2 marks a crucial change in what emerges from the non-aggregative approach. In Case 1, the process leads one to sympathize with each person's desire to press his claim. Moreover, after going through this process, one can see that there is no solution that does full justice to the personal perspective of each. By contrast, in Case 2, the process yields the conclusion that there is no claim that an individual will, from his permissible personal perspective, place into competition with the young man's claim to be saved from death. There is therefore a solution that is acceptable—indeed, *demande*d—from the permissible personal perspective of each.

The third reason to adopt this test of relevance is that, within the terms of the non-aggregative approach, it allows for a powerful explanation to a person whose claim is judged irrelevant. For one cannot complain that one's claim is not satisfied by an impartial third party when it would not even fall within one's personal prerogative to satisfy it oneself if no moral considerations besides the minimally required concern for the stranger's well-being stood in one's way.

We can now answer the question posed at the start of this section: Why should the non-aggregative approach be accommodated in the form of a constraint on the kinds of claims one can aggregate? The answer is this. The non-aggregative approach requires one to take up each person's permissible subjective perspective. From this perspective, one's own weaker claim may take priority over a stranger's competing stronger claim. But if one's claim is sufficiently weak, then the stronger claim will take priority. The point at which this happens has special significance. ARC treats this point as a boundary between relevant and irrelevant claims. When the non-aggregative approach yields unanimity, it holds that this unanimity ought to be respected. When, by contrast, this approach yields disagreement from each person's perspective, then ARC resolves this discord by an appeal to the aggregative approach.

This idea seems to fit with at least some people's everyday moral sensibility. On a number of occasions on which I have discussed cases involving a choice between either saving one young person from death or preventing a very minor harm to a multitude of others (who would have good lives in any case), an interlocutor has responded, unprompted and before being presented with the preceding analysis, along the following lines. "It would be wrong to allow the claim of the one to be saved from death to be outcompeted by any number of claims to be rid of the very minor harm, because given the difference in what is at stake for each person, someone

facing the very minor harm should withdraw her claim if she is aware of the situation.”<sup>19</sup> This thought is a rough approximation of the deep rationale I have here proposed for ARC.

In the remainder of this article, I shall consider a number of objections to this view.

**IV.** The first of these is due to John Broome. It runs as follows.<sup>20</sup> A key part of ARC’s appeal is that, unlike leading alternative principles of distributive ethics, it can accommodate the intuitive judgment in Case 2 that one ought to cure one young person’s terminal illness rather than save any number of otherwise well-off individuals from a very minor harm. However, this intuitive judgment is not reliable, because it ranges over all numbers, no matter how big, and we cannot have an intuitive grasp of every number.

Contra Broome, the preceding discussion establishes that the intuition in question need not rely on a grasp of every number. All that is required is an intuitive appreciation of the following reasoning. The gap between what is at stake for one individual who has a very weak claim and another individual who has a very strong claim may be so large that, in a one-to-one comparison, from the permissible personal perspective of a person with the very weak claim, it is more important that the stronger claim be satisfied. When this is the case, the person with the very weak claim is “silenced”—his claim must not play a role in determining whether the very weak claims or the competing very strong claim should be satisfied. What is true of one person with a very weak claim is true of every other such person. Therefore, for determining what one ought to do, it doesn’t matter how many people with such claims there are—no matter how large their number, one ought to satisfy the very strong claim. None of the premises in this reasoning relies on our intuitive grasp of the significance of large numbers. Our inability to intuitively apprehend such numbers is therefore irrelevant to the case for ARC.

V. I now turn to an objection raised by Derek Parfit.<sup>21</sup> It starts from the following series of choices (in each case, the choice is made from a feasible set of precisely two alternatives).

(i) Suppose that you can either cure one person of an illness that will kill him while young or, instead, cure a hundred thousand of a moderate lifelong mobility impairment. Also assume that, on ARC, the claim to be cured of the moderate impairment is relevant to the claim to be cured of the terminal illness, so that we should aggregate the claims of the hundred thousand to be cured of the former. If, as we can suppose, the sum of these hundred thousand claims is larger than the single competing claim to be cured of the terminal illness, it follows that, on ARC, you ought to cure the moderate impairment of the hundred thousand rather than the terminal illness of the one.

(ii) Now suppose that you can either cure a hundred thousand people of the moderate impairment or cure a billion of a very slight mobility impairment. Suppose also that the latter's claims are relevant to, and together outcompete, the former. It follows that, on ARC, you ought to cure the very minor impairment rather than the moderate impairment.

(iii) Finally, assume that, on ARC, the claims of the billion to be cured of the very slight impairment are irrelevant when competing with the claim of the one to be cured of the illness that will strike him down in his youth. You therefore ought to cure the one's terminal illness rather than save the billion from the very slight impairment.

Parfit then asks: what, according to ACR, ought you to choose when you have *three* feasible alternatives: save the one from death; save the hundred thousand from the moderate



impairment; or save the billion from the very slight impairment? (You can do only one of these things.) Using the aforementioned judgments from sets of two feasible alternatives, Parfit concludes that, on this view, from this feasible set of three alternatives you ought to save the hundred thousand from the moderate impairment rather than the one from death; you ought to save the billion from the very slight impairment rather than the hundred thousand from the moderate impairment; and you ought to save the one from death rather than the billion from the very slight impairment. He concludes that, “[o]n this view, [you] could not avoid acting wrongly. Whatever [you] do, [you] ought to have done something else instead.” Parfit adds that “[i]n cases of this kind, that is an unacceptable conclusion.”<sup>22</sup>

To answer this objection, we need to consider how the reasoning employed in choosing from these feasible sets of two alternatives applies to a choice from this set of three alternatives. By hypothesis, from the permissible personal perspective of someone with the claim to be cured of the very slight impairment, her cure does not take priority over curing a stranger of the terminal illness. Of course, the same will be true from the perspective of the person with the terminal illness. But, unlike in a choice between only two alternatives, we must also consider the perspective of another interested party: someone facing the moderate impairment. From her perspective, both the claim to be cured of the very slight impairment and the claim to be cured of the terminal illness belong to strangers. For her, these two claims therefore retain the relative significance they have from an objective point of view. Someone facing the moderate impairment will therefore also judge that, in a one-to-one comparison, one stranger’s terminal illness ought to take priority over another stranger’s very slight impairment. Within the terms of the non-aggregative approach, we therefore find the following agreement from the perspective of each: curing the terminal illness ought to take priority over curing the very slight impairment.

ARC respects this unanimous judgment by ruling that the very slight impairment is irrelevant. The remaining eligible alternatives are to cure the terminal illness or the moderate impairment. Since the latter is judged relevant to the former (because it can permissibly be pressed from the perspective of the person with the moderate impairment), this principle requires the choice of whichever of these alternatives satisfies the greatest sum of strength-weighted claims. By assumption, curing a hundred thousand people of the moderate impairment does this. ARC therefore mandates curing them.

It is straightforward to extend this reasoning to cases involving feasible sets with even more alternatives. From each feasible set, one simply assesses whether a claim is relevant by comparing it, one-to-one and from the permissible personal perspective of each, to the *strongest* competing claim in that set. A claim is irrelevant if and only if this strongest competing claim takes priority over it from *every* person's perspective. One then chooses an alternative that satisfies the greatest sum of strength-weighted, relevant claims.

**VI.** Our answer to the preceding objection shows that ARC always selects at least one alternative from a feasible set. But it may prompt a new objection: that ARC violates the following basic principle of rational choice.<sup>23</sup>

*Basic Contraction Consistency:* If an alternative is permissibly chosen from a given feasible set, then it is also permissibly chosen from any subset containing that alternative.<sup>24</sup>

To assess whether this condition is indeed violated by ARC, it is useful to use some shorthand. Let “DEATH” stand for curing one from death; “MODERATE” stand for curing a hundred thousand of the moderate impairment; and “VERY SLIGHT” for curing a billion of the very slight impairment. Given the discussion in the preceding section, it then appears that ARC mandates the following Basic Contraction Consistency-violating choices: from the feasible set {DEATH; MODERATE; VERY SLIGHT}, choose MODERATE; and from the set {MODERATE; VERY SLIGHT}, choose VERY SLIGHT.

In response, it is important to note that ARC offers a coherent explanation of this pattern of choices.<sup>25</sup> When one must choose from the set {MODERATE; VERY SLIGHT}, then the claims to the cure of the very slight impairment are relevant. By contrast, when one must choose from the set {DEATH; MODERATE; VERY SLIGHT}, then, on ARC, the claims to be free from the very slight impairment are not relevant. The distinction that ARC draws between “choosing VERY SLIGHT from {MODERATE; VERY SLIGHT}” and “choosing VERY SLIGHT from {DEATH; MODERATE; VERY SLIGHT}” is not *ad hoc*. Rather, it follows from the fact that, on this view, the relevance of a weaker claim depends on whether, from the permissible personal perspective of someone with the weaker claim, one gives this claim priority over the strongest competing claim in a one-to-one comparison. It therefore tracks the presence or absence of the type of unanimity from each person’s permissible perspective that ARC regards as crucially important. Since this pattern of choices tracks ARC’s rationale, it would be wrong to regard it as irrational in the sense in which this term is used in rational choice theory, viz. “not rationalisable by a coherent set of values.”

Ought one then to reject Basic Contraction Consistency as a necessary condition for rational choice? Some think so.<sup>26</sup> By contrast, I favour an approach outlined by Broome.<sup>27</sup> Basic Contraction Consistency is an internal consistency condition—it must be satisfied if choices are

to be explicable in terms of values that guide choice. When we judge whether it is violated by a moral view, we must therefore ensure that we identify alternatives by characteristics which that view takes to be pertinent grounds for choice. In other words, we must take one alternative to be identical to another if and only if it has, on the view in question, the same pertinent characteristics. Now, if we individuate alternatives in line with what, according to ARC, are their morally pertinent characteristics, we should therefore distinguish “VERY SLIGHT when chosen from {MODERATE; VERY SLIGHT}” and “VERY SLIGHT when chosen from {DEATH; MODERATE; VERY SLIGHT}.” We can do so by using the superscript “R” to denote an alternative with relevant claims and “¬R” to denote an alternative without relevant claims. Properly individuated, the feasible sets and choices mandated by ARC are therefore the following. From {DEATH<sup>R</sup>; MODERATE<sup>R</sup>; VERY SLIGHT<sup>¬R</sup>}, choose MODERATE<sup>R</sup> and from {MODERATE<sup>R</sup>; VERY SLIGHT<sup>R</sup>} choose VERY SLIGHT<sup>R</sup>. Since the latter set is not a sub-set of the former, these choices do not violate Basic Contraction Consistency.<sup>28</sup>

**VII.** I shall now address the following fundamental objection to my discussion of ARC. This principle merely reproduces, in an elaborate way, a set of pre-existing common judgments about when aggregation is required and when it is ruled out. ARC and its proposed rationale therefore do not offer anything new; nor do they do real work in justifying these judgments.<sup>29</sup>

The reply has several parts. First, one thing that is novel is the proposed connection between the answers to two questions that are typically discussed separately, namely: How much priority can one permissibly give oneself over a stranger? and Which claims aggregate? The proposed connection is controversial and testable. It is controversial, because leading thinkers disagree. John Taurek, for example, holds both that one can permissibly prioritise oneself over

others and that claims never aggregate.<sup>30</sup> And Parfit holds both that morality permits giving priority to oneself and that it requires aggregation of even comparatively very small claims.<sup>31</sup> It is also testable, since we have independent purchase on the correct answers to these questions. Consider, first, our judgments about the extent to which we are permitted to prioritise ourselves over strangers in situations of the kind outlined above. Of course, neither common-sense morality nor the moral theories that endorse such a personal prerogative precisely delineate the extent of this prerogative.<sup>32</sup> Nonetheless, there are, in everyday ethics, cases in which it is clearly true that the balance of one's interest and the interest of the stranger is such that one may favour oneself. Furthermore, there are cases in a "borderline zone," in which it is unclear whether or not it is permissible to favour oneself, but in which a moderate shift in the balance of interests would yield a more definite verdict. There are also cases in which this configuration is such that one must clearly aid the stranger. If the proposed version of ARC is correct, then these will correspond, respectively, to the following: cases in which aggregation is intuitively required; borderline cases for aggregation; and cases in which aggregation is intuitively prohibited.

In the harms involved in Case 1, we already have an example in which it is intuitively clear that one is permitted to save oneself from a lesser harm rather than a stranger from a greater harm and in which the lesser harm should be aggregated. In the harms involved in Case 2, we have an example in which it is clear that one is not permitted to save oneself from the lesser harm rather than a stranger from the greater harm and in which, intuitively, the lesser harm should not be aggregated. A further key test is therefore to consider a borderline case. The following seems to me such a case.

Suppose that you are a young adult and that a finger of yours has just been severed in an accident (for which no one is to blame). Its loss will modestly hinder your everyday activities

and be somewhat disfiguring, but in other respects, you will have a good life. Using your own resources, you can have it successfully reattached, or you can instead save a stranger's life. This seems to me a borderline case for the personal prerogative. It is unclear whether it is permissible to save yourself from this disfigurement. Moreover, if your loss were moderately greater—several fingers, say—it would be quite clear that you could permissibly save yourself, whereas if it were moderately smaller—a mere part of a finger—it would appear that you ought to save the stranger.

Now consider the following case.

**Case 3.** You, a morally-motivated stranger, can either cure one young man's terminal illness or prevent a number of people (who, overall, will lead good lives in any case) from losing a finger. Is there a number of people you can spare the latter harm for which you ought to spare them this harm instead of saving the one from death?

This seems to me to be a borderline case. Intuitively, at least, it is unclear whether there is a number of people who each face the loss of a single finger for which you ought to save that number from this loss rather than save a single young person's life. Moreover, if the former harm were moderately larger—say, each person faced the loss of several fingers—then it would seem one should save a very large number from this harm rather than one from death; whereas if the harm were significantly smaller—a mere part of a finger—it would be counterintuitive to judge that one ought to do so. It therefore seems that ARC passes this test.

I now turn to the second part of the reply to the objection that ARC offers nothing new. ARC tells us what to do in cases in which we may lack clear intuitive judgments. One such case

is the three-feasible-alternative case discussed in Sections V and VI. When you can either save one from death, or a hundred thousand from the moderate impairment, or a billion from a very slight impairment, it is, I submit, difficult to discern intuitively what you ought to do.<sup>33</sup> As we saw, ARC tells you to save the hundred thousand from the moderate impairment, and for a plausible reason. From the perspective of each person, in a one-to-one comparison of competing claims, one ought not to prioritise the very slight impairment over death. Within the terms of the non-aggregative approach, curing the very slight impairment is therefore unanimously judged ineligible. ARC therefore rules out curing the very slight impairment. Among the remaining alternatives, curing the moderate impairment satisfies the greatest sum of strength-weighted, relevant claims, and so ought to be chosen. Indeed, ARC selects at least one alternative from every set of alternatives. The collection of cases in which it tells us what to do therefore presumably includes many cases in which we do not have clear pre-theoretical judgments.

The final part of the response is this. The relevant intuitive judgments in Cases 1 and 2 stand in need of justification. The rationale I have proposed “does real work” in providing this justification because it shows that these judgments reflect important values. It also plays a crucial role in replying to objections, as we have seen in Sections IV, V, and VI. Against Broome’s ostensibly powerful debunking explanation of non-aggregative intuitions, I invoked ARC’s rationale to offer a vindicatory explanation of these intuitions. This rationale also helped explain why ARC can always choose at least one alternative from a feasible set and why it does not violate principles of rational choice. In the next section, I shall appeal to this rationale to rebut a final objection.

**VIII.** Consider the following case, which is a variation of a case formulated by Parfit.<sup>34</sup>

**Case 4.** Nine hundred young adults have contracted a disease. Before being affected, they all had life prospects equal to 1 on a cardinal well-being scale where 1 is a full life in good health and 0 is death at a young age. The disease will affect them differentially: if untreated, person One will experience a very slight loss in lifetime well-being (and so still lead a good life), person Two will experience a slightly greater loss, and so on in a linear fashion through to person Nine Hundred, who will end up with a very large loss and will lead a life only barely worth living. Fortunately, you have nine hundred doses of a medicine that can be consumed in either a one-dose form or a superdose form which uses up all nine hundred doses (no other dosages are possible). A single dose only slightly alleviates the disease's effects. In a person who can tolerate it, the superdose nearly completely eliminates its effects. Only Nine Hundred can tolerate the superdose. This means you have the following options.

EVERYONE ONE DOSE. On the aforementioned scale of well-being, One experiences a decline of  $1/1000$ , Two of  $2/1000$ , and so on until Nine Hundred, who suffers a loss of  $900/1000$ .

NINE HUNDRED THE SUPERDOSE. Nine Hundred has a decline of only  $1/1000$ , One has a decline of  $2/1000$ , Two a decline of  $3/1000$ , and so on until Eight Hundred Ninety-Nine, who has a decline of  $900/1000$ .

The distributions of losses for these options are listed in Table 1.



**Table 1. The distribution of losses of well-being in Case 4.**

Individual	EVERYONE ONE DOSE	NINE HUNDRED THE SUPERDOSE
One	1/1000	2/1000
Two	2/1000	3/1000
Three	3/1000	4/1000
...	...	...
Eight Hundred Ninety-Nine	899/1000	900/1000
Nine Hundred	900/1000	1/1000

In this case, since each individual’s well-being outcome is just 1 minus his loss, EVERYONE ONE DOSE leads to an anonymised distribution of well-being that is identical to the anonymised distribution of well-being under NINE HUNDRED THE SUPERDOSE. Put differently, the well-being associated with the worst anonymised position under EVERYONE ONE DOSE is identical to the well-being associated with the worst position under NINE HUNDRED THE SUPERDOSE and correspondingly for the well-being associated with every other position. All standard views of distributive ethics, from leximin through to utilitarianism, will therefore hold that these alternatives are equally good and that it is thus permissible to choose either one.

By contrast, it is plausible that ARC will demand that you give Nine Hundred the superdose. Because so much is at stake for Nine Hundred and he would be so badly off without it, he has a very strong claim to this dose. The largest of the competing claims is held by Eight Hundred Ninety-Nine. We can assume that even though she would be very badly off if she remained untreated, the small improvement in her well-being provided by a single dose means that her claim to that dose is too weak to be relevant to Nine Hundred’s competing claim. All

other claims to a single dose will therefore also be irrelevant, so that ARC will mandate treating Nine Hundred with the superdose.

In this case, Parfit agrees with standard distributive views that both courses of action are equally good and therefore both are permissible. He therefore sees this case as providing a counterexample to ARC.<sup>35</sup>

However, contrary to Parfit, it seems that in Case 4, only ARC arrives at the right conclusion. Consider what hangs on your choice for each person involved, taken separately. For Nine Hundred, what is at stake is the difference between a good life and a life barely worth living. By contrast, for every other person, all that depends on your choice is a marginal improvement in their quality of life. By assumption, this improvement is so small that none of them would be permitted to secure it for themselves alone rather than secure the larger improvement for Nine Hundred. Moreover, when you step back and take an impersonal view, it is clear that the value of these small potential improvements to everyone except Nine Hundred, is, in the aggregate, precisely equal to the value of the potential improvement in Nine Hundred's well-being. In sum, in this case, giving Nine Hundred the superdose is the only alternative that does justice to the perspective of each person taken separately, and it does so at no cost in terms of impersonal value. Rather than presenting an objection to ARC, Case 4 therefore reveals its advantages over familiar alternative views.<sup>36</sup>

However, Parfit also presents a case like the following, which is more challenging for ARC.

**Case 5.** This case is just like Case 4, except that if you give Nine Hundred the superdose, everyone's decline is very slightly greater than in Case 4, as follows.

NINE HUNDRED THE SUPERDOSE\*: Nine Hundred has a decline of 2/1000, One has a decline of 3/1000, Two a decline of 4/1000, and so on until Eight Hundred Ninety-Nine, who has a decline of 901/1000.

The distributions of losses for these options are listed in Table 2.

**Table 2. The distribution of losses of well-being in Case 5.**

Individual	EVERYONE ONE DOSE	NINE HUNDRED THE SUPERDOSE*
One	1/1000	3/1000
Two	2/1000	4/1000
Three	3/1000	5/1000
...	...	...
Eight Hundred Ninety-Nine	899/1000	901/1000
Nine Hundred	900/1000	2/1000

In Case 5, giving everyone one dose unquestionably leads to a better anonymised distribution of well-being: under this alternative, the well-being associated with the worst position is greater than the well-being associated with the worst position if you instead provide Nine Hundred with the superdose, and correspondingly for the well-being associated with every other position. Giving everyone one dose is therefore better on every standard measure of the goodness of distributions of well-being, from leximin through to utilitarianism.

How will ARC require you to choose? Because so much is at stake for Nine Hundred and he would be so badly off without it, he still has a very strong claim to this dose. As before, the

largest of the competing claims is held by Eight Hundred Ninety-Nine. Her claim to one dose is stronger than it was in Case 4, because this dose would do her more good. Nonetheless, because the improvement you can secure for her remains so small, her claim to a single dose is still very weak in comparison to Nine Hundred's Claim. Indeed, given how much is at stake for Nine Hundred, we can assume Eight Hundred Ninety-Nine's claim is too weak to be relevant to Nine Hundred's competing claim. All other claims to a single dose will therefore also be irrelevant, so that ARC will mandate treating Nine Hundred with the superdose.

Parfit regards this implication as unacceptable.<sup>37</sup> In his view, it is wrong to allow Nine Hundred's claim to trump everyone else's competing claims to receive aid, when satisfying his claim will unambiguously worsen the anonymised distribution of final well-being and lead another (i.e., Eight Hundred Ninety-Nine) to end up even worse off than Nine Hundred would be if everyone were given a single dose.

Case 5 highlights that ARC may mandate choices that have costs in terms of the impersonal goodness of the distribution of well-being. Parfit is right that this counts against ARC. However, as we saw in Case 4, the well-being attached to each position in the distribution of well-being is not the only thing of importance. What is at stake for each person and how this compares to what is at stake for another with a competing claim is also important. In Case 5, it remains true that Nine Hundred is the only person to whom you can make a large difference—the difference between a full, good life and a life barely worth living. If you forgo making this great difference to Nine Hundred's life, then all you can do is marginally improve the lives of a multitude of others. This provides a reason to give Nine Hundred the superdose, because that would best respect what matters most when we take up each person's position separately and compare her claim to the strongest competing claim. The question is, then, whether it is

reasonable to sacrifice some impersonal goodness of the distribution of well-being in order to do justice to the non-aggregative approach. If I have characterised this approach correctly, it is both deeply rooted in our moral psychology and reflects a form of respect for the separateness of persons. I therefore believe it is indeed reasonable to accept some loss of impersonal goodness for its sake. Of course, this sacrifice should not be too great. Fortunately, ARC limits the sacrifice of impersonal value it will demand in cases of this kind. To see why, imagine a choice between EVERYONE ONE DOSE and a new alternative, NINE HUNDRED THE SUPERDOSE\*\*, and progressively magnify the decline in the well-being attached to every anonymised position associated with the latter. As this decline grows larger, individuals' claims to a single dose grow stronger and become relevant. When enough of these smaller claims become relevant, they will together outcompete Nine Hundred's claim. ARC will then require the impersonally best alternative.<sup>38</sup>

In sum, Parfit's purported counterexamples are not decisive. Indeed, Case 4 *supports* ARC. And Case 5 merely reveals that we face a difficult choice. We can either do what every standard conception of the goodness of distributions says is best, or have the non-aggregative, sympathetic approach play a central role in our moral deliberations by giving special significance to the claims of those who have much more at stake than others. When faced with this dilemma, I have argued that it is reasonable to sacrifice some impersonal goodness to do justice to the non-aggregative approach.

**IX.** Many believe both that we ought to save a large number from being permanently bedridden rather than save one from death *and* that we ought to save one from death rather than save a multitude of people (who would be well off in any case) from very minor harms, no matter how

large this multitude. I have argued that these judgments reflect two conflicting responses, one aggregative and one non-aggregative, to the equal objective value of each person's well-being. I have also argued that ARC partially accommodates and arbitrates between these responses in a reasonable way. This principle holds that a person's weaker claim is irrelevant in the face of another's stronger competing claim just in case, in a one-to-one comparison of these claims, the stronger claim ought to take priority from *every* person's point of view. It thereby respects a form of unanimity that emerges from the non-aggregative approach. In the absence of such unanimity, this principle requires that we resolve conflicts of interest as the aggregative approach demands, by satisfying the greatest sum of strength-weighted claims.

I have argued that leading objections to this view are without force, with the following exception. ARC's accommodation of the non-aggregative approach is costly: in some cases, it mandates the selection of an alternative that worsens every anonymised position, from the worst to the best position, in a distribution of well-being. Such cases therefore reveal a way in which we may be forced to choose between pursuing what is impersonally best and giving a special role to the claim of a person who has much more at stake than anyone else. This choice is not an easy one. However, it seems reasonable to sacrifice some impersonal goodness in order to do justice to how things look when we take up each person's point of view, one person at a time, and see how what is at stake for him or her compares to what is at stake for someone with a competing claim.

---

\* This paper was presented to the Association for Legal and Social Philosophy, Philosophy for All, the Society for Applied Philosophy, and at the Joint Session of the Aristotelian Society and the Mind Association. It was also presented at Cambridge, Exeter, Glasgow, Harvard, the LSE, Münster, the NIH, Princeton, Stanford, and UC Berkeley. I thank those present for their responses. I am also grateful to Elizabeth Ashford, Luc Bovens, Richard Bradley, Susanne Burri, Simon Cotton, Marc Fleurbaey, Robert Goodin, Christopher Heathwood, R.J. Leyland,

---

Kasper Lippert-Rasmussen, William MacAskill, Michael Martin, Joseph Mazon, Veronique Munoz-Dardé, Michael Otsuka, Alan Patten, Philip Pettit, Daniel Ramöller, Henry Richardson, Peter Vallentyne, the Editors of this journal, and two anonymous referees for comments. Finally, I thank Princeton's University Center for Human Values for its Faculty Fellowship in 2012-13, during which this paper was written.

<sup>1</sup> I am here assuming that everyone with this condition has the well-being level that people typically assign to it in surveys. (See Erik Nord, Jose-Louis Pinto Prades, Jeff Richardson, Paul Menzel, and Peter Ubel, "Incorporating Societal Concerns for Fairness in Numerical Valuations of Health Programmes," *Health Economics* 8 (1999): 25–39.) I am of course not claiming that, in reality, all who are disabled in this way do or should regard their existence as only somewhat better than an early death.

<sup>2</sup> See David Brink, "The Separateness of Persons, Distributive Norms, and Moral Theory," in *Value, Welfare, and Morality*, ed. R.G. Frey and C.W. Morris (Cambridge: Cambridge University Press, 1993), pp. 252-289, on p. 270; Frances Kamm, *Morality, Mortality, Vol. I* (Oxford: Oxford University Press, 1993), Chapters 8-10 and *Intricate Ethics* (Oxford: Oxford University Press, 2007), pp. 297-298 and 484-486; T.M. Scanlon, *What We Owe to Each Other* (Cambridge, MA: Harvard University Press, 1998), pp. 238-241; Roger Crisp, "Equality, Priority, and Compassion," *Ethics* 113 (2003): 745-763, p. 754; Michael Otsuka, "Saving Lives, Moral Theory, and the Claims of Individuals," *Philosophy and Public Affairs* 34 (2006): 110-135; Larry Temkin, "Health Care Distribution and the Problem of Trade-offs," in *'Goodness' and 'Fairness': Ethical Issues in Health Resource Allocation*, ed. Daniel Wikler and Christopher Murray, World Health Organization, forthcoming; and Dale Dorsey, "Headaches, Lives, and Value," *Utilitas* 21 (2009): 36-58. For criticism of the "no" answer to Case 2, see Alistair Norcross, "Comparing Harms: Headaches and Human Lives," *Philosophy and Public Affairs* 26 (1997): 135-167; Derek Parfit, "Justifiability to Each Person," *Ratio* 16 (2003): 368–390; and John Broome, "A Comment on Temkin's Trade-offs," in *'Goodness' and 'Fairness' and Weighing Lives* (Oxford: Oxford University Press, 2004), pp. 56-58.

<sup>3</sup> See Frank Cowell, Marc Fleurbaey, and Bertil Tungodden, "The Tyranny Puzzle in Welfare Economics: An Empirical Investigation," STICERD working paper PEP-5 (2010).

<sup>4</sup> Scanlon, "Replies," *Ratio* 16 (2003): 424–439, p. 433; Fleurbaey, Tungodden, and Vallentyne, "On the Possibility of Nonaggregative Priority for the Worst Off," *Social Philosophy and Policy* 26 (2009): 258-285, p. 258.

<sup>5</sup> See Fleurbaey, et al., "On the Possibility."

---

<sup>6</sup> There are other non-standard theories that can accommodate these two judgments together. Perhaps the best-known of these is a “threshold view,” on which there is an absolute threshold of well-being, such that an improvement in well-being below the threshold, no matter how small and no matter how few people would receive it, always outweighs an improvement above the threshold, no matter how large and no matter how many people would receive it. When all improvements take place below the threshold, then the numbers count, as do their size and the level from which they take place. The same applies above the threshold. (See Tungodden, “The Value of Equality,” *Economics and Philosophy* 19 (2003): 1-44; Campbell Brown, “Priority or Sufficiency... or Both?” *Economics and Philosophy* 21 (2005): 199-220; and Matthew Adler, *Well-Being and Fair Distribution: Beyond Cost-Benefit Analysis* [Oxford: Oxford University Press, 2011], Chapter 5.) If we assume that life with the complete disability is below the threshold and that the harm of being bedridden for a day would occur to people who are above the threshold, then this view accounts for the aforementioned judgments in Cases 1 and 2. However, this view is open to the following objection. On the threshold view, an arbitrarily small gain for a person below the threshold can outweigh any gain above the threshold, no matter how large it is and no matter how many people would receive it. Brown and Adler rightly regard this as implausible. ARC is not vulnerable to this objection. ARC holds that a very small potential benefit to a badly off person gives rise to a weak claim. Because this claim is weak, the competing claim of someone above the purported threshold to a large benefit will be strong enough to be relevant. On ARC, you ought therefore to give a large benefit to each of a large number of people who are above the threshold rather than give a very small benefit to a single person below the threshold, when you cannot do both. (For objections to other non-standard views, see Fleurbaey et al., “On the Possibility.”)

<sup>7</sup> See Brink, “Separateness;” Kamm, *Morality, Mortality, Vol. I*, Chapters 8-10 and *Intricate*, pp. 297-298 and 484-486; Scanlon, *What We Owe*, pp. 238-41; and Temkin, “Health Care” and *Rethinking*.

<sup>8</sup> The proposal draws inspiration from Nagel, “Equality” and *Equality and Partiality* (Oxford University Press, 1991), and Kamm, *Morality, Mortality, Vol. I*, Chapters 8-10.

<sup>9</sup> The strongest claim is not necessarily had by the person who is worst off. Rather, as defined above, the strength of a person’s claim is a function both of how much a person could gain in well-being and how badly off he would be without this gain. The decision rule that best captures what is most important from this perspective is therefore *not* maximin, but rather: satisfy the largest claim. See also Scanlon, “Contractualism and Utilitarianism,” in *Utilitarianism and Beyond*, ed. Amartya Sen and Bernard Williams (Cambridge: Cambridge University Press,



1982), pp. 103-128, pp. 122-123 and Sophia Reibetanz, “Contractualism and Aggregation,” *Ethics* 108 (1998): 296-311.

<sup>10</sup> See his *The Theory of Moral Sentiments*, ed. D.D. Raphael and A.L. Macfie (Indianapolis: Liberty Fund, 1982).

<sup>11</sup> See, for example, Tehila Kogut and Ilana Ritov, “The ‘Identified Victim’ Effect: An Identified Group, or Just a Single Individual?” *Journal of Behavioral Decision Making* 18 (2005): 157-167. It is therefore no accident that this process of placing oneself in each person’s shoes and taking in what is at stake for each, considered separately, lies at the heart of John Taurek’s non-aggregative distributive morality. See his “Should the Numbers Count?” *Philosophy and Public Affairs* 6 (1977): 293-316.

<sup>12</sup> See Nagel, “Equality” and *Equality and Partiality*, pp. 65-69; Scanlon, *What We Owe*, pp. 238-241; and Temkin, *Rethinking*, p. 100.

<sup>13</sup> Wollner, “Egalitarianism, Numbers, and the Dreaded Conclusion,” *Ethical Perspectives* 19 (2012): 399-416, p. 409.

<sup>14</sup> *Morality, Mortality, Vol. I*, p. 154. Emphasis in original.

<sup>15</sup> For key contributions to the debate, see Shelly Kagan, *The Limits of Morality* (Oxford: Oxford University Press, 1991); Samuel Scheffler, *The Rejection of Consequentialism* (Oxford: Clarendon Press, 1994); and Kamm, *Morality, Mortality, Vol. II*, Chapter 8 and *Intricate*, Chapter 1.

<sup>16</sup> This is close to what Kamm, *Morality, Mortality, Vol. I*, Chapter 8, refers to as the “objectively permissible subjective perspective.”

<sup>17</sup> The notion of sympathy used here is therefore normative: when one imagines oneself in another’s position, one does not take on his actual self-concern. Rather, one takes on the concern he would maximally be permitted to have. In this, I follow Smith, who makes clear that when placing oneself in another’s shoes, one does not adopt all his preferences and feelings, but rather those that an idealized spectator could sympathize with—those that are proper. For example, on self-interest, he writes: “If [a man] would act so as that the impartial spectator would enter into the principles of his conduct, (...) he must (...) humble the arrogance of his self-love, and bring it down to something which other men can go along with” (*Moral Sentiments*, II.II.II; see also I.I.III-IV).

<sup>18</sup> Of course, the situation may change if A communicates that he does not want to be aided. I leave aside the role of individuals’ wishes of this kind.

<sup>19</sup> Those who responded in this way include an audience member at a Philosophy for All lecture and Kai Spiekermann.

<sup>20</sup> Broome, “Comment,” and *Weighing Lives*, pp. 56-58. Parfit, “Justifiability,” p. 385, raises a similar objection.

<sup>21</sup> Parfit, “Justifiability,” p. 384. I thank an anonymous referee for pressing me to discuss this objection.

<sup>22</sup> “Justifiability,” p. 384.

<sup>23</sup> This objection is found in Temkin, “Health Care” and *Rethinking*; Broome, “Comment;” and Wollner, “Egalitarianism.”

<sup>24</sup> This condition is also known under the name “Independence of Irrelevant Alternatives,” or “Alpha.” For a discussion of this condition, see Sen, *Collective Choice and Social Welfare* (San Francisco: Holden Day) and “Internal Consistency of Choice,” *Econometrica* 61 (1993): 495-521.

<sup>25</sup> Cf. Kamm, *Intricate*, pp. 484-487.

<sup>26</sup> Fleurbaey, et al., “On the Possibility,” pp. 283-284, rehearse the arguments for rejecting Basic Contraction Consistency.

<sup>27</sup> See Broome, *Weighing Goods* (Oxford: Blackwell, 1991), Chapter 5.

<sup>28</sup> For criticism of this “fine-grained individuation” approach, see Temkin, *Rethinking*, p. 465. For a response to this criticism, see my “Vaulting Intuition: Temkin’s Critique of Transitivity,” *Economics and Philosophy*, 29 (2013): 409-425.

<sup>29</sup> This objection was put to me by an anonymous referee.

<sup>30</sup> Taurek, “Numbers,” pp. 301 and 306-307.

<sup>31</sup> On permissible priority for oneself, see Parfit, *On What Matters, Vol. I* (Oxford: Oxford University Press, 2011) especially Section 59; on aggregation, see Parfit, “Justifiability,” pp. 384-385.

<sup>32</sup> See Taurek, “Numbers,” pp. 301 and 306-307; Scheffler, *Rejection*; Nagel, *Equality and Partiality*, Chapters 2 and 4; and Kamm, *Morality, Mortality, Vol. II*, Chapter 8, none of which precisely specify the boundaries of the prerogative. Parfit, *On What Matters, Vol. I*, esp. Chapter 6 and Section 59, argues that these boundaries are irremediably imprecise.

<sup>33</sup> Cf. Wollner, “Egalitarianism,” Section III and Temkin, *Rethinking*, pp. 45-52.

<sup>34</sup> Parfit, “Justifiability,” p. 383n16. See also Temkin, *Rethinking*, pp. 440-445.

<sup>35</sup> “Justifiability,” p. 383n16.

<sup>36</sup> Here, I am agreeing with Temkin, *Rethinking*, p. 440ff.

<sup>37</sup> Parfit, “Justifiability,” p. 383n16.

<sup>38</sup> As mentioned in the introduction, I have so far focused on cases in which the claims to an alternative are either all relevant or all irrelevant. The scenario I am here imagining, however, would involve some claims to an alternative (e.g. Eight Hundred Ninety-Nine’s claim to EVERYONE ONE DOSE) becoming relevant before others’ claims to that very same alternative reach the threshold of relevance defined by the test given in the main text. For such cases, the test I have offered may need adjustment—so that when, say Eight Hundred Ninety-Nine’s claim to EVERYONE ONE DOSE becomes relevant, other, weaker claims to that alternative also become relevant. For discussion of such cases, see Kamm, “Aggregation and Two Moral Methods,” *Utilitas* 17 (2005): 1-23.

### **Biographical note**

Alex Voorhoeve is an Associate Professor in the Department of Philosophy, Logic and Scientific Method at the London School of Economics. He works on distributive ethics and the theory of rational choice. His articles have appeared in *Philosophy and Public Affairs*, *The Journal of Political Philosophy*, and *Economics and Philosophy*, among other places. He is the author of *Conversations on Ethics* (Oxford: Oxford University Press, 2009).