

Aggregation, Balancing, and Respect for the Claims of Individuals

Bastian Steuwer

(Preprint, please cite published version)

Abstract. Most non-consequentialists “let the numbers count” when one can save either a lesser or greater number from equal or similar harm. But they are wary of doing so when one can save either a small number from grave harm or instead a very large number from minor harm. Limited aggregation is an approach that reconciles these two commitments. It is motivated by a powerful idea: our decision whom to save should respect each person who has a claim to our help, including those whom we fail to save. However, it has recently been argued that it is open to decisive objections. I develop a new limitedly aggregative view: *Hybrid Balance Relevant Claims*. This view is well-grounded in the reasons we have to be skeptical of aggregation and resolves all recent challenges by paying careful attention to the rationale for limited aggregation.

I. Introduction

Many non-consequentialists share a strong hesitation regarding the aggregation of harms across different individuals.¹ They object in particular to aggregation when trivial harms might thereby

¹ Examples include John Rawls, *A Theory of Justice* (Oxford: Oxford University Press, 1999), pp. 23-26; Robert Nozick, *Anarchy, State, and Utopia* (New York: Basic Books, 1974), pp. 32-33; John M. Taurek, Should the Numbers Count?, *Philosophy & Public Affairs* 6 (1977), pp. 293-316; Thomas Nagel, *Mortal Questions* (Cambridge: Cambridge University Press, 1979), ch. 8; T.M. Scanlon, *What We Owe to Each Other* (Cambridge, MA: The Belknap Press, 1998), pp. 238-41; and F.M. Kamm, *Intricate Ethics* (Oxford: Oxford University Press, 2007), pp. 31-40. Opposition to full aggregation and non-consequentialism nonetheless

outweigh significant harms. Such aggregation, a common complaint goes, fails to take seriously that individuals are separate persons. Most of these skeptics regarding aggregation would, however, still like the numbers to count when one can save either a lesser or greater number from equal or similar harm. What is therefore needed is a moral theory that allows the relative numbers to count sometimes but not always. Several philosophers have proposed theories of this kind.² The different approaches are motivated by a powerful idea: our decision whom to save should respect each person who has a claim to our help, including those whom we fail to save. Such views have been called *limited aggregation*.

The standard cases for limited aggregation are cases in which groups are homogeneous; i.e. groups in which everyone in the group has a claim that is as strong as the claim of every other member of the group. However, many cases are not like this. Often times we face decisions where the groups are heterogeneous; i.e. not all groups members face the same plight. Current proposals of limited aggregation have been shown to have devastating flaws when they are extended to cases with such heterogeneous groups. Patrick Tomlin has shown that each of two possible ways of extending a leading proposal of limited aggregation, Aggregate Relevant Claims, to heterogeneous

are distinct positions. For example, some non-consequentialists only depart from consequentialism by accepting either deontological constraints or agent-centered prerogatives or both. These forms of non-consequentialism, however, are only partial departures that tame a basically consequentialist outlook of morality with additional considerations. It is a half-hearted form of non-consequentialism; see Thomas Sinclair, Are We Conditionally Obligated to be Effective Altruists?, *Philosophy & Public Affairs* 46 (2018), pp. 36-59 (pp. 43-49). The opposition to full aggregation which is subject of this article is instead a hallmark of a more thorough form of non-consequentialism.

² F.M. Kamm, *Morality, Mortality*, vol. 1 (New York: Oxford University Press, 1993), pp. 156-61, and *Intricate Ethics*, pp. 31-77; Scanlon, *What We Owe to Each Other*, pp. 238-41; David Lefkowitz, On the Concept of a Morally Relevant Harm, *Utilitas* 20 (2008), pp. 409-423; and Alex Voorhoeve, How Should We Aggregate Competing Claims?, *Ethics* 125 (2014), pp. 64-87. In this article, I understand limited aggregation as a theory about what we ought to do and set aside a related view that interprets limited aggregation as part of our theory of the good. For the latter, see Larry Temkin, *Rethinking the Good* (Oxford: Oxford University Press, 2012), ch. 3; and Dale Dorsey, Headaches, Lives and Value, *Utilitas* 21 (2009), pp. 36-58.

groups violates one of two uncontroversial principles. On one extension of Aggregate Relevant Claims it violates a principle he calls Equal Consideration for Equal Claims which requires us to give all claims of equal strength equal weight in determining whom to save. On its alternative extension Aggregate Relevant Claims violates what he calls the Principle of Addition, which requires that adding a claim to a group cannot make saving this group less choiceworthy.³ In this article, I show how these problems can be resolved by a new theory of limited aggregation that is well-grounded in the reasons we have to be skeptical of aggregation and that meets this challenge set by Tomlin and related recent challenges. I propose the following theory:

Hybrid Balance Relevant Claims. Relevant individual claims ought to be balanced against one another, starting with the strongest claim(s) overall. If there are relevant claims on one side that are not balanced by claims on the other, then these unmatched claims will be decisive in determining what we ought to do. If the claims are evenly balanced, then we are permitted to save either group, or perhaps required to give equal chances. The relevance of claims is determined by two conditions:

- (1) *The local relevance condition:* A claim can be balanced against another claim only if the two claims are relevant to one another.
- (2) *The global relevance condition:* Every individual with a strong claim has a veto against the consideration of any type of claim that is irrelevant to her claim, if such consideration would lead to her claim not being satisfied (e.g. by not being saved).

³ Patrick Tomlin, On Limited Aggregation, *Philosophy & Public Affairs* 45 (2017), pp. 232-60. Tomlin's criticism has been extended by Joe Horton, Always Aggregate, *Philosophy & Public Affairs* 46 (2018), pp. 160-74. Like Tomlin and Horton, I will set aside well-acknowledged objections that limited aggregation violates axioms of rational choice, namely transitivity and the independence of irrelevant alternatives. See Derek Parfit, Justifiability to Each Person, *Ratio* 16 (2003), pp. 368-90 (pp. 384-85); and John Halstead, The Numbers Always Count, *Ethics* 126 (2016), pp. 789-802 (pp. 797-99). For responses to these criticisms see Kamm, *Intricate Ethics*, pp. 297-98, 484-87; Voorhoeve, How Should We Aggregate Competing Claims?, pp. 76-79; Alex Voorhoeve, Why One Should Count Only Claims with which One Can Sympathize, *Public Health Ethics* 10 (2017), pp. 148-56 (pp. 152-53); and Tomlin, On Limited Aggregation, p. 236 fn. 11.

Hybrid Balance Relevant Claims is a general theory for how to adjudicate between the competing claims of individuals. However, in this article my focus is on defending it in a narrower set of cases in which we are concerned only with saving people from harm as opposed to harming persons. In Section II, I will establish that the idea of relevance is key to any plausible theory of limited aggregation. Then in Section III, I will provide my defense of Hybrid Balance Relevant Claims and give a more precise understanding of what I mean by “balancing”. Section IV illustrates my view with some examples. Section V shows how my view escapes all recent challenges that have been raised against limited aggregation.

II. Relevance and Limited Aggregation

I have already mentioned that many philosophers are opposed to aggregation because it allows a large number of trivial claims to outweigh the significant claim of a single individual. A paradigm case for this phenomenon is *Life versus Headaches*. A fully aggregative view struggles to accommodate the intuition that we should not let a single person die in order to save any number of people from headaches. It seems that if we can aggregate the pain of the various headaches, that there will be a number of people suffering from headaches that outweighs the life of the one.

There are some views that allow aggregation in all cases but seek to avoid this conclusion. For example, a view could be fully aggregative but assign infinite disvalue to deaths. Such a view would then, however, struggle to accommodate the intuition that sometimes the relative numbers should count. The disvalue of two deaths would also be infinite and no number of a slightly lesser harm than death could outweigh single deaths.⁴ It would also struggle to account for cases in which harms other than death, for example paraplegia or excruciating pain, ought not be outweighed by

⁴ See also Michael Otsuka, Saving Lives, Moral Theory, and the Claims of Individuals, *Philosophy & Public Affairs* 34 (2006), pp. 109-135 (pp. 127-28).

any number of trivial harms, like mild transient headaches. Another attempt to accommodate the intuition in *Life versus Headaches* is to argue that we should accept full aggregation while adding that value functions are bounded. In such bounded value functions, the aggregate value of any number of a given harm has an upper bound. As the number of headaches approaches infinity, the value of saving these people approaches a fixed value lower than the value of saving a single person from death.⁵ Such views imply, implausibly, that our reasons for saving an additional person from serious harm diminish with the number of affected people. At some large number n , we have virtually no reason whatsoever to save additional people from serious harm. Such a view would therefore imply that we should rather save n people from a severe disability alongside one person with a headache instead of saving $n+1$ persons from a severe disability.

The better solution is, therefore, to adopt the idea that there are different kinds of claims. There is something that distinguishes headaches from deaths in a manner that bars us from trading off lives against headache relief. The idea of relevance can help here.⁶ Claims to headache relief are not relevant to claims to be saved from death. The same idea can explain why the numbers should count when the harms are more similar. Consider *Life versus Paraplegia*. If we save whichever group has the strongest individual claim, disregarding the numbers entirely, then we should save a single person from death regardless of the number of people that we could save from paraplegia.⁷ However, *Life versus Paraplegia* is different from *Life versus Headaches* insofar as the claims to be saved from paraplegia are, plausibly, relevant to the claims to be saved from death.

⁵ See Seth Lazar and Chad Lee-Stronach, *Axiological Absolutism and Risk*, *Noûs* 53 (2019), pp. 97-113. For objections similar to the one I raise see Otsuka, *Saving Lives, Moral Theory, and the Claims of Individuals*, p. 127 fn. 31 and Alex Voorhoeve, *Balancing small against large burdens*, *Behavioural Public Policy* 2 (2018), pp. 125-142 (pp. 132-34).

⁶ See Kamm, *Morality, Mortality*, vol. 1, chs. 8-10 and Lefkowitz, *On the Concept of a Morally Relevant Harm*.

⁷ Throughout the paper, I am using the term 'group' liberally in a manner akin to the use of 'set' and sometimes refer to single individuals as a group.

While the idea of relevance can give us a principle that can explain our intuitions in cases like *Life versus Paraplegia* and *Life versus Headaches*, the idea is incomplete in two ways. First, it is unclear how this idea can plausibly be extended to more complicated cases including those involving heterogeneous groups. Second, even if we had a decision procedure for such cases the question remains how to theoretically justify this decision procedure. While intuitive fit is an important part of a good moral theory, we also need to provide an underlying justification for the procedure. Otherwise our theory merely summarizes rather than justifies our immediate reactions about cases. This is the task I set out to do in the next section.

III. Justifying Hybrid Balance Relevant Claims

The starting point for my view is an idea that appears in the work of Thomas Nagel. Nagel writes about the reconciliation of two standpoints, the impartial and the partial standpoint. Impartial concern is, however, non-aggregative and should not be confused with impersonality which is aggregative.⁸ Impartiality is based on the recognition that everyone's life, including one's own, has objective importance and significance. Realizing this, we extend this significance to the lives of others. We imagine ourselves to be in their shoes and extend an impartial concern for them. This impartial concern is fragmented. Unlike fully aggregative theories that interpret impartiality as impersonality, we are not eroding the distinction between different viewpoints. This fragmented concern takes seriously the separateness of persons.⁹ The realization of the objective significance

⁸ Nagel, *Mortal Questions*, ch. 8, and *Equality and Partiality* (Oxford, 1991), chs. 2-8. Unlike me, Nagel sometimes speaks of impersonal concern as interchangeable with impartial concern. Impersonality might be one way to show impartiality but not the only one. Respect for the separateness of persons dictates that we should not equate impartiality with impersonality. For the contrast see Rawls, *A Theory of Justice*, pp. 165-68.

⁹ Nagel, *Mortal Questions*, pp. 126-27.

of everyone's life need not, however, lead us to abandon our own personal perspective in the world. Impartial concern goes along with legitimate partial concern for oneself.

The ideal that we are striving towards is unanimity. It is not unanimity in our rational self-interest, but rather unanimity among persons committed to finding common principles guiding our interactions. Our actions should be justifiable to each and every one who is affected. The kind of unanimity we are searching for is unanimity that is achieved by convergence from different standpoints.¹⁰

The two different perspectives, partial and impartial, can explain which claims are relevant.¹¹ From our first-personal perspective we can have a justified stronger concern for our own life than for the lives of others. When we imagine ourselves to be in the position of others, such imagination includes their self-favoring concerns. While individuals are entitled not to make use of their self-favoring concern, we cannot assume that individuals have volunteered to waive their moral claim to be aided. But in the absence of any such waiver, we need to consider the full strength with which individuals are entitled to press their claims. The maximum extent of such permissible concern thus determines when claims are relevant. A claim is relevant to another claim if and only if this claim can be preferred to the other claim from someone's point of view. A claim might be weaker and still relevant if the claim can be preferred from someone's self-favoring, partial perspective. For example, plausibly a person is not required to forgo saving herself from paraplegia to save another person from death. If this is the case, then it means that the claim to be saved from paraplegia can be preferred from this person's partial perspective. Hence, claims to be saved from paraplegia are relevant to claims to be saved from death.

The relevance test that invokes the partial perspective has the upshot that in a case where the claims of only one of the groups are relevant, unanimity naturally emerges. We can see this

¹⁰ Nagel, *Equality and Partiality*, pp. 33-40.

¹¹ I borrow this argument from Voorhoeve, *How Should We Aggregate Competing Claims?*.

with our paradigmatic case *Life versus Headaches*. From the perspective of the person about to lose her life, she should be favored. Both her impartial and partial concern favor this. From the perspective of each of the persons about to suffer a headache, the rescue of the life should be favored too. Their partial concern does not extend to the saving of themselves from a minor headache rather than another person's life. It would be unreasonable for these people to insist on their claim to be rescued. We can give a powerful justification to them for not rescuing them. Not even you with your partial concern would have been allowed to rescue yourself. So, how can you complain to me who does not have partial concern for you, for not rescuing you?¹²

With the theoretical framework in place, I can now turn to those cases which are central to my argument for a new view on limited aggregation and which are a challenge for limited aggregation. Let us start with cases in which claims of equal strength are at play, but the relative numbers differ. For example, we can save either A or B&C from equal harm. In this case every person has a claim that survives the test of partial concern. This means that we have a conflict of different standpoints. How should we resolve this conflict of standpoints? A leading proposal is *Aggregate Relevant Claims* according to which aggregating claims is allowed only in cases of such a conflict of standpoints. Only relevant claims can be aggregated.¹³ *Aggregate Relevant Claims* mediates between an aggregative approach and a nonaggregative approach to morality. Nonaggregation reigns in cases in which all relevant claims agree which group should be saved.

¹² Victor Tadros has provided a critique of this test for determining when claims are relevant which some philosophers writing on aggregation have found to be convincing. I, myself, remain unconvinced by Tadros's challenge, but addressing this disagreement would take me too far afield from my aim in this paper. I hope to engage with Tadros's challenge at another occasion. See Victor Tadros, *Localized Restricted Aggregation*, *Oxford Studies in Political Philosophy*. vol. 5, ed. David Sobel, Peter Vallentyne, and Steven Wall (Oxford: Oxford University Press, 2019), pp. 171-204, (pp. 176-78). Tadros's objection is an important reason for the development of Alec Walen's alternative view on aggregation that he dubs 'weak aggregation' in his *Risks and Weak Aggregation: Why Different Models of Risk Suit Different Types of Cases*, *Ethics* (forthcoming).

¹³ Voorhoeve, *How Should We Aggregate Competing Claims*.

Aggregation reigns in cases in which there are relevant claims for both sides. It is this feature of Aggregate Relevant Claims that makes it susceptible to Tomlin's criticism.¹⁴ When Tomlin extends Aggregate Relevant Claims to heterogeneous groups he asks how to identify which claims are relevant to see whether we should follow a nonaggregative or aggregative approach in the case.

I therefore endorse a different method that I call *Balance Relevant Claims* which does not distinguish at the outset between "aggregation cases" and "nonaggregation cases". Balance Relevant Claims resolves conflicts of standpoints by sequentially balancing individual claims against one another. Consider the following pair of cases. In the first case we can either save A or B, in the second case we can either save A or B&C. While in the first case the considerations for saving A and the considerations for saving B are equally strong, this is not the case in the second case. The fact that C is a third person means that C's claim breaks the tie and thus we should save B&C. We can explain the tie-breaking idea in terms of balancing claims. When we see that we can save A from death, we notice a strong claim to our assistance. If A's claim were the only thing to consider, we would be required to save A. But B's presence and B's claim balances the moral claim that A can raise. B's claim is just as strong as A's. Neither claim can ultimately determine what we ought to do. This is how I understand the term "balancing". If one claim is balanced, then this claim loses some of its moral force. It cannot ultimately determine that we should save the group to which the claim belongs. But the claim is not canceled; it retains some moral force.¹⁵ In the case of A and B this means that since we are required to save someone, saving either group is permissible, or perhaps we are required to give equal chances to both groups. By contrast, when

¹⁴ Tomlin, On Limited Aggregation.

¹⁵ In the first version of her argument for a balancing view Frances Kamm did speak about canceling (Frances Myrna Kamm, Equal Treatment and Equal Chances, *Philosophy & Public Affairs* 14 (1985), pp. 177-94). Later on, Kamm admits that the canceling metaphor is misleading (*Morality, Mortality*, vol. 1, pp. 116-17). Rahul Kumar used a neutralizing metaphor in later work (Contractualism on saving the many, *Analysis* 61 (2001), pp. 165-70). Kumar's argument was criticized on similar grounds to the ones presented here by Otsuka (Saving Lives, Moral Theory, and the Claims of Individuals, pp. 118-19).

C is also present, C's claim is not balanced. C's claim can then have the power to ultimately determine what we ought to do: namely, save B&C.¹⁶

The tie-breaking idea helps us to better understand how Balance Relevant Claims works, but it is not a good guide to justifying it. One unsuccessful justification for Balance Relevant Claims involves an appeal to the moral complaint that if we are not required to save B&C, then the additional presence of C does not make a moral difference. However, the additional presence of C can make a moral difference in other ways. For example, a weighted lottery that reflects the different numbers would ensure that the additional presence of C makes a moral difference by shifting the odds.¹⁷ For this reason, we should not rest our case for Balance Relevant Claims on the idea that tie-breaking is the only way to respect the fact of C's additional presence. Instead, Balance Relevant Claims is justified holistically. My aim in the following discussion is to make clear

¹⁶ The idea of tie-breaking brings out how balancing differs subtly from ordinary aggregation. We can see this with the help of the following analogy. There are two ways in mathematics to determine whether one set is larger than the other. One way counts the members of the set and then compares the number of elements in the set. Here the cardinal number, or sum-total, matters. Another way of comparing the size of sets does not require any numeracy skills or even knowledge of numbers. We can see whether there is a bijection, a one-to-one correspondence mapping between all elements of the two sets. In this method we only need to map individual members against one another. Balancing is like this second approach. In this sense, balancing employs a similar form of reasoning as the anonymous Pareto principle. B&C is anonymously Pareto superior to A because while B's claim can be matched by A's claim, C's claim cannot be matched. For an argument that Paretian reasoning departs from aggregative reasoning see Iwao Hirose, Saving the greater number without combining claims, *Analysis* 61 (2001), pp. 341-43.

¹⁷ Both Kamm and Scanlon rest their argument for views similar to Balance Relevant Claims on this moral difference argument. See Kamm, *Morality, Morality*, vol. 1, pp. 101, 114-19, *Intricate Ethics*, p. 33; and Scanlon, *What We Owe to Each Other*, pp. 231-35. The objection that one can make a moral difference in other ways is due to Michael Otsuka (Saving Lives, Moral Theory, and the Claims of Individuals, pp. 114-18). David Wasserman and Alan Strudler helpfully distinguish between what they call the marginal difference argument and the balancing argument. Kamm and Scanlon think that the marginal difference argument grounds the balancing argument. My point here is that we can retain balancing by giving it fresh foundations. David Wasserman and Alan Strudler, Can a Nonconsequentialist Count Lives?, *Philosophy & Public Affairs* 31 (2003), pp. 71-94 (pp. 82-89).

that Balance Relevant Claims can be developed into an overall compelling theory of limited aggregation.

The case of balancing highlights and expresses a different model of unanimity. In the case of *Life versus Headaches* we achieved unanimity by individuals withdrawing their claims because their partial perspective did not allow them to stake their claim. In the present case all claims are considered and weighed. We engage in a genuine confrontation of standpoints. But we can still say the following. It would be unreasonable for A to insist that his claim ultimately decides that she should be saved in the case of A versus B&C. Her claim was considered and balanced with the claim of someone else. She cannot insist that her claim has greater weight than the weight of balancing a single claim. She can only insist on a fair decision procedure that takes her moral claims and the severity of her plight into consideration. Balancing meets this demand. But she cannot insist on more. It would be unreasonable for her to insist on a particular outcome in which she is saved and to reject a principle requiring B&C to be saved.¹⁸

If we accept the balancing of claims, the question arises whether we can extend the idea of balancing to cases that are different from tie-breaking cases. For example, can several claims to be saved from paraplegia balance a single claim to be saved from death? It seems plausible that sometimes a greater number of weaker claims can balance a lesser number of stronger claims. Having a relevant claim means that one can maintain one's claim and insist that one's claim be considered. All cases involving opposing relevant claims thereby become cases involving a conflict

¹⁸ Might A insist on a weighted lottery that gives her a one third chance to be saved? I think not, since I believe a weighted lottery involves a failed attempt to reconcile two different approaches, which ends up abandoning a commitment to equal claims shared by these approaches. One way to reflect the equal claim of all three is to find a procedure for the distribution of the good that reflects this equality. This is what balancing does. Another way to reflect the equal claim of all three is to give everyone an equal chance of receiving the good. But a weighted lottery does not do this since both B and C will receive a two thirds chance of being saved. Insofar as we are moved by this model of fairness, we should be holding out for the equal chance of a fair coin toss as Taurek (Should the Numbers Count?) advocates. Of course, this model of fairness is disputed by limited aggregation.

of standpoints. While all can reasonably insist that their claims are considered, no one can insist that their claim has to be decisive. Balancing is a general method that allows us to resolve cases of such conflicts of standpoints. If one claim is outweighed by a multitude of weaker claims, this person's claim was outweighed by claims of people who were entitled to voice their claims and have their claims considered. If one has the strongest claim, but it has been considered and was outweighed, it is unreasonable for one to insist that one's claim has to have absolute priority over all weaker claims. If we accepted the absolute priority of stronger over weaker claims, there would be little point in distinguishing between relevant and irrelevant claims. However, the view that I advance here draws a distinction between cases in which claims cannot be staked because they are irrelevant, and cases in which claims are considered and outweighed.

The next extension of the idea of balancing concerns cases with heterogeneous groups. In homogenous group cases like *Life versus Headaches* or *Life versus Paraplegia*, the question of which claims are relevant and the order in which claims are balanced are straightforward. Either the claims that compete with the claims to be saved from death are relevant to the claim to be saved from death, or they are not. But now consider *Life and Broken Legs versus Paraplegia and Headaches*. How should we decide which claims are relevant in such cases? Are the claims against headaches relevant if they are relevant to claims against broken legs? Or are they irrelevant because they are irrelevant to claims to be saved from death? And how should we decide in which order to balance claims against one another? The key to thinking about these cases is, I believe, to apply the same principle that explains the difference between *Life versus Headaches* and *Life versus Paraplegia* to solve the questions of relevance and order.

Not saving a person with a strong claim to be saved requires a special justification to this person. Given that the person with the strongest claim is most likely to have grounds for grievance or complaint against our failure to save her, our justification for acting as we do must be primarily addressed to those who have the strongest claim on our aid. This gives an answer to the question of the order of balancing. It explains why balancing starts with the strongest claim and then works

its way down to less strong claims. The stronger the claim the greater is the urgency to give a justification to this person.

The same idea can also illuminate when claims are relevant. The justification we can give to this person has to respect the strength of the claim of the person to assistance. In *Life versus Paraplegia*, we can point out to the person dying that the claims of the many are all relevant to her plight. The other individuals are entitled to stake their claims on us. By contrast, it would be disrespectful to deny saving a person from death by pointing out that doing so would enable the agent to prevent some number of headaches. To consider the headaches as a reason not to save the person from death would trivialize her situation. The headaches are simply not relevant in comparison to the death.¹⁹ To say to the person: “I am sorry, but we cannot save your life because we are busy preventing many minor headaches” would trivialize what she stands to lose. This is not the case when we say: “I am sorry, but we cannot save your life because we are busy preventing many people from becoming paraplegic”. The “because” clause in the justification indicates that whether or not a justification is respectful depends on for whose sake we fail to save the person from death. This suggests the following general principle:

Respectful Failure to Save Principle. Every person that we fail to save is entitled to a respectful justification for our failure of save. It is disrespectful and impermissible to fail to save a person with a strong claim for the sake of persons whose claims are irrelevant to this strong claim.

The Respectful Failure to Save Principle tells us that what should guide our thinking about whether or not claims are relevant is determined by considering whether we can give a respectful justification to the person whom we fail to save. Applied to the case of balancing the claims of

¹⁹ For this see Kamm, *Intricate Ethics*, pp. 297-98, 484-87; Lefkowitz, On the Concept of a Morally Relevant Harm, pp. 421-23; and Voorhoeve, Why One Should Count Only Claims with which One Can Sympathize, pp. 152-53.

heterogeneous groups, this principle identifies two scenarios in which counting claims as relevant renders us unable to respectfully justify our refusal to save.

The first scenario concerns a local feature of the confrontation of two claims. Whether or not we can justify the balancing of claims will depend on the relation between the claim that is balanced and the balancing claims. These claims will need to be relevant to one another in order for the balancing to be respectful. Balancing a claim means depriving it of the force to ultimately and unilaterally decide what one ought to do. The claim still has force in deciding what ought to be done; in this sense it is not canceled. But it loses the force to ultimately and unilaterally “tip the balance”.²⁰ This is quite serious and deserves justification. We cannot justify it to a person that the moral force of her claim is diminished because of claims that are irrelevant to her claim. It is part of the meaning of an irrelevant claim that if a claim is irrelevant to claim X, then it can neither diminish nor override claim X.

What this means is that in each case of balancing, the claim must be relevant to the claim that is being balanced. However, what about claims that are relevant in this sense and would be ultimately decisive? Suppose, for the sake of argument, that claims to be saved from minor headaches would ultimately determine that we ought not save a person from death. This second scenario highlights the existence of what Frances Kamm has called “irrelevant utilities”.²¹ To consider a trivial harm like a headache or a sore throat as the ultimate reason to save a group would be trivializing the fate of a person whose life will not be saved and thus be disrespectful to her. The headache or sore throat is not relevant for this decision. This type of disrespect that flows

²⁰ Balancing, as I understand it here, is *not* a symmetric relation. For example, if A’s life is at stake and B and C have a claim against a severe impairment, then it is plausible that we should save A. In this case A would have balanced B and C, but B and C would not have balanced A. However, when two groups (e.g. A versus B suffering from equal conditions) that have equal claims oppose each other, then both groups balance each other. Thus, balancing is neither symmetric nor asymmetric.

²¹ Kamm, *Morality, Mortality*, vol. 1, pp. 144-64.

from the Respectful Failure to Save Principle is rather a global feature of the entire trade-off. The strongest claim in the opposing group fixes globally that trivial harms do not count.

This argument leads to the following two conditions for determining the relevance of claims:

(1) *The local relevance condition*: A claim can only be balanced with another claim if the two claims are relevant to one another.

(2) *The global relevance condition*: Every individual with a strong claim has a veto against the consideration of any type of claim that is irrelevant to her claim, if such consideration would lead to her claim not being satisfied (e.g. by not being saved).

The first scenario illustrates the need for a local relevance condition, while the second scenario illustrates the need for a global relevance condition. One concern about the hybrid character of my view is that it is gerrymandered and can explain our intuitions about cases only because it relies on divergent conditions that cannot be coherently defended. In response, I have argued that both conditions originate from a single overarching principle. This principle illuminates the idea of respect that proponents of limited aggregation have relied upon. The fact that there are two conditions is only the reflection of the fact that disrespect can manifest itself in various ways. I outlined two scenarios in which it can be disrespectful to fail to save a person, these two scenarios correspond to the two conditions of my hybrid view. A related concern is that the combination of the two conditions does not fit well the model of balancing. The concern is that the conditions license a form of double counting. Individuals with a claim that is balanced can nonetheless exercise their veto against the consideration of other claims. This claim would have balanced some claims *and* vetoed other claims. In response, we can see more clearly that this is not double counting by recalling the distinction between balancing and canceling. Balancing means that a claim is only deprived of some of its moral force and not fully canceled. In a confrontation between two equally strong claimants, each of the claims will be balanced yet nonetheless we are required to save one of them. These balanced claims retain some moral force, which is revealed by the fact that we cannot escape our duty to save on account of the fact that both are equally deserving of

our aid. The reason why a balanced claim retains this power is that failing to save anyone in the case of equally balanced claims would fail to respect their moral standing as persons who deserve to be saved. Similarly, without the veto condition, a decision-maker would be licensed to fail to save the person in a disrespectful manner. A person’s claim can never be deprived of the force to insist on respectful treatment.

IV. Illustrations of Hybrid Balance Relevant Claims

The view that I defend is complex. I will help to illustrate it with three examples. We can imagine the examples to be decisions about the allocation of scarce medical resources that can be used to save people from permanent medical conditions and restore them to full health. The medical conditions are specified by broad categories of severe impairments, moderate impairments, and mild impairments. While the mild impairments are relevant to the moderate impairments and the moderate impairments are relevant to the severe impairments, mild impairments are not relevant to severe impairments. For the sake of my illustrations, we can assume that ten claims against a moderate impairment are equal in strength to one claim against a severe impairment. Likewise, ten claims against a mild impairment are equal in strength to one claim against a moderate impairment. In my examples I illustrate claims that are balanced on either side by bracketing them.

Group A	Group B
Case One	
(1 Severe) 1 Moderate	(10 Moderate) 11 Mild

In Case One the claim to be saved from a severe impairment is balanced by the ten claims to be saved from a moderate impairment. If it was not for my global relevance condition, the claim against moderate impairment in favor of Group A could be outweighed by the claims against mild impairment. However, the global relevance condition blocks this. It would be disrespectful to the

person with the severe impairment not to save her for the sake of people afflicted with a mild impairment. Considering the mild impairments vitiates a respectful justification that we can give to the person losing out.

Group A	Group B
Case Two	
(1 Severe) (10 Mild) + 1 Mild	(10 Moderate) + (1 Moderate)

Case Two illustrates when claims that are irrelevant to the strongest overall claim nonetheless remain relevant. Here the claims against severe impairment are balanced by ten claims against a moderate impairment. The moderate impairment would then indicate that we should save Group B. But B's claim can be balanced in favor of A by taking into consideration the claims of the mildly impaired. Ignoring the claims of the mildly impaired here would not be disrespectful towards the one with the severe impairment. On the contrary, it is in her interest that these claims are to be counted.

Group A	Group B
Case Three	
(1 Severe) (10 Mild) + (1 Mild)	(10 Moderate) + (1 Moderate) (1 Mild) + 1 Mild

But what about Case Three then? Here the last unbalanced claim is one of a mild impairment and this claim advocates not saving the group with the most serious claim of severe impairment. Can we justify this to the person with the severe impairment? I believe we can. Had we not counted the persons with mild impairments in this case, we should still save Group B. The claims in Group B to be saved from a moderate impairment would outweigh the claim of the single person with a severe impairment in A. The person with the severe impairment cannot complain that we do not save her because of the claims to be freed from the mild impairment. Had we not counted these we would still not be permitted to save her. If only claims against severe and

moderate impairment count, we ought to save B. Accordingly, my global relevance condition does not block considering the claims of mild impairment in this case. There is no point for the person with the severe impairment to exercise her veto, since she would not be saved even if we fail to consider the claims of mild impairment. Consequently, we can justify our failure to save to the person with the severe impairment. Counting the persons with mild impairments does not vitiate a justification that we can give to this person.

Cases One and Two illustrate that the hybrid view has a certain asymmetry. We consider claims of mild impairment when they favor the person with the strongest claim, but not when they oppose the strongest claim. This asymmetry may seem suspect. The asymmetry does not, however, violate the principle that Tomlin has called Equal Consideration for Equal Claims. According to Equal Consideration for Equal Claims, we should give equal weight to claims of equal strength. My view does this. Whenever some claims of a certain relevance class become relevant, all claims of that class become relevant. If, in Case Two, there were other claims against mild impairments on the side of Group B, these would equally have to be counted. For every moral choice, it is either the case that the claims of a given class are relevant or irrelevant.

The asymmetry is more modest than a violation of Equal Consideration for Equal Claims in that the relevance of less strong claims depends, *inter alia*, on which group has the strongest claim in its favor. This modest asymmetry can be defended. My previous argument that illustrates the different justifications we can give to the person with the strongest claim does just this. Counting mild impairments in Case One would make the failure to save the person with the severe impairment disrespectful. Counting the mild impairment in Case Two, however, would not be disrespectful to the persons in Group B whom we would fail to save. Further, for the reasons I outlined above, counting mild impairments in Case Three does not vitiate the justification we can give to the person with a severe impairment. If this explanation succeeds, then the asymmetry is justified.

One other feature may seem to violate Equal Consideration for Equal Claims. Consider the following two cases.

Group A	Group B
Case Two	
(1 Severe) (10 Mild) + 1 Mild	(10 Moderate) + (1 Moderate)
Case Two*	
(1 Severe) + (1 Severe) 11 Mild	(1 Severe) (10 Moderate) + 1 Moderate

In Case Two, we ought to save Group A. The claims against mild impairment are locally relevant to the claims of moderate impairment in Group B. In Case Two*, we ought to save Group B. The claims against mild impairment are irrelevant according to my global relevance condition because of the presence of a claim against severe impairment in Group B. Even though we added equal claims to both sides, our verdict of permissibility changes. This feature does not violate Equal Consideration for Equal Claims. Both claims against severe impairment are equally considered. Indeed, it is because they are considered equally that claims against mild impairment become irrelevant on either side.

While it might seem counterintuitive that the addition of equally strong claims should change the permissibility of our decision, there is a rationale for this. With the addition of the claim against severe impairment in Group B something important changes. We can no longer justify considering the claims against mild impairment as ultimately decisive. In Case Two this was plausible since doing so would not result in a person not being saved whose claims are in a different class of relevance. But this changes in Case Two*. This change should lead us to consider the case differently and accept the verdict that Group B ought to be saved. The addition of the same claim on both sides can therefore change what one ought to do.

V. Hybrid Balance Relevant Claims and Objections to Limited Aggregation

So far, I have defended Hybrid Balance Relevant Claims as a theory of limited aggregation that can be applied both to homogeneous group cases and to heterogeneous group cases. In the remainder of the article I will show how the hybrid character of my view escapes challenges that other versions of limited aggregation face. I will start with challenges to views that rely only on a global relevance condition before moving on to challenges against purely local relevance. In a third step, I discuss and reject one challenge which all views of limited aggregation face.

A. *Problems with Global Relevance*

Consider first Tomlin's challenge to limited aggregation. Tomlin has pointed out that the idea of relevance in Aggregate Relevant Claims is ambiguous between two interpretations. Are claims relevant if they are relevant to the strongest claim with which they compete? Or are claims relevant if they are relevant to the strongest claim overall? Tomlin argues that on either of the two interpretations Aggregate Relevant Claims has deeply counterintuitive implications and violates plausible moral principles. Tomlin calls this the Anchoring Problem. The two alternatives that Tomlin outlines are both forms of global relevance that tell us that claims either are or are not relevant in a given choice situation. Since my view is a hybrid view that incorporates both a global and a local relevance condition, it avoids the Anchoring Problem.

The first possibility is that claims are relevant when they are relevant to the strongest claim with which they compete (*Anchor by Competition*).²² Tomlin provides the following counterexample.

²² Tomlin, *On Limited Aggregation*, pp. 240-44.

Group A	Group B
Anchor by Competition Case One	
1 Severe	10 Moderate
Anchor by Competition Case Two	
1 Severe	10 Moderate
1 Mild	1 Mild

This application of Anchor by Competition illustrates a violation of Equal Consideration for Equal Claims. For Anchor by Competition the mild impairments in Group A are counted while mild impairments in Group B are not. By contrast, my view resolves Case Two differently. In Case Two my global relevance condition treats the mild impairments as irrelevant utilities. This means that in the second case we are still in a tie. The global relevance condition that I adopt respects Equal Consideration for Equal Claims. If a person with a strong claim can veto other claims, the person vetoes all claims of the same type. Thus, claims are always either relevant or irrelevant in a given situation.

Consider now Tomlin's second alternative (*Anchor by Strength*).²³ This alternative holds that relevance is determined by reference to the strongest overall claim, regardless of which side this claim favors. The following example illustrates this.

Group A	Group B
Anchor by Strength Case One	
111 Mild	11 Moderate
Anchor by Strength Case Two	
1 Severe	11 Moderate
111 Mild	

²³ Tomlin, *On Limited Aggregation*, pp. 244-47.

Here is how Anchor by Strength resolves these cases: In Case One, the claims against mild impairment are relevant since the strongest overall claim is against moderate impairment. The claims against mild impairment can therefore outweigh the claims against moderate impairment. In Case Two, however, the claims against mild impairment are no longer relevant to the strongest overall claim according to Anchor by Strength. This makes it the case that the claims against moderate impairment can now outweigh the claim against severe impairment. Group B ought to be saved even though Group A has now one additional strong claim in its favor. This violates the Principle of Addition since Group A is less choiceworthy even though there is an additional claim in Group A present.

My hybrid view on the other hand allows that claims that are not relevant to the strongest claim can be relevant in balancing claims *provided that it is not to the disadvantage of any person with the strongest claim*. In the counterexample to Anchor by Strength it is in the interest of the person with the strongest claim that the weaker claims are counted. The person with the strongest claim cannot complain that considering these claims would be disrespectful to her. These claims are her “allies” and according to my view a strong claim cannot lose its “allies” by rendering them globally irrelevant. This means that the hybrid view, unlike Anchor by Strength, fulfils Tomlin’s Principle of Addition.

Since my view implies neither of the two problematic case judgments and also does not violate either of the two principles that Tomlin proposes, it avoids the Anchoring Problem. As I have just shown, part of the reason for this is the acceptance of some form of local relevance condition. Tomlin himself, together with Aart van Gils, has proposed a view that embodies a form of local relevance as a possible response to the Anchoring Problem.²⁴

²⁴ Aart van Gils and Patrick Tomlin, Relevance Rides Again? Aggregation and Local Relevance, *Oxford Studies in Political Philosophy*, vol. 6, ed. David Sobel, Peter Vallentyne, and Steven Wall (Oxford, 2020), pp. 221-55. Already in On Limited Aggregation Tomlin tentatively suggests a view of local relevance (pp. 259-

B. *Problems with Local Relevance*

Unlike Tomlin's (and van Gils's) own solution to the Anchoring Problem, I believe that while we should incorporate local relevance, we should not give up on global relevance altogether. To see why the hybrid character of my view is important consider Kamm's Sore Throat Case. In Kamm's case we have the choice between saving one life and saving another life alongside saving a person from a sore throat. But this additional sore throat should not tip the balance and make us to save the second person.²⁵ Not all versions of limited aggregation are able to accommodate this case. In particular, this case presents a challenge to versions of Balance Relevant Claims that lack a global relevance condition. Such views cannot tell us why it is impermissible to count the claim of the sore throat.

My view that combines a local and global relevance condition can do so. In the Sore Throat Case, the claim to be saved from a sore throat is globally irrelevant. It should not feature in our deliberation. The two claims to be saved from death would be evenly balanced against one another. Neither of the two claims retains its force in ultimately deciding what we ought to do. We are left with a tie. Either we have to give equal chances, or it is permissible to save either group.

The global relevance condition also allows my view to avoid a similar problem.²⁶ Imagine there is one claim against death in Group A that is balanced by enough claims in Group B so that a single claim against severe impairment remains. This claim against severe impairment is then balanced by enough claims in Group A so that a single claim against moderate impairment remains.

60). He credits Victor Tadros and Garrett Cullity for the idea. For Tadros's version see his *Localized Restricted Aggregation*.

²⁵ Kamm, *Morality, Mortality*, vol. 1, pp. 101-2, 146-47, *Intricate Ethics*, p. 34.

²⁶ Van Gils and Tomlin, *Relevance Rides Again?*, pp. 242-44. Kamm's own treatment of irrelevant utilities is susceptible to this related problem. See Kamm, *Intricate Ethics*, pp. 63-64.

This claim in turn is then balanced, and so on. In the end a single claim against a trivial inconvenience, like a sore throat, remains. If we did not accept a global relevance condition, we would have to accept that a sore throat could become decisive in this scenario. But it seems implausible that it really should. By accepting a global relevance condition, my view can appealingly explain why the sore throat should not be decisive. It would be disrespectful not to save a person from death (or severe impairment for that matter) because of the existence of a sore throat.

While my view incorporates and explains the intuition in the Sore Throat case, accepting that the sore throat makes a difference may not be a decisive objection for a proponent of purely local relevance. Some philosophers writing on aggregation have expressed doubt whether we should retain the intuition that it is permissible to save either of the two persons in Kamm’s case and believe that trivial harms can rightly be mandatory tie-breakers.²⁷ Versions of Balance Relevant Claims that contain only a local relevance condition are, however, subject to a different objection that my view avoids. Joe Horton provides this excellent criticism against views of purely local relevance. He devises the following case with two stages.²⁸

Group A	Group B
Local Relevance Stage One	
1 Severe	1,000 Mild
Local Relevance Stage Two	
1 Severe 11 Moderate	11 Moderate 1,000 Mild

²⁷ See e.g. Campbell Brown, Is close enough good enough?, *Economics & Philosophy* 36 (2020), pp. 29-59 (pp. 41-42); van Gils and Tomlin, Relevance Rides Again?, pp. 231-42; and Korbinian Rürger, Aggregation with Constraints, *Utilitas* (forthcoming).

²⁸ Horton, Always Aggregate, pp. 168-71. I changed the precise example so that it fits the stipulations about relevance that I have used throughout the paper.

At Stage One, we should save the person from the severe impairment. What about Stage Two? A form of Balance Relevant Claims with only a local relevance condition would resolve it as follows: The claim against severe impairment in Group A can be balanced by ten claims against moderate impairment in Group B. The remaining claim against moderate impairment in Group B can be balanced against one claim against moderate impairment in Group A. The remaining ten claims against moderate impairment can also be balanced and even outweighed by claims against mild impairment in Group B. Hence, we ought to save Group B.

Intuitively, the equal addition of claims should not make a difference in this case. My view can explain why in an intuitive fashion. The reason why we ought to save Group A in Stage Two as well is that the person with the severe impairment can complain and veto our failing to rescue her for the sake of people with mild impairments. This reason also sets this case apart from cases where the equal addition of claims does make a difference. In these cases, the addition of claims changes the justifications available to a group for not saving them. In the present case, however, the added claims do not enable us to give a different justification to the person with the severe impairment for failing to save her.

Horton explains that the failure of a view that only contains a local relevance condition is that it can allow for irrelevant claims to be “activated” in a manner that they can help outweighing the strongest claim. If there are any claims against moderate impairment in Group A, then, under the view Horton criticizes, it is always possible to balance the claims against moderate impairment in Group A against the claims against mild impairment in Group B.²⁹ The idea is that introducing claims against moderate impairment, even in Group A, makes mild impairments relevant, so that they can contribute to our reasons for saving Group B. The claims against mild impairment could here outweigh the remaining claims against moderate impairment. But this is what my global relevance condition blocks. The person with a severe impairment has a veto against claims of mild

²⁹ Horton, *Always Aggregate*, pp. 171-73.

impairment being considered. Mild impairments are treated as irrelevant utilities. My hybrid view therefore avoids the problem that Horton raises for views of pure local relevance.

C. Principle of Agglomeration

Horton, in his article, also raises another objection against limited aggregation. He describes the possibility that there are two separate moral decisions with one group we are required to save. But once we join the two decisions, it is possible that the group composed of those we ought to have saved no longer ought to be saved.³⁰ Horton might appeal here to something like the Principle of Agglomeration according to which combining groups that ought to be saved cannot render them a group that ought not to be saved.

The following case illustrates the problem.

Group A	Group B
Combined Case	
(1 Severe) (10 Mild) + (200 Mild)	(10 Moderate) + (1 Moderate) (200 Mild) + 800 Mild
Sub-Case One	
1 Severe	1,000 Mild
Sub-Case Two	
(110 Mild) + 100 Mild	(11 Moderate)

In both Sub-Cases we ought to save the sub-set of A, but in the combined case we ought to save the combination of the sub-sets of B. Notice, however, the following. We could also divide the case in the following sub-sets.

³⁰ See Horton, *Always Aggregate*, p. 173.

Group A	Group B
Sub-Case Three	
(1 Severe)	(10 Moderate) + 1 Moderate
Sub-Case Four	
(210 Mild)	(210 Mild) + 790 Mild

In these Sub-Cases it is obvious that any limited aggregation view would have to select B in both cases. What does it tell us? First, any view on limited aggregation needs to decide the four sub-cases the way I suggested. They follow straightforwardly from the stipulations I have made about relevance. If we accept the Principle of Agglomeration, we ought to save Group A because Sub-Case One and Two tell us to save sub-groups of A. But if we accept the Principle of Agglomeration, we also ought to save Group B because Sub-Case Three and Four tell us to save sub-groups of B. Yet we cannot save both A and B. In other words, *any* view of limited aggregation has to violate the Principle of Agglomeration.

Agglomeration may seem appealing at first blush. What difference can *mere composition* really make? Should it really matter in which groups we are finding people assembled? Agglomeration shares one feature with the independence of irrelevant alternatives. Of course, if alternatives really are irrelevant then they should not count. The problem is that sometimes alternatives can be relevant even if they do not appear to be at first sight. Whether or not one option is on the table influences, for example, the kinds of justifications we can give to those we do not save. Similarly, composition can be relevant if it is not *mere* composition but something else that changes.

A successful agglomeration argument should be one where mere scaling up results in a reversal of permissibility. Such an argument would object that a theory holds that it is permissible to act in a certain way towards A and B separately, but impermissible to act in this same way towards A&B combined.³¹ This agglomeration argument differs crucially from Horton's. It is

³¹ For an example of such an agglomeration argument (against ex post contractualism) see my Contractualism, Complaints, and Risk (unpublished manuscript).

symmetrical in the sense that everyone is equally affected in the sub-groups. Horton's agglomeration cases do not have this feature. They crucially rely on the fact that the groups that are combined are not symmetrical. This makes it plausible that it is not *mere* composition that changes. Joining the groups together changes the moral situation by changing the relations of individuals towards one another. For example, the existence of persons with moderate impairments enables us to give the person with a severe impairment a justification we were previously unable to give. These special relations between different persons only arise in the particular configuration of one case. They do not necessarily hold between sub-sets of these persons. Since limited aggregation is predicated on the idea that the moral relations of individuals to one another are important, we should not be surprised that changes in the relations of individuals will change permissibility verdicts. We can give a satisfactory explanation why agglomeration is not a mere change in composition.³²

VI. Conclusion

If we adopt Hybrid Balance Relevant Claims we can respond to the recent criticism of limited aggregation. Hybrid Balance Relevant Claims refines our ideas about which claims should be treated as relevant and which ones should not, as well as our ideas about how to count claims. The hybrid view it is motivated by a common motivation that our decisions whom to save should respect also those individuals whose claims we cannot fulfil. To solve the challenges limited

³² In his article Horton claims that his argument shows why intransitivity is not as innocuous as proponents of limited aggregation have claimed (*Always Aggregate*, pp. 170-71). What my response here shows is that Horton's agglomeration problem is not any more troubling to limited aggregationists than previous challenges were. My reply here invokes the importance of moral relations and relational properties, the very same considerations that are invoked to argue why transitivity and the independence of irrelevant alternatives fail.

aggregation faces, we only need to pay close attention to the guiding principle of respecting those whom we, sadly, cannot save.³³

³³ I presented a version of this paper at the London Workshop in Moral and Political Philosophy at University College London and the LSE Choice Group. For their comments, criticisms, and suggestions I thank the audiences as well as Nicolas Côté, Joe Horton, Todd Karhu, Michael Otsuka, Jonathan Parry, Thulasi K. Raj, Korbinian Rüger, Alex Voorhoeve, and multiple anonymous reviewers.