Utilitas (2020), **32**, 454–471
doi:10.1017/S095382082000014X

CAMBRIDGE
UNIVERSITY PRESS

ARTICLE

Aggregation with Constraints

Korbinian Rüger

Balliol College, Oxford
korbinian.rueger@balliol.ox.ac.uk

Abstract
Utilitarianism is often criticized because of its reliance on the interpersonal aggregation of harms and benefits. However, since the rejection of all forms of interpersonal aggregation strikes most people as implausible, some critics of utilitarianism have proposed theories of Limited Aggregation. These occupy the middle ground between fully aggregative and non-aggregative views. Recently, Limited Aggregation has been criticized for having counterintuitive implications that seem even worse than the counterintuitive implications of fully aggregative and non-aggregative views it tried to escape. I here propose a new view of Limited Aggregation that does better than existing accounts in this regard. It is more modest than existing accounts of Limited Aggregation, but it retains the view's core idea. This, I claim, is the thought that sometimes very strong individual claims stand in the way of realizing the best outcome.

Suppose that:

> *Death vs. Headaches.* We can either save Ann from a terminal illness or prevent some number of different people from suffering a mild headache.

Many people think that there is no number of people suffering from a mild headache for which we ought to help them rather than save Ann from death. Next, suppose that:

> *Death vs. Paraplegias.* We can either save Ann from a terminal illness or prevent some number of different people from suffering permanent paraplegia.

Many of the same people think that whether we ought to save Ann from death depends on the number of people we could save from paraplegia instead. These two judgements seem to be in tension. In *Death vs. Paraplegias* these people allow that a larger number of lesser harms (the paraplegias) together outweigh a smaller number of greater harms (the one death) in importance.[1] In *Death vs. Headaches* these same people deny that preventing the headaches can ever be more important than saving Ann from death, no matter how many people would actually suffer the headache. According to these two judgements, interpersonal aggregation of harms determines what we ought to do

[1]If you think that paraplegia is worse than death, then simply swap one for the other.

© The Author(s), 2020. Published by Cambridge University Press

https://doi.org/10.1017/S095382082000014X Published online by Cambridge University Press

in *Death vs. Paraplegias* but does not determine what we ought to do in *Death vs. Headaches*. Views exhibiting this pattern of judgements are committed to *Limited Aggregation*. Limited Aggregation occupies the middle ground between *Full Aggregation* and *Non-Aggregation*.[2]

> *Full Aggregation*. For any harms X and Y, where Y is worse than X, there is some *n*, such that we ought to save *n* people from X rather than one different person from Y.

> *Non-Aggregation*. For any harms X and Y, where Y is worse than X, there is no *n*, such that we ought to save *n* people from X rather than one different person from Y.

One fully aggregative view is utilitarianism. According to this view, we ought to do whatever leads to the best outcome, where the best outcome is the one maximizing overall well-being. The overall well-being of an outcome, in turn, is simply an additive, unweighted function of all the individual well-being levels in that outcome. Strictly aggregative views have no problem explaining the common intuition in *Death vs. Paraplegias*. For some number, the benefits to the people we can save from paraplegia, if we do save them, will together outweigh the benefit we can give to Ann. These views, however, cannot explain the common intuition in *Death vs. Headaches*.[3]

Non-aggregative views have the opposite problem. They cannot account for the judgement that it is required, or even permissible, to save some number of people from paraplegia, rather than Ann from death. One family of non-aggregative views holds that individuals we could benefit have "claims" to these benefits or "complaints" against not getting a benefit that do not lend themselves to the kind of combining together as required by aggregative views. Rather, we ought to respect the "separateness of persons" and, after conducting a series of pairwise comparisons, do whatever satisfies the strongest individual claim, or equivalently, minimizes the strongest individual complaint. On these views, we therefore ought to save Ann from death rather than any number of people from paraplegia. This judgement is at least as implausible as the strictly aggregative judgement that we ought to spare a large number of people from a headache rather than save Ann from death.

The fact that Limited Aggregation makes neither of these intuitively implausible judgements gives us some reason to believe that such a view is true. It is my aim in this article to spell out what I take to be the most plausible version of Limited

[2]For a recent systematic account of Limited Aggregation that sparked a lively debate see Alex Voorhoeve, How Should We Aggregate Competing Claims?, *Ethics*, 125.1 (2014), 64–87. For accounts in support of the Limited Aggregation intuitions preceding Voorhoeve's see e.g. Dale Dorsey, Headaches, Lives and Value, *Utilitas*, 21.1 (2009), 36–58, or F. M. Kamm, Nonconsequentialism, in *The Blackwell Guide to Ethical Theory*, ed. by Hugh LaFollette and Ingmar Persson, 2nd edn (Oxford: Blackwell, 2013), pp. 261–86. For rare examples of strictly non-aggregative views, see e.g. G. E. M. Anscombe, Who Is Wronged? Philippa Foot on Double Effect: One Point, in *Human Life, Action and Ethics: Essays* (Exeter: Imprint Academic, 2005), pp. 249–51. and John M. Taurek, Should the Numbers Count?, *Philosophy & Public Affairs*, 6.4 (1977), 293–316. For fully aggregative views see e.g. Alastair Norcross, Comparing Harms: Headaches and Human Lives, *Philosophy & Public Affairs*, 26.2 (1997), 135–67, and Derek Parfit, Justifiability to Each Person, *Ratio (New Series)*, 16.4 (2003), 368–90.

[3]This is not to say that there *could not* exist any aggregative view that could resist letting Ann die in Death vs. Headaches. We could, for example have a view with a fully aggregative, but lexicographic axiology, according to which well-being consists of, say, two elements, where no amount of one can ever outweigh any amount of the other. Such views shall not concern me here, however.

https://doi.org/10.1017/S095382082000014X Published online by Cambridge University Press

Aggregation. I am not offering a full justification of the view. Rather, I want to explore the most plausible way to hold on to the two intuitions in *Death vs. Headaches* and *Death vs. Paraplegias*, if we so want to. As I will argue, we then should adopt the view that if one person has much more at stake than any person whose interests conflict with hers, we ought to do whatever serves the interests of that person. Otherwise, we ought to do what realizes the best outcome, as determined by the most plausible axiology. This is a much more modest proposal than existing accounts of Limited Aggregation. But as we will see, it retains its core commitments without suffering from the same defects as existing accounts.

I will proceed in five steps. Building on a proposal by Alex Voorhoeve, in the first section I will introduce Limited Aggregation in more detail. In the second section I deal with an objection to Limited Aggregation that has recently been brought forward by Patrick Tomlin. This objection holds that the view faces insurmountable problems in cases where individuals are added to the groups of people we can help. In response to this objection, in the third section I introduce my own approach to Limited Aggregation, which I call *Aggregation with Constraints*, and argue that it does not suffer from the same problems. In the fourth section I deal with a recent objection to Limited Aggregation by Joe Horton. As I will reveal, Horton's objection is an instance of the more general, but hitherto unstated, worry that Limited Aggregation is implausible because it violates a certain *separability* condition. According to this condition, roughly, our choice between distributions should not depend on elements that are the same in these distributions. I will argue that this should not surprise us, as views of Limited Aggregation are by definition relational. Whether a given claim can determine what we ought to do depends on how strong this claim is compared to other claims under consideration. A violation of separability is thus a *central feature of*, rather than an *objection to*, the view. In the fifth and final section I deal briefly with the objection that my Aggregation with Constraints fails in cases where we have more than two options available.

I. Limited Aggregation

The intuitive idea behind Limited Aggregation is that interpersonal aggregation of harms and benefits should sometimes make a difference to what we ought to do and sometimes not. In *Death vs. Headaches* it should not make a difference, in *Death vs. Paraplegias* it should. Recent accounts of Limited Aggregation can be traced back to comments T. M. Scanlon made in *What We Owe to Each Other*. While justifying the "individualist restriction" of his Contractualism, which, in principle, forbids any form of interpersonal aggregation, he offers the following qualification:

> If one harm, though not as serious as another, is nonetheless serious enough to be morally "relevant" to it, then it is appropriate, in deciding whether to prevent more serious harms at the cost of not being able to prevent a greater number of less serious ones, to take into account the number of harms involved on each side. But if one harm is not only less serious than, but not even "relevant to", some greater one, then we do not need to take the number of people who would suffer these two harms into account in deciding which to prevent, but should always prevent the more serious harm.[4]

[4]T. M. Scanlon, *What We Owe to Each Other* (Cambridge, MA: London: Belknap Press of Harvard University Press, 1998), pp. 239–40. It is in fact very hard to allow for even Limited Aggregation within

https://doi.org/10.1017/S095382082000014X Published online by Cambridge University Press

Derek Parfit suggested a view that according to him captured the essence of Scanlon's remarks. He calls it the

> *Close Enough View*. When [and only when] burdens to different people are close enough in size, one greater burden could be morally outweighed by a sufficient number of lesser burdens.[5]

This view allows us to capture the aforementioned pair of intuitions about Deaths vs. Headaches and Death vs. Paraplegias. It was recently spelled out in a systematic way by Alex Voorhoeve.[6] On his approach, which he calls

> *Aggregate Relevant Claims* (*ARC*)
>
> 1. Each individual whose well-being is at stake has a claim on you to be helped. [. . .]
> 2. Individuals' claims compete just in case they cannot be jointly satisfied.
> 3. An individual's claim is stronger the more her well-being would be increased by being aided. [. . .]
> 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.[7]

Voorhoeve explains that, in a pairwise comparison between two claims, the weaker claim is relevant to the stronger claim if and only if the person holding that weaker claim would not be required to give it up, if she herself were the one to make the choice. The approach thus takes into account the subjective viewpoint of individuals from which it seems permissible to favour oneself to a certain degree. For example, it seems permissible for me to save myself from paraplegia rather than a stranger from death, but impermissible to spare myself a minor headache at the cost of the death of a stranger.

In *Death vs. Headaches,* ARC says that we ought to save Ann irrespective of the number of people we could spare the headache, because the individual claims to be spared a headache are not close enough to Ann's claim to be saved from death. We therefore ought to satisfy the single strongest individual claim. In *Death vs. Paraplegias* on the other hand, the individual claims to be saved from paraplegia presumably are close enough to Ann's claim. For some number of people facing paraplegia, we therefore ought to save them rather than Ann from death, since this maximizes the sum of strength-weighted, relevant claims.

Before pointing out the proposal's problems, I want to pause to emphasize that ARC, as I understand it, is a view about moral choice, and not about axiology. ARC will

Scanlon's Contractualism. See e.g. Michael Otsuka, Saving Lives, Moral Theory, and the Claims of Individuals, *Philosophy & Public Affairs*, 34.2 (2006), 109–35. This, however, will not be my topic here.

[5]Parfit, p. 278.

[6]See Voorhoeve, How Should We Aggregate Competing Claims?

[7]Voorhoeve, How Should We Aggregate Competing Claims?, p. 66. I have amended the view to reflect a simplifying restriction that the strength of a person's claim only turns on the size of a given benefit and not the absolute well-being level of a person.

https://doi.org/10.1017/S095382082000014X Published online by Cambridge University Press

sometimes select outcomes that are worse in terms of overall goodness than other available outcomes.[8] I take it that ARC and Limited Aggregation in general should only apply to the deontic assessment of choices and not to the evaluation of outcomes. As we will see later on, Limited Aggregation violates structural properties that, in my view, are inherent to the very concept of goodness.[9] To keep things simple, I will assume the utilitarian axiology according to which the best outcome is the one maximizing the overall unweighted sum of well-being. However, nothing hinges on this assumption, as long as we agree that axiology is a fully aggregative matter. If you prefer a prioritarian axiology, for example, then you can simply plug it in. Two further assumptions: First, I only consider choice under certainty, so each alternative act will be associated with exactly one outcome. And second, I only consider "same-people choices", i.e. all available outcomes contain the same set of people.

With these assumptions in place, we can now move on to the first set of objections against Limited Aggregation that I want to consider.

II. The anchoring ambiguity

The problem I deal with in this section is due to Patrick Tomlin.[10] It stems from an ambiguity in Voorhoeve's account that had not been appreciated before. The ambiguity is the following. There are (at least) two ways to understand the "strongest competing claim" (which Tomlin calls the "anchoring claim") to which any claim *c* must be close enough in order to be relevant. The first meaning is "the strongest claim with which *c* competes". The second meaning is "the strongest claim in the competition". These two meanings come apart when the two groups of claims that compete with each other are not homogeneous in the sense that it is not the case that all claims within each group are exactly equally strong.[11] Take the following example.

> *Death and Fingers vs. Legs.* We can either save Ann from a terminal illness and prevent a number of people from losing a finger or prevent a number of different people from losing a leg.

What in this case determines the relevance of the claims of those who stand to lose a finger? Ann's claim or the claims of those who stand to lose a leg? Ann's "death-claim" (i.e. her claim to be saved from death) is the strongest claim in the competition, but this claim doesn't compete with the finger-claims. If we suppose that Ann's claim is the anchoring claim, then in Tomlin's terminology we "anchor by strength". The strongest claims with which the finger-claims actually compete, however, are the leg-claims. If we take those claims to be the anchoring claims, then we "anchor by competition". If we now plausibly assume that finger-claims are close enough to leg-claims, but not to death-claims, then they are irrelevant under the first interpretation, but relevant under the second.

[8]Voorhoeve himself writes that ARC will sometimes "have costs in terms of *impersonal* goodness". See Voorhoeve, How Should We Aggregate Competing Claims?, p. 85 (my emphasis).

[9]Of course, some might disagree about this. See e.g. Larry S. Temkin, *Rethinking the Good : Moral Ideals and the Nature of Practical Reasoning* (New York: Oxford University Press, 2015).

[10]Patrick Tomlin, On Limited Aggregation, *Philosophy & Public Affairs*, 45.3 (2017), 232–60.

[11]I should note that Voorhoeve himself explicitly places such "diverse group cases" outside the scope of his original paper. However, it is clear that Voorhoeve doesn't think that ARC shouldn't apply beyond the "homogenous group cases" to which he restricts his discussion and none of his arguments would support such a restriction.

https://doi.org/10.1017/S095382082000014X Published online by Cambridge University Press

Tomlin shows that both interpretations give rise to implausible results. First consider Anchor by Competition and the following two-stage case.[12]

> *Two-Stage Case A*
> *Stage 1.* We can either save Ann from a terminal illness (A) or *n* different people from losing a leg (B).
> *Stage 2. m* people whom we can save from losing a finger are added to each side.

I will illustrate this and following cases with simple figures. Each column represents one group we can save. The numbers (1, *n*, *m* in this case) represent the number of people holding a claim of a given strength. The higher up the number is in a given column, the stronger the claim. To keep things simple, I will assume that there are only three strengths a claim can take. There will only be death-claims, leg-claims, and finger-claims. If claims are represented in parenthesis, that means that they are not present at the first stage but are added at the second stage. *Two-Stage Case A* can then be represented as follows.

Two-Stage Case A

A	B
1	
	n
(*m*)	(*m*)

Now suppose *n* is such that according to Limited Aggregation, at stage 1, we ought to save the people from losing a leg. Now we add *m* people who stand to lose a finger both to Ann's side and to the side of those people who stand to lose a leg. As before, we assume that finger-claims are close enough to leg-claims, but not to death-claims. According to Anchor by Competition, this means that the claims of the people we added to Ann's side are relevant, but the claims of the people we added to the other side are not. This seems counterintuitive, especially considering that at stage 2 we might have to reverse our judgement from stage 1 and should now save Ann from death and the *m* people on her side from losing a finger. This is because their relevant claims might turn the balance in Ann's favour, while the irrelevant claims on the other side carry no counterbalancing weight. The approach violates what Tomlin calls

> *Equal Consideration for Equal Claims*. All claims of equal strength ought to be given equal weight in determining which group to save.[13]

Consider now Anchor by Strength by way of the following case, also adapted from Tomlin.[14]

> *Two-Stage Case B*
> *Stage 1.* We can either save *n* people from losing a finger (A) or *m* people from

[12]See Tomlin, p. 240.
[13]Tomlin, p. 241.
[14]See Tomlin, p. 245.

https://doi.org/10.1017/S095382082000014X Published online by Cambridge University Press

losing a leg (B).

Stage 2. Ann, whom we can save from suffering a terminal disease, is added to A.

A	B
(1)	
	m
n	

Suppose, again, that finger-claims are relevant to leg-claims, but not to death-claims, and that *n* and *m* are such that at stage 1, we ought to save the people from losing a finger. It is obvious that adding Ann to this group ought to strengthen the case for saving them. Not only can we now save *n* people from losing a finger, but we can also save Ann's life and we obviously ought to do so. Anchor by Strength, however, might tell us otherwise. Since the strongest claims in the overall competition now no longer are the leg-claims, but Ann's death-claim, the finger-claims are no longer relevant and should play no role in our decision anymore. This means that we now have to weigh Ann's single claim against *m* leg-claims and depending on how large *m* is, ARC might tell us to satisfy those claims, which is clearly the wrong result. Anchor by Strength violates what Tomlin calls the

> *Principle of Addition.* Merely adding a claim to a group of claims cannot lessen that group's choiceworthiness compared with a fixed alternative.[15]

This principle seems intuitively plausible enough to reject any approach that violates it.[16] In his article Tomlin goes over a number of possible, more complicated anchoring rules, which, however, suffer from even more devastating counterexamples.[17] This suggests that the approach of determining a set of relevant claims and then restricting aggregation to these claims is doomed to fail. In the following section I therefore develop an alternative approach at fleshing out Limited Aggregation.

III. Aggregation with Constraints

I suggest the following specification of Limited Aggregation. In a choice situation with competing claims, we should look at the single strongest individual claim. If this claim is very much stronger than the strongest claim with which it competes, we do whatever satisfies this claim. If it is only slightly stronger, we do whatever is best in aggregate. Put

[15]Tomlin, p. 245

[16]I should note, however, that there is a phenomenon, familiar from social choice theory, which would also violate the Principle of Addition. This is the so-called "No Show Paradox". Many otherwise attractive voting procedures have the implication that a voter could be better off "not showing up", rather than casting the ballot for her favourite candidate in the sense that her voting for her favourite candidate will actually lead to this candidate not being elected when she would have been elected had the voter abstained. See e.g. Hervé Moulin, Condorcet's Principle Implies the No Show Paradox, *Journal of Economic Theory*, 45.1 (1988), 53–64.

[17]See Tomlin, pp. 247–50.

https://doi.org/10.1017/S095382082000014X Published online by Cambridge University Press

differently, there is a general rule to do whatever is best in aggregate, i.e. realize the best outcome. However, the strongest individual claim can raise a veto. We then need to see whether this veto is justified or not. If it is justified, we do whatever satisfies this claim. If not, we proceed to realize the best outcome. Call this

> *Aggregation with Constraints.* If one person has much more at stake than any person whose interests conflict with hers, we ought to do whatever serves the interests of that person. Otherwise we ought to do what realizes the best outcome.

This principle does better than ARC in Tomlin's cases. Reconsider

Two-Stage Case A

A	B
1	
	n
(*m*)	(*m*)

On Aggregation with Constraints, in stage 1, Ann's claim isn't much stronger (given our assumptions) than any other person's competing claim, so we ought to do what realizes the best outcome. This depends on how large *n* is. The added finger-claims at stage 2, do not change this assessment. Aggregation with Constraints does not violate Equal Consideration for Equal Claims, at least not in this case. One case, in which one might think that it looks like it does violate the principle, is the following.[18]

Two-Stage Case C
Stage 1. We can either save Ann from a terminal illness (A) or *n* different people from losing a finger (B).
Stage 2. m people whom we can save from losing a leg are added to each side.

A	B
1	
(*m*)	(*m*)
	n

At stage 1, Ann's claim is much stronger than any other claim, so by Aggregation with Constraints, we ought to save her. This judgement does not depend on how large *n* is. Now at stage 2 we may have to rethink our verdict. This is because the added leg-claim on the side of those threatened to lose a finger competes with Ann's death-claim and, by assumption, is close enough to her claim. On Aggregation with Constraints we therefore now ought to do what leads to the best outcome. If *n* is large enough, we will now have to let Ann die and satisfy the finger-claims and the leg-claims we added to their side.

[18]I thank Theron Pummer for prompting me to discuss such a case.

https://doi.org/10.1017/S095382082000014X Published online by Cambridge University Press

One might think that this violates Equal Consideration for Equal Claims. This is because, after all, we added the same number of leg-claims to each side and one could hold that since they carry equal weight, they should somehow "cancel each other out" and therefore together should do nothing to overturn our previous judgement at stage 1. The fact, however, that they do not cancel each out in this way does not imply that we do not give the two claims equal consideration or even that we disregard the leg-claims we added to Ann's side. To see this, consider an instance of *Two-Stage Case C*, where n is such that on aggregation, the finger-claims together with the one leg-claim would just barely outweigh Ann's single death claim. The leg-claim we added to Ann's side now turns the tables in her favour. On aggregation it adds just as much force to her side as the other leg-claim added to the other side. It is only when n is large enough that saving these people from losing a finger would add more value to the outcome than saving Ann could that we should overturn our judgement from stage 1. Only then is it the case that the added leg-claim to Ann's side can do nothing to help her. For all other n, it will turn out that we ought to help her and the person who was added to her side.

What is true, though, in cases like Two-Stage Case C, is that adding relevant claims to either side can "activate" previously irrelevant claims. It is like turning a switch that takes us from a non-aggregative to an overall aggregative mode. This implication may seem especially troublesome, when more claims are added to the group that will not be chosen than to the group that will be chosen. Take the following variation.

*Two-Stage Case C**

A	B
1	
(m)	(1)
	n

In this case, if n is very large, m could also be quite large, and we would still at stage 2 overturn our previous judgement. How problematic is this implication? As before, it does not violate Tomlin's Equal Consideration for Equal Claims. However, one might think that adding a claim to group B, while adding *more* claims of equal strength to an opposing group A, should not strengthen B*'s choiceworthiness compared to A.

I agree that at first blush this implication seems odd. However, on reflection, proponents of Limited Aggregation have a justification available. This is the thought that is at the very heart of Limited Aggregation. When we refuse to let a number of weaker claims outweigh one very strong claim, we do so because the person with the strong claim has so much more at stake than any person with a competing claim. Whenever one person (or a group of people) has much more at stake than any other person we could help instead, we ought to help this person irrespective of the numbers involved.

I take it that this is the driving intuition behind cases like *Death vs. Headaches* and it is fulfilled at the first stage, so we refuse the aggregation of the n finger-claims. The antecedent of this judgement no longer holds at the second stage, so by this judgement it is no longer true that we ought to help the person with the strongest claim, no matter what. There now is another person who has, by stipulation, a relevantly similar claim.

https://doi.org/10.1017/S095382082000014X Published online by Cambridge University Press

We are going to help this person alongside many others with weaker claims instead, and so realize the best outcome possible.

It makes a difference whether a claim competes against relevantly similar claims or not. If it does not, then we refuse to aggregate. If it does, then what we ought to do depends on the specific case and the specific version of Limited Aggregation. However, the fact that the existence of such a claim *could* make a difference to our mode of practical reasoning is not surprising.

The implication of Aggregation with Constraints in *Two-Stage Case C* is thus fundamentally different from the implication of ARC in Tomlin's *Two-Stage Case A*. In this case, it was the addition of finger-claims that, by stipulation, are irrelevant to Ann's death-claim, which changed things. In *Two-Stage Case C* it is the addition of relevant leg-claims that changes things. This is a crucial difference, as we should expect relevant claims to have the power to change the normative situation in which we find ourselves. We shouldn't expect irrelevant claims to have that power, however.[19] In essence, this is just a restatement of what it means to buy into Limited Aggregation. In certain situations, only some claims matter, others don't. Proponents of the view should therefore not be unnerved by this implication in *Two-Stage Case C*. Rather, this is what we should expect of a view of Limited Aggregation.

So far, I have shown that Aggregation with Constraints does not seem to run into the same problems as Anchor by Competition as it does not violate *Equal Consideration for Equal Claims*. To see whether it also does better than Anchor by Strength, consider now again Tomlin's second case,

Two-Stage Case B

A	B
(1)	
	m
n	

At stage 1, there is no single claim that is much stronger than any other competing claim, so we ought to do what realizes the best outcome and this depends on *n* and *m*. If we now at stage 2 add Ann's death-claim to the finger claim, no matter how large *n* and *m* are, we won't now have to switch from satisfying the finger-claims to satisfying the leg-claims. If at stage 1 we ought to satisfy the finger claims, then at stage 2, we will still do so. Unlike Anchor by Strength, Aggregation with Constraints does not violate the Principle of Addition and therefore does better. Adding a claim to one group of claims without at the same time also adding claims to the competing group can never lessen this group's choiceworthiness.

Here is the principled reason why my approach doesn't run into the problems raised by Tomlin. Whereas ARC aggregates some claims (the relevant ones) but not others (the irrelevant ones), my approach either aggregates all claims or none at all. It lacks the discriminating feature in ARC that gives rise to implausible results under both

[19]For a similar line of thought concerning this point see Victor Tadros, Localized Restricted Aggregation, *Oxford Studies in Political Philosophy*, 5 (2019), 171–204. Tadros proposes an otherwise very different view of Limited Aggregation.

https://doi.org/10.1017/S095382082000014X Published online by Cambridge University Press

ways of resolving the ambiguity in ARC. My approach is thus not subject to the same objections as both ways of resolving the ambiguity in ARC.

There is, however, a cost to Aggregation with Constraints that some proponents of Limited Aggregation will probably be hesitant to accept. This is the fact that it gives a different verdict from ARC in a well-known case by Frances Kamm.

> *Sore Throat.* We can either save Ann from death or save Bob from death and cure Carl's sore throat.[20]

Within our framework, this case can be represented as follows.

A	B
1	1
	1

Kamm writes that the benefit of curing Carl's sore throat should be disregarded. It is an "irrelevant utility" and should play no role in our decision. Many proponents of Limited Aggregation agree.[21] According to Kamm's approach, we ought to treat this case just as we would treat the case were Carl not present. That is, we should be indifferent between saving Ann or saving Bob. On my approach, this is not so. Since Ann's claim and Bob's competing claim are equally strong, and therefore obviously close enough, we ought to aggregate and do what realizes the best outcome. This means that we should save Bob and cure Carl's throat, which realizes the best outcome.

While I did share Kamm's intuition when I first encountered her case, I am prepared to give it up on reflection. If we *really* are indifferent between two alternatives, then any additional reason that speaks in favour of one of the alternatives, but not the other, should make it the case that we are no longer indifferent. On Aggregation with Constraints, this is what happens when we add Carl's very small claim to Bob's side. Before we added Carl's claim, the reasons speaking in favour of saving Ann were *exactly* equally as strong as the reasons in favour of saving Bob. This no longer is the case once we have added Carl. We now have more reason to save Bob and Carl than we have to save Ann and ought to do so.

As I said, this will come as a significant cost to some proponents of Limited Aggregation who are wedded to Kamm's intuition. Even so, this cost is, I think, much smaller than the costs of accepting the implications of ARC that Tomlin's cases brought out. So, Aggregation with Constraints still has the upper hand over ARC.

Before I conclude my exposition of Aggregation with Constraints, however, there is one more case I want to discuss. Consider the following variation of *Sore Throat*.

> *Sore Throat 2.* We can either save Ann and Dave from death or save Bob from death and cure Carl's and n other people's sore throat.

[20]See F. M. Kamm, *Morality, Mortality Volume I: Death and Whom to Save From It* (Oxford: Oxford University Press, 1998), p. 146.

[21]See e.g. Voorhoeve, How Should We Aggregate Competing Claims? and Tadros.

https://doi.org/10.1017/S095382082000014X Published online by Cambridge University Press

Since in this case, it is not true that one person has much more at stake than any other person whose claim competes with hers, on Aggregation with Constraints, we should still save Bob and cure the sore throats if there are enough of them. This is the case although Dave, also facing death, is now on Ann's side. One could argue that unlike the original Sore Throat Case, Sore Throat 2 is a simple variant of Death vs. Headaches and therefore a stronger objection to Aggregation with Constraints.[22]

This case does pose a stronger challenge than the original Sore Throat Case. However, as I shall now show, it is similar to the challenge that Two Stage Case C posed and we can give essentially the same answer. We can see this by cashing out Sore Throat 2 as a two-stage case as well.

Two-Stage Sore Throat 2
Stage 1. We can either save Ann from death (A) or cure *n* different people's sore throat (B).
Stage 2. Dave, whom we can save from death, is added to Ann's side. Bob, whom we can also save from death, is added to the other side.

A	B
1 (1)	(1)
	n

At stage 1, it is clear that we ought to save Ann. This is just a variant of Death vs. Headaches (or, rather, Death vs. Sore Throats). What we ought to do at stage 2, however, now depends on how large *n* is. This is because, after adding Bob's death-claim, it is no longer the case that one person has much more at stake than any other person whose interests conflict with hers. Thus, on Aggregation with Constraints, we now ought to do what is best in the aggregate. If *n* is large enough, we might have to overturn our judgement from stage 1.

This case has the same structure as Two Stage Case C, only that the added claims are of an even stronger kind, so by definition, they are of the kind that is relevant in this situation. As I write above, these sorts of claims have the power to change the normative situation we are in. When we refuse to let a number of weaker claims outweigh one very strong claim, we do so because the person with the strong claim has so much more at stake than any person with a competing claim. In this case, after we added Bob's claim opposing Ann's, this is no longer true. There is no one on either side who has much more at stake than any person with a competing claim. There are death-claims on both sides. This fundamentally changes the kind of justification we can give to the person whom we are not going to save. In a strictly pairwise comparison of the strongest claims, there would now be a tie where there wasn't one before. Therefore, the view that we should always satisfy the single strongest individual claim, which is the view that best embodies the non-aggregation sentiment, judges the two stages of Two-Stage Sore Throat 2 differently. Aggregation with Constraints does so as well.

This concludes my initial exposition of Aggregation with Constraints. We have seen that it does better with regard to some of the objections that have been levelled against

[22]I thank an anonymous referee for pressing me to discuss this case.

https://doi.org/10.1017/S095382082000014X Published online by Cambridge University Press

other accounts of Limited Aggregation. There are, however, possible objections to Limited Aggregation that also apply to the view I proposed. I will take these up in the remaining two sections.

IV. Separability and Horton's Dilemma

Joe Horton has recently brought forward a new challenge to Limited Aggregation.[23] He proposes a dilemma that, in his view, is faced by all views endorsing Limited Aggregation. His challenge is important. This is because both horns of his dilemma include a violation of a plausible separability condition, which he doesn't explicitly recognize. Separability says, roughly, that our choice between two distributions should not depend on elements that are the same in both distributions.

In this section, I show that Limited Aggregation does indeed violate this condition. This, however, is not actually surprising, but at the bedrock of the view's commitments. It brings to the surface the following: there is a tight connection between Limited Aggregation and relational factors as expressed by a violation of separability.

This section proceeds as follows. First, I will introduce the notion of separability as it applies here and show that it is tightly connected to aggregation. Second, I explain how Horton's objection amounts to the claim that Limited Aggregation violates separability. Third, I argue that this violation is not actually an objection to, but rather a central feature of, Limited Aggregation.

In general, separability is the idea that when evaluatively comparing two objects, we can ignore the objects' elements which are the same in both. It is a prima facie attractive principle and a powerful tool in ethics and decision theory. In decision theory, *separability of states of nature* is an important axiom of expected utility theory. It holds that when two alternative courses of action have the same outcome in some state of nature, then our choice should not depend on what happens in that state of nature.[24]

In ethics, we are more often concerned with the *separability of people*. This principle says that when comparing two distributions, we can ignore those people who fare the same in both distributions. In other words, our choice should be independent of the unaffected, where an individual i is unaffected just in case i has the same well-being level in both distributions. Separability is a practically convenient property, as it implies that when choosing between, say, two different policies, we can focus on those people actually affected by the policies. Here is a general formulation that serves our purpose.

> *Separability*: If two distributions have the same well-being level for one person, then our choice between them depends only on what they are like for other people.[25]

To illustrate the attractiveness of this principle, take the following case (Table 1), where the numbers represent the well-being levels of three people, i_1, i_2, and i_3, and we have to choose between two distributions, D1 and D2.

Our choice between D1 and D2 might be affected by many considerations, for example by the fact that total well-being is greater in D2 than in D1. But what our

[23]Joe Horton, Always Aggregate, *Philosophy & Public Affairs*, 46.2 (2018), 160–74.

[24]In decision theory, this is known as the "Sure Thing Principle". See e.g. Leonhard Savage, *The Foundations of Statistics*, 2nd edn (New York: Dover, 1972).

[25]For a precise formulation and discussion, see John Broome, *Weighing Goods* (Cambridge, MA: Basil Blackwell, 1991), ch. 4.

https://doi.org/10.1017/S095382082000014X Published online by Cambridge University Press

Table 1. D1 vs. D2

	i1	i2	i3
D1	2	2	1
D2	4	1	1

choice should *not* be affected by, separability holds, is how well i_3 is doing. Her well-being is the same either way and should not affect our choice. This means that on separability, D1 ought to be chosen rather than D2 if and only if (2,2) ought to be chosen rather than (4,1); and D2 ought to be chosen rather than D1 if and only if (4,1) ought to be chosen rather than (2,2).[26]

Separability plays an important role in some of the best positive arguments for utilitarianism and other fully aggregative theories, such as prioritarianism.[27] In fact, it can be proven that if a theory fulfils a number of prima facie plausible conditions and is separable, then it is fully aggregative, and vice versa.[28] So, there is a tight, well-known, connection between separability and aggregation. As I show below, Horton's argument exploits that connection.

However, before applying separability to the kinds of cases we are dealing with here, we need to take care of a slight complication. Since i_3 would not have a claim by the lights of our theory, she would not occur in the choices we have been considering thus far. In these choices we only considered the well-being of people who have claims on us. And since, by definition, only people whose well-being we can affect can have claims, i_3 would not be considered. We thus need a slightly stronger principle to show how Limited Aggregation violates separability. We need a principle that tells us that neither i_3 nor i_4 should affect our choice in a case like the following (Table 2).

Such a principle is a little harder to come by, as both i_3 and i_4 have a stake in our decision. i_3 has a claim to D3 and i_4 has a claim to D4. So, in some sense, we should not ignore them. However, it is clear that in the following case (Table 3), in which the decision only includes i_3 and i_4, we should be indifferent between D5 and D6.

We have no reason to favour i_3 over i_4 or vice versa and both have exactly the same stake in our decision, so it is clear that we should be indifferent between both options.[29]

[26]I am indebted to an anonymous referee for pressing me to make this point and for improving the following discussion.

[27]For an argument for utilitarianism based on separability, see e.g. Broome, *Weighing Goods*. For an argument for prioritarianism, see Matthew D. Adler, *Well-Being and Fair Distribution* (Oxford: Oxford University Press, 2012).

[28]See Broome, *Weighing Goods*, ch. 4. Broome proves that what he calls strong separability for value orderings implies their additivity, and vice versa. In other words, iff an ordering is separable then it is fully aggregative. This is Broome's "First Separability Theorem." Thus, rejecting full aggregation is equivalent to rejecting separability. However, we cannot readily rely on Broome's result here, since on Limited Aggregation, the "ought to be chosen over" relation is not transitive, and thus not an ordering. The "better than" relation with which Broome is concerned, on the other hand, *is* an ordering (pace e.g. Temkin). For Broome's method and results, see also W. M. Gorman, The Structure of Utility Functions, *Review of Economic Studies*, 35 (1968), 367–90, and David Kranz H. and others, *Foundations of Measurement: Volume 1* (Academic Press, 1971). However, there are positive arguments for a similar result that can largely do without transitivity. See e.g. Kacper Kowalzcyk, Yet Another Argument Against Anti-Aggregation (unpublished manuscript).

[29]This is also given by the following principle.

https://doi.org/10.1017/S095382082000014X Published online by Cambridge University Press

Table 2. D3 vs. D4

	i1	i2	i3	i4
D3	2	2	1	0
D4	4	1	0	1

Table 3. D5 vs. D6

	i3	i4
D5	1	0
D6	0	1

We might then think that attaching such indifferences to our choices should not change anything. In other words, when choosing between two distributions, our decision is unaffected by those elements between which we are indifferent. Call this

> *Separability**. If two distributions have subdistributions between which we are indifferent, then our choice between them depends only on what they are like in other subdistributions.

This principle implies that the presence of i_3 and i_4 should not affect our choice between D3 and D4 above. Prima facie, this principle will be attractive to many people. Like other separability principles, however, it is violated by some theories in distributive ethics, most notably egalitarian theories.[30] This is because, for these theories, it is important how individuals fare *relative to others*. It is also violated by Limited Aggregation and much for the same reason.

For example, to see that Aggregation with Constraints violates Separability*, we can just take another look at

Two-Stage Case C

A	B
1	
(*m*)	(*m*)
	n

Weak Anonymity. If two outcomes X and Y differ only in that the identities of two people have been permuted, we should be indifferent between X and Y.

I don't appeal to this general principle here, as it is likely to be rejected by many aggregation sceptics. However, even these sceptics will accept the indifference in this restricted case.

[30]In fact, a violation of separability is what distinguishes "true" egalitarian theories from other distribution sensitive theories. See John Broome, Equality versus Priority: A Useful Distinction, *Economics and Philosophy*, 31.02 (2015), 219–28. I agree with Broome on this.

https://doi.org/10.1017/S095382082000014X Published online by Cambridge University Press

A and B are equally choiceworthy inasmuch they both include the same number of people with leg-claims, *m*. Between these two subdistributions we are indifferent. According to Separability* this should then not affect our choice. However, we have seen that under Aggregation with Constraints, adding these claims *does* affect what we should do.

Now, Horton's dilemma is the following. The first horn: he shows that views of Limited Aggregation either have implications like my Aggregation with Constraints has in Two-Stage Case C above. So, a violation of Separability* is what gives rise to the first horn of Horton's dilemma. The second horn is the following: we can construct cases where, according to Limited Aggregation, we ought to save group A over B, and group C over D, but in a choice between the union of groups A and C and the union of groups B and D, we ought to choose B and D.[31] In other words, we are violating a principle we can call

> *Horton's Principle*. If we ought to choose distribution A over B and distribution C over D, then we ought to choose distribution (A&C) over (B&D).[32]

The specifics of Horton's case need not concern us here, but he is correct that views of Limited Aggregation can violate this principle. Underlying this implication, however, is, again, a violation of separability. This is not recognized by Horton.

Horton's principle follows from Separability* and transitivity. If we ought to choose A over B, then by Separability* we ought to choose (A&C) over (B&C). And if we ought to choose C over D, then by Separability* we ought to choose (B&C) over (B&D). If the transitivity of "ought to be chosen over" holds in this case,[33] it follows that we ought to choose (A&C) over (B&D), which is Horton's principle.[34] So, what gives rise to both horns of Horton's dilemma is the same thing: a violation of a separability. This means that Horton's objection amounts to the claim that all views of Limited Aggregation violate separability.

How surprising is this violation though? To see that it is not very surprising and actually a central part of the view, we just need to look at some of the different formulations of Limited Aggregation. According to Voorhoeve's ARC, "[a] claim is relevant if and only if it is sufficiently strong *relative to* the strongest competing claim."[35] And according to my Aggregation with Constraints, "if one person *has much more at stake than any person whose interests conflict with hers*, we ought to do whatever serves the interests of that person." On views of Limited Aggregation, what we should do depends in part on how people fare relative to others. That these views then violate separability is not a surprise or an unwelcome implication, but rather a central building block of the theory. Limited Aggregation is a *relational view*.

It is important to re-emphasize, however, that this does not imply anything about the underlying axiological structure. Aggregation with Constraints is a deontic view

[31]See Horton, p. 11.

[32]Horton does not explicitly mention this principle, but it is clear from his discussion that this is what he has in mind.

[33]I say "in this case", because I don't want to assume that the transitivity of this relation holds universally. See the next section.

[34]I am indebted to Kacper Kowalzcyk for discussion of this point in particular and the issues in this section more general.

[35]Voorhoeve, How Should We Aggregate Competing Claims?, p. 66 (my emphasis).

https://doi.org/10.1017/S095382082000014X Published online by Cambridge University Press

and the separability violations I have outlined occur only at that level. In fact, I am sceptical that any plausible ethical view could have this structure at the level of axiology.[36]

This concludes my discussion of Horton's objection to Limited Aggregation. It wasn't my aim to invalidate the objection. Rather I wanted to show that it is not so much an objection to Limited Aggregation as a restatement of a central part of the view. I argue elsewhere in more detail that there is a tight connection between this violation of separability and scepticism about aggregation. There I argue that aggregation sceptics should be egalitarians, thus violating separability.[37] For the remainder of this article, however, I now want to turn to a different set of objections that I need to address before Aggregation with Constraints can stand as a contender for the most plausible version of Limited Aggregation.

V. Cases with three or more alternatives

This objection is brought out when we apply my principle not to a binary choice but to a choice between three (or more) alternatives. Given the apparently intransitive structure of Limited Aggregation and the violation of plausible axioms of rational choice that follow from this structure, it is not surprising that Limited Aggregation faces problems when extended beyond binary choices. For Limited Aggregation, the relation "ought to be chosen over" is not always transitive. After all, as we have seen in many examples, it is the very idea of Limited Aggregation, that, if the numbers are right, we ought to save *n* people from losing an arm rather than Ann from death, and *m* people from losing a finger rather than *n* people from losing an arm, *but we should not* save *m* people from losing a finger rather than Ann from dying. This structure violates transitivity and in some cases in turn gives rise to a violation of other plausible principles of rational choice. Here, I do not want to add to these well-rehearsed general criticisms and responses to them.[38] Rather I want to focus on one case of this kind which seems to pose a challenge specifically for Aggregation with Constraints.

Death and Finger vs. Death vs. Fingers. We can either save Ann from a terminal illness and Bob from losing a finger, Annie from the same terminal illness, or *n* people from losing a finger.[39]

A	B	C
1	1	
1		*n*

In this case, Aggregation with Constraints seems to imply that we ought to choose C and save the people from losing a finger if there are enough of them. This is because,

[36]I tend to agree with Broome, who argues that, in his terminology, "good is coherent". See e.g. Broome, *Weighing Goods*, ch. 6.

[37]See Korbinian Rüger, Aggregation and Equality (unpublished manuscript).

[38]For some of the criticisms see e.g. John Halstead, The Numbers Always Count, *Ethics*, 126.3 (2016), 789–802; Alastair Norcross, Intransitivity and the Person-Affecting Principle, *Philosophy and Phenomenological Research*, 59.3 (1999), pp. 769–76, at 769. For responses see F. M. Kamm, *Intricate Ethics* (Oxford University Press, 2007). 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*, 2016, 1–9; Tomlin, p. 236 (fn. 11).

[39]Variations of this case were presented to me by Theron Pummer and Bastian Steuwer.

https://doi.org/10.1017/S095382082000014X Published online by Cambridge University Press

unlike, for example, *Death vs. Headaches*, it is no longer the case that "one person has much more at stake than any person whose interests conflict with hers" and so according to Aggregation with Constraints we ought to realize the best outcome, which by assumption means to satisfy the finger-claims, if n is large enough. One could plausibly argue that this is just as implausible as the claim that there is some n such that we ought to satisfy the finger-claims when the only other option is to save Ann, which would be equivalent to giving up our guiding intuition from *Death vs. Headaches*.

Another problem with a case like this is the following. If we consider pairwise comparisons of the three alternatives, we get a clear winner: A. Aggregation with Constraints would clearly choose A over C. It would also choose A over B. It is odd that the same view would then choose B out of all three options.

As I wrote above, because of their structural properties, such problems are bound to arise for any view of Limited Aggregation. However, if we are unsettled by this particular case, we could amend Aggregation with Constraints in a way that takes care of this case and that is in line with the rationale of the view.

This is first to eliminate all alternatives that do not include a claim that is close enough to the strongest claim. For *Death and Finger vs. Death vs. Fingers* this would mean eliminating C, and choosing A out of the remaining two alternatives. That is, we exclude all options that could not be chosen in a pairwise comparison with the alternative that includes the strongest claim(s), irrespective of the numbers involved. Only then do we apply Aggregation with Constraints as outlined above. This procedure is in line with the central idea that the strongest claim overall has special significance. The justification of our choice must first and foremost be directed at the person who has the greatest stakes in our decision.

VI. Conclusion

In this article I have clarified the notion of Limited Aggregation and have shown its limitations. I have shown where its most prominent account, Voorhoeve's Aggregate Relevant Claims, fails and have proposed an account that does better. I have called this Aggregation with Constraints. My proposal is a much more modest version of Limited Aggregation. If we are attracted to the initial idea, I take it this is the best we can hope for. However, the proposal manages to retain the central thought at the heart of Limited Aggregation. This is the thought that sometimes very strong individual claims stand in the way of doing what would be best in the aggregate and Ann gets to live in *Death vs. Headaches*. Of course, Aggregation with Constraints stands and falls with the idea that whether a person has much more at stake than any other person is very important. It thus puts even more pressure on the intuition in *Death vs. Headaches*. But we knew that all along. If this intuition weren't so robust, there would be no point in trying to come up with a theory of Limited Aggregation.[40]

[40]For comments and discussion, I would like to thank Ralf Bader, Jessica Fischer, Hilary Greaves, Joe Horton, Kacper Kowalzcyk, Michal Masny, Jeff McMahan, Theron Pummer, Tom Sinclair, Bastian Steuwer, Patrick Tomlin, Alex Voorhoeve, two anonymous referees of this journal, and audiences at University College London and at the University of St Andrews.

Cite this article: Rüger K (2020). Aggregation with Constraints. *Utilitas* **32**, 454–471. https://doi.org/10.1017/S095382082000014X

https://doi.org/10.1017/S095382082000014X Published online by Cambridge University Press