

Manuscript version: Author's Accepted Manuscript

The version presented in WRAP is the author's accepted manuscript and may differ from the published version or Version of Record.

Persistent WRAP URL:

<http://wrap.warwick.ac.uk/115239>

How to cite:

Please refer to published version for the most recent bibliographic citation information. If a published version is known of, the repository item page linked to above, will contain details on accessing it.

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions.

Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Publisher's statement:

Please refer to the repository item page, publisher's statement section, for further information.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk.

Relevance Rides Again? Aggregation and Local Relevance*

Aart van Gils and Patrick Tomlin

Forthcoming in Oxford Studies in Political Philosophy (Volume 6)

I Introduction

Often institutions or individuals are faced with decisions where not all claims can be satisfied. Sometimes, these claims will be of differing strength. In such cases, we must decide whether or not weaker claims can be aggregated in order to collectively defeat stronger claims. Some deny that such aggregation is ever permissible (Taurek, 1977; Munoz-Dardé, 2005; Thomas, 2012; Doggett, 2013). Call this *Anti-Aggregation*. However, this position seems unduly restrictive when claims are close in strength. For example, consider:

Case 1. You can save one person from death, or some larger number of people, N_1 , from paraplegia.

According to Anti-Aggregation, no matter how large N_1 gets, we should save one person from death. But imagine N_1 were one hundred, or one thousand, or one million people. It seems implausible that we should allow so many people to suffer paraplegia in order to save a single life. This may push us toward what we can call *Pure Aggregation*, a view under which there is no restriction on aggregating claims, no matter how weak they are. However, this kind of position is vulnerable to two famous kinds of examples, the most well-known and oft-discussed of which are T.M. Scanlon's 'World Cup Case' and F.M. Kamm's 'Sore Throat Case':

Case 2. *Scanlon's World Cup Case*: Suppose that Jones has suffered an accident in the transmitter room of a television station. Electrical equipment has fallen on his arm, and we cannot rescue him without turning off the transmitter for fifteen minutes. A World Cup match is in progress, watched by many people, and it will not be over for an hour. Jones's injury will not get

any worse if we wait, but his hand has been mashed and he is receiving extremely painful electrical shocks. Should we rescue him now or wait until the match is over? Does the right thing to do depend on how many people are watching—whether it is one million or five million or a hundred million? (Scanlon, 1998: 235).

Case 3. *Kamm's Sore Throat Case*: Suppose ... that we have a choice between saving A's life and saving B's, and alongside B is C who has a sore throat. Our drug that can save B's life can also in addition cure C's sore throat. (Kamm, 1998: 101).

Here is what Pure Aggregation would say about these cases. In Scanlon's case, provided the number of viewers were large enough, it would have us allow Jones to suffer the agonizing shocks in order to allow the many viewers to enjoy the football. In Kamm's case it would have us save B and C over A, since C's claim against a sore throat added to B's claim against death would outweigh A's claim against death. Both of these are entailments that many find counter-intuitive.

These two cases offer two different challenges to Pure Aggregation. Scanlon's case is an example of what we can call a *one vs. many* case – one strong claim (Jones's) is pitted against many weak claims (those of the viewers). Kamm's case is an example of what we can call a *tie break* case – a seemingly trivial claim is all that separates two otherwise equally matched groups of much stronger claims.

Many people are not attracted to either extreme view – Case 1 speaks against Anti-Aggregation, while Cases 2 and 3 speak against Pure Aggregation. Therefore, they seek a middle way, which we can call *Limited Aggregation*. The best-known version of Limited Aggregation, versions of which both Scanlon and Kamm endorse, is a view known as the *Relevance View*. The idea of the Relevance View is that weaker claims can be aggregated against stronger claims when they are sufficiently close in strength to be 'relevant' to the stronger claim, but not when they are 'irrelevant' to the stronger claim. So, for example, claims against paralysis can be aggregated against a claim against death,

but claims to watch the World Cup, or to avoid a sore throat, are ‘irrelevant’ and thus should not be counted.

The clearest articulation of this Relevance View is Alex Voorhoeve’s ‘Aggregate Relevant Claims’ (ARC), which states:

1. Each individual whose well-being is at stake has a claim on you to be helped. (An individual for whom nothing is at stake does not have a claim).
2. Individuals’ claims *compete* just in case they cannot be jointly satisfied.
3. An individual’s claim is *stronger*:
 - a. the more her well-being would be increased by being aided; and
 - b. the lower the level of well-being from which this increase would take place
4. A claim is *relevant* if and only if it is sufficiently strong relative to the strongest competing claim.
5. You should choose an alternative that satisfies the greatest sum of strength-weighted, relevant claims. (Voorhoeve, 2014: 66–67).

The Relevance View provides judgments in line with common intuitions across Cases 1, 2, and 3: provided N_1 were large enough, it would require saving the many against paraplegia in Case 1, but it would require saving Jones in Case 2, and would allow for either A or B and C to be saved in Case 3 (Kamm advocates tossing a coin).¹

Although the Relevance View vindicates our intuitions in these three cases, the Relevance View has always attracted criticism, often for the way it handles more complex cases (see, for example, Halstead, 2016). These complaints have been extensively discussed, and replied to (for example, Voorhoeve, 2014: 75–86). However, the Relevance View has recently come under renewed attack. In ‘On Limited Aggregation’ (hereafter OLA²), one of us (Tomlin, 2017) shows that when we consider complex cases involving groups of claims of diverse strength, the Relevance View suffers from an important ambiguity and seemingly fatal counter-intuitive entailments (OLA: 232–260).

In response to these problems, Victor Tadros has articulated a new version of the Relevance View, which he calls *Local Relevance* (Tadros, 2019). We will explain this view in depth in the next section. The key idea is that ‘relevance’ is not all-or-nothing: a claim is not relevant or irrelevant to some decision simpliciter. Rather, a claim can be relevant to some competing claims but not others, and therefore even if it is irrelevant to the strongest claims within a situation, it can still affect the decision as to which group to save in a limited way. The view is promising, and, as we will show, seems to get around the cases that OLA presents as objections to the original Relevance View, as articulated by ARC. However, the general idea of Local Relevance is vague – it isn’t clear exactly how to apply it to decisions about who to save.

In this paper, we want to explore and examine this view more closely. We will, first, introduce a more tightly-specified version of the Local Relevance view, called *Sequential Claims-Matching*. We will then show how Sequential Claims-Matching is able to meet the challenges presented in OLA. However, Sequential Claims-Matching faces difficulties and ambiguities of its own. In particular, we show that whilst it can deliver intuitive results in the OLA cases and Scanlon’s World Cup Case and other one vs. many cases, it struggles to capture the intuition that Kamm’s Sore Throat Case is designed to elicit and cannot necessarily rule out allowing tie-breaks in such cases. We will also show how there are two important ambiguities within the view that any would-be advocates will need to address. Finally, we briefly present an alternative version of Local Relevance, *Strongest Decides*.

One vs. many cases and tie break cases are important objections to Pure Aggregation, and the Relevance View allows us to vindicate both objections on the same grounds, giving a unified anti-aggregative rationale, whilst still allowing us to reject full Anti-Aggregation. However, in our view, the most promising account of Limited Aggregation, namely Local Relevance, will only justify Scanlon’s judgment in the World Cup Case, and *not* Kamm’s judgment in the Sore Throat Case. Therefore, we seek to pull apart the two most famous objections to purely aggregative views. It may be that our

intuitions about tie-break cases can be vindicated via some other route, but, we will argue, this cannot be based upon the best understanding of when claims are relevant.³ Before we begin our examination of Local Relevance and Sequential Claims-Matching, a word on method. In exploring Local Relevance, we will focus largely on whether it is able to vindicate core Limited Aggregation intuitions. Obviously, any full theory of Local Relevance will need to do more than this: for example, it will need an account of the deeper justification or rationale for the view, and to take a stand on how to distinguish which claims are relevant to which.⁴ While these are important tasks, it is clear from the wider literature that the search for a plausible position on aggregation which treads the middle ground between Anti-Aggregation and Pure Aggregation is *driven by* our intuitive responses to cases. Therefore, any defensible Limited Aggregation theory must both match our intuitions *and* have a firm theoretical grounding.

In this paper we seek to make progress on the first half of this coin – the search for an approach to aggregation that can make sense of common intuitions. This should help inform the search for a deeper justification: once we have a sense of which view looks intuitively plausible, we will have a better sense of what kind of deeper justification can be offered for it. Obviously, this must be an ongoing process of reflective equilibrium, but we cannot try to ‘do it all’ in this paper.

As a result, this paper deals with many hypothetical cases designed to show particular results, and they are quite complex. In order to be as clear as possible about what each case shows, we will both walk the reader through how a given view will handle the case and provide a brief summary of what we think the important upshot of the case is. For those happy to trust us, they can read the brief summary and ignore the detailed description of how and why we end up there.

II Local Relevance

In this section we will briefly outline the ambiguity which OLA identifies in the Relevance View, and the principles which OLA shows different versions of the Relevance View to violate. (We will use these principles to test Local Relevance and the

particular version of it that we will focus on, Sequential Claims-Matching, throughout the paper.) We will also introduce in greater detail Tadros' Local Relevance idea.

For the Relevance View to get off the ground, it needs to distinguish between those claims that are relevant within a choice situation and those that are not. For a claim to be relevant, it needs to be sufficiently strong in comparison with some other claim. Let us call this latter claim the 'anchoring claim'. In standard cases in the literature, it is always clear what the 'anchoring claim' is. For example, in Case 2, the 'anchoring claim' is clearly Jones's claim against suffering agonizing shocks, whilst in Cases 1 and 3 the 'anchoring claim' is the claim against death. OLA exposes an ambiguity in the Relevance View. There are two different anchoring rules we might endorse. In simple cases, like Cases 1, 2, and 3, these two anchoring rules identify the same anchoring claim. But when we consider more complex cases, these two rules suggest different answers to the question of which claim is the anchoring claim. These two anchoring rules are:

Anchor by Strength: in order to be relevant, a claim must be sufficiently strong relative to the *strongest overall* claim in the competition.

Anchor by Competition: in order to be relevant, a claim must be sufficiently strong relative to the strongest claim *with which it competes*.

OLA shows that these anchoring rules violate the following compelling principles. Anchor by Competition violates:

Equal Consideration for Equal Claims: all claims of equal strength ought to be given equal weight in determining which group to save.

Anchor by Strength violates:

The Principle of Addition: merely adding a claim to a group of claims cannot *lessen* that group's choice-worthiness, compared with a fixed alternative.

In addition, Anchor by Strength violates a similar though importantly different principle that OLA does not discuss:

The Principle of Strengthening: merely strengthening a claim within a group of claims cannot *lessen* that group's choice-worthiness, compared with a fixed alternative.

In his 'Localized Restricted Aggregation', Victor Tadros introduces an important distinction between 'Global Relevance' and 'Local Relevance' (Tadros, 2019). Under Global Relevance, claims that are judged to be irrelevant to an anchoring claim are irrelevant to the overall decision concerning which group to save. By contrast, under Local Relevance, if a weaker claim is judged irrelevant in comparison with an anchoring claim, while the weaker claim cannot counter-balance or outweigh the anchoring claim, it can counter-balance or outweigh other, weaker, claims with which it competes and thus remain part of the overall decision about which group to save. In sum, according to Global Relevance claims are either relevant or irrelevant simpliciter to a decision. According to Local Relevance, claims can be relevant to some competing claims but not others. Both Anchor by Competition and Anchor by Strength are Global Relevance views.

The core idea of Local Relevance is interesting, but it admits of many potential interpretations.⁵ We know that claims are not relevant or irrelevant simpliciter, but this does not tell us how to decide between two groups of claims in messy cases where the groups are made up of claims of diverse strength. It is clear from Tadros' paper that we must 'match up' claims, allowing them to counter-balance, or neutralize, claims to which they are relevant. However, there are many ways that we might go about doing this 'matching'. In the next section, we will introduce a more precise version of Local Relevance, which specifies how to match claims, before testing it against the OLA principles.⁶

III Sequential Claims-Matching

The central insight of Local Relevance is that claims are not to be judged relevant or irrelevant simpliciter. This insight is captured by this more precise view, Sequential Claims-Matching, suggested to us by Garrett Cullity.⁷ Sequential Claims-Matching provides a procedure through which we can decide which of two competing groups to save. Furthermore, this way of proceeding, in which we start with the strongest claim in the competition, as if it has a pro tanto claim to be saved which must be matched or defeated by claims relevant to it, seems like a natural extension of the Limited Aggregation approach. Limited Aggregation seeks to walk the middle ground between Anti-Aggregation and Pure Aggregation, and to combine personal and impersonal perspectives.⁸ According to the strongest claim a pro tanto claim to be saved seems to do justice to the personal, Anti-Aggregative, perspective.

Sequential Claims-Matching

- I. Identify the strongest claim-type T_1 . Does one group contain more individuals with claims of type T_1 than the other?
 - a. If not, eliminate all T_1 -claims from consideration.
- II. If so, match each T_1 -claim from the group with fewer T_1 -claims to a T_1 -claim from the group with more T_1 claims, and remove the matched claims from consideration.
- III. Now consider the remaining T_1 -claims. Does the other group contain claims of types that are relevant (i.e., sufficiently strong relative to) claim-type T_1 ?
 - a. If not, you should decide in favor of the group with the remaining T_1 claims.
- IV. If so, do the relevant competing claims outweigh the T_1 -claims?
 - a. If not, you should decide in favor of the group with the remaining T_1 claims.
- V. If so, match the set of remaining T_1 claims to a set of relevant competing claims with comparable weight, and remove the matched claims from consideration.
- VI. Now consider the remaining unmatched claims. Of these, identify the strongest claim-type T_2 . Repeat the above procedure.

- VII. Continue until either:
- a. one group contains unmatched claims, in which case you should decide in favour of that group; or
 - b. neither group contains unmatched claims. Then it is not the case that you should decide in favour of one group over the other (though you must save one).⁹

Sequential Claims-Matching is, as you can see, complicated.¹⁰ Nevertheless, the key idea is fairly simple: can the anchoring claim be matched by claims relevant to it? If not, we should meet the anchoring claim. If so, the anchoring claim is ‘matched’ to claims relevant to it, and all those claims are then set aside, since they counter-balance one another.¹¹ We now identify the strongest remaining claim, that becomes the new anchoring claim, and the process begins again.

Before we proceed to testing Sequential Claims-Matching against the OLA principles identified above, it may be helpful to show how it works in a simple case. This will also allow us to introduce our way of presenting cases. In the following table, on the left-hand side is the strength level of the claim, Level 1 being the strongest. Under headings Group A and Group B are the numbers of claims at each level within each group, each claim being held by a distinct individual. For shorthand, we will refer to claims at Level 3 in Group B as B₃ claims.

Suppose that relevance extends two levels up. That is, Level 3 claims are relevant to Level 1 claims, but Level 4 or Level 5 claims are *not* relevant to Level 1 claims. Furthermore, imagine that two claims at Level *x* precisely match one claim at the level above. So, for example, one claim at Level 1 will be precisely matched by two claims at Level 2, and those will be precisely matched by four claims at Level 3. Unless we state otherwise, these stipulations also apply to all cases hereafter.

Here is our initial case:

Case 4.

Level	Group A	Group B
1	1	
3		5
5	3	

Here is how Sequential Claims-Matching would handle this case.

First, it would identify A_1 as the initial anchoring claim. There are no B_1 claims to match A_1 with, so we see whether there are weaker but relevant claims in Group B. There are – the five B_3 claims. Four of these B_3 claims will precisely match the A_1 claims, so we take these four B_3 claims and the A_1 claim which are then set aside, as they counter-balance one another. The one remaining B_3 claim then becomes the new anchoring claim. There are no A_3 claims with which to match the remaining B_3 claim, and so we look to see if there are weaker but relevant claims in Group A. There are – the three A_5 claims. However, three A_5 claims do not match one B_3 claim (since four Level 5 claims precisely match one Level 3 claim), and so we would save Group B in this case.

Sequential Claims-Matching has attractive implications in key cases that OLA uses to undermine Anchor by Competition and Anchor by Strength. Against Anchor by Competition, OLA (240–241) shows that adding equal numbers of equally strong claims to two groups will force us to switch groups. Even stranger, when one additional claim is added to one group (Group A), and a billion equally strong claims are added to another (Group B), Anchor by Competition may require us to switch from saving Group B to saving Group A (OLA: 242). Both of these counter-intuitive entailments result because Anchor by Competition violates Equal Consideration for Equal Claims: it allows that claims in one group will be relevant, while claims of equal strength in another will not, as they compete with different ‘anchoring claims.’

Sequential Claims-Matching does not have these odd implications in the cases that OLA considers. This is because it allows that as soon as claims of strength X in one group are relevant, all such claims are relevant.

OLA (244–247) shows that Anchor by Strength has even odder implications: adding a very strong claim (such a claim against death) to a group of weaker claims would force us to switch *away* from saving that group, even if the competing group is not altered at all. This violates the Principle of Addition. Again, Sequential Claims-Matching would avoid this implication. Since relevance is determined by whether or not a claim is relevant to claims it competes with, adding a single claim could not suddenly rule as irrelevant claims within its own group.

Thus far we see that Sequential Claims-Matching is a better-specified version of Tadros' Local Relevance view, and that it is preferable to both Anchor by Competition and Anchor by Strength in that it provides more intuitively attractive judgments in the key cases presented in OLA. In the rest of the paper, however, we will show that Sequential Claims-Matching faces difficulties and ambiguities of its own.

IV Sequential Claims-Matching and Sore Throats

Recall that there are two famous objections to Pure Aggregation. The first, exemplified by Scanlon's World Cup Case, is the *one vs. many* case. As well as handling the more complex cases introduced in OLA and Section IV of the present paper, Sequential Claims-Matching would have no problem with such cases. The strongest claim (the 'one') would be the initial anchoring claim, and it would not be matched by any relevant claims (the 'many' being irrelevant) and so we would meet the strongest claim (e.g., we would save Jones in Case 2).

The second objection, exemplified by Kamm's Sore Throat Case, is the *tie-break* case. Here, however, we will show that Sequential Claims-Matching as we have defined it thus far would not be able to vindicate Kamm's intuition and would allow the sore throat to break the tie. We will then consider four possible amendments to Sequential Claims-

Matching which will allow us to vindicate the intuition, but which are found wanting in other, more serious, ways. Therefore, we recommend rejecting the Sore Throat intuition, as the lowest caliber bullet of those on offer.

Case 3 Upshot: Sequential Claims-Matching will recommend saving B and C (a life and a sore throat) over A (a life).

Consider how Sequential Claims-Matching would handle Case 3 (Kamm's Sore Throat Case). Step I of Sequential Claims-Matching requires that we identify the strongest claims, and Step I.a requires that where there are equal numbers of such claims, they are matched to each other and set aside. So, we would put aside A's claim against death, and B's claim against death. Step VI of Sequential Claims-Matching states that we should then identify the next strongest claim, and that becomes the new anchoring claim. In Case 3, this would be C's claim against a sore throat. That claim is unmatched and so we would save B and C over A according to Step VII.a.

Therefore, as it stands, Sequential Claims-Matching only meets one of the two prominent objections to Pure Aggregation that in turn motivate the desire to find a plausible version of Limited Aggregation. The original Relevance View, which relied on Global Relevance, ruled out considering small claims such as sore throats, when much more serious claims were on the table. However, as a Local Relevance view, Sequential Claims-Matching never takes claims *fully* off the table. It therefore allows C's sore throat to act as a tie break in Case 3.¹²

Could Sequential Claims-Matching be altered so that it can vindicate common intuitions in the Sore Throat case? We will now consider four ways to alter or supplement Sequential Claims-Matching so that it can avoid having this counter-intuitive implication in the Sore Throat case.

Response 1

The first, seemingly obvious, way that we could try to fix this problem would be to amend Step VI. The problem arises because the two claims against death counter-balance each

other out, leaving us with an unmatched claim against a sore throat, which becomes the new anchor. But sore throats are irrelevant to claims against death, and so it seems troubling, from a Limited Aggregation perspective, to allow a sore throat to determine who lives and who dies. Claims that have been ‘set aside’ should still be allowed to determine whether remaining claims are relevant. We could alter Sequential Claims-Matching so that when it comes to considering ‘new anchoring claims’, we only consider those claims which are relevant to the weakest of those claims that have been previously matched and set aside. For example, in Case 3, A and B’s claims may have been set aside, but since C’s claim is relevant to neither of these, it should not be allowed to become a ‘new anchor’. Consider this replacement for Step VI:

Step VI*: Now consider the remaining unmatched claims. Are any of them relevant to the weakest claim that has previously been matched?

- a. If so, identify the strongest claim-type T₂. Repeat the above procedure.
- b. If not, then it is not the case that you should decide in favour of one group over the other (though you must save one).

In Case 3, Step VI* would not allow C’s sore throat to condemn A to death. However, while Step VI* can save the Sore Throat Case intuition, it comes at a price too high to pay: Sequential Claims-Matching using Step VI* would violate the Principle of Addition and the Principle of Strengthening.

To see this, consider Case 5. Case 5 is a two-stage case. Initially we have a comparison between two groups. Then some additional claims are added to one or both groups. (See OLA and Horton, 2018 for use of this method). From here onward, numbers in brackets represent claims added at the second stage.

Case 5.

Level	Group A	Group B
1	1	

2		2
3		(1)
5	5	

Case 5 Upshot: Sequential Claims-Matching (incorporating Step VI*), would initially not require us to save one group over the other. By adding a claim to Group B, we would then be required to save Group A.

According to Sequential Claims-Matching using Step VI*, we would handle this case in the following way. At Stage 1, A₁ would be the initial anchoring claim. That would be precisely matched by the two B₂ claims. So, the A₁ claim and the B₂ claims would be set aside. The remaining five A₅ claims are not relevant to the B₂ claims, and so we would end in a tie.

At Stage 2, we add in the single B₃ claim. This would fundamentally alter the structure of the case. For now, after having matched A₁ to the B₂ claims, we *would* have a new anchoring claim, namely the B₃ claim. This would qualify as a new anchoring claim as it is relevant to the previously matched B₂ claims. This B₃ claim, however, would ‘activate’ the previously irrelevant A₅ claims, which would not only match but outweigh the B₃ claim, so we would save Group A. Therefore, by adding a claim to Group B we would switch from a tie, and viewing both groups as equally choice-worthy, at Stage 1, to saving Group A at Stage 2. This would clearly violate the Principle of Addition, which, recall, states that:

The Principle of Addition: merely adding a claim to a group of claims cannot *lessen* that group’s choice-worthiness, compared with a fixed alternative.

A similar case (call it Case 5*) shows that Sequential Claims-Matching would also violate the Principle of Strengthening. Imagine that instead of adding the B₃ claim, it was initially a much weaker claim (for example, a Level 15 claim) that became much more serious at Stage 2. If it becomes a B₃ claim at Stage 2, the same process described above would go through, and so the claim moving from B₁₅ to B₃ would force us to switch

from viewing both groups as equally choice-worthy to choosing Group A, in violation of the Principle of Strengthening.¹³

Given that this way of vindicating the Sore Throat intuition comes at such a high price, we are back to allowing sore throats to break ties between equally-matched claims against death.

Response 2

Here is a second way in which Sequential Claims-Matching might be amended to avoid it having to allow the sore throat to make the difference in Case 3, without violating the Principles of Addition and Strengthening.¹⁴ This amendment retains and builds upon the amendment examined in Response 1 (i.e., Step VI* is retained). Return to Case 5. What goes wrong there, for Sequential Claims-Matching (amended to incorporate Step VI*), is that the addition of the B₃ claim ‘activates’ the five A₅ claims, which not only match the B₃ claim, but *outweigh* it. This means that we go from a situation in which both groups were regarded as equally choice-worthy to a situation in which Group A is favored, even though a claim was added to Group B.

This could be avoided if we refuse to allow the A₅ claims to *outweigh* the B₃ claim, but rather only allow them to ‘match’ or ‘disable’ the B₃ claim. If the role of the A₅ claims is restricted in this way, then all they can do is counter-balance the B₃ claim, preventing it from tipping things in Group B’s favor, and so the addition of the B₃ claim would leave us where we were after Stage 1, namely in a tie.

In his original formulation of the idea of Local Relevance Tadros states that while a claim may lack force against a stronger claim to which it is not relevant, it can still ‘retain such force for other aspects of the overall decision [about which group to save]. For example, it may counterbalance other [claims].’ (Tadros, 2019). Tadros leaves it open here what kind of roles weaker claims might play but gives the *example* of counter-balancing. The amendment to Sequential Claims-Matching we are considering here would restrict weaker claims to *only* counter-balancing stronger claims, therefore not allowing them to

outweigh them. So, in Case 5, at Stage 2 all the A₅ claims can do is match the added B₃ claim (even though there is enough of them to outweigh the B₃ claim in a ‘straight fight’).

Note that the claim here cannot be that weaker claims are *always* restricted to counterbalancing rather than outweighing stronger claims to which they are relevant. If that were the case, Sequential Claims-Matching would deliver the implausible recommendation in Case 1 that no matter how large the number of people facing paralysis got, we would only ever consider them to have matched the claim against death. So, we need a more precise and complex articulation of this limited role for weaker claims. We propose this:

Restricted Role: Claims can only counter-balance, but not outweigh, claims to which they are relevant if they are also competing with claims with which they are not relevant.

This may sound confusing, but again a simple example will illustrate the point of Restricted Role. Restricted Role would allow the claims against paralysis to outweigh the claim against death in Case 1, but would only allow the A₅ claims to match, but not outweigh, the B₃ claim in Case 5, because the A₅ claims also are in competition with the B₂ claims. This makes sense within the Local Relevance view, because while the A₅ claims are relevant to the B₃ claim, they are not relevant to the competition between the A₁ claim and B₂ claims, and so, we might think, should only be able to neutralize the B₃ claim, without affecting the competition between the stronger claims.

However, we have concerns about this view. Consider, first, the following case:

Case 6.

Level	Group A	Group B
1	1	
3		5
5	(10)	(10)

Case 6 Upshot: At Stage 1, Sequential Claims-Matching with Restricted Role would instruct us to Save Group B. We then add equal numbers of equally strong claims to both groups. This leads to a tie at Stage 2.

At Stage 1, the five B₃ claims are relevant to, and outweigh, the single A₁ claim. We then add, at Stage 2, ten Level 5 claims to each side. We would expect adding equal claims like this to leave everything else the same. Indeed, it is precisely this intuition on which Tadros relies in order to show how Local Relevance performs better than Global Relevance (in a case he calls ‘Adding People’ (Tadros, 2019: section III)). However, on the version of Sequential Claims-Matching we are considering here, adding these equal claims would result in a tie.

The A₁ claim is the initial anchoring claim. This is matched by four B₃ claims, leaving one B₃ claim as the new anchoring claim. This is matched by four A₅ claims, leaving six A₅ claims as the new anchoring claims. These are matched by six of the ten B₅ claims, leaving four B₅ claims unmatched. Under Sequential Claims-Matching as we originally considered it, these four B₅ claims would carry the day in favor of Group B.

However, under the current proposal, they are irrelevant, and will not be counted. Since B₅ claims compete with an A₁ claim to which they are irrelevant, they can only disable or match claims to which they are relevant, and so beyond the six B₅ claims required to match the A₅ claims, B₅ claims can make no difference. Thus, we would have a tie, in violation of the key intuition that has been used to show why some version of Local Relevance is necessary, that adding equally strong claims of equal numbers to the groups should not alter which group is to be saved.¹⁵

In addition, this violates Equal Consideration for Equal Claims. *All* of the A₅ claims matter – the more A₅ claims we have, the more that speaks in favor of saving Group A. A₅ claims can not only match but *outweigh* the B₃ claims. B₅ claims, however, matter only insofar as they match A₅ claims, and so six of them are not counted.

Some have declared themselves happy to let go of Equal Consideration for Equal Claims. It is surely a far weaker bullet to bite than, say, the Principle of Addition. But it isn't just the principle that should cause us to reject Response 2. The counter-intuitive cases, independently, weigh heavily against it. As we have already noted, Case 6 is a variant of the kind of case that was used as a rationale for Local Relevance. Further, imagine, at Stage 2, that we added five A₅ claims and one million B₅ claims (call this Case 6*). The five B₅ claims would be treated equally to the five A₅ claims, since it would match it. But the 999,995 other B₅ claims would be superfluous. We would go from saving Group B to a tie, even though we had added equally strong claims to both sides, but with numbers heavily in favor of Group B. This is counter-intuitive.

This view would also leave us with a lot of ties: in going down the chain of matching claims to claims, as soon as we got past the point at which the claims were relevant to the strongest claim, we would be in 'matching only' territory. Complex, real-world, cases, would therefore almost always end in a tie.

Response 3

A third way of amending Sequential Claims-Matching so as to capture the Sore Throat intuition, is suggested by the way that Tadros attempts to escape this conclusion. Tadros tries to use the fair procedure of tossing a coin between A and B to exclude or outweigh C's claim. It is plausible that the value of a fair procedure would explain our intuitions in the Sore Throat case. After all, for many, the intuition is not merely that you are not required to save B and C over A, it is that you are required to toss a coin.¹⁶

Tadros tries to bring in this kind of consideration via two routes. The first is to claim that A has an interest in a chance of being saved, and C's claim is irrelevant to A's interest in having that chance. The second route is to claim that tossing a coin is a fair procedure that has independent value, and that C's claim is not sufficiently strong to outweigh that value.

These two routes look importantly different. The first incorporates interests in chances to be saved *within* the Sequential Claims-Matching framework: A's interest in having a

chance at survival is strong enough to deem C's claim *irrelevant*.¹⁷ The second claims that fair procedure is an independent value, and the difference between the groups (i.e., C's claim) is not sufficient to warrant abandoning that value. This does not rely on C's claim being irrelevant.

Incorporating a concern for chances within the Sequential Claims-Matching procedure is undermotivated and leads to counter-intuitive results. Meanwhile, relying on fair procedure, or the distribution of chances, as an independent value is not so much an amendment of Sequential Claims-Matching as downgrading it to only pro-tanto guidance, in competition with other considerations.

If this is right, then our intuitions in tie-break cases can only be explained by entirely different considerations from those that explain our intuitions in one vs. many cases. Scanlon's World Cup viewers have irrelevant claims; whilst Kamm's sore throat victims have relevant claims, but those claims (or the difference between the two groups) are outweighed by some alternative value. Recall that the original Relevance View appeared to explain both intuitions. Sequential Claims-Matching is only able to vindicate the one vs. many intuitions – something else, on this third response, must take care of the tie-break cases.

So, only Tadros' first route is an *amendment* to Sequential Claims-Matching that could see it vindicate the Sore Throat intuition from the inside, by ruling C's claim to be irrelevant. There seem to be two ways to understand A's claim to holding a lottery – via A's interest in the chance of being saved, or via A's moral claim to a fair procedure.

If A's claim to a chance to be saved is based on the cost of going from a 50% chance to be saved to a 0% chance of being saved, this would surely be counter-balanced by B's interest in going from a 50% chance to be saved to a 100% chance to be saved. B has an equally strong interest in *avoiding* the lottery as A has in holding the lottery. This would still leave C as a tie-breaker.

If A's claim to a 50% chance of being saved is based on that being *fair*, then it must be that fairness declares that a fair coin is the correct procedure *prior to C's claim being dismissed as irrelevant*. But this is far from uncontroversial. In addition, once we *moralize* claims, allowing them to depend on fairness for example, we have moved a long way from the Relevance View and the account of claims with which we began.

Even if A's interest in a 50% chance can be explained, however, this amendment has counter-intuitive implications:

Case 7.

Level	Group A	Group B
1	1	1
2		
3		(1)
4	3	

Case 7 Upshot: At Stage 1, the groups are equally choice-worthy. We then add a claim to Group B, and this addition requires us to save Group A.

At Stage 1, A₁ and B₁ have equal claims. According to Tadros, the A₄ claims are irrelevant to B₁'s claim *and* (let us stipulate), B₁'s claim to a fair procedure. So, we would view both groups as equally choice-worthy and hold a lottery. However, at Stage 2, within the Sequential Claims-Matching framework, the addition of the B₃ claim changes things. It is relevant to A₁'s claims, and so, absent other considerations, would act as a tie-breaker. However, it would also 'activate' the A₄ claims, as they are relevant to it, and we would then move to the usual 'back and forth' of Sequential Claims-Matching. The A₄ claims would not only be relevant but would outweigh the B₃ claim, and so we would save Group A. So, adding a claim to Group B at Stage 2 would move us from a lottery to saving Group A. This violates the Principle of Addition.

Response 4

An alternative way to save Kamm's Sore Throat intuition within Sequential Claims-Matching would be to endorse the idea of 'rough equality' (see Hirose 2015: 35). That is the idea that the scales of justice should not be finely tuned. As we go down the Sequential Claims-Matching process, so long as the groups are held to be 'roughly equal', then it is not the case that we must save one group over the other. If we consider Case 3, we can see that the claims of A and B are precisely matched, and so the addition of C's sore throat doesn't take us out of the zone of 'rough equality', and so it is not true that we must save B and C over A.

We have several concerns about this line of thinking. One problem that advocates of 'rough equality' will surely come up against is that in specifying this view they will need to decide whether what counts as 'rough equality' is a constant or contextual matter. To explain: in Case 3, if the 'rough equality' view is to save the Sore Throat intuition, one sore throat cannot tip the scales. But is this because of the broader context (where two lives are also at stake) or because sore throats never tip scales? On the latter view, under which sore throats cannot be difference-makers, we would end up viewing two sore throats as 'roughly equal' to one. This seems implausible. But on the other view, under which what counts as 'rough equality' is contextual, or proportional, to what is at stake, then we will end up viewing 1001 lives as 'roughly equal' 1000. We find this counter-intuitive, though some have endorsed this claim (Kamm, 1998: 103 though see Kamm, 2007: 33).

Our second response to 'rough equality' is to wonder *why* (other than it vindicating the tie-break intuition) we should prefer 'scales of justice' that ignore seemingly morally-relevant information. When pressed on this, advocates of the view seem to rely on the following kind of case. Imagine you could save the life of a twenty-year-old or a twenty-one-year-old. Both have the same life expectancy. Therefore, one is deprived of one more year of life than the other.¹⁸ If this translates into a stronger claim for the twenty-year-old, then, without rough equality, we should save the twenty-year-old over the twenty-one-year-old. Many find this counter-intuitive.

While we agree with the intuition that we are not required to save the twenty-year-old in this case, we do not think this case provides much backing for the idea of ‘rough equality’ in cases of multiple claims. In essence, the issue is that the two cases raise two importantly different questions. Case 3 raises the question of how we should balance different people’s claims of different strengths against one another. The question raised when choosing between two people of roughly similar ages, however, raises the question of how to establish the strength of a claim. All those who reject Pure Aggregation will endorse the idea that the ‘separateness of persons’ matters in some way, and so we should be careful in reading straight across from a case in which what is at stake is *how strong* a claim is to a case in which what is at stake is whether, or to what extent, a claim *counts*. For example, we don’t think the choice between saving a twenty-year-old or a twenty-year-old is the *same* as the choice between saving a twenty-year-old or a twenty-year-old *and* a person with one year of quality of life left (if they are saved).

However, even if we accept ‘rough equality’, problems remain. One, which we will describe in detail below as ‘A Further Problem’, is that ‘rough equality’ will still sometimes allow one headache to determine who lives and who dies. Another is that even if ‘rough equality’ explains why we shouldn’t let the headache break the tie in Case 3, then the *rationale* for tossing a coin in Case 3 is *entirely distinct* from the rationale for saving Jones in Case 2. In other words, even if Sequential Claims-Matching combined with a commitment to ‘rough equality’ saves the anti-aggregative intuitions in Cases 2 and 3, there is no single underlying rationale that explains both.

The original Relevance View, if accepted, would explain both intuitions. But ‘rough equality’ and Sequential Claims-Matching are wholly distinct. One can accept Sequential Claims-Matching without accepting ‘rough equality’, and one can accept ‘rough equality’ without accepting Sequential Claims-Matching, or indeed any Limited Aggregation view – ‘rough equality’ could be attached to Pure Aggregation, for example. Indeed, not only are Sequential Claims-Matching and ‘rough equality’ distinct positions, one might even say that they are somewhat in tension with one another. Sequential Claims-Matching seems to offer us a precise method for determining which group to save. To throw the vague idea of ‘rough equality’ over the top of that procedure may undermine its merits.

A Further Problem

Thus far in this section we have been considering whether Sequential Claims-Matching can accommodate the intuition that a sore throat cannot decide who lives and who dies in a tie break case. We have considered four ways of doing so, and there are, at the least, concerns with each. However, even if Sequential Claims-Matching can be modified or supplemented such that we can avoid this entailment in tie break cases, no version of Sequential Claims-Matching can escape the entailment that a single sore throat will sometimes decide who lives and who dies. Consider, for example, Case 8.

Case 8.

Level	Group A	Group B
1	1	
2		3
3	3	
4		3
5	3	
...
99	3	
100		2 (+1)

Case 8 Upshot: At Stage 1, we have a tie. At Stage 2, through the addition of another B100 claim (a very weak claim), the tie is broken and we save Group B. As such one very weak claim determines whether the strongest claim in the competition, A1, is met.

Imagine in between Level 5 and Level 100, the pattern of three claims in Group A at each odd-numbered level, and three claims in Group B at each even-numbered level continues. Sequential Claims-Matching (as originally conceived) would then handle this case as follows. The A1 claim would be the initial anchoring claim. Two B2 claims would match it, leaving one B2 claim as the new anchoring claim. That would be matched by

two A₃ claims, leaving one A₃ claim as the new anchoring claim. This process would continue all the way down, until one A₉₉ claim was left as the anchoring claim. This would be matched by the two B₁₀₀ claims, leaving us with a tie. However, at Stage 2, if we add another B₁₀₀ claim – an extremely weak claim – this would tip things in favor of Group B, meaning that that extremely weak claim is decisive. Therefore, whether or not that additional B₁₀₀ claim is present could decide whether A₁ lives or dies.

This entailment could be avoided if we endorse the Restricted Role for weaker claims, considered in Response 2, since then the additional B₁₀₀ claim added at Stage 2 could only counter-balance the A₉₉ claims, which were already counter-balanced. However, consider a modified version of the case (Case 8*), in which we start with only one B₁₀₀ claim. This single B₁₀₀ claim would fail to counterbalance the A₉₉ anchoring claim, and so we would save Group A at Stage 1. Adding in a second B₁₀₀ claim at Stage 2 would counter-balance the A₉₉ claim, and so take us to a tie. So, even if we accept Restricted Role, the single B₁₀₀ claim would again be decisive, even though it is extremely weak in comparison with the strongest claims in play.

Of course, ‘rough equality’ would still rule out a single headache making the difference in Case 8, but all ‘rough equality’ can do is to kick this can further down the road. If we keep adding headaches to Group B, at *some point* we will emerge out of the zone of ‘rough equality’ and find ourselves in the zone of inequality. Therefore, while it may be difficult to find the exact place where it occurs, a single headache will still take us from a tie to saving Group B.

Summary of Section V

In this section we have shown how Sequential Claims-Matching as we originally formulated it cannot capture Limited Aggregation intuitions in tie break cases. We then considered four ways in which Sequential Claims-Matching could be amended in order to get around this problem, but each seemed to generate further difficulties for the view, causing us to violate the very principles which OLA had shown the Relevance View to violate. In addition, in non-tie-break cases such as Case 8, we still wouldn’t be able to avoid a single very small claim making all the difference. Our sense here, is that the

advocate of Limited Aggregation should reject the tie break intuition, though we accept that not all advocates of Limited Aggregation will feel able to do so.

Even if Sequential Claims-Matching is accepted as we originally formulated it, or some alternative version is put forward, however, two key ambiguities about the view remain, which we will now describe.

VI Ambiguity 1: How Should We Match Claims in Sequential Claims-Matching?

We have noted how Tadros' explanation of Local Relevance view seems to require us to 'match up' claims with one another, so that they can counter-balance other claims to which they are relevant. Sequential Claims-Matching is a clear advance on the loosely-specified Local Relevance view because it clearly outlines a process by which we should match up claims. Nevertheless, it still contains an ambiguity concerning how to 'match up' claims. The relevant steps within the Sequential Claims-Matching view are as follows:

III. Now consider the remaining T_1 -claims. Does the other group contain claims of types that are relevant (i.e., sufficiently strong relative to) claim-type T_1 ?

- a. If not, you should decide in favor of the group with the remaining T_1 claims.

IV. If so, do the relevant competing claims outweigh the T_1 -claims?

These steps show how anchoring claims must be compared with claims that are relevant to them. However, more than one kind of claim can be relevant to a single anchor, and in such a case we need a procedure to decide which of the weaker claims to compare with the anchor. Consider this case:

Case 9.

Level	Group A	Group B
1	1	
2		3
3	2	
4	4	
5		
6		20

Case 9 Upshot: A B2 claim can be matched with *either* A3 or A4 claims. Whether the B2 claim is matched with A3 or A4 claims determines which group is saved. Sequential Claims-Matching provides no guidance as to which is preferable.

The A1 claim would be the initial anchoring claim. This claim is matched by two B2 claims, leaving the remaining B2 claim as the new anchoring claim. Here is where the ambiguity concerning how to match up claims comes in. Both the A3 and A4 claims are relevant to the B2 claim. Either the two A3 claims or the four A4 claims would exactly counterbalance the one remaining B2 claim. However, whether we match the B2 claim either with the two A2 claims or the four A4 claims will decide which group we will save.

Suppose, first, that we match the B2 claim with the two A3 claims. These are all the A3 claims there are, and so we move one level down in order to identify the next anchoring claim. The A4 claims become the new anchoring claims. What is crucial here is that the twenty B6 claims are relevant to the A4 claims and outweigh them. So, matching in this way means that we would save Group B.

But now let's suppose that when we have the remaining B2 claim as the anchoring claim that we match it with the four A4 claims. That would mean that the only claims that are left in the competition are two A3 claims and twenty B6 claims. But B6 claims are not relevant to A3 claims. There could be any number of B6 claims but they could never

outweigh a single A₃ claim. So, the two A₃ claims are unmatched. Therefore, were we to match claims in this way, we would save Group A.

This case shows that there is an ambiguity in how we match claims on Sequential Claims-Matching that could be decisive in deciding which group to save. What Sequential Claims-Matching requires, therefore, is a principled way in which to determine how we match claims in this kind of case. We have been able to conceive of three rival possibilities, all of which have some plausibility, but which would give conflicting advice in Case 9.

The first possibility is the following:

Match to the Strongest Competing Claim: Anchoring claims should be matched with the strongest unmatched claims with which they compete.

This rule has some plausibility. Sequential Claims-Matching in general has a structure which encourages us to look at the ‘next level down’ – that is, after all, the next claim in the sequence. It also seems a non-arbitrary and clear way to decide how to match up the claims. In Case 9, Match to the Strongest Competing Claim would have us match the B₂ claim to the A₃ claims, leaving the A₄ claims to be outweighed by the B₆ claims. Therefore, we would save Group B.

The second possibility is the following:

Match in the Interest of the Overall Strongest Claimant: Anchoring Claims should be matched in whatever way is in the interest of strongest overall claimant.

Limited Aggregation generally has, of course, much in common with Anti-Aggregation. Anti-Aggregation is concerned, above all else, with the person who possesses the strongest claim. Limited Aggregation doesn’t hold this concern above all else – it would allow the weaker claims to outweigh the strongest claim in Case 1 – but it is plausible to

think that it would inherit this general priority to the worst off. Match in the Interest of the Overall Strongest Claimant says that when there is a plurality of ways of matching up claims within the confines of Sequential Claims-Matching, and where different ways of matching would produce different outcomes (as in Case 9) we should, in essence, allow the person possessing the overall strongest claim in the competition to choose how we should match up the claims. In Case 9, A₁ is the strongest claim. If we match the B₂ claim with the A₄ claims, this leaves the A₃ claims unmatched, since they are 'out of reach' of the B₆ claims. And so, Group A would win, which is, of course, what A₁ would want.

The third possibility is as follows:

One or the Other: when there is an ambiguity over how to match up claims, and different ways of doing so would require saving different groups, then it is not the case that you should decide in favour of one group over the other (though you must save one).

This view takes the ambiguity we have noted in this section not as something to be rectified, but rather as to reveal that we are not required to save one group over the other. Since the Sequential Claims-Matching process can be used to justify saving either group, we cannot say that we are required to save one group over the other. In such circumstances, it may be that we are permitted to choose, or are required to toss a coin.

Advocates of Local Relevance owe us an account of how to match up claims with one another. Sequential Claims-Matching provides such an account. However, there remains an ambiguity in that even within the strictures of that more precise theory, there can be more than one way to match up claims. So, advocates of Sequential Claims-Matching will need to further specify the view.

VII Ambiguity 2: Partial Leftovers

Thus far we have been considering cases in which two claims at one level *precisely* match one claim at the level above. However, it is unlikely that things will always be quite so neat, especially given how small the differences between claims can be. For example, even

if one accepts that two claims at Level 2 will precisely match one claim at Level 1, imagine if one of the Level 2 claims becomes ever so slightly stronger (for example, if it is a claim against partial paralysis, the loss of use of one more finger). Does this remain a Level 2 claim with no additional weight? Surely not. But equally implausible is that it would become a Level 1 claim, and double in weight.

To put things more concretely, many health systems use Quality Adjusted Life Years (QALYs) or Disability Adjusted Life Years (DALYs) to measure the strength of claims. Imagine a case in which a single person facing the loss of 50 QALYs faces three people facing the loss of 20 QALYs each. The single strongest claim outweighs two of the weaker claims and is defeated by three of the weaker claims. 2.5 of the weaker claims precisely match the stronger claim.¹⁹

If whole numbers of weaker claims do not precisely match stronger claims, we need to decide how to handle the ‘partial claims’ that are left over once the matching has taken place. To make things clearer, consider this case. In this case 2.5 claims at Level x precisely match one claim at the level above.

Case 10.

Level	Group A	Group B
1	1	
2		3
3		
4		
5	20	

Case 10 Upshot: Following Sequential Claims-Matching, 0.5 of a B2 claim is left over from the first round of ‘matching’ and becomes the new anchor. But it is ambiguous in Sequential Claims-Matching whether the twenty A5 claims need to be relevant to that 0.5 of a B2 claim or to a full B2 claim. Which route

we take, determines whether the twenty A₅ claims are relevant and so which group we save.

A₁ is the initial anchoring claim. It is matched by 2.5 B₂ claims, seemingly leaving half a B₂ claim as the new anchoring claim. There is an ambiguity about what kinds of claims should be considered relevant to this partial B₂ claim. The key question is whether the twenty A₅ claims are relevant to the half B₂ claim. If it were a full B₂ claim, then there would be no ambiguity, since Level 5 claims are not relevant to Level 2 claims. However, in this case, there is half a Level 2 claim remaining. In order to be relevant, should we demand that the A₅ claims are relevant to a full B₂ claim or a half B₂ claim? This suggests two possible principles.

Full Claim Relevance: To be relevant, a claim must be relevant to the claim with which it competes. (For example, the full B₂ claim.)

Partial Claim Relevance: To be relevant, a claim must be relevant to the proportion of the claim with which it is in competition. (For example, the half B₂ claim.)

On the one hand, we could imagine one of the B₂ claimants objecting to small A₅ claims being considered relevant to her far weightier claim. On the other hand, it seems unfair to expect A₅ claims to have to be relevant to a full B₂ claim when they are only in competition with half a B₂ claim.

In our view, Partial Claim Relevance seems the right view. However, this again raises the issue of sore throats. Imagine a case in which there is a partial leftover of 1/1000 of a very serious claim. Very weak claims, which would only need to be relevant to (and not necessarily equal to) the leftover, could then tip the balance. This adds to the case for advocates of Limited Aggregation letting go of the tie-break intuition.

VIII An Alternative Version of Local Relevance

We began by highlighting how Tadros' Local Relevance was promising. We sought to provide a more concrete version through Sequential Claims-Matching and have discussed at length the merits and problems of this view. One ambiguity we highlighted (in Section VI) concerns how to match up claims with one another. One suggestion we offered was to Match in the Interest of the Overall Strongest Claimant. However, we have only considered this within the confines of Sequential Claims-Matching. But why not match in the interest of the overall strongest claimant more generally? So, whenever we face cases with groups of claims of diverse strength, if there is any way to match the claims such that the strongest claimant would have her claim met, then we should allow this. The basic idea here is that we take from Anti-Aggregation the priority for the worst off and we allow that to dictate the way in which claims are matched. Call this view *Strongest Decides*. Here is an example where this view would deliver a different verdict from Sequential Claims-Matching:²⁰

Case 11.

Level	Group A	Group B
1	1	
3	3	3
5		10

Case 11 Upshot: Sequential Claims-Matching would tell us to save Group B.

Strongest Decides, by contrast, would tell us to save Group A.

Sequential Claims-Matching would first match two B₃ claims with the A₁ claim, leaving one B₃ claim as the new anchoring claim. That would be matched by one A₃ claim, leaving two A₃ claims. Finally, these two A₃ claims would be outweighed by the ten B₅ claims. So, we would save Group B.

Strongest Decides, by contrast, would have the A₁ claimant decide how we match claims in this case. It is in A₁'s interest to have the three A₃ claims be matched with the three B₃

claims, as such taking all these claims out of consideration. This would leave A1's own claim and ten B5 claims. Level 5 claims are not relevant to a Level 1 claim, and so we would save Group A.

Here is a potential problem with Strongest Decides: it may violate Equal Consideration for Equal Claims.²¹ It is clear that, in some cases, there will be more than one way of 'matching up' claims, and the way that would be favored by the person with the strongest claim may be one in which some, say, Level 5 claims are matched with claims to which they are relevant, while other Level 5 claims are matched with claims to which they are not relevant. Therefore, in the 'matching scheme' the strongest claim would select, some Level 5 claims will be counted (as they're relevant to the claims they're matched with) and others will not (as they're not relevant to the claims they're matched with).

While that much is clear, it is not clear whether this in fact violates Equal Consideration for Equal Claims. OLA (241) makes it clear that *any* plausible view will not advocate Equal Treatment for Equal Claims, for when we have two groups with equal claims on both sides, we will end up saving one group and not the other. What OLA objected to, and Equal Consideration for Equal Claims attempted to articulate, was that *at the outset* some views ruled that some claims of a certain strength mattered, and other claims, of the same strength, did not. Strongest Decides does not do this.

Strongest Decides only really comes into play when there is a plurality of ways in which to 'match up' claims with one another. Since equal claims can always be matched to each other, claims of equal strength will never be ruled as irrelevant *at the outset*: there is always a potential 'match up' in claims in which all claims of the same strength are relevant. The strongest claim is able to dictate which 'match up' is chosen *as a tie breaker*. This may involve choosing a 'match up' in which some Level 5 claims are relevant and others are not. Is this like accepting that we will save some Level 5 claims and not others, which is surely acceptable? Or is it like ruling out some Level 5 claims but not others at the outset? Or, is it like neither? Equal Consideration for Equal Claims may need to be further refined. Or perhaps it must be abandoned. Perhaps it didn't properly articulate

what Anchor by Competition was getting wrong. At any rate, this is an issue that Strongest Decides must confront.

X Concluding Remarks

This paper is an exploration. Many of us would like to avoid biting the bullets of the extreme Anti-Aggregation and Pure Aggregation positions. Limited Aggregation therefore seems our best hope. But as OLA shows, the Relevance View (or what Tadros calls ‘Global Relevance’) is open to potentially devastating objections. Tadros’ version of Limited Aggregation, ‘Local Relevance’, is attractive, but can be fleshed out in a variety of ways. It is this idea that we have sought to precisify and explore. We have done this by, first, providing a clearer and more tightly-specified version of the idea, Sequential Claims-Matching, and, second, by subjecting that view to scrutiny. We have found that Sequential Claims-Matching stands up well against alternative versions of Limited Aggregation in terms of the OLA cases. But it suffers from problems and ambiguities of its own.

First, it does not seem able to vindicate some of the key intuitions that have motivated the search for Limited Aggregation cases (those elicited from ‘tie break’ cases). We considered several possible refinements to Sequential Claims-Matching which might allow it to account for these intuitions, but each was found to either be wanting, or to rest the vindication of those intuitions on alternative values from outside the Sequential Claims-Matching procedure. This seems like a considerable cost when compared with Global Relevance Views, like ARC, which were able to explain our judgments about tie-break and one vs. many cases in the same way. One option is for the advocate of Local Relevance to let go of the tie-break intuitions, and, on reflection, this is the route that we ourselves would recommend. What is distinctive about Local Relevance is that no claim is ever ‘ruled out’ altogether. This means that one tiny claim can always make the difference between who lives and who dies. This seems counter to some of the intuitions and commitments that drive the Relevance View but is, in our view, the smallest caliber bullet to bite having precisified the view.

Second, while it is clearer than Tadros' view on how to match claims, there are two important ambiguities within Sequential Claims-Matching: an ambiguity about how to match up claims, and a further ambiguity about how to handle what we have called partial left-overs. Thus, we have shown that Sequential Claims-Matching, at best, needs to be further specified. In addition, we have articulated an alternative version of Local Relevance: Strongest Decides. This view tracks more closely to the Anti-Aggregation view, giving the strongest claim under consideration a defeasible right to be saved. However, this view potentially conflicts with the Equal Consideration for Equal Claims principle. We recommend further scrutiny of this view, and the principle it seems to violate.

In conclusion, we offer an interim report on Local Relevance. The general idea admits of several possible interpretations. We need a clear sense of which versions appear to be the most plausible, both in terms of how they handle key cases, and the sort of deeper justifications which can be offered on their behalf. Even within these versions, there are ambiguities, for example, on how to handle partial leftovers. Second, while Local Relevance offers intuitively plausible responses to some cases, since it does not fully rule out any claims, no matter how weak, as completely irrelevant, it will struggle to handle cases, like tie-break cases, in which one small claim 'tips the balance'.²² Our recommendation is to bite this bullet. The alternative seems to be to downgrade Local Relevance to only being a *part* of the aggregative story. Further work is needed to see whether these challenges can be met, or whether some alternative version of Limited Aggregation, abandoning altogether the focus on 'relevance' is the best hope for those who wish to tread the tightrope between the extremes of Anti-Aggregation and Pure Aggregation.

References

Broome, J. (2004). *Weighing lives*. Oxford: Oxford University Press.

Doggett, T. (2013). Saving the Few. *Noûs*, 47(2), 302–315.

Halstead, J. (2016). The Numbers Always Count. *Ethics*, 126(3), 789–802.

<https://doi.org/10.1086/684707>

- Hirose, I. (2015). *Moral Aggregation*. Oxford ; New York: Oxford University Press.
- Horton, J. (2018). Always Aggregate. *Philosophy & Public Affairs*, 46(2), 160–174.
<https://doi.org/10.1111/papa.12116>
- Kamm, F. M. (1998). *Morality, Mortality Volume I: Death and Whom to Save From It*. Oxford University Press. <https://doi.org/10.1093/0195119118.001.0001>
- Kamm, F. M. (2005). Aggregation and Two Moral Methods. *Utilitas*, 17(01), 1–23.
<https://doi.org/10.1017/S0953820804001372>
- Kamm, F. M. (2007). *Intricate ethics: rights, responsibilities, and permissible harm*. Oxford; New York: Oxford University Press. Retrieved from
<http://public.eblib.com/choice/publicfullrecord.aspx?p=415447>
- Kamm, Frances M. (2015). Bioethical Prescriptions. *Journal of Medical Ethics*, 41(6), 493–495. <https://doi.org/10.1136/medethics-2014-102543>
- Lang, G., & Lawlor, R. (2015). Numbers scepticism, equal chances and pluralism: Taurek revisited. *Politics, Philosophy & Economics*, 15(3), 298–315.
<https://doi.org/10.1177/1470594X15618967>
- Lawlor, R. (2006). Taurek, Numbers and Probabilities. *Ethical Theory and Moral Practice*, 9(2), 149–166. <https://doi.org/10.1007/s10677-005-9004-4>
- MacMahan, J. (2003). *The Ethics of Killing: Problems at the Margins of Life*. New York: Oxford Univ. Press.
- Munoz-Dardé, V. (2005). The Distribution of Numbers and the Comprehensiveness of Reasons. *Proceedings of the Aristotelian Society*, 105, 191–217.
- Nagel, T. (1989). *The View from Nowhere* (First issued as an Oxford University Press paperback). New York, NY: Oxford Univ. Press.

- Nagel, T. (1995). *Equality and Partiality*. Oxford University Press.
<https://doi.org/10.1093/0195098390.001.0001>
- Peterson, M. (2009). The Mixed Solution to the Number Problem. *Journal of Moral Philosophy*, 6(2), 166–177. <https://doi.org/10.1163/174552409X402331>
- Peterson, M. (2010). Some Versions of the Number Problem Have No Solution. *Ethical Theory and Moral Practice*, 13(4), 439–451.
- Ross, D. (2002). *The Right and the Good*. (P. Stratton-Lake, Ed.). Oxford University Press. <https://doi.org/10.1093/0199252653.001.0001>
- Scanlon, T. M. (1998). *What We Owe to Each Other*. Belknap Press of Harvard University Press.
- Tadros, V. (2019). Localized Restricted Aggregation. *Oxford Studies in Political Philosophy*, 5.
- Taurek, J. M. (1977). Should the Numbers Count? *Philosophy & Public Affairs*, 6(4), 293–316.
- Thomas, A. (2012). Giving Each Person Her Due: Taurek Cases and Non-Comparative Justice. *Ethical Theory and Moral Practice*, 15(5), 661–676.
<https://doi.org/10.1007/s10677-012-9358-3>
- Tomlin, P. (2017). On Limited Aggregation. *Philosophy & Public Affairs*, 45(3), 232–260.
<https://doi.org/10.1111/papa.12097>
- Voorhoeve, A. (2014). How Should We Aggregate Competing Claims? *Ethics*, 125(1), 64–87. <https://doi.org/10.1086/677022>

* We are grateful to Mike Otsuka, audiences at LSE and Pavia, two anonymous referees, and Peter Vallentyne for useful feedback. We are especially grateful to Garrett Cullity for the suggestion of Sequential Claims-Matching, and subsequent comments.

¹ Although Voorhoeve's view has this attractive implication in Kamm's case, strictly speaking Voorhoeve (2014: 67) restricts his theory to cases in which each member of a group has the same strength of claim.

² We will refer to this paper as OLA throughout to save Tomlin from having to refer to himself in the third person. Or at least from having to do so again after this note.

³ Alternative ways might involve appealing to a general form of pluralism (Ross, 2002); (Lang & Lawlor, 2015), or specifically by balancing aggregate goodness of satisfied claims with the unfairness of unsatisfied claims Lawlor 2006; Hirose, 2015; Peterson 2009; Peterson 2010).

⁴ Voorhoeve (2014: 79–82) proposes a test based on the personal prerogative. Kamm (2015) rejects this approach.

⁵ In response to Horton (2018), Tadros (2019, Section IV) refined his own version somewhat.

⁶ For an alternative OLA-inspired response to Tadros, see Horton (2018).

⁷ OLA (252–257) suggested a similar 'matching' procedure for 'quasi-competition' cases (cases where some subset of one large group must be saved) but did not consider this kind of procedure for 'competition cases' (where we must save one group or the other). OLA argues that quasi-competition cases are actually more analogous to many policy decisions.

⁸ For similar approaches in the literature that look to combine and/or balance personal and impersonal perspectives, see Kamm's (five different versions of) "Subjectivity" (1998: Chapters 8–10; 2007: 34–36, 39, 50) and Thomas Nagel's (1989; 1995) "View from Nowhere" vs. "View from Somewhere."

⁹ We leave it open here whether in such a situation an agent has a free choice, or whether they must choose via a fair procedure such as a coin toss.

¹⁰ As formulated here, Sequential Claims-Matching only applies to two-group cases. It may become more complicated still for three-group cases. The following seems like the most plausible procedure for three- or more-group cases: first, eliminate all groups with *no* claims relevant to the strongest overall claim. Second, conduct a series of pairwise comparisons using Sequential Claims-Matching. If a single group emerges, save that group. If no single group emerges, then it isn't the case that you must save one group over the others. We are grateful to Garrett Cullity for comments here.

¹¹ 'Set aside' here does *not* mean 'cancelled'. If the claims were cancelled, then two equally strong claims would cancel each other and nobody would have a claim to be saved. We're grateful to Mike Otsuka for useful comments here.

¹² We are grateful to an anonymous referee for helping us to describe the structural differences between Kamm's Irrelevant Utilities view and Sequential Claims-Matching.

¹³ We are grateful to Anna Mahtani and Alex Voorhoeve for useful discussion here.

¹⁴ This has been proposed in discussed by both Victor Tadros and Alec Walen.

¹⁵ In response to Horton (2018), Tadros accepts that sometimes adding claims of unequal numbers but equal strength should cause us to switch to saving the group to which fewer claims have been added. However, as we have stressed here, he accepts the intuition that this is unacceptable in cases such as Case 6 and 6*. Tadros (2019, Sections III-IV).

¹⁶ We are grateful for comments from an anonymous referee here.

¹⁷ Kamm sometimes appears to adopt this reasoning – that the sore throat is too weak a claim compared with the claim to a *chance of life*. At other times she appears to claim that the sore throat is too weak compared to the claim against death. See Kamm (2005: 13), where both types of reasoning are on show.

¹⁸ For competing accounts of the moral relevance of this, see McMahan (2003: 165–174); Broome (2004: 241–253).

¹⁹ For related discussion see Kamm (2005), section VII (15–18). Kamm's discussion clearly indicates that lesser claims must compete with left-overs but is ambiguous between our Full Claim Relevance and Partial Relevance.

²⁰ We draw here on cases presented in Horton (2018), though we put them to different ends here.

²¹ We are grateful to Bastian Steuwer for pressing this worry.

²² There are further potential problems with Local Relevance. In particular Joe Horton (2018) has shown cases in which it appears to violate this principle:

The Principle of Net Addition: Adding claims of equal strength but differential numbers cannot make the group to which more claims are added less choice-worthy compared with the group to which fewer claims are added.

To be clear: this is our attempt (with help from Mike Otsuka) to articulate a principle which Horton's cases show Local Relevance to violate. Horton doesn't try to articulate such a principle.