

Tie-breaks and two types of relevance

Article

Published Version

Creative Commons: Attribution 4.0 (CC-BY)

Open Access

Hart, J. (2022) Tie-breaks and two types of relevance. *Ethical Theory and Moral Practice*, 25. pp. 314-334. ISSN 1572-8447
doi: <https://doi.org/10.1007/s10677-022-10270-3> Available at
<https://centaur.reading.ac.uk/102794/>

It is advisable to refer to the publisher's version if you intend to cite from the work. See [Guidance on citing](#).

To link to this article DOI: <http://dx.doi.org/10.1007/s10677-022-10270-3>

Publisher: Springer

All outputs in CentAUR are protected by Intellectual Property Rights law, including copyright law. Copyright and IPR is retained by the creators or other copyright holders. Terms and conditions for use of this material are defined in the [End User Agreement](#).

www.reading.ac.uk/centaur

CentAUR

Central Archive at the University of Reading

Reading's research outputs online



Tie-breaks and Two Types of Relevance

James Hart¹ 

Accepted: 20 January 2022
© The Author(s) 2022

Abstract

Sometimes we must choose between competing claims to aid or assistance, and sometimes those competing claims differ in strength and quantity. In such cases, we must decide whether the claims on each opposing side can be aggregated. Relevance views argue that a set of claims can be aggregated only if they are sufficiently strong (compared to the claims with which they compete) to be morally relevant to the decision. Relevance views come in two flavours: Local Relevance and Global Relevance. This paper presents a trilemma for both. Namely, that neither view can capture our intuition in tie-break cases, without forfeiting our intuitions in other important cases.

The paper then presents a way to salvage relevance views and capture all our intuitions using a Hybrid view. By distinguishing between two types of relevance we can combine the strengths of Local and Global Relevance views such that we can hold all our intuitions, consistently and in a non-ad-hoc manner. Building on this, the paper demonstrates how we might amend the strongest formulation of a Relevance view, into a Hybrid account.

The main focus of this paper is to develop a relevance view that can capture all the intuitions we have; however deeper justifications will still be needed for a full account of any relevance view. Thus, the end of the paper briefly considers what deeper justification might support the Hybrid view, indicating the direction such literature might go. Lastly, the paper considers the advantages of this view over a rival Hybrid view.

Keywords Aggregation · Limited aggregation · Global relevance · Local relevance

1 Introduction

Questions about aggregation arise when we need to adjudicate between conflicting claims of different strengths. For instance, should a charity spend its limited funds on giving one

✉ James Hart
james.hart@pgr.reading.ac.uk

¹ Philosophy Department, University of Reading, Reading, Berkshire, England

individual life-saving dialysis treatment or some large number of individuals life altering cataract surgery? Does the combined weight of reasons to treat the cataract patients outweigh the reason to save the dying person's life?

Two basic responses have been levied to this issue. The first, Pure Aggregation, states that we should always aggregate, no matter how small the claim. Thus, assuming the number of cataract patients is large enough, Pure Aggregation tells us to treat the cataract patients. The second response, Anti-Aggregation, states that we should not aggregate under any circumstances. Instead, we should make individual pairwise comparisons and act on the strongest claim. Thus, Anti-Aggregation tells us to give the treatment to the individual requiring dialysis. However, both basic responses run into serious problems in different cases.

Firstly, Anti-Aggregation. Consider the following:

'Case 1. You can save one person from death, or some larger number of people, N_1 , from paraplegia.' (van Gils and Tomlin 2019, p. 221).

Because saving someone from paraplegia is such a significant benefit, it seems we should save the people from paraplegia if N_1 is large enough. For instance, if there are a hundred people threatened with paraplegia, then we should save the paraplegics. However, Anti-Aggregation holds that we should never aggregate weaker claims to overpower a single stronger claim. In this way, Anti-Aggregation tells us to save the dying person, even if the number of people you could save from paraplegia was in the thousands, millions or even billions. This conclusion is unacceptable.

Secondly, Pure Aggregation. Consider the following two cases:

'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, p. 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, p. 101)' (van Gils and Tomlin 2019, p. 222).

Because fifteen minutes of missed football is such an insignificant cost and saving Jones would be such a significant benefit, it seems we ought to save Jones from agony in Case 2. The same reasoning explains why it seems we ought to ignore C's sore throat in Case 3. Yet, Pure Aggregation tells us to leave Jones until the match is over, and that C's sore throat is a good enough reason to decide in favour of B. Pure Aggregation has these consequences even if the small costs were fleeting itches or mildly uncomfortable thoughts.

For these reasons there is a growing literature of Limited Aggregation views, which try to straddle the gap between the more extreme Pure Aggregation and Anti-Aggregation views. Such views attempt to find a principled way to capture our intuitions in these cases. The ideal Limited Aggregation view would save the paraplegics in Case 1, Jones in Case 2, and flip a coin in Case 3.

Limited Aggregation usually involves an appeal to relevance. A claim can only be considered if it is relevant, i.e., if it is sufficiently close in size to the claim with which it competes. In Case 1 then the claims to be saved from paraplegia are strong enough to be relevant to the claim to be saved from death, and so we must aggregate. In Case 2, the claims to watch the World Cup uninterrupted are not relevant to Jones' claim to be saved from agonising electrical shocks, so we cannot aggregate. Lastly, in Case 3 the sore-throat claim¹ is not strong enough to be relevant to the decision of whose life to save.

Now, there is significant disagreement as to how to flesh out and justify relevance. This is a large and complex issue which I take to go beyond the scope of the paper.² To avoid committing myself to any particular view, I shall only speak in the most basic terms of respect and leave it to the reader to substitute in their preferred justification. I do this because my view is unchanged by which justification one prefers and couching it in the most basic terms allows us to see the generality of my conclusions. I shall take a claim to be relevant when it would be disrespectful to ignore it, and irrelevant when it would be disrespectful to consider it. To elucidate this further, it would be disrespectful to consider the weaker claim if doing so would trivialise what the stronger claimant would lose and if one cannot provide an acceptable justification to the stronger claimant as to why she is not being saved. Similarly, it would be disrespectful to ignore the weaker claim, if doing so would fail to recognise what is at stake for him and if no acceptable justification can be given as to why his claim is not being considered.

In Case 1 then, because there is a lot at stake for all those with paraplegia claims it would be disrespectful of those claimants if we did not consider attending to their claims. So, respect permits us to aggregate the respective claims, and the paraplegia claims are stronger on aggregate, so we save the paraplegics. In Cases 2 and 3, the opposite is true. What is at stake for the football fans is insignificant compared to what Jones is facing, and so it would be disrespectful to consider attending to the football fans instead of attending to Jones. Similarly, it would be disrespectful to consider attending to a single sore throat rather than giving A an equal opportunity to be saved.

¹ For simplicity, in the rest of the paper I shall speak of 'sore-throat claims' instead of 'claims to be saved from sore throats' etc.

² Voorhoeve parses relevance out primarily from the personal perspective of the weaker claim, if the stronger claim is so much stronger that she could not press her claim, then it would be disrespectful to consider the weaker claim. If, on the other hand, the weaker claim could licitly choose to save herself over the stronger claim then the weaker claim is relevant, and it would be disrespectful to ignore it (Voorhoeve 2014, pp. 70–75). Kamm rejects understanding relevance from the personal perspective of the weaker claim (Kamm 1998, pp. 170–171) (Kamm, 2015), and offers instead a number of different ways to incorporate the personal perspective (Kamm 1998, Chapters 8–10). Kamm's approaches put a greater emphasis on the personal perspective of the stronger claim, but also on what is salient to a disinterested third party. Tadros denies that partial concern provides the right explanation for these respect conditions at all. Instead, he argues that such respect conditions are 'morally basic' and constitute 'a plausible instance of the more general phenomenon of contextual interaction' (Tadros 2019, pp. 177–178).

Thus, we can see that if it is respectful to compare competing claims, then said claims are relevant to each other, but if it is disrespectful to make such a comparison then the weaker claim is not relevant to the stronger claim.

1.1 Global and Local Relevance

Importantly for our purposes, there are two distinct types of relevance views: Global Relevance and Local Relevance (Tadros 2019). Both views handle homogenous cases, cases where each group only contains one type of claim, in the same way. But the two views differ in heterogenous cases, where the strength of claims differ within groups as well as across groups. Global Relevance views consider whether it is respectful to consider a claim in a competition by comparing that claim to the strongest claim with which it competes. If a claim is not relevant to the strongest claim with which it competes then it can play no part in which group to save. In this way Global Relevance decides whether a claim is relevant or irrelevant simpliciter.

Local Relevance views, on the other hand, consider whether a claim is respectful to *each and every* type of claim with which it competes. So, a claim might be irrelevant regarding the strongest claim in the competition, but relevant to a weaker claim in the competition. That a claim is irrelevant to the strongest claim, does not mean it cannot play any role in which Group to save. Thus, if the strongest claims are counterbalanced, then weaker claims can play a role in deciding which group to save.

To make this clear consider how the different views handle the following case:

Case 4. We can save Group A or Group B. In Group A we have a death claim, in Group B we have a death claim and a hundred lost-hand claims.

Let us suppose that our respect-based reasons tell us that lost-hand claims are too weak to be relevant to death claims - no number of lost-hand claims should ever be saved instead of a life. In this case Global Relevance will tell us to flip a coin. This is because the lost-hand claims are irrelevant to the death claim with which they compete, and therefore cannot play any part in deciding who wins.

But this seems wrong. It seems that we should not completely ignore the lost-hand claims in this case. The lost-hand claims, whilst not strong enough to compete and outweigh a death claim, are strong enough to choose between two differing death claims. It seems disrespectful to the people whose hands are in danger, to ignore their claims when all else is equal. Lost hands are unlike sore throats in this way. Sore throats seem to lose all moral relevance when compared to what is at stake for a dying person, but lost hands do not.³ I will return to the question of why this might be in the Deeper Justification section.

Firstly though, the intuition that Case 4 gets at might not be immediately clear. So, let us test it. Suppose you are a doctor deciding whether to send your limited resources to Ward A or Ward B. In Ward A there is a dying man, and in Ward B there are a hundred people with strong, yet (compared to death) irrelevant claims. In this paper I use lost-hand claims, but

³ To put it in Voorhoeve's terms, the partial concern that lost-hand claimants have for their claim does not permit them to prefer their own rescue over the rescue of a dying person, but partial concern does permit them to prefer their own rescue over an equal opportunity of rescue for a competing death claim.

they could be lost-finger claims, lost-limb claims or some other strong claim that is slightly too weak to ever aggregate against a death claim.⁴ Let us call this:

Case 5. We can save Ward A or Ward B. In Ward A we have a death claim, in Ward B we have a hundred lost-hand claims.

Suppose that after some deliberation, you work out that the lost-hand claims are irrelevant and so send the resources to Ward A. However, if the hundred people in Ward B had slightly stronger claims, then their claims would have been relevant and by sheer weight of number significantly outweighed the death claim. Now, just before you send off those resources to Ward A, you are informed that a dying person has been wheeled into Ward B, and you could save her life as well as the hundred. Who should you send your resources to? My strong intuition is that we ought to send the resources to Ward B rather than flip a coin. Whilst the hundred do not have claims strong enough to aggregate against the dying person, they do have claims strong enough to be decisive in a tie-break situation. In fact, it seems that it would be disrespectful to not include the lost-hand claims in our decision when all other features are balanced. Global Relevance cannot capture this intuition - on any Global Relevance view claims are either always or never taken into consideration, there is no gradation.⁵

Local Relevance solves this problem. According to local relevance, the lost-hand claims, whilst irrelevant to the death claims, can remain relevant to what we ought to do all things considered. Because the death claims counterbalance each other, the lost-hand claims are not thrown into direct comparison with the death claims. Thus, we ought to take the lost-hand claims into consideration, and given that everything else is equal we should save Group B. This is a much better result.

Local Relevance views however have a problem with capturing Kamm's sore throat intuition. We can see why when we consider again Case 3, where we can save A's life, or B's life and C's sore throat:

'According to Local Relevance, C's sore throat cannot outweigh A's death, but it may be relevant more generally to what I am permitted or required to do. A's claim to be cured from the lethal illness is exactly counterbalanced by B's identical claim. Therefore, there is nothing left for C's claim to outweigh. Therefore, Local Relevance implies that I ought to cure B and C rather than flipping a coin, which seems wrong.' (Tadros 2019, pp. 184–185).

Some limited aggregationists are willing to bite the bullet on this as the 'the smallest caliber bullet' (van Gils and Tomlin 2019, p. 253) (Rüger 2020, pp. 464–465) and some may not share Kamm's intuition at all. However, I find Kamm's case compelling. Further examples are more powerful still. It seems clear that neither a single speck of dust in an eye nor a fleeting itch should decide between two lives. Thus, any theory that can capture these intuitions has a significant *prima facie* advantage over one that cannot.

⁴ Readers are encouraged to replace lost-hand claims throughout the paper with the strongest claims they can imagine that nevertheless remain irrelevant to claims against death.

⁵ Global Relevance has several other fatal flaws that I will not go into here. For more on this, see Tomlin (2017).

The rest of this paper will be spent exploring how we might amend Relevance views to accommodate the Sore Throat Intuition and the intuition uncovered by Case 4. By the end of the paper, I will have demonstrated that the only way to accommodate them is to bring back Global Relevance and combine the insights of both types of relevance in a Hybrid view.

2 A Trilemma

First though, let us strengthen the problem. Consider Cases 3, 4 and 5 again:

Case 3

| Group A | Group B |
|-----------------------|--|
| 1 person facing death | 1 person facing death 1 person facing a sore throat |

Case 4

| Group A | Group B |
|-----------------------|--|
| 1 person facing death | 1 person facing death 100 people facing a lost hand |

Case 5

| Group A | Group B |
|-----------------------|-------------------------------|
| 1 person facing death | 100 people facing a lost hand |

Intuitively, we should flip a coin in Case 3, save Group B in Case 4 and save Group A in Case 5. Yet both Global Relevance and Local Relevance struggle to capture what we want to say in all three of these cases.

Consider Global Relevance first. Global Relevance gets the right answer in Cases 3 and 5. In Case 3, the sore throat is dismissed from consideration because it is irrelevant to the death claim, and thus we have equal claims on both sides and must flip a coin. In Case 5 the lost-hand claims are irrelevant to the person facing death so we dismiss them from consideration, leaving only the death claim, and so Group A is saved.

However, as we have already seen, in Case 4, Global Relevance tells us that the lost-hand claims are irrelevant to the death claim and so must be dismissed entirely from consideration. Thus, Global Relevance tells us to flip a coin, when it seems we ought to save Group B.

Now let us consider Local Relevance. Local Relevance gets the right answer in Cases 4 and 5. Firstly, in Case 5, Local Relevance tells us that the lost-hand claims are irrelevant to the death claim, and so cannot outweigh it, and therefore we must save Group A. Secondly, in Case 4, whilst the lost-hand claims are irrelevant to the death claims, they are not needed to outweigh the death claim, and thus are still relevant to our deliberation. Given that all else is equal in Case 4, the lost-hand claims tie-break in favour of Group B.

However, Local Relevance gets the wrong answer in Case 3. In Case 3, Local Relevance tells us that whilst the sore throat is irrelevant to the death claim, it can remain relevant to our deliberation. Given that the death claims counterbalance each other, the sore throat is the decisive claim, and Global Relevance tell us to save Group B. But this, of course, violates the intuition Kamm identifies.

Amendments to the Local Relevance view have attempted to capture our intuition in Case 3 - Tadros argues for the following:

‘if I choose to save B because I can cure C’s sore throat, C’s sore throat grounds a permission not to save A from death where otherwise I would be required to give him an equal chance of being saved. Friends of Local Relevance might claim that this is objectionable—C’s sore throat cannot ground a permission not to give A a chance of being saved from death, either on its own, or in conjunction with other permission-grounding types of facts. It can only outweigh duty-grounding types of facts that do not disable its duty-grounding force.’ (Tadros 2019, p. 185).

But if Tadros takes this position he is unable to capture our intuition in one of either Case 4 or Case 5. Suppose lost hands can ground a permission not to give another claimant a chance of being saved from death - such that the lost hand claims are relevant. If so, Tadros’ view will get the wrong answer in Case 5 - the lost-hand claims would be able to outweigh the death claim. Conversely, if lost hands cannot ground such a permission, then in Case 4 the lost-hand claims cannot be decisive in favour of Group B, again getting the wrong answer.⁶

Similarly, van Gils and Tomlin suggest a possible sore throat amendment to their version of Local Relevance, the Sequential Claims Matching view, which requires an understanding of anchoring claims.

The anchoring claim is the strongest claim under consideration (or claims if there are multiple claims of the same strength); in the cases above, the death claims are the original anchoring claims. It is the anchoring claim we compare other claims to, to determine whether they are sufficiently strong to be relevant. In this way the person with the anchoring claim has a *pro tanto* claim to be saved - unless the claim is silenced or outweighed by other claims relevant to it, we must save the anchoring claimant. On a Local Relevance view, when an anchoring claim is matched, the strongest unmatched claim becomes the new anchoring claim. This is how claims, which were not relevant to the original anchoring claim, might still be relevant to the overall decision by being relevant to new anchoring claims.

Van Gils and Tomlin then suggest the following amendment to their account:

“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.” (van Gils and Tomlin 2019, p. 233).

This amendment allows their version of Local Relevance to capture our intuition in Case 3. Because the sore throat is not relevant to the weakest previously matched claim (i.e., the death claim) it cannot become the new anchoring claim and so we must flip a coin. Naturally, however, this amendment also loses our intuition in Case 4. Before we could say that, whilst the lost-hand claims were irrelevant to the death claims, they could still play a role in our deliberation. However, with the sore throat amendment, the lost-hand claims can no longer play a role as they are not relevant to the weakest previously matched claim.

⁶ If Tadros were to recognise that the strength a claim needs to be relevant to a chance to be saved differed from the strength needed to be relevant to a claim to be saved, then he could capture all these cases. To do so he would need to distinguish between two types of relevance, in the same way that I do in the next section. Still, Tadros’ approach here helpfully supports the deeper justification for my view.

So, no account so far considered can capture all these intuitions about relevance. Perhaps, there are good reasons for biting the bullet in one or more of these cases. However, a view that could capture our intuition in all three cases would be *prima facie* stronger and deserves serious consideration.

3 Two Types of Relevance

The problem above seems to be that there is a tension as to how far we want relevance to extend. In Case 5 we want to say that only very similar claims are relevant, whilst in Case 4 we want to say that more distant claims are relevant (whilst still allowing even more distant claims to remain irrelevant). We are being pulled in two directions.

This tension suggests that what we have previously treated as a single phenomenon might in fact be two. We have been wrong to consider only one type of relevance, instead we need to differentiate between types of relevance more finely. We will explore the deeper rationale behind such a distinction in a later section. First let us consider how exactly we might distinguish between types of relevance, and whether the distinction allows us to untangle the tension between these cases.

The tension in Cases 4 and 5 is that relevance is needed to play two distinct roles in the different cases. In Case 5 relevance is used to determine whether a claim can compete with another claim, i.e., whether the hundred lost-hand claims can match with the death claim. In Case 4, relevance is used to determine whether a claim can become the anchoring claim, i.e., whether the lost-hand claims can be decisive between two death claims. These two roles we might call the *Matching Role* and the *Anchoring Role*:

Matching Role: the role relevance plays in determining whether a claim can match or outweigh a competing claim.

Anchoring Role: the role relevance plays in determining whether a claim can become the new anchoring claim.

Up until now, we have had only one type of relevance to play both these roles, and this is what leads us to the tension in the view. If we want to capture all our intuitions, then we need to separate these roles from each other and distinguish between two types of relevance that can play the different roles.

Fortuitously, we have already distinguished between two types of relevance - Global Relevance and Local Relevance. Let us reintroduce the two types of relevance in more formal terms:

Global Relevance: A claim, X, is Globally Relevant if X is strong enough compared to the strongest claim with which it competes.

Local Relevance: A claim, X, is Locally Relevant to another claim, Y, if X is strong enough compared to Y.

With this in view, we can see that there are four notions of relevance in play. We have Global Relevance regarding the Matching Role, Global Relevance regarding the Anchoring Role, Local Relevance regarding the Matching Role and Local Relevance regarding the Anchoring Role.

Global Relevance and Local Relevance have previously been considered rival understandings of relevance, with proponents of either view ascribing both the Matching and Anchoring roles to their preferred type of relevance. However, there is no reason to suppose we must choose one form of relevance for both roles. Instead, we could choose one type of relevance for one role, and the other type of relevance for the other role - I call this the Hybrid Approach.

Recognising this, Local Relevance seems most appropriate for the Matching Role, and Global Relevance most appropriate for the Anchoring Role. Let us quickly consider why.

Firstly, the Matching Role. If we used Global Relevance for the Matching Role, whether a claim can be aggregated against a competing claim would not be determined by whether it is relevant to that particular claim, but whether it is relevant to the strongest claim in the competition. This seems overzealous - the strongest claim in the competition should not determine whether two smaller sets of claims are comparable! However, if we use Local Relevance, we determine whether a claim can be aggregated against a competing claim, by looking to see if it is relevant to that particular claim. This seems much more intuitive.

Secondly, the Anchoring Role. Where the Anchoring Role determines whether a claim can become the new anchoring claim, it also determines whether the last remaining claim can tie-break between the already matched claims. It seems obvious that Global Relevance is more appropriate for this role. If a claim is going to tie-break between claims, not only is it tie-breaking between smaller claims, but it is also, of course, tie-breaking between the strongest claims. Thus, whether it is strong enough to become the new anchoring claim, should be determined in comparison with the strongest overall claim.

I posit we formalise this by combining Local Relevance and the Matching Role into Competing Relevance, and Global Relevance and the Anchoring Role into Anchoring Relevance:

Competing Relevance: the relevance a claim needs to be able to match or outweigh a competing claim. A claim, X, has competing relevance with another claim, Y, if X is sufficiently close in strength to Y such that X can match with Y.

Anchoring Relevance: the relevance a claim needs to be able to become the new anchoring claim. A claim, X, has anchoring relevance if X is sufficiently close in strength to the strongest claim (matched or unmatched) with which X is in competition, such that X can become the new anchoring claim when all stronger claims are matched.

Once we recognise this distinction, we can set Competing Relevance and Anchoring Relevance to extend different distances. For instance, we can set Competing Relevance to extend only to very similar claims, such that no claim weaker than a lost-limb claim can match with a death claim. We can do this, whilst simultaneously setting Anchoring Relevance to extend further, such that lost-hand claims can become the new anchoring claims when all death claims are matched.

Consider how this might work in Cases 3, 4 and 5:

Case 3

| Group A | Group B |
|-----------------------|--|
| 1 person facing death | 1 person facing death 1 person facing a sore throat |

Case 4

| Group A | Group B |
|-----------------------|--|
| 1 person facing death | 1 person facing death 100 people facing a lost hand |

Case 5

| Group A | Group B |
|-----------------------|-------------------------------|
| 1 person facing death | 100 people facing a lost hand |

With Competing Relevance extending only as far as lost-limb claims, in Case 5 the hundred lost-hand claims cannot be matched against the death claim. Thus, Group A wins in Case 5. Simultaneously, with Anchoring Relevance extending further, in Case 4 the lost-hand claims can become the new anchoring claims, and therefore be decisive in favour of Group B. Yet, Anchoring Relevance would not extend so far as to include sore throats and so we must flip a coin in Case 3. So, we get the right answer in all three cases.

4 Amending Sequential Claims Matching by Closeness

So far, I have spoken in generalities about Relevance views and how the Hybrid Approach might be used. But now I want to turn to a particular way of capturing the Hybrid Approach. This will work as a proof of concept, of how we might adapt Relevance views to include the Hybrid Approach, and significantly strengthen my favoured Relevance view.

To do so I am going to build on the Sequential Claims Matching by Closeness (SCMC) view. This view is a version of Sequential Claims Matching, initially established by van Gils and Tomlin (van Gils and Tomlin 2019), strengthened against a variety of objections, and designed to better capture the idea of respect outlined earlier. It does so by virtue of its matching principle:

Match by Closeness: when matching claims, where possible, match claims with closest relevant claims.

So, in SCMC we have a procedure which compares competing claims according to their closeness. First, we compare identical claims, then relevantly similar claims, and then wider relevant claims, until we reach the limit of relevance.⁷ SCMC does so by splitting claims into 'levels' of claims sorted by their strength, with the relative strength of claims between levels remaining the same. Let Level 1 Claims be death claims, Level 2 Claims be paraplegia claims, Level 3 Claims be lost-limb claims and so on. Each level has half the weight of

⁷ It is this insight that means this account also avoids Tomlin's (2017) and Horton's (2018) criticisms of limited aggregation. For more on this, and the arguments for the Matching by Closeness principle, see Hart's (2022).

the next level up, such that two Level 2 claims, match with exactly one Level 1 claim, and four Level 3 claims match with exactly two Level 2 claims or one Level 1 claim.

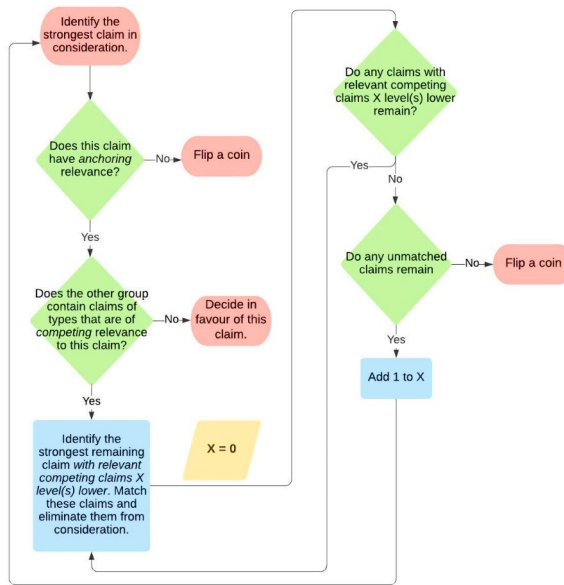
The following is an amended version of SCMC which distinguishes between the two types of relevance, to capture the Hybrid Approach:

Amended Sequential Claims Matching by Closeness (ASCMC).

- I. Identify the strongest claim still in consideration and ask, does it have *anchoring* relevance?
 - a. If so, go to Step II
 - 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)
- II. Does the other group contain claims of types that are of *competing* relevance to this claim?
 - a. If so, go to Step III
 - b. If not, decide in favour of the group with the strongest claim still in consideration
- III. Identify the strongest remaining claim *with* competing claims [zero] levels lower. Match these claims and eliminate them from consideration.
- IV. Continue step III until there are no claims with competing claims [zero] levels lower.
- V. Repeat steps I - IV, with competing claims one level lower, then two levels lower, and so on.
- VI. Continue until either:
 - a. One group contains unmatched claims with anchoring relevance, 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).

The below is a flow chart which helps to illuminate the sequence. We loop through the sequence starting by matching identical claims, then very close claims, and lastly more distant claims until we either have a tie or claims that cannot be matched.

Amended Sequential Claims-Matching by Closeness – Flow Chart



Once we adapt the sequence in this way, we fully capture the two types of relevance within the sequence. Therefore, we can redeem the sore throat intuition, whilst also getting the right answer in Cases 4 and 5. Consider again Case 3:

Case 3

| Level | Group A | Group B |
|-------------------|---------|---------|
| 1 - Death | 1 | 1 |
| 2 - Paraplegia | | |
| 3 - Lost Limb | | |
| 4 - Lost hand | | |
| ... | | |
| 100 - Sore throat | | 1 |

Suppose that the levels refer to claims against death, paraplegia, lost limbs, lost hands, and sore throats, respectively. Thus, Group A is made up of a death claim, whilst Group B is made up of a death claim and a sore-throat claim. For ease of reference, we will call the Level 1 claims in Groups A and B, the A1 and B1 claims, and the Level 100 claims in Group B, B100 claims.

Now we can set the extent of relevance. For the purposes of this paper, let us say that Competing Relevance is set at two levels, such that Level 3 claims have Competing Relevance to Level 1 claims, but Level 4 and below claims do not. Similarly, let us set the extent of Anchoring Relevance at four levels, such that Level 5 claims have Anchoring Relevance to Level 1 claims, but Level 6 and below claims do not. With this in view, let us go through Case 3 slowly to demonstrate how the sequence works:

Step I: First we identify the strongest claim, which is A1 (or B1). As the overall strongest claim in the competition it, trivially, has Anchoring Relevance.

Step II: We check if the A1 claim has claims of Competing Relevance to it - it does with the B1 claim, so we move to Step III.

Step III: We identify the strongest claim with claims zero levels lower (i.e., claims on the same level). Again, this is the A1 claim, and we match it with the B1 claim. I have demonstrated this on the below table by crossing through the A1 and B1 Claim.

Case 3

| Level | Group A | Group B |
|-------|---------|---------|
| 1 | † | † |
| 2 | | |
| 3 | | |
| 4 | | |
| ... | | |
| 100 | | 1 |

Step IV and V: We check to see if there are any unmatched claims, which there are so we loop back round to Step I.

Step I (Round 2): On the second pass through the sequence, we identify the B100 claim as the strongest claim still in consideration, we check to see if it has Anchoring Relevance. Anchoring Relevance only extends four levels, and thus we can see that the B100 claim does not have Anchoring Relevance, and so we must flip a coin.

Thus, by amending the sequence in this way we have managed to capture the sore throat intuition!

Let us also, quickly consider Case 4 and Case 5:

Case 4

| Level | Group A | Group B |
|-------|---------|---------|
| 1 | 1 | 1 |
| 2 | | |
| 3 | | |
| 4 | | 100 |

On the first pass through the sequence, we match the A1 claim and the B1 claim. On the second pass, at Step I, we identify the B4 claims as the strongest remaining claims, and we check to see if they have anchoring relevance - they do so we move onto Step II. At Step II, we look to see if there are any remaining claims with competing relevance to the B4 claims and, since there are none, Step II tells us to choose in favour of Group B. Our intuition is captured here too.

Case 5

| Level | Group A | Group B |
|-------|---------|---------|
| 1 | 1 | |
| 2 | | |
| 3 | | |
| 4 | | 100 |

Now Case 5. At Step I we identify the A1 claim as the strongest claim, recognise that it has anchoring relevance and move on to Step II. At Step II, we look to see if there are any

claims with competing relevance to the A1 claim. The B4 claims are too distant to have Competing Relevance, and so Step II tells us to decide in favour of Group A. Our intuition is captured once again!

A further positive of fleshing out the Anchoring and Competing Relevance distinction in this way, is that doing so allows further gradation. A claim without anchoring relevance might not be able to be decisive in favour of a Group, but it can still play a role in our overall decision by preventing other claims from becoming decisive. Consider the following case:

Case 6

| Level | Group A | Group B |
|-------|---------|---------|
| 1 | 1 | 1 |
| 2 | | |
| 3 | | |
| 4 | | |
| 5 | 10 | |
| 6 | | 30 |

In this case, anchoring relevance is set at four levels. Thus, the ten A5 claims in Group A do have anchoring relevance, but the thirty B6 claims do not. Nonetheless, the B6 claims can still play a role in the outcome. This is because they still have competing relevance with the A5 claims, and thus can counterbalance the A5 claims and prevent them from being decisive. This means we do not leap straight from claims with relevance to those with no relevance but have a more natural weakening of relevance.

5 Deeper Justification

Now the primary purpose of this paper is to vindicate Limited Aggregation intuitions and defend Limited Aggregation against the charge of inconsistency. As I have shown, the Hybrid Approach is fully able to capture our intuitions and thus the main task of the paper is complete.

What is novel about my account is the way in which I apply the different types of relevance to the different roles. I do not introduce any distinction in relevance that has not been previously discussed, I make use of two already plausible conceptions of relevance and simply show that they are compatible. The distinction I do introduce, between the Matching and Anchoring role, simply describes how we already use relevance in two subtly different ways.

As such, there seems to be little reason to think that both roles must be fulfilled by one type of relevance. Hence, no special or extra explanation is needed as to why we should adopt the Hybrid Approach as compared to a non-hybrid approach. That it gets the right answers in the above cases is itself a significant part of the reason to prefer this account over either Local or Global Relevance by itself.

Nonetheless, any Limited Aggregation account is going to need a deeper theoretical justification if it is to provide a complete explanation, and this account is no exception. Thus, it would be prudent if we spent a little time indicating what theoretical grounding we might have for distinguishing between the two types of relevance. Again, to demonstrate the gen-

erality of my account I will couch this discussion in the most basic terms of respect and allow readers to insert their preferred rationale behind relevance to fill in the detail.

Here I think two distinctions might be helpful in explaining the two types of relevance:

1. Act disrespect vs. procedural disrespect.
2. Claims to be saved vs. claims to an equal opportunity to be saved.

Firstly, let us handle the simpler distinction between act disrespect and procedural disrespect, and how this distinction might elucidate the two types of relevance. Act respect is about whether an action is respectful or not. For instance, if I choose to save hands over a life, that act is disrespectful because the outcome is disrespectful: it is disrespectful to attend to hands when someone's life is on the line, as it fails to recognise the strength of the dying person's claim, and what is at stake for them. On the other hand, if I am faced with a choice between saving two lives, then, whichever life I save, the action is respectful.

Procedural respect, on the other hand, is about the process that leads to the action. For instance, choosing between two life claims by flipping a coin shows procedural respect, because doing so recognises what is at stake for both claimants and gives them an equal chance of being saved. On the other hand, if I choose to save the person with blue eyes then I am showing procedural disrespect even though the act of saving that claimant's life is not itself disrespectful. Similarly, if we were choosing between a life claim and a lost-hand claim, and I choose to save the life because the dying person had blue eyes, I would be equally procedurally disrespectful whilst doing the respectful act.

We can now see how this distinction might explain the two types of relevance. Competing Relevance is concerned with action respect. Competing Relevance rules out comparisons between competing claims of different strengths, if such a comparison could lead to a disrespectful action. For instance, it would be a disrespectful action to save a multitude of claimants from lost hands instead of saving a single life, and so we must not match lost-hand claims with death claims.

Anchoring Relevance on the other hand is concerned with procedural respect. Certain claims are so weak that it would be disrespectful to even take them into consideration when deciding between stronger claims. Thus, we dismiss considerations such as eye colour, but we will also dismiss very small considerations such as sore throats. In the same way that it is disrespectful to choose between two death claims because of eye colour it is disrespectful to choose between two death claims because of a sore throat.

The second distinction makes use of the difference between a claimant's claim to be saved, and a claimant's claim to an equal opportunity to be saved.⁸ It is important to note that a claimant's claim to be saved from X is significantly stronger than another claimant's claim to an equal opportunity to be saved from X. This is why a claim to be saved from death will outweigh a claim to an equal opportunity to be saved from death. It is why we should choose to save the group with two dying people, rather than flip a coin between that group and the group with one dying person.

Now, a claim might be disabled (or silenced) in one of two ways. It might be disabled by the strength of a competing claim to be saved, or it might be disabled by the strength of a competing claim to an equal opportunity to be saved. Given that the claim to be saved is stronger than a claim to an equal opportunity to be saved, it would be appropriate if a claim-

⁸ As we have seen above, Tadros says something like this (2019, pp. 185–186).

ant's claim to be saved had stronger disabling power than a claimant's claim to an equal opportunity to be saved. Thus, a claimant's claim to be saved disables stronger competing claims than a claimant's claim to an equal opportunity to be saved does.

Let us consider this idea in a little more detail. Someone who is dying has a very strong claim to be saved. The strength of their claim is such that it disables lost-hand and sore-throat claims; it would be disrespectful to attend to sore throats in light of the dying person's claim to be saved. Thus, lost-hand and sore-throat claims, cannot be aggregated against the dying person's claim to be saved.

However, if the dying person's claim to be saved is matched by another dying person's claim to be saved, the dying person no longer has an unmatched claim to be saved; their claim to be saved has been taken out of consideration. Thus, their claim to be saved can no longer disable the lost-hand and sore-throat claims. Hence, it is no longer disrespectful to consider attending to the smaller claims because they are no longer in competition with the dying person's claim to be saved.

Nonetheless, the dying person still has a claim to an equal opportunity to be saved, even though this claim is much weaker than his initial claim to be saved. His claim to an equal opportunity to be saved can still disable the sore-throat claims, but it cannot disable the lost-hand claims. It is disrespectful to the dying person to consider attending to sore-throat claims in light of his claim to an equal opportunity to be saved. However, it is not disrespectful to consider attending to lost-hand claims in light of his claim to an equal opportunity to be saved.⁹

We can see then how this distinction might explain the distinction between Anchoring and Competing Relevance and why they might extend different distances. Competing Relevance is based around what a claimant's claim to be saved can disable, and Anchoring Relevance is based around what a claimant's claim to an equal opportunity to be saved can disable. In this way, we might say a claim has Anchoring Relevance if it is strong enough to respectfully outweigh the strongest claimant's claim to an equal opportunity to be saved. And we might say a claim has Competing Relevance if it is strong enough to respectfully compete with another claimant's claim to be saved.

So, I have identified two potential ways in which we might begin to understand the deeper justification behind the differing notions of relevance. There is certainly a debate still to be had about the best rationale for this distinction, and which wider account of relevance gives the best explanation. Unfortunately, there is no space to develop these arguments further in this paper. Nevertheless, what I have shown here is that the Hybrid Approach is not ad hoc but points to a genuine moral difference and makes sufficient reference to the respect-based reasons that all relevance theorists agree on. By couching the Hybrid View in these terms, I hope the generality of my conclusions to all relevance views should be evident. I leave it to the reader to flesh the distinction out further with their own preferred rationale.

⁹ In Voorhoeve's terms, the permissible personal perspective will not allow a claimant to insist upon their own rescue over the rescue of another's significantly stronger claim. However, the permissible personal perspective might allow them to prefer their own rescue over an equal opportunity of rescue for another's significantly stronger claim.

6 A Rival Hybrid View

I am not the first to attempt a hybrid relevance view. Bastian Steuwer has very recently offered his own hybrid view that also attempts to avoid the pitfalls of both the Global and Local Relevance views.¹⁰ His view is simpler and involves only two conditions:

- ‘(1) The local relevance condition: A claim can only be balanced with another claim if the two claims are relevant to one another.
 (2) The global relevance condition: Every individual with a strong claim has a veto against the consideration of any type of claim that is irrelevant to her claim, if such consideration would lead to her claim not being satisfied (e.g. by not being saved).’
 (Steuwer 2021, p. 19).

This view is interesting and certainly avoids many of the issues that plague non-hybrid views. Nevertheless, to finish this paper, I want to sketch why this view is less attractive than my view.

First, because Steuwer does not distinguish between the matching and anchoring role, he cannot get the right answer in Case 4. My view allows a claim, such as a lost-hand claim, to tie-break between two death claims, even though no number of lost-hand claims could match or outweigh a single death claim. Steuwer’s view cannot capture this because the lost-hand claims (being irrelevant to a death claim) would be vetoed.

Steuwer’s view also suffers from some wider problems. Consider the following case:

Case 7

| Level | Group A | Group B |
|-------|---------|---------|
| 1 | | |
| 2 | (1) | (2) |
| 3 | 10 | |
| 4 | | |
| 5 | | 100 |

At the first stage, (indicated by the claims without brackets) Group B wins – the hundred B5 claims outweigh the A3 claims. But at the second stage we add stronger Level 2 claims (indicated by the claims within brackets) – but we add more to Group B than to Group A. Nevertheless, on Steuwer’s view, we now swap to saving Group A, because the A2 claim has a veto against the inclusion of the B5 claims. So Steuwer’s view violates the Principle of Net Addition:

‘*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.’ (van Gils and Tomlin 2019, p. 253).

Now consider another similar case:

¹⁰ I am grateful to the second anonymous reviewer for bringing this paper to my attention.

Case 8

| Level | Group A | Group B |
|-------|---------|---------|
| 1 | | (2) |
| 2 | (2) | |
| 3 | 10 | |
| 4 | | |
| 5 | | 100 |

In this case, at stage 1, Group B wins again. But at stage two we add new claims, this time an equal number of claims, but stronger claims in Group B. Nevertheless, Group A now wins. Again, the B5 claims are vetoed by the A2 claims and so the A3 claims are decisive. Thus, Steuwer's view also violates a principle I call:

Greater Consideration for Stronger Claims: adding claims of differential strength and equal numbers cannot make the group to which stronger claims are added less choice-worthy compared with the group to which weaker claims are added.

My view, in contrast, does not violate either principle. Whilst the B5 claims might not be able to compete with the A2 claims, they still have anchoring relevance because anchoring relevance is set at four levels. Thus, the B5 claims can match with the remaining A3 claims, and then tie-break between the balanced claims. These cases give us further reason to think that any successful limited aggregation view must distinguish between the two types of relevance: if we do not, even the best relevance views will violate very commonly held principles.

Lastly, I am also sceptical of Steuwer's use of the veto. On his view, the strongest claimant has a veto against any irrelevant claims to it if those claims do not benefit her chance to be saved. This seems to me too arbitrary and unfair to those weaker claims. Consider the following case:

Case 9

| Level | Group A | Group B |
|-------|---------|---------|
| 1 | 1 | |
| 2 | | 2 |
| 3 | | |
| 4 | 20 | 19 (+2) |

At the first stage the Level 4 claims matter. They are decisive in favour of Group A: the strongest claimant has decided that the Level 4 claimants deserve sympathy and should be included in the moral decision. But at the second stage when two extra Level 4 claims are added to Group B, the strongest claimant does an about-face and decides that Level 4 claimants no longer matter, leaving us with a tie. This seems very morally inconsistent, and I think the Level 4 claimants would rightfully feel aggrieved by this change. In this case the Level 4 claimants do not seem to be treated with respect or as ends in themselves, but merely as instrumentally useful to the strongest claimant.

My view avoids this. A claim's relevance is fixed from the start. There is no arbitrariness, and all claimants are treated with the respect they deserve as ends in themselves. Now my view might get the same answer as Steuwer's in similar cases. However, the reasoning in those cases is known in advance, such that whether a claim matters does not change depending on whether it is helpful for another claim.

To conclude, my Hybrid View combines the positives of Global and Local Relevance and thus gets the right answer in both the heterogeneous and the sore throat cases which each Relevance view respectively struggles with. By distinguishing between the Anchoring and Competing Relevance we also avoid the negative consequences of Steuwer's hybrid account and capture a more gradual weakening of relevance as per Case 4. Nevertheless, whilst I have shown the hybrid view to not be entirely ad hoc, more needs to be said to connect the view to the rationale of relevance views in the wider literature and to provide a deeper justification for the distinctions uncovered.

Acknowledgements Special thanks must be extended to my supervisors Philip Stratton-Lake, Jonathan Way and Brad Hooker, whose insight, patience, and guidance at every step of this paper's development has been invaluable. I also want to thank all the participants of the Graduate Research Seminars in the Philosophy Departments at the University of Reading and the University of Southampton for their helpful comments. In particular, I want to thank Steven Wu for extensive written comments and long discussions, and Aart van Gils who initially triggered my interest in this topic. The paper has also been significantly strengthened thanks to comments from the two anonymous reviewers at Ethical Theory and Moral Practice.

Funding This work was supported by the SWW DTP [Grant Number: AH/R012776/1]. The SWW DTP is funded by the Arts and Humanities Research Council.

Declaration

Conflict of Interest/Competing Interests None.

Availability of Data and Material Not applicable.

Code Availability Not applicable.

Authors' Contributions Not applicable.

Ethical Approval Not applicable.

Consent to Participate Not applicable.

Consent for Publication Not applicable.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Hart J (2022) Local Relevance and Match by Closeness. Unpublished Manuscript. University of Reading.
Horton J (2018) Always Aggregate. *Philos Public Affairs* 46(2):160–175
Kamm F (1998) *Morality, Mortality Volume I: Death and Whom to Save From It*. Oxford University Press, Oxford. doi:<https://doi.org/10.1093/0195119118.001.0001>

- Kamm F (2015) Bioethical Prescriptions. *J Med Ethics* 41(6):493–495
- Rüger K (2020) Aggregation with Constraints. *Utilitas* 32:454–471
- Scanlon T (1998) *What We Owe to Each Other*. Harvard University Press, Cambridge, MA. doi:<https://doi.org/10.2307/j.ctv134vmrn>
- Steuwer B (2021) Aggregation, Balancing, and Respect for the Claims of Individuals. *Utilitas* 33:17–34
- Tadros V (2019) Localized Restricted Aggregation. *Oxford Studies in Political Philosophy* 5:172–204. doi:<https://doi.org/10.1093/oso/9780198841425.001.0001>
- Tomlin P (2017) On Limited Aggregation. *Philosophy & Public Affairs*, 232–260
- van Gils A, Tomlin P (2019) Relevance Rides Again? Aggregation and local relevance. *Oxford Studies in Political Philosophy* 6:221–255. doi:<https://doi.org/10.1093/oso/9780198852636.003.0008>
- Voorhoeve A (2014) How should we aggregate competing claims? *Ethics* 125(1):64–87

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.