

On limited aggregation

Article

Accepted Version

Tomlin, P. (2017) On limited aggregation. *Philosophy & Public Affairs*, 45 (3). pp. 232-260. ISSN 1088-4963 doi: <https://doi.org/10.1111/papa.12097> Available at <https://centaur.reading.ac.uk/72162/>

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/10.1111/papa.12097>

Publisher: Wiley

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

PATRICK TOMLIN

On Limited Aggregation*

Forthcoming in Philosophy & Public Affairs

I Introduction: Limited Aggregation and the Relevance View

Consider the following pair of cases:

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

Case 2. You can save one person from death, or some larger number of people, N_2 , from a mild headache.

On one view, let's call it *Pure Aggregation*, in both cases we have to see how large N is before deciding what to do. If N_1 gets large enough, we should save the people from paralysis. And if N_2 gets large enough, we should prevent the mild headaches. According to another view, *Anti-Aggregation*, in both cases we should save the one person from death: we should simply satisfy the strongest claim, no matter how large the number of people possessing competing weaker claims gets.

Many people believe that we should take prima facie inconsistent stances on these two cases.¹ That is, we should accept aggregation of the lesser claims in Case 1, but not in Case 2: there is some number of people we should save from

* I am grateful to Aart van Gils for sparking my interest in this topic. The paper was presented and discussed at Oxford, York, and Goethe Universität's Forschungskolleg Humanwissenschaften. For further help in improving the paper, I am grateful to Aart van Gils, John Halstead, Robert Jubb, Thomas Sinclair, Zofia Stemplowska, Alex Voorhoeve, and the editors of *Philosophy & Public Affairs*. I am especially grateful to Christian Barry, Brad Hooker, Jeff McMahan, Kieran Oberman, and Victor Tadros for detailed discussion and advice. Whilst writing the paper, I was supported by the University of Reading 2020 Fellowship, and Goethe Universität's Justitia Amplificata centre and Forschungskolleg.

¹ For relevant survey data, see Alex Voorhoeve, 'Balancing Small against Large Burdens' in *Behavioural Public Policy* (forthcoming); and 'Healthy Nails versus Long Lives: An Analysis of a Dutch Priority-Setting Proposal' (unpublished m/s, available at: <http://personal.lse.ac.uk/voorhoev/Healthy%20Nails%20versus%20Long%20Lives.pdf>)

paralysis over saving a single person from death, but there is no number of mild headaches which could outweigh, in the relevant sense, the importance of saving someone's life. These people favour *Limited Aggregation*.

The challenge for the Limited Aggregation view is to explain why aggregation is acceptable in Case 1 but not in Case 2. The most influential answer to this challenge is to posit a distinction between the way the two sets of claims relate to one another in the two cases: in Case 1, the claim against death and the claims against paralysis are sufficiently close in strength to be, in T.M. Scanlon's words, 'relevant' to one another; whilst in Case 2, the claims against mild headaches possessed by the many are not close enough in strength to the claim against death for them to be considered 'relevant'. Aggregation is then limited to weaker claims that are 'relevant'.

This kind of approach has been advocated by some of the most influential moral and political philosophers of the last thirty years. For example, in an oft-cited passage, Scanlon argues:

If one harm, though not as serious as another, is nonetheless serious enough to be morally 'relevant' to it, then it is appropriate, in deciding whether to prevent more serious harms at the cost of not being able to prevent a greater number of less serious ones, to take into account the number of harms involved on each side. But if one harm is not only less serious than, but not even 'relevant to,' some greater one, then we do not need to take the number of people who would suffer these two harms into account in deciding which to prevent, but should always prevent the more serious harm.²

² T.M. Scanlon, *What We Owe to Each Other* (Cambridge, MA: Harvard University Press, 1998), pp. 239-240. See also: T.M. Scanlon, 'Replies' in *Ratio* 16 (2003): 424-439. It is a challenge for Scanlon to incorporate this position within his contractualist framework. I will not take up that issue here.

F.M. Kamm defends a similar view. According to Kamm's Principle of Irrelevant Utilities, some gains or losses are too small to be included in a decision when much more is at stake for others.³ Michael Otsuka agrees, appealing to Kamm's principle to explain why we should not aggregate the complaints of many people subjected to a low risk of death, when compared with a certainty of serious harm for one person.⁴

Alex Voorhoeve has recently defended a variant of this view. Voorhoeve calls his account Aggregate Relevant Claims (ARC). Voorhoeve gives a very clear account of how ARC will work, at least in a limited set of cases:⁵

ARC:

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

³ F.M. Kamm, *Morality, Mortality Volume 1* (New York: Oxford University Press, 1993), chs. 8-10; and *Intricate Ethics* (Oxford: Oxford University Press, 2007), pp. 297-298, 484-486. For discussion, see Iwao Hirose, *Moral Aggregation* (Oxford: Oxford University Press, 2014), ch. 8.

⁴ Michael Otsuka, 'Risking Life and Limb: how to discount harms by their probability' in I. Glenn Cohen, Norman Daniels and Nir Eyal, *Identified versus Statistical Lives: An Interdisciplinary Perspective* (Oxford: Oxford University Press, 2015), pp. 79-80. Elsewhere, Otsuka has argued for the need for a theory that treads the line between Anti-Aggregation and Pure Aggregation. See his 'Saving Lives, Moral Theory, and the Claims of Individuals' in *Philosophy & Public Affairs* 34 (2006): 109-135.

⁵ Alex Voorhoeve, 'How Should We Aggregate Competing Claims?' in *Ethics* 125: 64-87, at pp. 66-67.

Larry Temkin claims that trade-offs between stronger and weaker claims are sometimes desirable, but that ‘virtually all agree that, other things equal, it would be worse if fifty people suffered from AIDS, quadriplegia, severe psychosis, or being deaf, dumb, and blind, than if virtually *any* number of people suffered from a minor nosebleed, a slight cold, a sprained finger, or a short mild headache.’⁶ Temkin argues, with Kamm, Scanlon, and Voorhoeve, that these judgments are best explained by the relationship between the two sets of claims, namely their relative strength.⁷ Temkin finds such a view ‘compellingly plausible.’⁸ Temkin’s view, however, differs in an important way from Scanlon’s, Kamm’s, and Voorhoeve’s. Temkin’s view concerns the evaluative question of which outcomes are better or worse, whilst those of Scanlon, Kamm, and Voorhoeve concern the normative question of whom we ought to save. In what follows, I will primarily focus on the normative version of this approach, but what I say applies, *mutatis mutandis*, to the evaluative version as well.

While many proponents of Limited Aggregation are non-consequentialists and many proponents of Pure Aggregation are consequentialists, the two sets of positions do not map on to one another as neatly as one might suppose. First, we might hold a Limited Aggregation view about goodness, like Temkin’s, but nevertheless be consequentialists in believing that we are required to maximize the good. Second, we might favour Pure Aggregation about *whom to save*, but employ all sorts of other paradigmatically non-consequentialist distinctions, views, and principles, such as the doing/allowing distinction, the asymmetry between harm and benefit, the means principle, the moral relevance of intentions, fundamental rights, and the personal prerogative. For example, it doesn’t follow from saying that we ought to save the many in Case 2 that we

⁶ Larry S. Temkin, *Rethinking the Good* (Oxford: Oxford University Press, 2012), pp. 33. Emphasis in original.

⁷ *Ibid.*, p. 32.

⁸ *Ibid.*, p. 37.

can permissibly *kill* the one to achieve the same end, or that the one must kill herself, or allow herself to die, to achieve the same end.

Let us call this version of the Limited Aggregation position concerning the ‘relevance’ between claims the *Relevance View*. The Relevance View certainly delivers judgments in line with common intuitions about Cases 1 and 2, and is attractive for that reason. But Cases 1 and 2 both have three key features which cases used in discussions of these issues often share. First, there are only two groups of potential beneficiaries. Second, the groups are homogenous. That is, within each group, everyone has precisely the same claim. Third, the groups are in competition – you can only save all of one group, or all of another. Although the cases used to motivate the intuitions behind the Relevance View share these stylised features, the position is, of course, supposed to have application beyond these specific examples.⁹ In particular, proponents have argued that healthcare decisions ought to be made in accordance with, and healthcare systems organised along the lines suggested by, the Relevance View.¹⁰

Some well-known criticisms of the Relevance View focus on the first key feature, and present cases in which more than two homogenous groups are in

⁹ It is worth noting that in ‘How Should we Aggregate Competing Claims?’ Voorhoeve explicitly limits his discussion, for simplicity, to cases in which ‘claims to an alternative are either all irrelevant or relevant’ (p. 67) (i.e., roughly-speaking, homogenous group cases). However, Voorhoeve doesn’t say that ARC *doesn’t* apply beyond such cases, and he doesn’t articulate an independent view for ‘diverse group’ cases, nor a rationale for why they should be treated differently. Kamm is clear that her Principle of Irrelevant Utilities *does* apply to ‘diverse group’ cases, and key intuitions in support of the principle are derived from such cases (such as one in which one claim against death is pitted against one claim against death plus one claim against a sore throat. See: *Morality, Mortality Volume 1*, p. 146.) If the Relevance View were restricted to homogenous group cases, it would not apply to, or draw support from, Kamm’s tie-break cases. It would also lack the real-world relevance its adherents claim for it. Furthermore, such a restriction would seem ad hoc, and none of the arguments given for the Relevance View seem to support restricting it in this way. At the very least, I take it that adherents of the Relevance View think that the notion of ‘relevance’ should apply beyond homogenous group cases, and so my criticisms about the ambiguity of relevance, and how various anchoring rules would classify claims as relevant or irrelevant, would apply to any such view.

¹⁰ See, for example, the exchange between Voorhoeve and Kamm in *Journal of Medical Ethics* 41 (2015): 492-495; Temkin’s discussion of WHO decisions in *Rethinking the Good*, pp. 79-80; Voorhoeve, ‘Healthy Nails versus Long Lives’; Voorhoeve, ‘Balancing Small against Large Burdens’.

competition. Once we move from two- to three-group cases, the Relevance View seems to breach two plausible axioms of rationality: transitivity, and the Principle of Irrelevant Alternatives.¹¹ These criticisms are well-known, and the advocates of the Relevance View have discussed and replied to them, convincingly in my view.¹²

In this paper I seek to advance some new and, I think, more damning criticisms of the view, both in terms of its theoretical plausibility and in terms of its practical relevance when it comes to making decisions about healthcare resources. Instead of focussing on cases with more than two groups, I focus on cases which do not share the second and third common features: homogenous groups, and competition between groups. Once we look at these cases, some serious, hitherto unexplored, problems with the Relevance View come to light. These problems concern the Relevance View's plausibility in cases with internally diverse competing groups, and its scope.

¹¹ In assessing the accusation of intransitivity, it is important to keep the two versions of the Relevance View separate. In order to endorse the evaluative version, we must deny that 'all things considered better than' is transitive. This is a significant bullet to bite, though some are prepared to bite it, or at least consider it (Temkin, *Rethinking the Good*). But the accusation that 'ought to be chosen over' is not transitive across different pairwise comparisons is much less troubling: 'ought to be chosen over' is not even the sort of thing we would *expect* to be transitive, and therefore much less is at stake in rejecting transitivity in this context. Indeed, even utilitarians should reject transitivity here. Imagine I must choose whom to give some good to, candidate A, B, or C. Considering only the happiness of non-candidates, we would rank them as follows: A>B>C. But it will make C very very unhappy to lose out to A in particular, so much so that this outweighs the non-candidates' preference for A over C. Therefore, in pairwise competitions, according to utilitarianism, I should choose C over A, A over B, and B over C. The utilitarian may deny that 'Choosing A and so upsetting C' is *the same option* as 'Choosing A over B'. But this is to invoke a moralized way of individuating options, and the proponent of the Relevance View can do the same, thereby preserving transitivity. For example, this is Voorhoeve's response ('How Should We Aggregate Competing Claims', pp. 78-79). The disagreement therefore comes down to what is morally relevant, not whether one theory violates transitivity whilst another does not. On transitivity and deontic concepts more generally, see: Temkin, *Rethinking the Good*, pp. 195-201; F.M. Kamm, 'Supererogation and Obligation' in *Journal of Philosophy* 82 (1985): 118-138.

¹² Variants of these criticisms can be found in: Derek Parfit, 'Justifiability to Each Person' in *Ratio* 16 (2003): 368-390, at p. 384; Alistair Norcross, 'Contractualism and Aggregation' in *Social Theory and Practice* 28 (2002): 303-314, at pp. 306-309; John Halstead, 'The Numbers Always Count' in *Ethics* 126 (2016): 789-802, at pp. 797-799. For responses, see: Kamm, *Intricate Ethics*, pp. 298 and 484-487; Voorhoeve, 'How Should We Aggregate Competing Claims', pp. 76-79, and 'Why One should Count only Claims with which One Can Sympathize' in *Public Health Ethics* 10:2 (2017): 148-56 ; Temkin, *Rethinking the Good*.

Consideration of cases involving groups of claims of diverse strength reveals two problems. First, there is a key ambiguity in the position. I call this the ‘anchoring problem’. Second, any way of resolving the problem entails deeply counter-intuitive judgments that violate core normative principles. The critique of the Relevance View’s scope shows that it simply *does not apply* to many decisions about healthcare resource allocation. Furthermore, there is no simple or obvious way to extend or amend the Relevance View so that it will apply to healthcare choices, since obvious extensions have counter-intuitive results, or pose more questions than they answer. I conclude by discussing what other options we have, if we reject the Relevance View.

II The anchoring problem

Cases 1 and 2 involve homogenous groups. That is, we can help either some number of people with condition x, or some number of people with condition y. Within each group, each person is identical in the relevant respects: they will be equally badly off if left alone, and equally well off if assisted, and so their claims are of identical strength. But the Relevance View is surely supposed to apply beyond these cases: if it applied *only* to cases involving homogenous groups, it would have no relevance outside philosophical examples. In the real world, even when we save many people with the same condition their experiences of that condition will vary, often markedly, such that they will have claims of different strengths.

The focus on homogenous group cases is, of course, the kind of simplifying move that philosophers often make, in order to help us get clear on the issues. However, concentrating on homogenous group cases hides an important ambiguity in the Relevance View.¹³ We can call this the ‘anchoring problem’. Once the ambiguity is revealed, there are two obvious interpretations. But

¹³ Both Derek Parfit and Frances Kamm discuss cases in which the groups are not homogenous, but both use cases in which the strongest claim competes against a group of diverse claims. Therefore, these cases neither show up the ambiguity I explore here, nor produce the counter-intuitive results I later show. See: Parfit, ‘Justifiability’, p. 383 n.16; F.M. Kamm, *Mortality, Mortality Volume I*, chs. 8-10.

neither is attractive in diverse group cases. Nor are more complex anchoring principles.

In order 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. In order 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 Cases 1 and 2, it is obvious what the anchoring claim is – it is the claim against death. All other claims must be measured against this claim to see if they are relevant. This is also true of all of the examples in the literature which I have come across: there is always some claim – the strongest claim under consideration – which is clearly the anchoring claim. Voorhoeve captures this idea with his fourth principle:

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

However, this principle is open to two interpretations. We could either Anchor by Strength, or Anchor by Competition:

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*.

In order to see how these two interpretations come apart, consider Case 3 below. In this case, and in all cases that follow, the following relevance relations hold:

Severe impairment is relevant to death. Saving ten people from severe impairment outweighs saving one person from death.

Mild impairment is relevant to severe impairment. Saving ten people from mild impairment outweighs saving one person from severe impairment.

Mild impairment is not relevant to death.¹⁴

Case 3. You can save Group A or Group B. Group A consists of some people who are facing death and some who are facing mild impairment. Group B consists of some people who are facing severe impairment.

We need to decide whether those in Group A who face mild impairment have relevant claims. According to Anchor by Strength they do not: their claim is not relevant to the strongest claims *in the competition* (the claims against death, held by other members of Group A). In contrast, according to Anchor by Competition, those facing mild impairment *do* have a relevant claim, since the strongest claims *with which they compete* are the claims against severe impairment.

This is the ‘anchoring problem’. We do not see this problem in homogenous group cases, since in such cases all claims either (a) are the strongest claim in the competition, or (b) compete with the strongest claim in the competition, and so the question of whether we Anchor by Strength or Anchor by Competition is never raised, let alone explicitly answered. The focus on homogenous group cases has hidden the distinction between the two interpretations, and so nobody has thus far articulated the distinction or defended one interpretation over the other (even if those advancing the Relevance View may have had one in mind).

III Assessing Anchor by Competition

¹⁴ I have chosen the generic terms ‘severe impairment’ and ‘mild impairment’ in order to bracket disagreement over exactly what is ‘relevant to’ what, since my claims in this paper do not depend upon any particular view about which conditions are relevant to which.

I now want to show how diverse group cases show *both* interpretations of the anchoring rule have seriously counter-intuitive implications. I shall also highlight how the respect-based reply that some proponents of the Relevance View have given to other supposed counter-examples does not apply to my cases.

Let's examine Anchor by Competition first. Consider this case:

Case 4.

Stage 1: you can save Group A or Group B. Group A consists of one dying person. Group B consists of 10 people facing severe impairment. The Relevance View requires you to save Group B.

Stage 2: 10 people facing mild impairment are added to both Groups.

Since we had a clear conclusion at Stage 1, and then added equal claims to both sides, we might expect everything to remain the same. Either the claims of those facing mild impairment are relevant, in which case the two sets of claims counter-balance one another, or they are not relevant, in which case they would have no bearing on our choice.

However, according to Anchor by Competition, the claims of those facing mild impairment who were added to Group A are relevant (since they compete with those in Group B facing severe impairment), whilst the claims of those added to Group B are not relevant (since they compete with the person in Group A facing death). This is despite the fact that all these people have the exact same claim. This seems to violate a key principle that proponents of the Relevance View would most likely accept, a principle we can call *Equal Consideration for Equal Claims*.

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

To be clear, any plausible view concerning whom we should save between competing groups will allow that people holding equal claims may be *treated* differentially. Unless you are prepared to save nobody, in any case with equal claims on both sides you will save some but not others with the same strength of claim. But Anchor by Competition introduces a different form of inequality between equal claims, since it treats equally strong claims differently not only in terms of outcomes, but also in terms of how they affect the decision, since it recommends *considering* some claims against mild impairment as relevant to the decision, but not others. In rejecting a principle which does not require us to save the largest of two groups of people facing death, Scanlon objects that such a principle does ‘not require agents to treat the claims of each person who could be saved as having the same moral force.’¹⁵ Surely exactly the same objection applies when some claims against mild impairment are counted, and some are not.

Anchor by Competition would even allow weaker claims to be considered relevant, whilst stronger ones are considered irrelevant. Imagine a condition, mild impairment+, which is more serious than mild impairment, but not relevant to death. In Case 4, if claims against mild impairment+ are added to Group B at Stage 2, their claims are irrelevant, whilst claims against mild impairment added to Group A are relevant.

Aside from these theoretical issues, the biggest problem with Anchor by Competition is that it is vulnerable to, to my mind, devastating counter-examples, once we look at diverse group cases. For example, consider this case:

Case 5.

¹⁵ Scanlon, *What We Owe to Each Other*, p. 232. Scanlon’s application of this principle to the particular case has been criticized, but not the principle itself. See, for example, Michael Otsuka, ‘Scanlon and the Claims of the Many versus the One’ in *Analysis* 60 (2000): 288-293; and ‘Saving Lives, Moral Theory, and the Claims of Individuals.’

Stage 1: you can save Group A or Group B. Group A consists of one dying person. Group B consists of 10 people facing severe impairment. The Relevance View requires you to save Group B.

Stage 2: one person facing mild impairment is added to Group A. One billion people facing mild impairment are added to Group B.

Imagine that the decision to save Group B at Stage 1 was a very close one – the claims of the 10 facing severe impairment only *just* outweigh the claim of the dying person. Therefore, adding one relevant claim to Group A could tip the balance in favour of Group A. If that happens here, then at Stage 2 we should save Group A, because the lone claim added to Group A is relevant, whilst the one billion claims added to Group B are not.

This implication of Anchor by Competition is counter-intuitive in two ways. First, the recommendation at Stage 2 that we should save one person from death and one person from mild impairment over ten people from severe impairment and one billion people from mild impairment is counter-intuitive by itself. Second, the comparison between the recommendations given at Stage 1 and Stage 2 is also counter-intuitive. It is bizarre that at Stage 1 we ought to save Group B, and then by adding claims of the same strength to both groups, but with numbers heavily in favour of Group B, we ought to switch to Group A.

It is important to note that this counter-intuitive implication cannot be deflected by appealing to respect, in the way that supporters of the Relevance View seek to deflect other potentially counter-intuitive cases involving the addition of claims, which appear to violate the Principle of Irrelevant Alternatives. In these cases, we start with two homogenous groups, and then add a third, separate, competing group. Adding a new group forces a switch of who we save, even though we don't save the new group. This kind of case is a typical objection to the Relevance View:

Case 6.

Stage 1: You can save Group B or Group C. Group B consists of 10 people facing severe impairment. Group C consists of 100 people facing mild impairment.

Stage 2: Group A is added. This consists of one person facing death.

According to the Relevance View, we ought to save Group C over Group B at Stage 1. But when we add Group A, we ought to save Group B (since Group C's claims are no longer relevant). This is said to violate the Principle of Irrelevant Alternatives, since the addition of Group A causes us to switch, but not to A.

Advocates of the Relevance View have responded to this by arguing that introducing Group A changes the choice situation in a morally relevant way. They argue that it would be disrespectful to allow the person in Group A to die in order to save Group C, but it is not disrespectful to that person to allow them to die in order to save Group B, since Group B hold claims relevant to their own. Voorhoeve states that 'ARC rules that it would be disrespectful to the person whose life is at stake to spare a multitude [from mild impairment] rather than save his life. It would not be disrespectful to instead save many from [severe impairment].'¹⁶ Kamm also argues that respect explains why we should choose in this way. She says it would be disrespectful to attend to people facing paraplegia over someone who will die, but not disrespectful to attend to people facing quadriplegia.¹⁷

However, this line cannot be run against our Case 5 counter-example. In that case, we were all set to allow the person in Group A to die in order to save the people facing severe impairment. This is not disrespectful. It is hard, therefore, to see why it would be disrespectful to allow them to die in order to save those exact same people *plus some others*. As for the additional person added to

¹⁶ Voorhoeve, 'Why One should Count only Claims with which One Can Sympathize'.

¹⁷ Kamm, *Intricate Ethics*, p. 298.

Group A (the person whose claim against mild impairment tips things in favour of Group A) it is hard to see how they could have any respect-based complaint: competing against them is a larger number of people with a stronger claim, and a far larger number of people with the same claim as they have.

If we break this down into sub-groups, the claims of those facing severe impairment trump the dying person's claim, and the claims of the one billion trump the claim of the lone person facing mild impairment. It seems to be some kind of moral alchemy, therefore, under which the claims of the dying person and the lone person facing mild impairment somehow combine to defeat the others, and there seems to be no respect-based rationale for allowing this alchemy to stand.

IV Assessing Anchor by Strength

Let's now examine Anchor by Strength, according to which we find out whether a claim is relevant by comparing it to the strongest claim in the competition. Regardless of which anchoring rule proponents of the Relevance View originally had in mind, this variant of the Relevance View may seem attractive, as it would avoid the problems of Anchor by Competition: since all claims are compared with a single anchor (the strongest claim), claims of equal strength will either all be relevant, or all be irrelevant. Therefore, Equal Consideration for Equal Claims is respected.

However, this view is also subject to devastating counter-examples.

Case 7.

Stage 1: You can save Group A or Group B. Group A consists of 100 people facing mild impairment. Group B consists of 10 people facing severe impairment.

Stage 2: One person facing death is added to Group A.

The Relevance View coupled with Anchor by Strength would give seriously counter-intuitive judgments in this case. Imagine a life boat crew faces a situation like Case 7. At Stage 1, they radio their ethics advisor. In accordance with Anchor by Strength, she tells them to save Group A – the 100 claims against mild impairment are relevant to, and outweigh, the ten against severe impairment. At Stage 2, they radio in with new information. There is an additional person in Group A and they are dying. The lifeboat crew now press on toward Group A with renewed vigour – not only are they going to save 100 people from mild impairment, they are *also* going to save someone's life. The ethics advisor radios back: 'Turn the boat around.'

Why does the ethics advisor tell them to do this? She does so because now that there is someone facing death involved, the bar against which 'relevant claims' must be assessed shifts. According to Anchor by Strength, whilst the claims of the 100 in Group A were relevant to the strongest claims in play at Stage 1 (those of the 10 in Group B), they are not relevant to the person dying in their midst at Stage 2, and so their claims no longer count. The new calculation pits the one dying person against the 10 facing severe impairment, and so the lifeboat must be turned around, as Group B's aggregate relevant claims defeat the claim of the sole dying person.

This is a counter-intuitive implication. Anchor by Strength says that we should save the 100 over the ten at Stage 1, but the ten over the 100 *and* the dying person at Stage 2.¹⁸ It is worth spelling out how this differs from counter-examples that show that the Relevance View violates the Principle of Irrelevant Alternatives, such as Case 6. In those cases, adding a *new group* changes which of two groups we ought to save. In this case, however, it isn't the addition of *another group* that forces us to change our minds; it is an addition *to* a group –

¹⁸ If the ethics advisor favours a version of the Relevance View under which we should hold weighted lotteries that only take into account relevant claims, the addition of the dying person should change the weighting of the lottery from being in Group A's favour to being in Group B's favour. I find this just as counter-intuitive.

the addition of a dying person to Group A means that we ought to *turn away* from Group A. The dying person isn't an 'alternative' (irrelevant or not). Instead, they act as a moral repellent from Group A.

Therefore, this counter-example shows that Anchor by Strength violates another principle that we intuitively endorse:

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

For those of us who are at least attracted to the Relevance View it is quite plausible that adding a claim to a group of claims would not make it *more* choice-worthy – irrelevant claims are neutral. But the idea that adding a claim (indeed, a very weighty claim) to a group could *lessen* its choice-worthiness is seriously counter-intuitive.

Perhaps some will be prepared to bite the bullet, and both deny the Principle of Addition and agree with Anchor by Strength's recommendation in Case 7. After all, the intuition underlying both our rejection of the Case 7 recommendation and our affirmation of the Principle of Addition may be based on assuming that 'ought to be chosen over' is transitive: Saving Group A with the dying person (call this A*) ought to be chosen over (or is better than) saving Group A. Saving A ought to be chosen over (or is better than) saving B. Therefore, saving A* ought to be chosen over (or is better than) saving B. But we already know that advocates of the Relevance View reject transitivity of whatever they take to be the relevant relation, be it normative or evaluative.¹⁹ Indeed, it was a case of

¹⁹ Or, at least, they reject transitivity for simple descriptions of the options (e.g., 'Save A'). It is possible to preserve transitivity if we allow options to be described in moralized terms. See n. 11 above; and Voorhoeve, 'How Should We Aggregate Competing Claims', pp. 78-79. Voorhoeve's response draws on John Broome, *Weighing Goods* (Oxford: Blackwell, 1991), ch. 5.

‘mere addition’ that led Larry Temkin toward rejecting transitivity for ‘all-things-considered-better-than’ in the first place.²⁰

However, while the Principle of Addition appears to draw support from transitivity, rejecting the transitivity of ‘ought to be chosen over’ or ‘all-things-considered-better-than’ does not entail rejecting the Principle of Addition. It is consistent with denying that A ought to be chosen over B and B ought to be chosen over C *always* leads to the conclusion that A ought to be chosen over C, that it *sometimes* or even *often* does. And while I constructed a defence of the Principle of Addition based on transitivity above, the Principle does not explicitly rely on it. The Principle of Addition is a very restricted claim. It only applies to cases in which everything is left exactly the same, but one further claim (from an already existing person²¹) is added to one of the groups.

So, even if we reject transitivity, the Principle of Addition at the least represents a further bullet, and a pretty significant one, that must be bitten in order to affirm Anchor by Strength. Furthermore, even if we reject the Principle of Addition, Case 7 puts independent pressure on the Relevance View: whatever principles we hold in the background, the advice in Case 7 is implausible.

The respect-based response to Irrelevant Alternative cases will, again, not help us here. It may well be disrespectful to save those facing mild impairment *over* the dying person in Case 6. But the same does not ring true with our case here. Here the presence of the dying person causes us to turn away from a group *which contains them*, when the alternative is to save them *alongside* those facing mild impairment. It is hard to portray this as a form of respect. Indeed, it seems to suggest that the dying person’s life has negative moral value.

²⁰ Larry S. Temkin, ‘Intransitivity and the Mere Addition Paradox’ in *Philosophy & Public Affairs* 16 (1987): 138-187; and *Rethinking the Good*, pp. vii-ix.

²¹ Some of the key cases in which transitivity seems to fail are cases in which additional people (and not just additional claims) are added. See Temkin, ‘Intransitivity and the Mere Addition Paradox’; Derek Parfit, ‘Future Generations: Further Problems’ in *Philosophy & Public Affairs* 11 (1982): 113-172, at pp. 158-169.

V Alternative anchoring principles

Anchor by Competition and Anchor by Strength are the two most plausible interpretations of the anchoring rule in the literature on the Relevance View (and in particular of Voorhoeve's principle 4 – his key statement of the anchoring rule). Both Anchor by Competition and Anchor by Strength produce highly counter-intuitive results and violate key plausible principles in diverse group cases. Perhaps some other anchoring rule can do a better job? Although not attributable to any of the proponents of Limited Aggregation in the literature, the following rule might be preferable²²:

Anchor by Competition (Transfer): in order to be relevant, a C-type claim must be sufficiently strong relative to the strongest claim with which any C-token completes.

In other words, in order for a claim against mild impairment to be relevant, it must be sufficiently strong relative to the strongest claim with which *any* claim against mild impairment competes. If any claim against mild impairment is *irrelevant*, then this irrelevance transfers across groups. Notice how this anchoring rule would avoid the counter-intuitive implications of Anchor by Competition and Anchor by Strength, exposed in Cases 5 and 7 respectively. In Case 5, since the claims of the one billion facing mild impairment are not relevant to death, then the claim of the one person facing mild impairment is also not relevant. Therefore, adding the additional claims at Stage 2 would not force us to switch groups, and we would not violate the principle of Equal Consideration for Equal Claims. In Case 7, adding the dying person to Group A would not render the claims of those facing mild impairment in Group A irrelevant. Therefore, we would not switch to Group B at Stage 2, and therefore would not violate the Principle of Addition.

²² I am grateful to an editor of *Philosophy & Public Affairs* for suggesting this anchoring rule to me.

However, Anchor by Competition (Transfer) faces other problems. One is that it only transfers irrelevance to claims of *equal* strength. This means that, as with Anchor by Competition, claims might be considered irrelevant even though *weaker* claims are considered relevant. There are fixes for this problem. But Anchor by Competition (Transfer) also violates the Principle of Addition. Consider this case:

Case 8.

Stage 1: You can save Group A or Group B. Group A consists of 10 people facing severe impairment, and one person facing mild impairment.

Group B consists of 101 people facing mild impairment. All claims are relevant, and the Relevance View directs you to save Group B (101 people facing mild impairment).

Stage 2: one person facing death is added to Group B.

At Stage 2, claims against mild impairment are no longer relevant according to Anchor by Competition (Transfer). Since the single claim against mild impairment in Group A is now competing with a claim against death, it is not relevant, and so – due to the transfer of irrelevance – no claim against mild impairment is relevant. This means that the solo claim against death is now competing with the ten claims against severe impairment, which outweigh the claim against death. Therefore, the Relevance View tells us, at Stage 2, to switch from saving Group B to Group A: the effect of adding the claim against death to Group B is that we switch away from Group B. This is the same counter-intuitive implication as in Case 7.

These two problems with Anchor by Competition (Transfer) stem from two key restrictions. First, relevance is heavily restricted: C-claims are only relevant when they are sufficiently strong in comparison to the strongest claim with which *any* C-claim competes. *Irrelevance* is transferred. This is why all claims against mild impairment are suddenly ruled out in Case 8, Stage 2. Second,

transfer of irrelevance is restricted to claims of exactly the same strength, which is why, for example, claims against mild impairment+ might be considered irrelevant, even though weaker claims against mild impairment are considered relevant.

Both these issues can be addressed by this anchoring principle, which builds on the idea of ‘transferring’, but transfers *relevance* instead of *irrelevance*:

Anchor by Competition (Transfer 2): in order to be relevant a claim must be either: (a) sufficiently strong relative to the strongest claim with which it competes (i.e., pass the Anchor by Competition test); or (b) at least as strong as another relevant claim.

The effect of Anchor by Competition (Transfer 2) is basically to take the weakest claim that would pass the Anchor by Competition test, and then make all claims that are as strong as, or stronger than, that claim relevant. This anchoring rule would avoid the counter-intuitive implications of Anchor by Competition (Transfer). First, since it transfers relevance to all claims *as strong or stronger*, it will never allow a situation in which a weaker claim is relevant whilst a stronger claim is irrelevant. Second, in Case 8, Stage 2, since the claims against mild impairment in Group B continue to be relevant to the claims against severe impairment in Group A, all claims against mild impairment continue to be relevant. Therefore, adding a dying person to Group B in Case 8 does not force us to switch to saving Group A, and therefore Anchor by Competition (Transfer 2) does not violate the Principle of Addition in that case. In addition, like Anchor by Competition (Transfer), this principle would not fall foul of Cases 5 or 7.

However, Anchor by Competition (Transfer 2) does violate the Principle of Addition in other diverse group cases. For example, this case:

Case 9.

Stage 1: You can save Group A or Group B. Group A consists of one person facing death. Group B consists of a large number of people, N_9 , facing mild impairment. Since the claims of Group B are not relevant to claims against death, you should save Group A.

Stage 2: A single person facing mild impairment is added to Group A.

At Stage 2, the claims of Group B suddenly become relevant. Group B's claims become relevant because the claim against mild impairment that is added to Group A is relevant, since it is sufficiently strong in comparison with the strongest claims with which it competes, namely the equally strong claims in Group B. Because this one claim against mild impairment is relevant, relevance is transferred to all such claims. If N_9 is large enough, then at Stage 2, the claims of Group B will outweigh the two claims in Group A, with the result that at Stage 2 we should switch to saving Group B, even though all that changes is that a single claim is added to Group A. The Principle of Addition is therefore violated.

This may seem less alarming than previous violations of the Principle of Addition – after all, the claim that is added, but treated as if it has negative weight, is a claim that, in the context of a choice involving a claim against death, is not that important. But that the added claim is a relatively weak one is itself important, when we look at things from the perspective of the person facing death who is not saved. First, a person who is dying is not saved in order to save many people facing mild impairment. Therefore, not only does Anchor by Competition (Transfer 2) violate the Principle of Addition, it violates the core intuitions behind the Relevance View. Second, the reason this person is allowed to die whilst claims irrelevant to their own are satisfied is because they happen to find themselves in need alongside a person *whose claim is also irrelevant to their own*.

We could continue to search for an anchoring principle which will avoid violating both Equal Consideration for Equal Claims and the Principle of Addition. I fear this is impossible. Any principle (such as Anchor by Competition) which only allows anchoring claims to affect the relevance of claims in other groups will allow for a plurality of anchors, and therefore violate Equal Consideration for Equal Claims. But any principle which corrects for this by insisting that there is only one anchor for all claims will violate the Principle of Addition.

VI Healthcare cases are not cases of competing claims

In addition to the theoretical issues identified above, the Relevance View has less practical relevance than its proponents presume. In this section and the following one, I will show that the Relevance View doesn't apply to many cases of healthcare resource allocation, and that it is difficult to see how its core claims can be modified such that it does apply to such cases.

The key cases used to discuss the Relevance View are cases of competition – we can either save this person or group, or some other group. Alex Voorhoeve explicitly limits his version of the Relevance View (Aggregate Relevant Claims) to cases of competing claims. Competing claims, recall, are defined like this:

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

Therefore, ARC applies only when there are individual claims that cannot be jointly satisfied. Proponents of the Relevance View present decisions relating to healthcare resources as (complex) cases of competing claims: various groups of claims compete for scarce resources, so we can save some people from death, or others from paralysis, or some other group from chronic pain, and so on. My central claim in this section is that many decisions regarding the allocation of healthcare resources are *not* cases of competing claims in Voorhoeve's sense.

Some healthcare decisions might be cases of competing claims – for example when we can only fund research into treatments for condition x or condition y. But decisions about which *treatments* (once available) to provide will rarely, if ever, look like this.²³ Consider this case:

Case 10. A hospital trust has \$1m at its disposal. Two people require life-saving operations, costing \$500,000 each. 1 billion people require a headache pill, costing \$0.001 each.

It is tempting to put these two sets of claims into groups, and to say that *either* we save the dying *or* we save the people with headaches (both of these courses of action use up the full \$1m budget). These, after all, look like the plausible alternatives: either the claims of the dying outweigh the claims of the headache sufferers, or they don't, or we refuse to contemplate aggregation in this case. Each of these three alternatives recommends saving all the people with headaches or all the people who are dying. And this is how the literature tends to discuss such cases.²⁴ But separating them out into two groups, one of which is to be saved, makes this *seem* like a case of competing claims when in fact it is not. There are no two *individual* claims in this case that cannot be jointly satisfied. We can jointly satisfy both claims against death, or any two claims against headaches, or any claim against death alongside any claim against a headache. And this is true of healthcare resource allocation more generally: of all the claims we might meet, there will usually be various (perhaps morally implausible) possible distributions which will allow us to satisfy any two of the individual claims. In other words, healthcare resource decisions are not cases of people with condition x competing with people with condition y. Rather, there is one large group of people with a diverse range of claims. Various subsets of

²³ Kamm advocates slightly different positions (both of which employ the notion of relevant claims) for individual medical decisions and decisions about which treatments to fund. See her discussion of Sob₃ and Sob₄ in *Morality, Mortality Volume 1*, ch. 10.

²⁴ See, for example, the exchange between Voorhoeve and Kamm in *Journal of Medical Ethics* 41 (2015): 492-495; Temkin's discussion of WHO decisions in *Rethinking the Good*, pp. 79-80; Voorhoeve, 'Healthy Nails versus Long Lives'; Voorhoeve, 'Balancing Small against Large Burdens'.

this large group can be saved. Each person is both in competition with every other person, and potentially can be saved alongside every other person. Call this ‘quasi-competition’.

VII Applying the Relevance View to quasi-competition cases

Strictly speaking, therefore, Voorhoeve’s ARC, as it stands, has *nothing* to say about Case 10, and any healthcare decision like it.²⁵ And, as I will show in this section, it is unclear how we should apply the Relevance View more generally, an argument developed with cases of competing claims (such as Cases 1 and 2) in mind, to cases in which claims quasi-compete. We can try to amend or clarify the Relevance View so that it does apply to healthcare cases involving quasi-competition. In this section I consider some possible amendments. But these amendments are problematic – some entail counter-intuitive results, and others raise a series of difficulties with interpreting the requirements of respect in quasi-competition cases. Since advocates of the Relevance View have focused on competition cases, in which these issues do not arise, they have not addressed these difficulties.

The simplest way to amend or clarify the Relevance View so that it applies to quasi-competition cases would be as follows:

Relevance View (quasi-competition):

1. Individuals’ claims *quasi-compete* when they are both members of a set of claims which cannot all be satisfied.

²⁵ There are further important differences between healthcare *systems* and individual treatments. Healthcare decisions are iterative. This is relevant because whilst any individual minor ailment may not be relevant to death, a lifetime of untreated minor ailments may be (Temkin, *Rethinking the Good*, pp. 79-80; Scanlon, *What We Owe to Each Other*, pp. 236-238). In such cases, it seems better to apply to Relevance View to the policy or sequence of decisions, rather than each individual decision (this would also counter problems identified in Johan E. Gustafsson, ‘Sequential Dominance and the Anti-Aggregation Principle’ in *Philosophical Studies* 172 (2015): 1593-1601.) Furthermore, national healthcare systems paid for by taxation are often seen not as benevolent distributors, but as social insurance schemes, and if it is rational to insure against minor ailments, then on that model they should be covered (Voorhoeve, ‘Why One should Count only Claims with which One Can Sympathize’ and ‘Healthy Nails versus Long Lives’).

In cases of quasi-competition:

2. A claim is relevant if it is strong enough in comparison with the strongest claim under consideration.
3. You should choose an alternative that satisfies the greatest sum of strength-weighted, relevant claims.

So far as I can see, this is what Relevance View advocates have in mind when they apply their view to healthcare systems. For example, this is roughly how a proposed Dutch system works, which Voorhoeve considers to be an instantiation of ARC.²⁶ To test this view, consider:

Case 11. A hospital trust has \$1m at its disposal. There are three groups of patients in need. Group A is a dying person who can be saved at minimal cost. Group B is a group of people facing severe impairment who are extremely expensive to save. Group C is a very large group of people facing mild impairment who can each be saved at minimal cost.

In this case, the Relevance View (quasi-competition) would have us do the following. First, eliminate from consideration Group C, as their claims are irrelevant to the strongest claim. Second, satisfy the greatest sum of strength-weighted claims from those in Groups A and B. We will save Group A (the dying person), and then as many of Group B as we can. But this is an odd thing to do. Group A is being saved, so we don't need to worry about disrespecting them through failing to save them in order to satisfy claims irrelevant to their own. And Group C's very cheap claims are relevant to, and much more efficient to meet than, Group B's. In a competition between Groups B and C, we would save Group C. So even if we accept the respect-based rationale for the Relevance View, it gives us no reason to save Group B here. Under the Relevance View (quasi-competition) we seem to be saving Group B and not Group C as a form of respect to Group A *even when Group A is to be saved*.

²⁶ Voorhoeve, 'Healthy Nails versus Long Lives'.

There is another way to amend the Relevance View so that it can be applied to cases of quasi-competition. Call this the Relevance View (sub-competitions). This would involve seeing the overall quasi-competition as divided into a series of sub-competitions over the money.

Relevance View (sub-competitions):

In cases of quasi-competition:

1. Sub-Competition 1: Establish whether to meet the costs of satisfying the strongest claim, or to use that money to satisfy weaker relevant claims.
2. If the strongest claim is not to be satisfied, set it aside.
3. Sub-Competition 2: Establish whether to meet the costs of satisfying the strongest remaining claim (i.e., claims that have not been met or set aside), or to use that money to satisfy weaker relevant claims.
4. If the strongest remaining claim is not to be satisfied, set it aside.
5. Repeat 3 and 4 until resources run out.

In Case 11, we would first see whether it was worth spending the small amount on Group A or on the (expensive but relevant) Group B claims. We would decide in favour of Group A. With that sub-competition settled, we would then see whether Group B or C should be saved with the remaining funds, with Group B as the anchoring claim for this second sub-competition. We would decide in favour of Group C. This gets the right result in Case 11, as it doesn't continue to try to respect Group A even after it's been determined that they will be saved.

But the Relevance View (sub-competitions) may over-respect the dying in another way. Consider this case:

Case 12. A hospital trust has \$1m at its disposal. There are three groups of patients in need: Group A is one person who needs a life-saving treatment which costs \$500,000; Group B is 40 people who face severe

impairment – their treatment costs \$25,000 each; and Group C is 10,000 people who face mild impairment – their treatment costs \$100 each.

Here is how the Relevance View (sub-competitions) would handle Case 12. The strongest claim is the claim against death. The first sub-competition is therefore a competition over whether we should spend \$500,000 on saving that person, or half of Group B (since Group C's claims are irrelevant to death). We would save half of Group B. With that sub-competition settled, we now must decide how to spend the remaining \$500,000. Of the remaining 'live' claims, the remaining half of Group B's claims are the strongest. So, we take each member of Group B in turn, and decide whether to spend \$25,000 saving them, or 250 members of Group C. Group C's claims are relevant to these claims, and 250 of Group C's claims outweigh one of Group B's, and so the members of Group C would win each of these sub-competitions, and we would end up saving half of Group C.

Here is a potential problem with this view. In the first sub-competition, the Relevance View (sub-competitions) instructs us to save 20 people facing severe impairment over the dying person, in order to ensure that the dying person is not disrespected. But recall that ten claims against severe impairment outweigh death – ten would be *enough* to ensure respect. It is therefore arguable that we *also* wouldn't have disrespected the dying person in Case 12 by saving only ten people from severe impairment (at a cost of \$250,000) and using the remaining money elsewhere. Once ten people have been saved from severe impairment, we have ensured respect to the dying person, and so to save *more than* ten potentially over-respects the dying. This is *especially* problematic when, as in Case 12, the additional ten are an inefficient use of the remaining resources (since some claims relevant to their claims, and with greater aggregate weight, could be met, namely those held by Group C).

If this is indeed a problem, we can alter the Relevance View (sub-competitions) so that if we decide not to satisfy a stronger claim, we only satisfy *enough*

weaker, relevant claims so as to meet the minimum threshold for respect, allocating the rest of the money for further sub-competitions. But this is also potentially problematic: now we are using money that we might have used to save the dying person to satisfy claims irrelevant to hers. We save some people from mild impairment *instead of* saving her. Perhaps this is the kind of disrespectful action we must avoid.

These two variants of the Relevance View (sub-competitions) raise questions about how to understand the requirements of respect in quasi-competition cases.²⁷ I do not have the space to try to resolve these issues here. What I want to make clear is that these are difficulties with applying the Relevance View to quasi-competition cases, and therefore healthcare allocation, with which its proponents have not dealt, since they have assumed that the Relevance View unproblematically extends to healthcare allocation. Those who wish to apply the Relevance View to healthcare cases therefore must address these issues.

There are two main questions posed by Case 12. First, when we have decided not to meet a stronger claim (such as death), does respect require that *enough* relevant claims are satisfied, or that *only* relevant claims are met using the resources that might have been spent on the strongest claim? The two variants of the Relevance View (sub-competitions) come apart on this issue. It is unclear to me which is a better understanding of respect in these circumstances.²⁸

Second, must respect be shown by *the sub-competition of which the strongest claim is part*, or by *each sub-competition*? If the former, we can ensure respect to the unmet strongest claim (whether this is done by meeting *enough* relevant claims, or *only* relevant claims) at the first sub-competition, and then move on,

²⁷ I am grateful to two editors of *Philosophy & Public Affairs* and an audience in York (and Martin O'Neill in particular) for useful comments and discussion here.

²⁸ In discussing this work with others, I have found them to have divergent intuitions on this, and related, questions. Relatedly, the literature is unclear on what counts as disrespect in competition cases. For example, Kamm is inconsistent on whether (dis)respect is shown by *whom you save*, or whose claims you even *consider* – within one paragraph she suggests that it would be disrespectful to ‘attend to’, ‘save’, and ‘consider’ the wrong people. *Intricate Ethics*, p. 298.

putting the strongest claim aside, safe in the knowledge that respect has been achieved. If the latter, however, both variants of the Relevance View (sub-competitions) will disrespect the unmet strongest claim at the next sub-competition, since at that point we will be considering, and potentially meeting, claims irrelevant to the strongest claim.

These difficulties about how to understand the requirements of respect in quasi-competition cases show that we cannot read directly across from competition cases to quasi-competition cases. In the standard cases of competition among homogenous groups, these questions simply don't arise. In such cases, like Case 6, you either save the people facing severe impairment, or you don't: there is no question of whether saving some but not all (and saving some from mild impairment as well) would show sufficient respect. And there is no question as to whether respect must be shown through one stage of the process, or at each stage in a decision-making process. Applying the Relevance View to quasi-competition cases is not a simple matter, and so those wishing to apply the view to healthcare cases must show how this is to be done.

VIII Concluding remarks

Advocates of the Relevance View think that it is a position that works well for cases involving competing claims, and present healthcare decisions as these kinds of cases. I have shown here that healthcare decisions are not, in general, cases of competing claims, they are cases of quasi-competition, and that it is not obvious how to apply the Relevance View to cases of quasi-competition. I have also shown that in cases of diverse competing groups, the Relevance View is shown to have a hitherto unacknowledged ambiguity, and to lead to seriously counter-intuitive implications.

Where does this leave us? First, the distinction between cases of competing claims and cases of quasi-competing claims is an important one. Too often the literature suggests that we must save the people with condition x or the people with condition y. But in quasi-competition cases we can save some of the

people with condition x and some of the people with condition y. At the least, proponents of the Relevance View must explain how their view handles such cases, since, as I have shown, it is not obvious how it would or should do so. It is, however, questionable whether we should seek to apply the Relevance View beyond competition cases, given the flaws I exposed with the view as applied to competition cases in sections II-V.

Second, the anchoring problem is one that proponents of the Relevance View must address: if they wish to defend the view, they must defend a particular anchoring rule. In doing so, however, they must be prepared to bite bullets and jettison one of the normatively compelling principles I have outlined.

Third, by showing these powerful counter-examples to the Relevance View, and in showing what principles we would need to reject in order to endorse the counter-intuitive implications, I have put further pressure on, and presented new bullets to bite, for advocates of the Relevance View. This doesn't show us that we *must* reject the Relevance View – after all, both Pure Aggregation and Anti-Aggregation are hard to swallow (recall Cases 1 and 2). Each position has counter-intuitive implications, and individuals will decide differently on which implications they are and are not prepared to accept. But my cases, and the principles they show the Relevance View to violate, are very significant bullets. I was prepared to accept the Relevance View until I developed the arguments presented in sections II-V. I can no longer do so. While I find some of the implications of the other views counter-intuitive, they at least make internal sense. The counter-examples I present here are not only counter-intuitive in terms of what they recommend, the various comparisons between Stage 1 and Stage 2 in my cases also show a bizarre internal logic to the Relevance View that it is very hard to accept. Adding or subtracting claims to and from groups does not have such odd and normatively disturbing effects in either Pure Aggregation or Anti-Aggregation: neither violates Equal Consideration for Equal Claims, or the Principle of Addition. I have been impressed by the replies that advocates of the Relevance View have been able to make to criticisms from

the Principle of Irrelevant Alternatives and transitivity, but I cannot imagine rejecting either Equal Consideration for Equal Claims, or the Principle of Addition, or accepting the recommendations of the view in the cases presented here. At the least, these are concerns with the view that must be addressed.

If, with me, you reject the Relevance View, we are left with the following three options. First, we could accept Anti-Aggregation. I find this the least plausible of our options. Imagine N_1 in Case 1 were one billion people. It is implausible that we should allow so many people to suffer such a serious harm in order to save one life. Second, we could accept some form of Pure Aggregation (this could involve giving some additional weight to the worse off in a prioritarian or egalitarian manner, rather than simply maximizing the satisfaction of claims). However, no matter how far we stack the deck in favour of those with more serious claims, this position does force us to acknowledge that if N_2 gets large enough in Case 2, we ought to prevent the mild headaches. This is a sizeable bullet to bite.²⁹

The final option is to try to find an alternative version of Limited Aggregation, or at least a position that explains our judgments in Cases 1 and 2, without running into the problems of Anti-Aggregation, Pure Aggregation, or the Relevance View. One version of this option would be to articulate a plurality of values, which must be traded off against one another in hard cases. Another would be to continue to use the notion of the ‘relevance’ of claims, but not to use it to establish whether claims are relevant or irrelevant to the decision *simpliciter*. Instead, we could ask whether a claim is relevant to a series of sub-competitions or questions. For example, we could ask whether the strongest claim is ‘matched’ by enough claims relevant to it. If not, we meet the strongest claim. If it is matched, we look at the next strongest claim, and see if that is

²⁹ For an attempt to show that it is not as significant a bullet as we might be tempted to think, see Alistair Norcross, ‘Comparing Harms: Headaches and Human Lives’ in *Philosophy & Public Affairs* 26 (1997): 135-167.

matched by enough claims relevant to it, and so on.³⁰ As things stand, I think we are forced to choose between Pure Aggregation, or finding another version of Limited Aggregation.

³⁰ Variants of this view were suggested to me independently by an editor of *Philosophy & Public Affairs* and Victor Tadros. Tadros' version, which takes the arguments of the present paper as its starting point, is outlined in his 'Restricted Aggregation: the Local Relevance View' (unpublished m/s). The view, though about competition cases, has much in common with the second variant of the Relevance View (sub-competitions), outlined in Section VII. As well as the concerns about the demands of respect expressed there, I have additional concerns with this view, but I certainly think it is worthy of exploration. One concern is that there may be a plurality of ways to 'match up' relevant claims, which would deliver different verdicts. My chief concern is that the view allows that whether or not we should save a life (or indeed hundreds of thousands of lives) could turn on whether or not one individual has a mild headache. Because the headache is ruled as irrelevant to whether the deaths are 'matched', but not ruled out as irrelevant to the larger question of which group we should save, it could end up being the determining factor. This goes against many Limited Aggregation intuitions – when people are dying, these tiny claims just *don't matter*. See, for example, Kamm on tie-breaks in *Mortality, Mortality Volume 1*, ch. 8.