Always Aggregate

1. Introduction

According to some moral views, at least when other things are equal, we should maximize the sum of well-being.[1] These views are *fully aggregative*. The most influential fully aggregative view is *Utilitarianism*, but there are many others.[2]

Fully aggregative views tend to have very counterintuitive implications in cases in which we can save either a few people from very large burdens or a huge number of people from very small burdens.[3] Suppose that in

> *Death v Headaches:* You can save either X from death or a huge number of people from headaches.[4]

This paper began as comments on Victor Tadros's 'Localised Restricted Aggregation', which I gave at the University of St Andrews Centre for Ethics, Philosophy, and Public Affairs in October 2017. I am grateful to Theron Pummer for inviting me to give those comments, and to Victor, Theron, and everyone in attendance for a helpful discussion. I also benefited from the comments of two anonymous editors at this journal, and from discussions with Alexander Dietz, Abelard Podgorski, Jonathan Quong, Alexander Sarch, Victor Tadros, Patrick Tomlin, Alec Walen, and audiences at the UCL Philosophy Departmental Symposium and the LSE/UCL Workshop on Risk and Aggregation in Ethics.

[1] The 'other things equal' clause allows these views to accommodate moral constraints and agent-centered prerogatives. See Robert Nozick, *Anarchy, State, and Utopia* (Blackwell Publishing, 2006): 28–30; and Samuel Scheffler, *The Rejection of Consequentialism* (Oxford University Press, 1994).

[2] For example, *Prioritarianism*. See Derek Parfit, 'Equality and Priority', *Ratio* 10 (1997): 202–221.

[3] There are ways to avoid these implications, but they lead to even greater problems. See James Griffin, *Well-Being: Its Meaning, Measurement, and Moral Importance* (Oxford University Press, 1989): 75–92; Alastair Norcross, 'Comparing Harms: Headaches and Human Lives', *Philosophy & Public Affairs* 26 (1997): 135–167; Larry Temkin, *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning* (Oxford University Press, 2012): 45–52; and Gustaf Arrhenius and Wlodek Rabinowicz, 'Value Superiority', in Iwao Hirose and Jonas Olson (eds.), *The Oxford Handbook of Value Theory* (Oxford University Press, 2015): 225–248.

[4] I take this example from Norcross, 'Comparing Harms'. For another influential example, see T. M. Scanlon, *What We Owe to Each Other* (Harvard University Press, 1998): 235.

Intuitively, you should save X, no matter how many people face headaches. But fully aggregative views imply that, if enough people face headaches, you should instead save them, for the well-being generated by saving X from death will be less than the sum of well-being generated by saving a large enough number of people from headaches.

Some philosophers have suggested that what matters in cases like Death v Headaches is not the sum of well-being that you could bring about, but rather the strength of the individual complaints that could be made against your act.[5] X could make a much stronger individual complaint against you saving the people facing headaches than any of them could make against you saving X. That seems to be why you are morally required to save X, no matter how many people face headaches.

This line of reasoning supports a *non-aggregative* moral view.[6] According to non-aggregative moral views, at least when other things are equal, we should minimize the strongest individual complaint. These views have much more plausible implications in cases like Death v Headaches. But they have very counterintuitive implications in cases in which we can save either a few people from very large burdens or a huge number of people from burdens that are just slightly smaller. Suppose that in

> *Death v Quadriplegia:* You can save either X from death or a huge number of people from quadriplegia.

[5] For example, Scanlon, *What We Owe to Each Other*, 235; and John M. Taurek, 'Should the Numbers Count?', *Philosophy & Public Affairs* 6 (1977): 293–316.

[6] For a defence of non-aggregative views, see Elizabeth Anscombe, 'Who is Wronged?', *The Oxford Review* 5 (1967): 16–17; Taurek, 'Should the Numbers Count?'; Véronique Munoz-Dardé, 'The Distribution of Numbers and the Comprehensiveness of Reasons', *Proceedings of the Aristotelian Society* 105 (2005): 191–217; and Tyler Doggett, 'Saving the Few', *Noûs* 47 (2013): 302–315.

Since X has a stronger individual complaint against death than any of the others has against quadriplegia, non-aggregative views have the very counterintuitive implication that you should save X, no matter how many people face quadriplegia.

These problems with fully aggregative and non-aggregative views have inclined many philosophers to accept some kind of *partially aggregative* moral view.[7] According to partially aggregative moral views, at least when other things are equal, we should minimize the sum of strength-weighted, *relevant* complaints. I will get precise about when complaints are relevant shortly, but here is how these views apply to our cases. In Death v Headaches, the complaints of the people facing headaches are not relevant, because they are too weak relative to the complaint that X has against being left to die. So partially aggregative views imply that you should save X, no matter how many people face headaches. In Death v Quadriplegia, the complaints of the people facing quadriplegia are relevant, because they are sufficiently strong relative to the complaint that X has against being left to die. So partially aggregative views imply that, if enough people face quadriplegia, you should save them.

Partially aggregative views are not without problems. Most notably, they seem to give rise to either deontic cycling or violations of a principle known as 'the independence of irrelevant alternatives'.[8] However, these problems have not dissuaded many proponents of these views. They have argued that deontic cycling can be avoided in a principled way, and that either the violations of the independence of irrelevant alternatives are merely apparent or

[7] For example, F. M. Kamm, 'Nonconseqeuntialism', in Hugh LaFollette (ed.), *The Blackwell Guide to Ethical Theory* (Blackwell Publishing, 2013): 278–284; Scanlon, *What We Owe to Each Other*, 239–240; and Alex Voorhoeve, 'How Should We Aggregate Competing Claims?', *Ethics* 125 (2014): 64–87.

[8] See Derek Parfit, 'Justifiability to Each Person', *Ratio* 16 (2003): 368–390, at 384–385; and John Halstead, 'The Numbers Always Count', *Ethics* 126 (2016): 789–802, at 797–799.

this principle is less plausible than it initially seems.[9] I will not rehearse this debate here. I want to instead focus on a more recent challenge to partially aggregative views.

Patrick Tomlin has recently argued that the most promising partially aggregative views in the literature have implausible implications in certain cases in which there are additions or subtractions to the groups of people that we can save.[10] He concludes that we must either develop a new kind of partially aggregative view or, because the counterintuitive implications of non-aggregative views seem to him much worse than the counterintuitive implications of fully aggregative views, accept a fully aggregative view.

Several philosophers have already begun responding to Tomlin's argument by developing partially aggregative views that avoid the relevant implications. Victor Tadros has made particularly quick and impressive progress.[11] But I believe these efforts will ultimately be in vain. In this paper, I extend Tomlin's argument to create a dilemma for partially aggregative views that is both fatal and unavoidable. Since I agree with Tomlin that the counterintuitive implications of non-aggregative views are much worse than the counterintuitive implications of fully aggregative views, I conclude that we should accept a fully aggregative view.

In Section 2, I summarise Tomlin's argument. In Section 3, I present the new kind of partially aggregative view that Tadros thinks can overcome this argument. In Section 4, I present an initial problem for this new kind of partially aggregative view, and then show that

[9] For example, F. M. Kamm, *Intricate Ethics: Rights, Responsibilities, and Permissible Harm* (Oxford University Press, 2007): 484–486; Voorhoeve, 'How Should We Aggregate Competing Claims?', 76–79; and Alex Voorhoeve, 'Why One Should Count Only Claims with Which One Can Sympathize', *Public Health Ethics* 10 (2017): 148–156, at 152–153.

[10] Patrick Tomlin, 'On Limited Aggregation', *Philosophy & Public Affairs* 45 (2017): 232–260.

[11] Victor Tadros, 'Localised Restricted Aggregation', *Oxford Studies in Political Philosophy, Volume 5* (forthcoming). Also see Aart van Gils and Patrick Tomlin, 'Relevance Rides Again? Aggregation and Local Relevance', unpublished.

this problem is the seed of a fatal dilemma that no partially aggregative view can avoid. In Section 5, I explain why it would be a mistake to embrace either horn of this dilemma.

2. Tomlin's Argument

2.1. Competitive Relevance and Broad Relevance

According to partially aggregative views, at least when other things are equal, we should minimize the sum of strength-weighted, relevant complaints. When is a complaint relevant? Tomlin identifies two possible answers.[12] According to *Competitive Relevance*, a complaint is relevant if and only if it is sufficiently strong relative to the strongest complaint with which it competes. According to *Broad Relevance*, a complaint is relevant if and only if it is sufficiently strong relative to the strongest complaint in the competition.

Let me illustrate this distinction. Suppose that a complaint against a lost finger is sufficiently strong relative to a complaint against a lost arm, and a complaint against a lost arm is sufficiently strong relative to a complaint against death, but a complaint against a lost finger is not sufficiently strong relative to a complaint against death. Now consider the following two cases.

> *Case 1:* You can save either group A, which contains 1 person facing death and 1 person facing a lost finger, or group B, which contains 1 person facing a lost arm.

> *Case 2:* You can save either group A, which contains 1 person facing death, or group B, which contains 1 person facing a lost arm and 1 person facing a lost finger.

[12] Tomlin refers to these answers respectively as 'Anchor by Competition' and 'Anchor by Strength'. See 'On Limited Aggregation', 239.

Focus on the potential complaint of the person facing a lost finger. On Competitive Relevance, this complaint is relevant in Case 1, for it is sufficiently strong relative to a complaint against a lost arm, which is the strongest complaint with which it competes, but it is not relevant in Case 2, for it is not sufficiently strong relative to a complaint against death, which is the strongest complaint with which it competes. On Broad Relevance, this complaint is not relevant in either case, for it is not sufficiently strong relative to a complaint against death, which is the strongest complaint in both competitions.

2.2. Against Competitive Relevance

Tomlin argues that, regardless of which view about relevance we adopt, partially aggregative views have implausible implications in certain cases in which there are additions or subtractions to the groups of people that we can save.[13] Suppose that 10 complaints against a lost arm equals 1 complaint against death (that is, the combined moral weight of the former is equal to the moral weight of the latter). Now consider *Case 3*, which has two stages.

> *Stage 1:* You can save either group A, which contains 1 person facing death, or group B, which contains 10 people facing a lost arm.
>
> *Stage 2:* 1 person facing a lost finger is added to group A, and 1,000,000 people facing a lost finger are added to group B.

Partially aggregative views imply that it is permissible to save either group at Stage 1. If it is permissible to save either group at Stage 1, it must be at least permissible to save B at Stage 2. But on Competitive Relevance, partially aggregative views imply otherwise. Since the

[13] Tomlin, 'On Limited Aggregation', 240–247.

1,000,000 complaints against a lost finger added to B are in competition with a complaint against death, they are not relevant. But the 1 complaint against a lost finger added to A is relevant, for it is in competition only with a complaint against a lost arm. This complaint must tip the balance. So partially aggregative views imply that you should save A.

Of course, we could reject the suppositions about how these complaints compare. But that would not solve the problem. If partially aggregative views are to capture the judgments that motivate them, there must be some set of complaints C1, C2, and C3 such that C1 is relevant to C2, C2 is relevant to C3, and C1 is not relevant to C3. So there must be some set of complaints that gives rise to this problem.

2.3. Against Broad Relevance

We can avoid the above problem by adopting Broad Relevance. On this view, none of the complaints added at Stage 2 are relevant, for they are not sufficiently strong relative to the strongest complaint in the competition, which is a complaint against death. So partially aggregative views imply that Stage 2 makes no moral difference.

However, as Tomlin argues, Broad Relevance faces an even worse problem. Suppose that 4,000 complaints against a lost finger outweigh 20 complaints against a lost arm, and 20 complaints against a lost arm outweigh 1 complaint against death. Now consider *Case 4*, which has two stages.

> *Stage 1:* You can save either group A, which contains 4,000 people facing a lost finger, or group B, which contains 20 people facing a lost arm.
>
> *Stage 2:* 1 person facing death is added to group A.

Partially aggregative views imply that you should save A at Stage 1. If you should save A at Stage 1, you should clearly save A at Stage 2. But on Broad Relevance, partially aggregative views imply otherwise. At Stage 2, the 4,000 complaints against a lost finger in A become irrelevant, for they are not sufficiently strong relative to the strongest complaint in the competition, which is now a complaint against death. Since the only relevant complaints are the 1 complaint against death and the 20 complaints against a lost arm, and the latter outweighs the former, partially aggregative views imply that you should save B.

Thus, regardless of whether we accept Competitive Relevance or Broad Relevance, partially aggregative views have implausible implications in certain cases in which there are additions or subtractions to the groups of people that we can save. Tomlin concludes that we must either develop a new kind of partially aggregative view or accept a fully aggregative view.[14]

3. Tadros's Response

Tadros argues that we should respond to Tomlin by accepting a new kind of partially aggregative view that he calls *Local Relevance*.[15] He does not formulate this view precisely, but the basic idea is this: if a complaint C1 is not sufficiently strong relative to a competing complaint C2, C1 cannot contribute to counterbalancing C2, but C1 can still contribute to counterbalancing other competing complaints relative to which it is sufficiently strong.

Recall Case 3.

Stage 1: You can save either group A, which contains 1 person facing death, or group B, which contains 10 people facing a lost arm.

[14] Tomlin, 'On Limited Aggregation', 259–260.

[15] Tadros, 'Localised Restricted Aggregation'.

Stage 2: 1 person facing a lost finger is added to group A, and 1,000,000 people facing a lost finger are added to group B.

On Competitive Relevance, partially aggregative views have the implausible implication that, though it is permissible to save either group at Stage 1, you should save A at Stage 2. On Local Relevance, partially aggregative views still imply that it is permissible to save either group at Stage 1, but they need not imply that you should save A at Stage 2. Though the complaints against a lost finger added to B cannot contribute to counterbalancing the complaint against death in A, they can still contribute to counterbalancing the complaint against a lost finger added to A. Since they more than counterbalance this complaint, this complaint cannot tip the balance in favour of A. So Stage 2 makes no difference.

Recall Case 4.

Stage 1: You can save either group A, which contains 4,000 people facing a lost finger, or group B, which contains 20 people facing a lost arm.

Stage 2: 1 person facing death is added to group A.

On Broad Relevance, partially aggregative views have the implausible implication that, though you should save A at Stage 1, you should save B at Stage 2. On Local Relevance, partially aggregative views still imply that you should save A at Stage 1, but they need not imply that you should save B at Stage 2, for the complaints against a lost finger in A can always contribute to counterbalancing the complaints against a lost arm in B.

Local Relevance therefore avoids the problems that Tomlin presents for both Competitive Relevance and Broad Relevance.

4. A Fatal Dilemma

In this section, I present an initial problem for Local Relevance, and then show that this problem is the seed of a fatal dilemma for all partially aggregative views.

4.1. A Path Dependence Problem

Local Relevance seems more promising than Competitive Relevance and Broad Relevance. But it faces a path dependence problem. Suppose that 2,000 complaints against a lost finger equal 20 complaints against a lost arm, and 20 complaints against a lost arm outweigh 1 complaint against death. Now consider *Case 5*, which has two stages.

> *Stage 1:* You can save either group A, which contains 1 person facing death, or group B, which contains 20 people facing a lost arm.
>
> *Stage 2:* 2,000 people facing a lost finger are added to each group.

Partially aggregative views imply that you should save B at Stage 1. If you should save B at Stage 1, it must be at least permissible to save B at Stage 2. Local Relevance can capture this judgment if we apply it as follows.

Group A	Group B
1 person facing death	20 people facing a lost arm
~~2,000 people facing a lost finger~~	~~2,000 people facing a lost finger~~

Here, we allow the complaints added to each group to counterbalance each other, so that Stage 2 makes no difference. However, we could instead apply Local Relevance as follows.

Group A	Group B
1 person facing death	~~20 people facing a lost arm~~
~~2,000 people facing a lost finger~~	2,000 people facing a lost finger

Here, we allow the complaints added to A to counterbalance the complaints against a lost arm in B. If we apply Local Relevance in this way, it has the implausible implication that, though you should save B at Stage 1, you should save A at Stage 2. Tadros give us no reason to apply Local Relevance in the former way rather than the latter.

4.2. A Fatal Dilemma

Perhaps Tadros can give us a reason to apply Local Relevance in the former way rather than the latter. As I will now show, Local Relevance would then face a different problem.

We begin with a second illustration of the path dependence problem. Consider *Case 6*, which has two stages.

> *Stage 1:* You can save either group A, which contains 1 person facing death, or group B, which contains 4,000 people facing a lost finger.
>
> *Stage 2:* 20 people facing a lost arm are added to each group.

Partially aggregative views imply that you should save A at Stage 1. If you should save A at Stage 1, it must be at least permissible to save A at Stage 2. Local Relevance can capture this judgment if we apply it as follows.

Group A	Group B
1 person facing death	4,000 people facing a lost finger
~~20 people facing a lost arm~~	~~20 people facing a lost arm~~

However, we could instead apply Local Relevance as follows.

Group A	Group B
1 person facing death	~~2,000 people facing a lost finger~~
~~20 people facing a lost arm~~	2,000 people facing a lost finger
	20 people facing a lost arm

If we apply Local Relevance in this second way, it has the implausible implication that, though you should save A at Stage 1, you should save B at Stage 2, for the 20 complaints against a lost arm added to B outweigh the 1 complaint against death in A.

Suppose that Tadros gives us a reason to apply Local Relevance in the former way rather than the latter. Now consider three more cases, each of which has only one stage.

Case 7: You can save either group C, which contains 1 person facing death, or group D, which contains 20 people facing a lost arm.

Case 8: You can save either group E, which contains 20 people facing a lost arm, or group F, which contains 4,000 people facing a lost finger.

Case 9: You can save either group C + E or group D + F.

Since 20 complaints against a lost arm outweigh 1 complaint against death, you should save D in Case 7. Since 4,000 complaints against a lost finger outweigh 20 complaints against a lost arm, you should save F in Case 8. Since you should save D rather than C in Case 7, and F rather than E in Case 8, it seems clear that you should save D + F in Case 9. Local Relevance can capture this judgment only if we apply it in the latter of the following two ways.

Group C + E	Group D + F
1 person facing death	4,000 people facing a lost finger
~~20 people facing a lost arm~~	~~20 people facing a lost arm~~

Group C + E	Group D + F
1 person facing death	~~2,000 people facing a lost finger~~
~~20 people facing a lost arm~~	2,000 people facing a lost finger
	20 people facing a lost arm

But C + E is equivalent to A in Case 6, and D + F is equivalent to B in Case 6. So it would be inconsistent to apply Local Relevance in the former way in Case 6 and the latter way in Case 9. So Local Relevance must have an implausible implication in one of these cases.

This dilemma for Local Relevance is really a dilemma for all partially aggregative views. In order to avoid an implausible implication in Case 6 (this is the first horn of the

dilemma), a partially aggregative view must balance claims in the former way. In order to avoid an implausible implication in Case 9 (this is the second horn of the dilemma), a partially aggregative view must balance claims in the latter way. But it is inconsistent to balance claims in the former way in Case 6 and the latter way in Case 9. So all partially aggregative views must have an implausible implication in one of these cases.

On reflection, it is unsurprising that partially aggregative views face this dilemma. These views imply that, given a straight choice between C and D, you should save D, and given a straight choice between D and F, you should save F, and yet given a straight choice between C and F, you should save C. Proponents of these views have argued that this kind of cycle is unimportant, because it disappears when all three options are available simultaneously (they hold that the presence of C makes F irrelevant, so that D can triumph).[16] But this cycle is not unimportant, for it is clearly what gives rise to the dilemma. Since partially aggregative views imply that you should save C rather than F, and D is identical to E, it seems these views should direct you to save C + E rather than D + F. But because they imply that you should save D rather than C and that you should save F rather than E, it also seems they should direct you to save D + F rather than C + E. They cannot do both.[17]

5. Sharpening the Horns

Proponents of partially aggregative views have defended the cyclic implications noted above by offering justifications that make these implications easier to accept.[18] They might be

[16] Since partially aggregative views generate this cycle, they imply that 'ought to do rather than' is an intransitive relation. But, importantly, they do not have the much more controversial implication that 'all-things-considered better than' is intransitive, for claims about what we ought to do need not imply claims about what is better. See Kamm, *Intricate Ethics*, 484–486; Voorhoeve, 'How Should We Aggregate Competing Claims?', 76–79; and Voorhoeve, 'Why One Should Count Only Claims with Which One Can Sympathize', 152–153.

[17] I am grateful to an anonymous editor for suggesting this way of presenting the dilemma.

[18] See the references in footnote 16.

tempted to respond to my dilemma in a similar way, embracing one of the horns and offering a mitigating justification. In this section, I explain why that would be a mistake.

5.1. Against Embracing the First Horn

How bad would it be to embrace the first horn? Consider a modified version of Case 6.

> *Stage 1:* You can save either group A, which contains 1 person facing death, or group B, which contains a ginormous number of people facing a lost finger.
>
> *Stage 2:* 1,000,000 people facing a lost arm are added to group A, and 20 people facing a lost arm are added to group B.

Partially aggregative views imply that you should save A at Stage 1. If you should save A at Stage 1, it is clear that you should save A at Stage 2. Since the number of people facing a lost arm added to A is vastly greater than the number added to B, this is even clearer than in the original case. If we embrace the first horn of my dilemma, however, we must allow the 1,000,000 complaints against a lost arm added to A to be neutralized by the complaints against a lost finger in B, and that leaves the 20 complaints against a lost arm added to B free to outweigh the 1 complaint against death in A. We must then accept the bizarre implication that, though you should save A at Stage 1, you should save B at Stage 2.

Could reflecting on possible justifications for partially aggregative views make this implication easier to accept? These views are typically justified by appeal to respect. Here is a respect-based justification for the implication that we are considering: It is disrespectful to press a complaint against a harm when this complaint is intended to compete with a complaint against a much stronger harm, but it is not disrespectful to press a complaint

against a harm when this complaint is intended to compete only with a complaint against a harm that is sufficiently close in size. Thus, at Stage 1, the people facing a lost finger cannot press a relevant complaint, for their complaints would be intended to compete with the complaint against death, and that would be disrespectful. At Stage 2, however, the people facing a lost finger can press a relevant complaint, for their complaints can be intended to compete only with the complaints against a lost arm added to A. The complaints against a lost finger in B defeat the complaints of the people facing a lost arm in A, and that leaves the complaints against a lost arm added to B free to defeat the complaint against death in A.

If this respect-based justification helps, it does not help enough. It remains bizarre to hold that you should save A at Stage 1 and then, after we add complaints of equal strength to both groups with the numbers heavily in favour of A, you should save B at Stage 2. Furthermore, respect-based justifications for the opposite result are easy to find, making it dubious to place much weight on any of these justifications. We could reason as follows: It is disrespectful to press a complaint against a harm whenever doing so would result in someone not being saved from a much stronger harm. At Stage 1, the people facing a lost finger cannot press a relevant complaint, for doing so would result in someone not being saved from death, and so would be disrespectful. At Stage 2, the people facing a lost finger still cannot press a relevant complaint, for doing so would still result in someone not being saved from death. Since there is nothing to defeat the complaints against a lost arm added to A, and these complaints defeat the complaints against a lost arm added to B, we have a respect-based justification for holding that you should save A at both Stage 1 and Stage 2.[19]

[19] This alternative justification could be used to support the bizarre implication that Competitive Relevance has in Case 3: the 1,000,000 people facing a lost finger added to B cannot make a relevant complaint, for doing so would result in someone not being saved from death, and so would be disrespectful, whereas the person facing a lost finger added to A can make a relevant complaint, for doing so would not result in anyone not being saved from death. Thus, Tomlin is wrong when he says that there is 'no respect-based rationale' for the bizarre implication that Competitive Relevance has in Case 3. See Tomlin, 'On Limited Aggregation', 244.

5.2. Against Embracing the Second Horn

How bad would it be to embrace the second horn? We would then accept that, though you should save D rather than C in Case 7, and F rather than E in Case 8, you should save C + E rather than D + F in Case 9. To reveal how bizarre this implication is, we can combine the cases and make things more vivid. Suppose that on your left are two buttons marked C and D, and on your right are two buttons marked E and F. If you press a button, that will save the corresponding group. But you can press only one button in each pair. Suppose next that your arm span is just slightly too short for you to reach both pairs of buttons simultaneously. It follows, on the view that we are considering, that you should press button D on your left and then button F on your right, even though, had your arm span been just slightly longer, it would have been permissible for you to simultaneously press buttons C and E.

We could again provide a respect-based justification for this bizarre implication, by appealing to the latter justification described above. But that would not be enough to make the implication plausible. And the former justification yields the opposite result, making appeal to either justification very dubious.

6. Summing Up (Pun Intended)

Fully aggregative views tend to have counterintuitive implications in cases in which we can save either a few people from very large burdens or a huge number of people from very small burdens. They imply, for example, that you should save an enormous number of people from headaches rather than saving one person from death. We have just seen that all partially aggregative views have implausible implications in either Case 6 or Case 9. Both implications seem to me much worse than the counterintuitive implications of fully aggregative views. Since the counterintuitive implications of non-aggregative views also

seem to me much worse than the counterintuitive implications of fully aggregative views, I conclude that we should accept a fully aggregative view.

I will end by noting that, in other work, I have argued that all partially aggregative views have implausible implications in certain cases involving risk.[20] To my mind, even considered on their own, these implications are sufficient reason to reject partially aggregative views in favour of fully aggregative views. But we should not consider these implications on their own. We should add them to the implausible implications that partially aggregative views have in either Case 6 or Case 9. When we do that, it is even clearer that partially aggregative views should be rejected, and that we should always aggregate.[21]

Bibliography

Anscombe, Elizabeth. 'Who is Wronged?', *The Oxford Review* 5 (1967): 16–17

Arrhenius, Gustaf and Rabinowicz, Wlodek. 'Value Superiority', in Iwao Hirose and Jonas Olson (eds.), *The Oxford Handbook of Value Theory* (Oxford University Press, 2015): 225–248

Broome, John. *Weighing Lives* (Oxford University Press, 2004)

Doggett, Tyler. 'Saving the Few', *Noûs* 47 (2013): 302–315

van Gils, Aart and Patrick Tomlin, 'Relevance Rides Again? Aggregation and Local Relevance', unpublished

Griffin, James. *Well-Being: Its Meaning, Measurement, and Moral Importance* (Oxford University Press, 1989)

Halstead, John. 'The Numbers Always Count', *Ethics* 126 (2016): 789–802

Horton, Joe. 'Aggregation, Complaints, and Risk', *Philosophy & Public Affairs* 45 (2017): 54–81

Huemer, Michael. 'In Defence of Repugnance', *Mind* 117 (2008), 899–933

Kamm, F. M. *Intricate Ethics: Rights, Responsibilities, and Permissible Harm* (Oxford University Press, 2007)

——'Nonconseqeuntialism', in Hugh LaFollette (ed.), *The Blackwell Guide to Ethical Theory* (Blackwell Publishing, 2013): 261–286

Munoz-Dardé, Véronique. 'The Distribution of Numbers and the Comprehensiveness of Reasons', *Proceedings of the Aristotelian Society* 105 (2005): 191–217

[20] Joe Horton, 'Aggregation, Complaints, and Risk', *Philosophy & Public Affairs* 45 (2017): 54–81.

[21] It is also worth noting that several philosophers have argued that our intuitions about cases like Death v Headaches are not reliable, because these cases require us to imagine very large numbers of people and we cannot accurately imagine very large quantities. See Norcross, 'Comparing Harms', 146–147; Parfit, 'Justifiability to Each Person', 385– 386; John Broome, *Weighing Lives* (Oxford University Press, 2004): 56–57; Michael Huemer, 'In Defence of Repugnance', *Mind* 117 (2008): 899–933, at 907–910; Halstead, 'The Numbers Always Count', 797; and Horton, 'Aggregation, Complaints, and Risk', 72–73.

Norcross, Alastair. 'Comparing Harms: Headaches and Human Lives', in *Philosophy & Public Affairs* 26 (1997): 135–167

Nozick, Robert. *Anarchy, State, and Utopia* (Basic Books, 2013)

Parfit, Derek. 'Equality and Priority', *Ratio* 10 (1997) 202–221

——'Justifiability to Each Person', *Ratio* 16 (2003): 368–390

Scanlon, T. M. *What We Owe to Each Other* (Harvard University Press, 1998)

Scheffler, Samuel. *The Rejection of Consequentialism* (Oxford University Press, 1982)

Tadros, Victor. 'Localised Restricted Aggregation', *Oxford Studies in Political Philosophy, Volume 5* (forthcoming)

Taurek, John M. 'Should the Numbers Count?', *Philosophy & Public Affairs* 6 (1977): 293–316

Temkin, Larry. *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning* (Oxford University Press, 2012)

Tomlin, Patrick. 'On Limited Aggregation', *Philosophy & Public Affairs*: 45 (2017): 232–260

Voorhoeve, Alex. 'How Should We Aggregate Competing Claims?', *Ethics* 125 (2014), 64–87

——'Why One Should Count Only Claims with Which One Can Sympathize', *Public Health Ethics* 10 (2017): 148–156