

Partial Aggregation in Ethics

(forthcoming in *Philosophy Compass*)

1. Introduction

According to the moral view known as *Act Utilitarianism*, we should always maximize aggregate well-being. This view has plausible implications in a wide range of cases.

However, it has very counterintuitive implications in cases in which we can save either a few people from severe burdens or many people from minor burdens. Suppose that in

Death v Migraines: We can save either X from death or a huge number of people from migraines.

Intuitively, we should save X, no matter how many people face migraines. But Act Utilitarianism implies that, if enough people face migraines, we should instead save them, for the well-being at stake for them will sum to more than the well-being at stake for X.

There are other moral views that have this implication in *Death v Migraines*, and that have corresponding implications in equivalent cases. We can say that these views are *fully aggregative*. Act Utilitarianism is probably the most influential fully aggregative view, but these views can take many other forms. They can endorse moral options, which permit us to give extra weight to our own well-being and the well-being of those we love; require us to respect moral constraints, such as constraints against killing and breaking promises; and direct us to give extra weight to the well-being of the badly off.¹

¹ For discussion of moral options and moral constraints, see Shelly Kagan, *Normative Ethics* (London: Routledge, 2018). For discussion of equality and priority for the badly off, see Derek Parfit, 'Equality and Priority', *Ratio* 10 (1997): 202–221.

Some philosophers have suggested that what matters in cases like Death v Migraines is not the aggregate well-being that we could bring about but rather the strength of the individual complaints that could be made against our act.² X could make a much stronger individual complaint against us saving the people facing migraines than any of them could make against us saving X. That seems to be why we are morally required to save X, no matter how many people face migraines.

This line of reasoning supports a moral view that we can call *Minimax Complaint*. According to this view, we should always minimize the strongest individual complaint. This view has intuitive implications in cases like Death v Migraines. However, it has very counterintuitive implications in cases in which we can save either a few people from severe burdens or many people from burdens that are just somewhat less severe. Suppose that in

Death v Paraplegia: We can save either X from death or a huge number of people from paraplegia.

Since X has a stronger individual complaint against death than any of the others has against paraplegia, Minimax Complaint has the very counterintuitive implication that we should save X, no matter how many people face paraplegia.

There are other moral views that have this implication in Death v Paraplegia, and that have corresponding implications in equivalent cases. We can say that these views are *non-aggregative*.³ Minimax Complaint is probably the most influential non-aggregative view, but

² See John M. Taurek, 'Should the Numbers Count?', *Philosophy & Public Affairs* 6 (1977): 293–316; and T. M. Scanlon, *What We Owe to Each Other* (Cambridge, MA: Harvard University Press, 1998): Chapter 5.

³ For discussion, see Elizabeth Anscombe, 'Who is Wronged? Philippa Foot on Double Effect: One Point', *The Oxford Review* 5 (1967): 16–17; Taurek, 'Should the Numbers Count?'; Scanlon, *What We Owe to Each Other*, Chapter 5; Véronique Munoz-Dardé, 'The Distribution of Numbers and the Comprehensiveness of Reasons',

again these views can take many other forms. They can endorse moral options, moral constraints, and priority for the badly off.

We might hope to find a moral view that forbids aggregation in cases like *Death v Migraines* while allowing aggregation in cases like *Death v Paraplegia*. We can say that such a view would be *partially aggregative*.⁴ Is there any systematic and plausible partially aggregative view? If not, should we accept a fully aggregative or a non-aggregative view? These questions have central importance in moral and political philosophy, for our answers are likely to influence the positions we take in a wide range of moral and political debates. But these questions also have considerable practical importance. For example, they have implications for how governments should allocate health care funding, and for how philanthropists should choose the charities to which they donate.

In this article, I review the most natural and influential ways of developing partially aggregative views and explain the main problems they face. Since the article is intended as a primer, I do not aim to defend any thesis. But it might be helpful for the reader to know that, in my opinion, we should accept a fully aggregative view. These views seem to me more plausible than non-aggregative views, and their most counterintuitive implications seem to me less bad than the worst implications of the best partially aggregative views.⁵

Proceedings of the Aristotelian Society 105 (2005): 191–217; Tyler Doggett, ‘Saving the Few’, *Noûs* 47 (2013): 302–315; and Kieran Setiya, ‘Love and the Value of a Life’, *Philosophical Review* 123 (2014): 251–280.

⁴ The most influential partially aggregative views are those defended by F. M. Kamm, T. M. Scanlon, and Alex Voorhoeve. For the view defended by Kamm, see her *Morality, Mortality: Death and Whom to Save From It* (Oxford: Oxford University Press, 1993), chapters 9 and 10; *Intricate Ethics: Rights, Responsibilities, and Permissible Harm* (Oxford: Oxford University Press, 2007): chapters 1 and 16; and ‘Nonconsequentialism’, in Hugh LaFollette (ed.), *The Blackwell Guide to Ethical Theory* (Oxford: Blackwell Publishing, 2013): 278–284. For the view defended by Scanlon, see his *What We Owe to Each Other*, Chapter 5. For the view defended by Voorhoeve, see his ‘How Should We Aggregate Competing Claims?’, *Ethics* 125 (2014): 64–87; and ‘Why One Should Count Only Claims with Which One Can Sympathize’, *Public Health Ethics* 10 (2017): 148–156.

⁵ Some philosophers argue that our anti-aggregative intuitions are systematically unreliable, because the cases used to prompt them involve very large numbers of people and we cannot accurately imagine very large

In Section 2, I focus on *axiological* partially aggregative views. These are views on which partially aggregative intuitions are justified by appeal to axiological claims—claims about what is good or what is bad. In Section 3, I focus on *purely deontic* partially aggregative views. These are views on which partially aggregative intuitions are justified without appeal to axiological claims. In Section 4, I briefly flag a further problem for partially aggregative views, which has to do with how these views handle risk.

2. Axiological Partially Aggregative Views

A partially aggregative view is axiological if it endorses the following claims:⁶

- (1) We should do what makes things go best, other things equal.
- (2) It is worse if a sufficiently large number of people suffer a moderate burden, such as paralysis, than if one person suffers a severe burden, such as death.

quantities. I am not sure how much weight to place on this argument. For versions of the argument, see Alastair Norcross, ‘Comparing Harms: Headaches and Human Lives’, *Philosophy & Public Affairs* 26 (1997): 135–167, 146–147; Derek Parfit, ‘Justifiability to Each Person’, *Ratio* 16 (2003): 368–390, 385–386; John Broome, *Weighing Lives* (Oxford: Oxford University Press, 2004): 56–57; Michael Huemer, ‘In Defence of Repugnance’, *Mind* 117 (2008): 899–933, 907–910; John Halstead, ‘The Numbers Always Count’, *Ethics* 126 (2016): 789–802, 797; and Joe Horton, ‘Aggregation, Complaints, and Risk’, *Philosophy & Public Affairs* 45 (2017): 54–81, 72–73. For responses, see Theron Pummer, ‘Intuitions about Large Number Cases’, *Analysis* 74 (2013): 37–46; Voorhoeve, ‘How Should We Aggregate Competing Claims?’, 75–76; and Victor Tadros, ‘Localized Restricted Aggregation’, *Oxford Studies in Political Philosophy* 5 (2019): 171–204, 175–176.

⁶ The arguments in this section are based on discussion of axiological partially aggregative views in Norcross, ‘Comparing Harms’; Erik Carlson, ‘Aggregating Harms—Should We Kill to Avoid Headaches?’, *Theoria* 66 (2000): 246–255; Michael Otsuka, ‘Saving Lives, Moral Theory, and the Claims of Individuals’, *Philosophy & Public Affairs* 34 (2006): 109–135, 126–128; Dale Dorsey, ‘Headaches, Lives, and Value’, *Utilitas* 21 (2009): 36–58; Larry Temkin, *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning* (Oxford: Oxford University Press, 2012); Gustaf Arrhenius and Wlodek Rabinowicz, ‘Value Superiority’, in Iwao Hirose and Jonas Olson (eds.), *The Oxford Handbook of Value Theory* (Oxford: Oxford University Press, 2015): 225–248; and Seth Lazar and Chad Lee-Stronach, ‘Axiological Absolutism and Risk’, *Noûs* 53 (2019): 97–113.

(3) It is worse if one person suffers a severe burden, such as death, than if any number of people suffer a minor burden, such as a migraine.

(2) and (3) are axiological claims, and (1)–(3) have implications that match partially aggregative intuitions. Since (1)–(3) are each plausible, they provide some justification for those intuitions.

It might seem that (2) and (3) are too similar to partially aggregative intuitions to provide them with much independent justification. But all justifications must end somewhere, and appealing to the plausibility of these claims might be as deep as we can go.

There is a more serious problem for axiological partially aggregative views. (3) can be challenged as follows:

If someone suffers a migraine, that is bad. If two people suffer migraines, that is twice as bad. More generally, the badness of n people suffering migraines is n times that of one person suffering a migraine. So, for it to be worse if one person dies than if any number of people suffer migraines, the badness of one person dying must be infinite. But the badness of one person dying is not infinite, for the badness of n people dying is n times that of one person dying.⁷ So, there is some number such that, if that number of people suffer migraines, that is worse than one person dying.

To defend (3) against this challenge, we could deny that the badness of n people suffering migraines is n times the badness of one person suffering a migraine. There are two main ways to do this.

⁷ Though infinity might come in different sizes, same-sized infinities do not sum to larger infinities. See Otsuka, ‘Saving Lives’, 127–128.

First, we could hold that the badness of n people suffering migraines is n times the badness of one person suffering a migraine only up to a certain value for n , and that once this limit is reached, adding another person suffering a migraine does not make things any worse. It would follow that there is an upper bound on the badness of people suffering migraines, and hence that one person dying could be worse than any number of people suffering migraines even though the badness of death is not infinite.

This position seems to me pretty implausible. It is difficult to believe that, once enough people are suffering migraines, it is not bad for others to suffer migraines. Furthermore, this position at least suggests that, once enough people are suffering migraines, there is no moral reason to prevent others from suffering migraines as well.

Second, we could hold that, at least once a certain number of people are suffering migraines, every additional person suffering a migraine adds less badness to the total, with the badness added by each additional person decreasing asymptotically towards zero. It would again follow that there is an upper bound on the badness of people suffering migraines, and hence that one person dying could be worse than any number of people suffering migraines even though the badness of death is not infinite.

This position avoids the troubling implications of the former position. But it has other, very similar implications. It implies that the badness of a person suffering a migraine depends on how many others are suffering migraines alongside them, and it at least suggests that, as the number of people suffering migraines increases, the strength of our moral reason to prevent others from suffering migraines decreases asymptotically towards zero.

There is another way to defend (3) against this challenge. We could hold that there are multiple scales of badness, with migraines ranked on one scale and deaths ranked on another, and with the latter scale 'trumping' the former. We could model this position using ordinal

numbers instead of cardinal numbers, associating one migraine with 1, two migraines with 2, and so on, and associating one death with $\omega + 1$, two deaths with $\omega + 2$, and so on.

An obvious problem with this position is that it seems as much in need of explanation and justification as the partially aggregative intuitions with which we began. A subtler problem is that, if we rank migraines and deaths on different scales, we cannot capture the intuition that the difference in badness between one migraine and two migraines is less than that between one death and two deaths. For the same reason, we cannot use expected value theory to infer what we should do in cases where we can either certainly prevent some number of migraines or take a chance of preventing some number of deaths.

Though none of the positions described above seems to me very attractive, others might be willing to embrace these positions to avoid the counterintuitive implications of fully aggregative and non-aggregative views. Erik Carlson, Seth Lazar, and Chad-Lee Stronach seem tempted by the position on which the badness of people suffering migraines decreases asymptotically towards zero, and James Griffin, Dale Dorsey, and Larry Temkin seem tempted by the position on which there are multiple scales of badness.⁸

Since the preceding arguments might not be decisive, it is worth discussing another challenge to (3):

Consider the following spectrum of possible outcomes.

- A single person dies.
- A large number of people suffer a burden fractionally less severe than death.

⁸ Carlson, 'Aggregating Harms', 248–249; Lazar and Lee-Stronach, 'Axiological Absolutism and Risk', 100; James Griffin, *Well-Being: Its Meaning, Measurement, and Moral Importance* (Oxford: Oxford University Press, 1989): Chapter 5; Dorsey, 'Headaches'; and Temkin, *Rethinking the Good*, 35–37 and 52–57.

➤ A much larger number of people than in the preceding outcome suffer a burden fractionally less severe than the burden in the preceding outcome.

➤ A much larger number of people than in the preceding outcome suffer a burden fractionally less severe than the burden in the preceding outcome.

...

➤ A much larger number of people than in the preceding outcome suffer migraines.

It seems clear that, for any two adjacent outcomes in this spectrum, given the choice between preventing the former and preventing the latter, we should prevent the latter.

An axiological partially aggregative view can yield this result only if, for any two adjacent outcomes, the latter is worse than the former. But if the second outcome is worse than the first, and the third is worse than the second, and so on, then, by the transitivity of ‘worse than’, the last is worse than the first. It follows that it is worse if a sufficiently large number of people suffer migraines than if one person dies.

There are three main ways to defend (3) against this challenge.

First, we could deny that it is possible for there to be a spectrum of possible outcomes like that on which the challenge depends. We could do this by holding that, if all the possible burdens that we could experience were arranged in a spectrum from most severe to least severe, there would be at least one pair of adjacent burdens that were sufficiently different in severity for it to be plausible that a sufficient number of people suffering the more severe burden would be worse than any number of people suffering the less severe burden.⁹

This position seems to me pretty implausible. It is difficult to believe that the spectrum of burdens we could experience would be as discontinuous as this position holds.

⁹ Carlson seems to endorse a version of this response in ‘Aggregating Harms’, 250–252.

Second, we could deny that ‘worse than’ is a transitive relation. This would allow us to accept that the second outcome in the spectrum is worse than the first, the third is worse than the second, and so on, without accepting that the last is worse than the first.

Temkin offers several arguments in support of this position, but even he holds back from endorsing it.¹⁰ If we say ‘B is worse than A’, that just seems to mean ‘B ranks lower on the scale of goodness than A’. If B ranks lower on a scale than A, and C ranks lower on this scale than B, C must rank lower on this scale than A. So, it seems to be a conceptual truth that ‘worse than’ is a transitive relation.

Third, we could reject the judgment that, for any two adjacent outcomes in the spectrum, given the choice between preventing the former and preventing the latter, we should prevent the latter. We would hold instead that, for at least one pair of adjacent outcomes, given the choice between preventing the former and preventing the latter, either we should prevent the former, or it is vague what we should do.¹¹

I think it would be implausible to insist that there is a pair of adjacent outcomes in the spectrum for which, given the choice between preventing the former and preventing the latter, we should prevent the former.¹² But it might seem more plausible to hold that, for at least one pair of adjacent outcomes, it is vague which we should prevent. This position might seem especially promising because the spectrum used to challenge (3) closely resembles the

¹⁰ Temkin, *Rethinking the Good*. See also Larry Temkin, ‘A Continuum Argument for Intransitivity’, *Philosophy & Public Affairs* 23 (1996): 175–210. For responses to Temkin, see Norcross, ‘Comparing Harms’; Broome, *Weighing Lives*, 50–63; Michael Huemer, ‘Transitivity, Comparative Value, and the Methods of Ethics’, *Ethics* 123 (2013): 318–345; and Alex Voorhoeve, ‘Vaulting Intuition: Temkin’s Critique of Transitivity’, *Economics and Philosophy* 29 (2013): 409–423.

¹¹ For a good discussion of vagueness, see Timothy Williamson, *Vagueness* (London: Routledge, 1994).

¹² For a defence of this position, see Dorsey, ‘Headaches’.

spectrums standardly used to illustrate the infamous Sorites Paradox, and it is widely accepted that vagueness will feature in the correct resolution of that paradox.¹³

Though this appeal to vagueness seems to me a more promising way to defend (3) against the spectrum challenge, the resulting position still seems to me unattractive. For any two adjacent outcomes in the spectrum, it seems to me clear, not vague, that we should prevent the former rather than the latter, given the choice.

But, again, I do not aim to defend any thesis. I hope instead to have given the reader a good sense of what axiological partially aggregative views are and what problems they face. By way of summary, axiological partially aggregative views must either deny that the badness of n people suffering a minor burden is n times that of one person suffering a minor burden, or hold that there are multiple scales of badness with some scales trumping others. They face significant problems either way. They also have counterintuitive implications when applied to the spectrum described above. Though these problems seem to me decisive against axiological partially aggregative views, others might be willing to live with them to avoid the counterintuitive implications of fully aggregative and non-aggregative views.

3. Purely Deontic Partially Aggregative Views

A partially aggregative view is purely deontic if it attempts to justify partially aggregative intuitions without appeal to axiological claims. Frances Kamm and T. M. Scanlon have both

¹³ An influential example of the Sorites Paradox posits a spectrum of men where the first has a very large amount of hair, every subsequent man has one hair less than the preceding man, and the last man has no hair. Given the very plausible premise that a single hair cannot make the difference between being bald and being not bald, it follows that either the first man in the spectrum is bald or the last man in the spectrum is not bald. For detailed comparison of the Sorites Paradox and spectrums like that used to challenge (3), see Temkin, *Rethinking the Good*, Chapter 9 and appendices C and D; Jacob Nebel, 'The Good, the Bad, and the Transitivity of *Better Than*', *Noûs* 52 (2018): 874–899; and Theron Pummer, 'Sorites on What Matters', in Tim Campbell, Jeff McMahan, and Ketan Ramakrishnan (eds.), *Festschrift in Honor of Derek Parfit* (Oxford: Oxford University Press, forthcoming).

proposed influential purely deontic partially aggregative views, but the most thorough elaboration and defence of such a view comes from Alex Voorhoeve.¹⁴

Voorhoeve calls his view *Aggregate Relevant Claims* (ARC). According to ARC, at least when other things are equal, we should satisfy the greatest sum of strength-weighted, *relevant* claims. A person has a claim to our acting in some way if and only if it would benefit her, and the strength of her claim increases both the more she stands to gain and the worse off she will be if her claim is not satisfied. A claim is relevant if and only if it is sufficiently strong relative to the strongest claim with which it competes.

ARC captures partially aggregative intuitions as follows. Suppose you can save either one person from death or a huge number of people from paralysis. Each of these people has a claim to your help, and since a claim to be saved from paralysis is very close in strength to a claim to be saved from death, all these claims are relevant. So, on the plausible assumption that the strength-weighted sum of a huge number of claims to be saved from paralysis is greater than the strength-weighted sum of one claim to be saved from death, ARC implies you should save the people facing paralysis. Suppose next that you can save either one person from death or a huge number of people from a migraine. Each of these people has a claim to your help, but since a claim to be saved from a migraine is very weak relative to a claim to be saved from death, the claims to be saved from a migraine are not relevant. So, ARC implies you should save the person facing death, no matter how many people face migraines.

Why does a claim become irrelevant when it competes with a much stronger claim? Voorhoeve suggests the following answer.¹⁵ Suppose A must choose between saving herself from a minor burden, such as a migraine, and saving B from a severe burden, such as death. Plausibly, A should save B. So, if you must choose between saving A from a minor burden

¹⁴ See the references in footnote 4.

¹⁵ A similar answer is presented by Kamm. See the references in footnote 4.

and saving B from a severe burden, A should withdraw her claim to your help. So, if you must choose between saving many people from a minor burden and saving one person from a severe burden, each of the people facing the minor burden should withdraw their claim to your help. So, a claim is irrelevant when it competes with a much stronger claim.

This rationale for ARC is attractive, but it is importantly incomplete. Though it seems clear that A would be obliged to save B from a severe burden rather than saving herself from a minor burden, it does not necessarily follow that A would be obliged to save B from a severe burden rather than saving herself *and a large number of others* from a minor burden.¹⁶ Furthermore, though the latter claim might be intuitive, it cannot be assumed in a rationale for ARC, for it is too similar to the intuitions that the rationale is meant to explain.

Voorhoeve is aware of this problem with the rationale, and of another problem that we will come to shortly. He therefore modifies and extends the rationale.¹⁷ I do not have space to discuss his extended rationale here, but it pays to read his article with these problems in mind.

Return now to the spectrum used to challenge axiological partially aggregative views in the previous section. ARC has very plausible implications when applied to this spectrum. Since the burdens in any two adjacent outcomes are very close in strength, ARC implies that, for any two adjacent outcomes, given the choice between preventing the former and preventing the latter, we should prevent the latter. But because the burdens in the first and last outcomes are far apart in strength, ARC also implies that, given the choice between preventing the former and preventing the latter, we should prevent the former.

¹⁶ This problem is developed in greater detail in Halstead, ‘The Numbers Always Count’; and Tadros, ‘Localized Restricted Aggregation’, 176–178.

¹⁷ Voorhoeve, ‘How Should We Aggregate Competing Claims?’, 68–75. John Halstead argues that the extended rationale does not succeed in overcoming the first problem. See his ‘The Numbers Always Count’. For an alternative rationale for a view like ARC, see Tadros, ‘Localized Restricted Aggregation’, 172–178.

Because ARC has these implications, it implies that the relation ‘ought to be done rather than’ is intransitive. This might seem to be a problem, for it might seem that ‘A ought to be done rather than B’ just means ‘A ranks higher on the scale of permissibility than B’. However, ‘A ought to be done rather than B’ plausibly instead means ‘A ranks higher on the scale of permissibility than B *given the options that are currently salient*’. If that is the correct interpretation, ‘ought to be done rather than’ need not be transitive.

If we accept ARC, what should we think about the axiological relationship between the outcomes in the spectrum, and in particular about the axiological relationship between the first outcome and the last? There seem to me three main options.

First, we could deny that the last outcome in the spectrum is better than the first by holding that, for at least one pair of adjacent outcomes, it is vague whether the latter is worse than the former. This option comes with familiar baggage, however, for it means either denying that the badness of n people suffering minor burdens is n times that of one person suffering a minor burden, or holding that there are multiple scales of badness.

Second, we could accept that every outcome in the spectrum is worse than the immediately preceding outcome, accept that ‘worse than’ is transitive, and hence accept that the last outcome is worse than the first. It might seem strange to accept that the last outcome is worse than the first while also holding that we should save one person from death rather than saving any number of people from migraines. However, many philosophers already accept that we should sometimes bring about the worse of two available outcomes.

Third, we could follow philosophers such as Peter Geach, Philippa Foot, and Judith Jarvis Thomson in holding that, though outcomes can be good or bad *for particular people*, they cannot be good or bad *simpliciter*.¹⁸ We would then hold that none of the outcomes in

¹⁸ Peter Geach, ‘Good and Evil’, *Analysis* 17 (1956): 33–42; Philippa Foot, ‘Utilitarianism and the Virtues’, *Mind* 94 (1985): 196–209; and Judith Jarvis Thomson, *Normativity* (Chicago: Open Court, 2008).

the spectrum are better than, worse than, or even equally as good as any of the other outcomes.

We can now turn to consider some problems for ARC. We have to build up to these problems gradually, so bear with me...

Recall that, according to ARC, a claim is relevant if and only if it is sufficiently strong relative to the strongest claim with which it competes. Let a *minor claim* be a claim to be saved from a minor burden. Let a *moderate claim* be the strongest claim such that, when it competes with a minor claim, the minor claim is relevant. Let a *strong claim* be the strongest claim such that, when it competes with a moderate claim, the moderate claim is relevant.

Note that, given these definitions, a minor claim is not relevant when it competes with a strong claim. So, while there is some number N of moderate claims that we should satisfy rather satisfying one strong claim, and while there is some number $N+$ of minor claims that we should satisfy rather than satisfying N moderate claims, there is no number of minor claims that we should satisfy rather than satisfying one strong claim.

Suppose you can save either one person with a strong claim, N people with a moderate claim, or $N+$ people with a minor claim. ARC implies you should save the N people with moderate claims, for the moderate claims sum to more than the strong claim, and the minor claims are made irrelevant by being in competition with the strong claim.

We can now outline two problems for ARC.

First, the implication that ARC has in this three-option case is not supported by the simple rationale for ARC outlined earlier. Though that rationale might explain why the minor claims should be withdrawn *relative to the strong claims*, it does not explain why they should

be withdrawn *simpliciter*. For the minor claims are sufficiently strong relative to, and together sum to more than, the moderate claims.¹⁹

If the minor claims should not be withdrawn simpliciter, and should instead be withdrawn only relative to the strong claims, it seems you should satisfy the moderate claims over the strong claim, satisfy the minor claims over the moderate claims, and satisfy the strong claim over the minor claims. We then have a *deontic cycle*—a situation in which you should do B rather than A, do C rather than B, and yet do A rather than C. It would follow that, in the three-option case, you cannot avoid acting wrongly, for whichever option you choose, there is another option that you should have chosen instead. This result seems implausible.²⁰

As mentioned earlier, Voorhoeve is aware of this further problem with the simple rationale for ARC. His modified and extended rationale is meant to overcome it.²¹ Again, it pays to read his article with this problem in mind.

Second, because ARC avoids deontic cycling in the three-option case and instead implies that you should satisfy the moderate claims, it violates a principle known as *the independence of irrelevant alternatives* (also known as *basic contraction consistency*). According to this principle, if it is permissible to choose an option O from an option set S, it is permissible to choose O from any subset of S containing O. ARC violates this principle by implying that it is permissible for you to satisfy the moderate claims given the option set {one strong claim, N moderate claims, N+ minor claims}, but wrong for you to satisfy the moderate claims given the subset {N moderate claims, N+ minor claims}.²²

¹⁹ Johanna Privitera develops this problem in greater detail in ‘Aggregate Relevant Claims in Rescue Cases?’, *Utilitas* 30 (2017): 228–236.

²⁰ Parfit presents the same problem for the partially aggregative view defended by Scanlon. See Parfit, ‘Justifiability to Each Person’, 384.

²¹ Privitera argues that it does not succeed. See her ‘Aggregate Relevant Claims in Rescue Cases?’.

²² Kamm also discusses this problem with respect to her view. See her *Intricate Ethics*, 484–486.

Voorhoeve denies that ARC violates the independence of irrelevant alternatives in this kind of case. He takes it to be an implication of ARC that the option of *satisfying N moderate claims when you could instead satisfy either one strong claim or N+ minor claims* is different from the option of *satisfying N moderate claims when your only alternative is satisfying N+ minor claims*. It follows that the option set {N moderate claims, N+ minor claims} is not a subset of {one strong claim, N moderate claims, N+ minor claims}.

This response does not seem to me compelling. If we generally individuate options according to the option sets of which they are a part, it becomes impossible for one option set to be a subset of another, so the independence of irrelevant alternatives becomes trivial.²³ Furthermore, as John Halstead argues, what ultimately matters is not whether ARC violates the independence of irrelevant alternatives, but rather whether it has the kind of counterintuitive implications that are generally associated with violations of this principle.²⁴ We can show that it does by dramatizing the implications of ARC as follows:

Doctor 1: “There is one person with a strong claim in ward A, N people with moderate claims in ward B, and N+ people with minor claims in ward C. We have only enough time to get to one ward, so we should head to ward B.”

Doctor 2: “Actually, there is not enough time to get to ward A anyway.”

Doctor 1: “In that case, we should head to ward C.”

²³ For discussion of this point, see Tina Rulli and Alex Worsnip, ‘IIA, Rationality, and the Individuation of Options’, *Philosophical Studies* 173 (2016): 205–221.

²⁴ Halstead, ‘The Numbers Always Count’.

ARC implies that Doctor 1 responds appropriately, but her response seems very odd.

Though this implication of ARC seems to me very counterintuitive, Voorhoeve and a significant number of other philosophers are willing to embrace it.²⁵ They argue that the counterintuitiveness of this implication is significantly mitigated when the rationale for ARC is kept in mind, and also that accepting this implication is a price worth paying to avoid the counterintuitive implications of fully aggregative and non-aggregative views.

Since the preceding problems with ARC may not be decisive, it is worth discussing a further problem. This problem was originally noticed by Patrick Tomlin.²⁶ Roughly, the problem is that ARC has very counterintuitive implications in cases in which there are additions or subtractions to the groups of people we can save.

Let $N\#$ be the number of moderate claims that is needed to exactly balance (or as nearly balance as possible) one strong claim. Now, consider the following, two-stage case.

Stage 1: You can save either group A, which contains one person with a strong claim, or group B, which contains $N\#$ people with moderate claims.

Stage 2: A single person with a minor claim is added to group A, and a gazillion people with minor claims are added to group B.

ARC implies that it is permissible to save either group at Stage 1. If it is permissible to save either group at Stage 1, it seems it must be at least permissible to save B at Stage 2. But ARC

²⁵ See Voorhoeve, 'How Should We Aggregate Competing Claims?', 76–79; Voorhoeve 'Why One Should Count Only Claims with Which One Can Sympathize', 152–153; Kamm, *Intricate Ethics*, 484–486; Patrick Tomlin, 'On Limited Aggregation', *Philosophy & Public Affairs*: 45 (2017): 232–260, 236–237; Tadros, 'Localized Restricted Aggregation', 185–191; and Seth Lazar, 'Limited Aggregation and Risk', *Philosophy & Public Affairs* 46 (2018): 117–159, 126–130.

²⁶ Tomlin, 'On Limited Aggregation'.

implies otherwise. Since the minor claims added to B compete with a strong claim, they are not relevant. But the one minor claim added to A is relevant, for the strongest claims with which it competes are moderate. This claim must tip the balance. So, ARC implies that, though it is permissible to save either group at Stage 1, you should save A at Stage 2.

It seems to me that, if the rationale for ARC were really sufficient mitigation for the counterintuitive implications associated with violations of the independence of irrelevant alternatives, it would also be sufficient mitigation for the counterintuitive implication identified by Tomlin.²⁷ But, surprisingly, proponents of partially aggregative views have been much less willing to embrace the latter implication. They have instead attempted to develop alternatives to ARC that avoid that implication. However, as they acknowledge, these alternatives have other implications that are not obviously less counterintuitive.²⁸

We have been focusing on ARC because it is the most developed purely deontic partially aggregative view in the literature. But the problems faced by ARC will be faced by other purely deontic partially aggregative views, and by any partially aggregative view that has plausible implications when applied to the spectrum used to challenge axiological partially aggregative views. We are now in a position to understand why.

Let PA be any partially aggregative view that has plausible implications when applied to the spectrum used to challenge axiological partially aggregative views. PA must imply that, given the choice between preventing the first outcome in the spectrum and preventing the last, we should prevent the former. Presumably, it also implies that, given the choice, we

²⁷ Tomlin argues that the rationale for ARC does not explain this implication in a way that mitigates its counterintuitiveness. See his ‘On Limited Aggregation’, 242–244. I think he is mistaken about that. See my ‘Always Aggregate’, *Philosophy & Public Affairs* 46 (2018): 160–174, 172–173 and footnote 19.

²⁸ See Tadros, ‘Localized Restricted Aggregation’; Aart van Gils and Patrick Tomlin, ‘Relevance Rides Again? Aggregation and Local Relevance’, *Oxford Studies in Political Philosophy* 6 (2020): 221–256; Korbinian Rürger, ‘Aggregation with Constraints’, *Utilitas* (forthcoming); and Bastian Steuwer, ‘Aggregation, Balancing, and Respect for the Claims of Individuals’, *Utilitas* (forthcoming). See also Horton, ‘Always Aggregate’.

should prevent the second outcome rather than the last, and the third outcome rather than the last. But, as we move along the spectrum, we must eventually come to the first outcome that is such that, given the choice between preventing this outcome and preventing the last, we should prevent the latter. Let us call this *the middle outcome*.

PA must also imply that there is least one outcome preceding the middle outcome that is such that, given the choice between preventing this outcome and preventing the middle outcome, we should prevent the latter. Presumably, it implies that there are many such outcomes. Let the first outcome with this property be *the upper outcome*.

It follows from these definitions that, though we should prevent the last outcome rather than the middle outcome, and though we should prevent the middle outcome rather than the upper outcome, we should prevent the upper outcome rather than the last outcome.

Which outcome should we prevent given a choice between the upper outcome, the middle outcome, and the last outcome? PA now faces a dilemma: either it implies that none of these options is permissible, because it generates a deontic cycle, or it violates the independence of irrelevant alternatives. If it violates the independence of irrelevant alternatives, it is likely to face versions of the problem that Tomlin identified for ARC.

By way of summary, though purely deontic partially aggregative views can have plausible implications when applied to the spectrum used to challenge axiological partially aggregative views, those that do face other important problems. They either generate deontic cycles or violate the independence of irrelevant alternatives. If they violate the independence of irrelevant alternatives, they are likely to face versions of the problem that Tomlin identified for ARC. These problems seem to me too high a price to pay for avoiding the counterintuitive implications of fully aggregative views, but others might disagree.

4. Partially Aggregative Views and Risk

There is a further problem that seems to afflict all partially aggregative views, whether axiological or purely deontic, and which also afflicts non-aggregative views. It can be shown that, in cases involving risk, these views either have very counterintuitive implications, have implications that are in tension with the intuitions that motivate rejection of fully aggregative views, or draw sharp distinctions between cases in which we face single decisions and seemingly morally equivalent cases in which we face repeated decisions. I do not have space to discuss this trilemma here, but I have discussed it elsewhere, and the literature on it has been growing quickly.²⁹ I hope to have left the reader feeling ready to dive in.

Acknowledgments

For helpful comments and discussion, I am grateful to Alexander Dietz, Kerah Gordon-Solmon, Nathan Howard, Nicholas Laskowski, Véronique Munoz-Dardé, Abelard Podgorski, Theron Pummer, Anthony Taylor, Patrick Tomlin, and Alec Walen.

²⁹ I develop this trilemma in ‘Aggregation, Risk, and Reductio’, *Ethics* 130 (2020): 514–529. For related challenges to partially aggregative views, see Norcross, ‘Comparing Harms’; Sophia Reibetanz, ‘Contractualism and Aggregation’, *Ethics* 108 (1998): 296–311; Alastair Norcross, ‘Great Harms from Small Benefits Grow: How Death Can Be Outweighed by Headaches’, *Analysis* 58 (1998): 152–158; Tom Dougherty, ‘Aggregation, Beneficence, and Chance’, *Journal of Ethics and Social Philosophy* 7 (2013): 1–19; and Horton, ‘Aggregation, Complaints, and Risk’. For responses to these challenges, see Michael Otsuka, ‘Risking Life and Limb’, in I. Glenn Cohen, Norman Daniels, and Nir Eyal (eds.), *Identified Versus Statistical Lives* (Oxford: Oxford University Press, 2015): 77–93; Johann Frick, ‘Contractualism and Social Risk’, *Philosophy & Public Affairs* 43 (2015): 175–223; Lazar, ‘Limited Aggregation and Risk’; Korbinian Rürger, ‘On *Ex Ante* Contractualism’, *Journal of Ethics and Social Philosophy* 13 (2018): 240–258; Lazar and Lee-Stronach, ‘Axiological Absolutism and Risk’; Tadros, ‘Localized Restricted Aggregation’; Kerah Gordon-Solmon, ‘Should Contractualists Decompose?’, *Philosophy & Public Affairs* 47 (2019): 259–287; and Alec Walen, ‘Risks and Weak Aggregation: Why Different Models of Risk Suit Different Types of Cases’, *Ethics* 131 (2020): 62–86.

Bibliography

- Anscombe, Elizabeth. 'Who is Wronged? Philippa Foot on Double Effect: One Point', *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: Oxford University Press, 2015): 225–248
- Broome, John. *Weighing Lives* (Oxford: Oxford University Press, 2004)
- Carlson, Erik. 'Aggregating Harms—Should We Kill to Avoid Headaches?', *Theoria* 66 (2000): 246–255
- Doggett, Tyler. 'Saving the Few', *Noûs* 47 (2013): 302–315
- Dorsey, Dale. 'Headaches, Lives, and Value', *Utilitas* 21 (2009): 36–58
- Dougherty, Tom. 'Aggregation, Beneficence, and Chance', *Journal of Ethics and Social Philosophy* 7 (2013): 1–19
- Foot, Philippa. 'Utilitarianism and the Virtues', *Mind* 94 (1985): 196–209
- Frick, Johann. 'Contractualism and Social Risk', *Philosophy & Public Affairs* 43 (2015): 175–223
- Geach, Peter. 'Good and Evil', *Analysis* 17 (1956): 33–42
- van Gils, Aart and Patrick Tomlin, 'Relevance Rides Again? Aggregation and Local Relevance', *Oxford Studies in Political Philosophy* 6 (2020): 221–256
- Gordon-Solmon, Kerah. 'Should Contractualists Decompose?', *Philosophy & Public Affairs* 47 (2019): 259–287
- James Griffin, *Well-Being: Its Meaning, Measurement, and Moral Importance* (Oxford: 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
- 'Aggregation, Risk, and Reductio', *Ethics* 130 (2020): 514–529
- 'Always Aggregate', *Philosophy & Public Affairs* 46 (2018): 160–174
- Huemer, Michael. 'In Defence of Repugnance', *Mind* 117 (2008): 899–933
- 'Transitivity, Comparative Value, and the Methods of Ethics', *Ethics* 123 (2013): 318–345
- Kagan, Shelly. *Normative Ethics* (London: Routledge, 2018)
- Kamm, F. M. *Intricate Ethics: Rights, Responsibilities, and Permissible Harm* (Oxford: Oxford University Press, 2007)
- Morality, Mortality: Death and Whom to Save From It* (Oxford: Oxford University Press, 1993)
- 'Nonconsequentialism', in Hugh LaFollette (ed.), *The Blackwell Guide to Ethical Theory* (Oxford: Blackwell Publishing, 2013): 261–286
- Lazar, Seth. 'Limited Aggregation and Risk', *Philosophy & Public Affairs* 46 (2018): 117–159
- Lazar, Seth and Lee-Stronach, Chad. 'Axiological Absolutism and Risk', *Noûs* 53 (2019): 97–113
- Munoz-Dardé, Véronique. 'The Distribution of Numbers and the Comprehensiveness of Reasons', *Proceedings of the Aristotelian Society* 105 (2005): 191–217
- Nebel, Jacob. 'The Good, the Bad, and the Transitivity of *Better Than*', *Noûs* 52 (2018): 874–899
- Norcross, Alastair. 'Comparing Harms: Headaches and Human Lives', *Philosophy & Public Affairs* 26 (1997): 135–167

- ‘Great Harms from Small Benefits Grow: How Death Can Be Outweighed by Headaches’, *Analysis* 58 (1998): 152–158
- Otsuka, Michael. ‘Risking Life and Limb’, in I. Glenn Cohen, Norman Daniels, and Nir Eyal (eds.), *Identified Versus Statistical Lives* (Oxford: Oxford University Press, 2015): 77–93
- ‘Saving Lives, Moral Theory, and the Claims of Individuals’, *Philosophy & Public Affairs* 34 (2006): 109–135
- Parfit, Derek. ‘Equality and Priority’, *Ratio* 10 (1997): 202–221
- ‘Justifiability to Each Person’, *Ratio* 16 (2003): 368–390
- Privitera, Johanna. ‘Aggregate Relevant Claims in Rescue Cases?’, *Utilitas* 30 (2017): 228–236
- Pummer, Theron. ‘Intuitions about Large Number Cases’, *Analysis* 74 (2013): 37–46
- ‘Sorites on What Matters’, in Tim Campbell, Jeff McMahan, and Ketan Ramakrishnan (eds.), *Festschrift in Honor of Derek Parfit* (Oxford: Oxford University Press, forthcoming)
- Reibetanz, Sophia. ‘Contractualism and Aggregation’, *Ethics* 108 (1998): 296–311
- Rüger, Korbinian. ‘Aggregation with Constraints’, *Utilitas* 32 (2020):
- ‘On *Ex Ante* Contractualism’, *Journal of Ethics and Social Philosophy* 13 (2018): 240–258
- Rulli, Tina and Alex Worsnip, ‘IIA, Rationality, and the Individuation of Options’, *Philosophical Studies* 173 (2016): 205–221
- Scanlon, T. M. *What We Owe to Each Other* (Cambridge, MA: Harvard University Press, 1998)
- Setiya, Kieran. ‘Love and the Value of a Life’, *Philosophical Review* 123 (2014): 251–280
- Steuwer, Bastian. ‘Aggregation, Balancing, and Respect for the Claims of Individuals’, *Utilitas* (forthcoming)
- Tadros, Victor. ‘Localised Restricted Aggregation’, *Oxford Studies in Political Philosophy* 5 (2019): 171–204
- Taurek, John M. ‘Should the Numbers Count?’, *Philosophy & Public Affairs* 6 (1977): 293–316
- Temkin, Larry. ‘A Continuum Argument for Intransitivity’, *Philosophy & Public Affairs* 23 (1996): 175–210
- Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning* (Oxford: Oxford University Press, 2012)
- Thomson, Judith Jarvis. *Normativity* (Chicago: Open Court, 2008)
- 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
- ‘Vaulting Intuition: Temkin’s Critique of Transitivity’, *Economics and Philosophy* 29 (2013): 409–423
- ‘Why One Should Count Only Claims with Which One Can Sympathize’, *Public Health Ethics* 10 (2017): 148–156
- Walen, Alec. ‘Risks and Weak Aggregation: Why Different Models of Risk Suit Different Types Cases’, *Ethics* 131 (2020): 62–86
- Williamson, Timothy. *Vagueness* (London: Routledge, 1994)