Aggregation, Risk, and Reductio*

Joe Horton

Is there any number of people you should save from paralysis rather than saving one person from death? Is there any number of people you should save from a migraine rather than saving one person from death? Many people answer "yes" and "no," respectively. The aim of *partially aggregative* moral views is to capture and justify combinations of intuitions like these. In this article, I develop a risk-based reductio argument that shows that there can be no adequate partially aggregative view. I then argue that the only plausible response to this reductio is to accept a *fully aggregative* view.

I. INTRODUCTION

Is there any number of people you should save from a substantial burden, such as paralysis, rather than saving one person from a severe burden, such as death? Is there any number of people you should save from a minor burden, such as a migraine, rather than saving one person from a severe burden, such as death? Many people answer "yes" and "no," respectively. The aim of *partially aggregative* moral views is to capture and justify combinations of intuitions like these.[1] These views contrast with *fully aggregative* moral views, which imply that the answer to both questions is "yes," and

* For helpful comments and discussion, I am grateful to Nilanjan Das, Alexander Dietz, Jessica Fischer, Nathan Howard, Nicholas Laskowski, Véronique Munoz-Dardé, Abelard Podgorski, Anthony Taylor, Han van Wietmarschen, three anonymous editors at this journal, and an audience at Cardiff University.

1. For influential partially aggregative views, see F. M. Kamm, "Nonconsequentialism," in *The Blackwell Guide to Ethical Theory*, ed. Hugh LaFollette (Oxford: Blackwell, 2013), 278–84; and Alex Voorhoeve, "How Should We Aggregate Competing Claims?," *Ethics* 125 (2014): 64–87. For criticisms of these views, see Derek Parfit, "Justifiability to Each Person," *Ratio* 16 (2003): 368–90, 384–85; John Halstead, "The Numbers Always Count," *Ethics* 126 (2016): 789–802, 797–99; Joe Horton, "Aggregation, Complaints, and Risk," *Philosophy and Public Affairs* 45 (2017): 54–81; Patrick Tomlin, "On Limited Aggregation," *Philosophy and Public Affairs* 45 (2017): 232–69; and Joe Horton, "Always Aggregate," *Philosophy and Public Affairs* 46 (2018): 160–74. For responses to these criticisms, see F. M. Kamm, *Intricate Ethics:*

Ethics 130 (July 2020): 514–529

with *nonaggregative* moral views, which imply that the answer to both questions is "no."[2]

Though partially aggregative views have much intuitive appeal, I believe they should be rejected in favor of fully aggregative views. I here explain why. In Section II, I develop a risk-based reductio argument that shows that there can be no adequate partially aggregative view. I then argue that the only plausible response to this reductio is to accept a fully aggregative view. In Section III, I apply this reductio to a range of partially aggregative views from the literature, including some that were specifically designed with risk in mind.

Since there has been much discussion in the literature of partially aggregative views and risk, it is worth highlighting what is new about my reductio.[3] Though it has been shown that influential partially aggregative views have problematic implications in certain cases involving risk, it has not been shown that risk is a decisive problem for all partially aggregative views. Furthermore, several philosophers have recently argued that more sophisticated partially aggregative views can avoid problems with risk.[4] If my reductio is sound, it shows that even these more sophisticated views should be rejected.

Rights, Responsibilities, and Permissible Harm (Oxford: Oxford University Press, 2007), 484–86; Voorhoeve, "How Should We Aggregate"; Alex Voorhoeve, "Why One Should Count Only Claims with Which One Can Sympathise," *Public Health Ethics* 10 (2017): 148–56; Victor Tadros, "Localized Restricted Aggregation," *Oxford Studies in Political Philosophy* 5 (2019): 171–204; and Aart Van Gils and Patrick Tomlin, "Relevance Rides Again? Aggregation and Local Relevance," *Oxford Studies in Political Philosophy* 6 (2020): 221–56.

2. For discussion of fully aggregative views, see Samuel Scheffler, *The Rejection of Consequentialism* (Oxford: Oxford University Press, 1994); Alastair Norcross, "Comparing Harms: Headaches and Human Lives," *Philosophy and Public Affairs* 26 (1997): 135–67; and Parfit, "Justifiability to Each Person." For discussion of nonaggregative views, see Elizabeth Anscombe, "Who Is Wronged? Philippa Foot on Double Effect: One Point," *Oxford Review* 5 (1967): 16–17; John M. Taurek, "Should the Numbers Count?," *Philosophy and Public Affairs* 6 (1977): 293–316; T. M. Scanlon, *What We Owe to Each Other* (Cambridge, MA: Harvard University Press, 1998), chap. 5; Véronique Munoz-Dardé, "The Distribution of Numbers and the Comprehensiveness of Reasons," *Proceedings of the Aristotelian Society* 105 (2005): 191–217; Tyler Doggett, "Saving the Few," *Noûs* 47 (2013): 302–15; and Kieran Setiya, "Love and the Value of a Life," *Philosophical Review* 123 (2014): 251–80.

3. See Sophia Reibetanz, "Contractualism and Aggregation," *Ethics* 108 (1998): 296–311; Michael Otsuka, "Risking Life and Limb," in *Identified versus Statistical Lives*, ed. I. Glenn Cohen, Norman Daniels, and Nir Eyal (Oxford: Oxford University Press, 2015), 77–93; Johann Frick, "Contractualism and Social Risk," *Philosophy and Public Affairs* 43 (2015): 175–223; Horton, "Aggregation, Complaints, and Risk"; Seth Lazar, "Limited Aggregation and Risk," *Philosophy and Public Affairs* 46 (2018): 117–59; Tadros, "Localized Restricted Aggregation," 192–202; Seth Lazar and Chad Lee-Stronach, "Axiological Absolutism and Risk," *Noûs* 53 (2019): 97–113; and Kerah Gordon-Solmon, "Should Contractualists Decompose?," *Philosophy and Public Affairs* 47 (2019): 259–87.

4. Lazar, "Limited Aggregation and Risk"; Tadros, "Localized Restricted Aggregation," 192–202; Lazar and Lee-Stronach, "Axiological Absolutism and Risk"; and Gordon-Solmon, "Should Contractualists Decompose?"

II. THE REDUCTIO

In this section, I develop a risk-based reductio argument that shows that there can be no adequate partially aggregative view. I then respond to a challenge to the reductio, and I argue that the only plausible response to the reductio is to accept a fully aggregative view. Though the reductio is new in applying to all possible partially aggregative views, it draws heavily on some of my earlier work, and it builds on work by Alastair Norcross and Tom Dougherty.[5]

A. The Reductio

Suppose for reductio that PA is an adequate partially aggregative view, meaning it avoids implausible implications in individual cases, avoids implications that are in tension with the intuitions that incline people toward partially aggregative views, and avoids inconsistent implications across cases that are in all morally relevant respects equivalent.

Suppose that in

> *Villain 1*: A villain has kidnapped A and B. He will either (1) inflict a migraine on A, or (2) inflict a one-in-a-zillion chance of death on B. You must choose which.[6]

PA will imply that you should (or at least that it is permissible to) choose (2). We frequently impose tiny chances of death on some people as a side effect of sparing others from minor burdens, and this behavior seems clearly permissible. For example, if someone is suffering a migraine, it is

5. See Horton, "Aggregation, Complaints, and Risk"; Alastair Norcross, "Great Harms from Small Benefits Grow: How Death Can Be Outweighed by Headaches," *Analysis* 58 (1998): 152–58; and Tom Dougherty, "Aggregation, Beneficence, and Chance," *Journal of Ethics and Social Philosophy* 7 (2013): 1–19. There are two important differences between the argument presented here and the argument presented in my earlier article. First, the argument presented in my earlier article distinguishes between "ex ante" and "ex post" partially aggregative views and raises objections to both, but it merely suggests, and does not show, that all hybrid partially aggregative views face similar objections. Second, the argument presented here uses cases that differ from the corresponding cases used in my earlier article, in that the people you can expose to lesser burdens are different from the people you can expose to tiny chances of death. Because the cases used in my earlier article stipulate that these people are the same, the argument in that article fails to apply to all possible partially aggregative views. For example, it fails to demonstrate any problem for the partially aggregative view subsequently developed in Lazar, "Limited Aggregation and Risk." This deficiency is shared by the argument presented by Dougherty, while the argument presented by Norcross applies only to partially aggregative views on which what you ought to do is determined by the value of the outcomes you can bring about. I also here respond to three challenges (see Secs. II.*B*, II.*C*, and II.*E*) that are not considered in my earlier article.

6. Why must you choose? Why not keep your hands clean? If you refuse to choose, the villain will destroy the world.

permissible to drive them to a pharmacy, even though this exposes pedestrians—even several pedestrians—to a tiny chance of death.[7]

Suppose next that in

> *Villain 2*: A villain has kidnapped ten zillion X people and ten zillion Y people. He will either (1) inflict a migraine on each X person, or (2) randomly select and kill ten Y people. You must choose which.

PA will imply that you should choose (1). It would be bizarre to hold that, in cases involving certainty, there is no number of people you should save from migraines rather than saving one person from death, and yet also hold that, in some cases involving risk, you should save a huge number of people from migraines rather than saving ten people from death.[8]

Suppose finally that in

> *Villain 3*: A villain has kidnapped ten zillion X people and ten zillion Y people. He pairs each X person with a Y person. For each pair, the villain will either (1) inflict a migraine on the X person, or (2) give the Y person a ticket for a lottery with ten zillion tickets. You must choose between these options for each pair in turn. You know that, after you have chosen for each pair, the villain will randomly select ten tickets and kill anyone who has a corresponding ticket.

Here, you face a series of choices where each choice is exactly like the choice in Villain 1. Since PA will imply that you should (or at least that it is permissible to) choose (2) in Villain 1, it will imply that you should (or at least that it is permissible to) choose (2) for each pair in Villain 3. But choosing (2) for each pair in Villain 3 is choosing that the villain randomly select and kill ten Y people rather than inflicting a migraine on each X person. And that is the same choice PA will condemn in Villain 2. So, PA will be inconsistent, in the sense that it will have different implications across cases that are in all morally relevant respects equivalent.

It follows that, if PA avoids implausible implications in individual cases (Villain 1), and avoids implications that are in tension with the intuitions that incline people toward partially aggregative views (Villain 2), it has inconsistent implications across cases that are in all morally relevant respects equivalent (Villain 2 and Villain 3). This contradicts our definition of PA. So, there can be no adequate partially aggregative view.

B. Sequences of Acts

The final stage of this reductio might be challenged. I claimed that, because PA will imply that you should choose (2) in Villain 1, it will imply

7. If you remain skeptical that you should choose (2) in Villain 1, hold fire until Sec. II.*E*.

8. Lazar might deny that this position is bizarre. See Lazar, "Limited Aggregation and Risk," 132, 157. I provide another reason to reject this position below, in Sec. III.*C*.

that you should choose (2) for each pair in Villain 3. But, as some proponents of partially aggregative views have suggested, PA could be a view that applies not to individual acts, but rather to sequences of acts.[9] In Villain 3, if you choose (2) for each pair, you perform a sequence of acts that you know will result in ten deaths. If PA applies to sequences of acts, it could condemn this sequence.

Suppose PA does apply to sequences of acts, and that it forbids you from choosing (2) for each pair in Villain 3. What will it imply you should do in this case? There are two possibilities. PA will imply either that you should choose (1) for every pair, or that you should choose (2) for some number of pairs and then choose (1) for the others.

If PA is applied in this way, it has other implausible implications. Suppose that in

> *Long Life*: You will live an extremely long time—zillions of years. As you look ahead at your long life, you know there will be frequent opportunities to spare some people from minor burdens, or give them minor benefits, by acting in ways that expose others to tiny chances of death. Given the extreme length of your life, it is a statistical certainty that, if you take these opportunities, eventually you will kill someone.

If PA applies to sequences of acts, it will forbid you from taking these risky opportunities. It will imply either that you should never take these opportunities, or that you should take some and then refuse to take any more. Both implications are implausible.

It might be objected that Long Life differs from Villain 3 in two respects to which PA could be sensitive. First, as you give out the lottery tickets, the chance of your having caused a death increases steadily until it becomes certain, whereas each time you take one of the risky opportunities, the chance of your having killed someone remains extremely low. Second, if you give out all the lottery tickets, you are certain to cause a death, whereas if you always take the risky opportunities, you are only statistically certain to cause a death, for there is an infinitesimally small chance that you will never kill anyone.

These distinguishing features of Villain 3 are inessential to the reductio. We can replace Villain 2 and Villain 3 with the following cases.

> *Villain 2**: A villain has kidnapped ten zillion X people and ten zillion Y people. He will either (1) inflict a migraine on each X person, or (2) inflict a one-in-a-zillion chance of death on each Y person. You must choose which.

9. Lazar, "Limited Aggregation and Risk," 141–42, 158; Tadros, "Localized Restricted Aggregation," 197–202; and Lazar and Lee-Stronach, "Axiological Absolutism and Risk," 104–8.

> *Villain 3**: A villain has kidnapped ten zillion X people and ten zillion Y people. He pairs each X person with a Y person. For each pair, the villain will either (1) inflict a migraine on the X person, or (2) inflict a one-in-a-zillion chance of death on the Y person. You must choose between these options for each pair in turn.

If you choose (2) in Villain 2*, it is almost certain that at least one person will die, and there is a good chance that more will die. So, PA will imply that you should choose (1). But PA will also imply that you should choose (2) for each pair in Villain 3*, for each choice is exactly like the choice in Villain 1. So, PA will have inconsistent implications across Villain 2* and Villain 3*, which contradicts our definition of PA. And Villain 3* lacks the features that distinguish Villain 3 from Long Life. If you choose (2) for each pair, it is only statistically certain that someone will die, and as you make each choice, the chance of your having caused a death remains extremely low, at least until someone actually dies.

It might next be objected that the risky opportunities you would be taking in Long Life are special, because they are licensed or even required by social norms to which everyone has tacitly consented, or to which everyone would consent under appropriate conditions.[10] I think this claim is false, but rather than arguing against it, we can shift our focus to a case on which our social norms have no bearing. Suppose that in

> *Migraine Magic*: You have limited magical powers that allow you to cure migraines with a wave of your hand. These powers work almost perfectly, but there is always a one-in-a-zillion chance of them misfiring and causing a bystander to drop dead.

If PA applies to sequences of acts, it has the implausible implication that whether it is permissible for you to use your powers to help migraine sufferers depends on how long you will live.

C. Moral Equivalence

I claimed earlier that choosing (2) for each pair in Villain 3 is "in all morally relevant respects equivalent" to choosing (2) in Villain 2. This claim might also be challenged. It might be argued that, because the permissibility of our acts is determined by more than just their outcomes, there could be an important moral difference between bringing about an

10. I am grateful to an anonymous reviewer for suggesting this objection. The reviewer also pointed out that refusing to take the risky opportunities in Long Life might impose a significant social cost on you, since you would likely be seen as a jerk. Notice that refusing to use your magical powers in Migraine Magic need impose no similar social cost on you, at least if we imagine that you are able to keep your powers a secret.

outcome via a sequence of choices and bringing about the same outcome via a single choice.[11]

But even if this difference is sometimes morally important, it is very implausible that it makes a decisive moral difference in these villain cases. If it did, it would follow that, although it is permissible to sequentially choose (2) for each pair in Villain 3, it is wrong to cut to the chase, by making a single choice for the villain to give a lottery ticket to each Y person.

D. Embracing Full Aggregation

I have argued that, if we accept that you should choose (2) in Villain 1, we must accept that you should choose (2) for each pair in Villain 3. If we accept that you should choose (2) for each pair in Villain 3, then, to avoid inconsistency across cases, we must accept that you should choose (2) in Villain 2. So, we must either reject the claim that you should choose (2) in Villain 1, or accept that you should choose (2) in Villain 2.

We cannot plausibly reject the claim that you should (or at least that it is permissible to) choose (2) in Villain 1. Again, if someone is suffering a migraine, it is permissible to drive them to a pharmacy, even though this exposes pedestrians to a tiny chance of death. So, we must accept that you should choose (2) in Villain 2. That is, we must accept that you should choose that the villain randomly select and kill ten of the Y people rather than inflicting a migraine on each of the X people. This claim is counterintuitive, but it is much less counterintuitive than embracing inconsistency across Villain 2 and Villain 3.

The most natural explanation of why you should choose (2) in Villain 2 is that the aggregate of ten zillion migraines morally outweighs ten deaths. If we accept this explanation, we should accept a fully aggregative view.

However, before we accept this explanation, we should consider an alternative. It is better to be subjected to a tiny chance of death than to suffer a migraine. That is why it is rational to take aspirin, despite the tiny chance of suffering a fatal allergic reaction (or choking to death on the pill). So, each of the X people has a stronger individual claim to be spared from a migraine than each of the Y people has to be spared from a tiny chance of death. So, the reason you should choose (2) in Villain 2 could be that this satisfies the strongest individual claim.

If we accept this alternative explanation, we accept a particular nonaggregative view. According to this view, at least when other things are

11. I am grateful to an anonymous reviewer for pressing me to address this challenge. The reviewer suggested that the moral importance of this difference might be defended by appeal to the intentions or other mental states with which you act. But, as I explain in the main text, holding that this difference makes a decisive moral difference in these villain cases would be implausible regardless of how it is defended.

equal, you should satisfy the strongest individual claim.[12] As well as implying that you should choose (2) in Villain 2, this view correctly implies that you should choose (2) in Villain 1.

However, this view, like all nonaggregative views, implies that there is no number of people you should save from paralysis rather than saving one person from death. That seems to me very implausible. And there is another reason to reject this view. Suppose that in

> *Villain 4*: A villain has kidnapped A and ten zillion others. He will either (1) inflict a migraine on A, or (2) randomly select and kill ten of the others. You must choose which.

Since it is better to be subjected to a tiny chance of death than to suffer a migraine, A has a stronger individual claim to your choosing (2) than any of the others has to your choosing (1). So, the view we are considering has the very implausible implication that you should choose (2)—it implies that you should choose for ten people to die rather than for one person to suffer a migraine.[13] More generally, this view has the very implausible implication that you should do what will be very bad for many people rather than what will be much less bad for a few people whenever you have sufficiently low credence about who the former people are.

I conclude that we should accept the aggregative explanation of why you should choose (2) in Villain 2. We should therefore accept a fully aggregative view. It is worth emphasizing that this does not mean accepting act utilitarianism or any form of act consequentialism. Though these are the most influential fully aggregative views, there are many other forms that fully aggregative views can take. For example, there are fully aggregative views on which, at least when other things are equal, you should satisfy the greatest sum of strength-weighted claims. Nor does accepting a fully aggregative view mean rejecting moral options or moral constraints. Like nonaggregative and partially aggregative views, fully aggregative views can permit us to pursue our projects and forbid us from harming people and breaking promises, even when this reduces total well-being.[14]

E. Strengthening the Reductio

Some people might still be tempted to reject the claim that you should (or that it is permissible to) choose (2) in Villain 1. It is therefore worth

12. See Scanlon, *What We Owe to Each Other*, chap. 5.

13. A similar objection is presented in Reibetanz, "Contractualism and Aggregation," 302–3.

14. For discussion of moral options and moral constraints, see Robert Nozick, *Anarchy, State, and Utopia* (New York: Blackwell, 2006), 28–30; Scheffler, *Rejection of Consequentialism*; and Shelly Kagan, *Normative Ethics* (New York: Routledge, 2018).

showing that the reductio can be strengthened in a way that undermines this response.

We begin by replacing Villain 1 with the following case.

> *Villain 1*$^{\#}$: A villain has kidnapped a million A people and B. He will either (1) inflict a migraine on each A person, or (2) inflict a one-in-a-zillion chance of death on B. You must choose which.

Clearly, PA will imply that you should choose (2). It would be very implausible to allow any number of migraines, however large, to avoid any chance of a death, however small.

We next replace Villain 2 and Villain 3 with the following cases.

> *Villain 2*$^{\#}$: A villain has kidnapped ten million zillion X people and ten zillion Y people. He will either (1) inflict a migraine on each X person, or (2) randomly select and kill ten Y people. You must choose which.

> *Villain 3*$^{\#}$: A villain has kidnapped ten million zillion X people and ten zillion Y people. He divides the X people into groups of one million and then pairs each group with a different Y person. For each pair, the villain will either (1) inflict a migraine on the million X people, or (2) give the Y person a ticket for a lottery with ten zillion tickets. You must choose between these options for each pair in turn. You know that, after you have chosen for each pair, the villain will randomly select ten tickets and kill anyone who has a corresponding ticket.

PA will imply that you should choose (1) in Villain 2$^{\#}$. But it will also imply that you should choose (2) for each pair in Villain 3$^{\#}$, for each choice is exactly like the choice in Villain 1$^{\#}$. So, PA will have inconsistent implications across Villain 2$^{\#}$ and Villain 3$^{\#}$, which contradicts our definition.

III. APPLYING THE REDUCTIO

In this section, I apply my reductio to a range of partially aggregative views from the literature, including some that were specifically designed with risk in mind.[15]

15. I focus on the partially aggregative views that have been most influential in the literature. I do not discuss what might seem the most natural partially aggregative view, which holds, roughly, that the disvalue of a severe burden is greater than the disvalue of any number of minor burdens. This view has not had much uptake, probably because it is more difficult to make plausible than it initially seems. For discussion, see Norcross, "Comparing Harms"; Gustaf Arrhenius and Wlodek Rabinowicz, "Value Superiority," in *The Oxford Handbook of Value Theory*, ed. Iwao Hirose and Jonas Olson (Oxford: Oxford University Press, 2015), 225–48; and Lazar and Lee-Stronach, "Axiological Absolutism and Risk."

A. Aggregate Relevant Claims

Alex Voorhoeve has developed a partially aggregative view that he calls *Aggregate Relevant Claims* (ARC).[16] According to ARC, at least when other things are equal, you should satisfy the greatest sum of strength-weighted, relevant claims. A person has a claim to your 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 the intuitions that motivate partially aggregative views. Suppose that 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, ARC implies that 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 migraine claims are not relevant. So, ARC implies that 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.[17] Suppose that 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 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.

ARC faces a problem that is common among partially aggregative views. It gives rise to either deontic cycling or violations of a principle known as "the independence of irrelevant alternatives."[18] However,

16. Voorhoeve, "How Should We Aggregate." A similar view is presented in Kamm, "Nonconsequentialism," 278–84.

17. A similar answer is presented in Kamm, "Nonconsequentialism," 278–84. Voorhoeve calls this answer "a rough approximation" of his rationale for ARC. His actual rationale avoids some problems with this approximation, but it faces other difficulties. For discussion, see Johanna Privitera, "Aggregate Relevant Claims in Rescue Cases?," *Utilitas* 30 (2017): 228–36; and Tadros, "Localized Restricted Aggregation," 176–78.

18. See Reibetanz, "Contractualism and Aggregation," 308–11; Parfit, "Justifiability to Each Person," 384–85; and Halstead, "Numbers Always Count," 797–99.

Voorhoeve argues that this problem is not decisive, and many philosophers seem to agree.[19] I will therefore set it aside.

Recall that, according to my reductio, no partially aggregative view can simultaneously avoid implausible implications in individual cases, avoid implications that are in tension with the intuitions that incline people toward partially aggregative views, and avoid inconsistent implications across cases that are in all morally relevant respects equivalent. ARC falls on the second horn of this trilemma. Recall Villain 2.

> *Villain 2*: A villain has kidnapped ten zillion X people and ten zillion Y people. He will either (1) inflict a migraine on each X person, or (2) randomly select and kill ten Y people. You must choose which.

Again, it is better to be subjected to a tiny chance of death than to suffer a migraine. That is why it is rational to take aspirin, despite the tiny chance of suffering a fatal allergic reaction. So, a claim to be saved from a migraine is stronger than a claim to be saved from a tiny chance of death. So, ARC implies that you should choose (2). This implication is in tension with the intuitions that incline people toward partially aggregative views.

There are other cases involving risk in which ARC has even more implausible implications. Suppose that in

> *Villain 5*: A villain has kidnapped A and ten zillion others. He will either (1) paralyze A, or (2) randomly select and kill ten of the others. You must choose which.

A has a claim to be saved from paralysis, and each of the ten zillion has a claim to be saved from a tiny chance of death. Since the latter claim is very weak relative to the former, the claims of the ten zillion are not relevant. So, ARC has the very implausible implication that you should choose (2)—it implies that you should choose for ten people to die rather than for one person to suffer paralysis. More generally, ARC implies that you should do what will be very bad for many people rather than what will be less bad for a few people whenever you have sufficiently low credence about who the former people are.[20]

There is a natural way to modify ARC to get the right implication in Villain 5. Though the claims of the ten zillion are very weak relative to

19. See Voorhoeve, "How Should We Aggregate," 76–79; Voorhoeve, "Why One Should Count," 152–53; Kamm, *Intricate Ethics*, 484–86; Tomlin, "On Limited Aggregation," 236–37; Lazar, "Limited Aggregation and Risk," 126–30; and Tadros, "Localized Restricted Aggregation," 185–91.

20. I also present this objection to ARC in Horton, "Aggregation, Complaints, and Risk," 57–58.

the claim of A, the burden that grounds these claims—death—is comparable to the burden that grounds the claim of A. We could argue that the relevance of competing claims is determined not by their relative strength, but rather by the relative size of the burdens that ground them. The claims of the ten zillion would then be relevant, and they would together outweigh the claim of A.

We can call this modified view *ARC 2.0*. The relevance criterion employed by this view cannot be explained in the way Voorhoeve explains his relevance criterion, but perhaps it could be explained in some other way. I will not pursue this issue here.

ARC fell on the second horn of my reductio trilemma. ARC 2.0 avoids this horn. In Villain 2, each of the X people has a claim to be saved from a migraine, and each of the Y people has a claim to be saved from a tiny chance of death. Since the burden that grounds the claims of the X people is very small relative to the burden that grounds the claims of the Y people, the modified relevance criterion implies that the claims of the X people are not relevant. So, ARC 2.0 implies that you should choose (1).

However, ARC 2.0 falls on the first horn of my reductio trilemma. Recall Villain 1.

> *Villain 1*: A villain has kidnapped A and B. He will either (1) inflict a migraine on A, or (2) inflict a one-in-a-zillion chance of death on B. You must choose which.

Since the burden that grounds the claim of A is very weak relative to the burden that grounds the claim of B, the modified relevance criterion implies that the former claim is not relevant. So, ARC 2.0 implies that you should choose (1). Again, that is not plausible. Worse still, by the same reasoning, ARC 2.0 implies that you should choose (1) in Villain 1#.

B. Minimize Relevant Complaints

Some philosophers have suggested an alternative partially aggregative view, specifically designed with risk in mind.[21] This view focuses not on satisfying the relevant claims that people have to your assistance, but rather on minimizing the relevant complaints that people could make against your act once the full effects of your act are known.

We can call this alternative view *Minimize Relevant Complaints* (MRC). According to MRC, at least when other things are equal, you should minimize the expected sum of strength-weighted, relevant complaints. A person has a complaint against your act if and only if it actually burdens her,

21. See Reibetanz, "Contractualism and Aggregation," 304; Otsuka, "Risking Life and Limb," 88; Horton, "Aggregation, Complaints, and Risk," 65–66; and Lazar, "Limited Aggregation and Risk," 136–42.

and the strength of her complaint increases both the worse off she ends up and the greater the difference between her final well-being and her expected well-being had you performed the act that was best for her. A complaint is relevant if and only if it is sufficiently strong relative to the expected strongest complaint associated with every alternative act. To calculate the expected strongest complaint associated with an act, we consider each outcome that might result from the act, take the strongest complaint in each outcome, multiply the strength of this complaint by the probability of this outcome obtaining, and then sum the results. To calculate the expected sum of strength-weighted, relevant complaints associated with an act, we consider each outcome that might result from the act, sum the strengths of the relevant complaints in each outcome, multiply the result by the probability of this outcome obtaining, and then sum the results.

MRC is complicated. However, when there is no risk, MRC is simple to apply, and it has the same implications as ARC. Suppose that you can save either one person from death or a huge number of people from paralysis. If you do the former, there will be a huge number of complaints about paralysis. If you do the latter, there will be one complaint about death. Since these complaints are close in strength, they are all relevant. So, MRC implies that you should save the huge number of people from paralysis. Suppose next that you can save either one person from death or a huge number of people from a migraine. If you do the former, there will be a huge number of complaints about a migraine. If you do the latter, there will be one complaint about death. Since the migraine complaints are not sufficiently strong relative to the death complaint, the migraine complaints are not relevant. So, MRC implies that you should save the one person from death, no matter how many people face migraines.

MRC avoids the first horn of my reductio trilemma. Recall Villain 1.

> *Villain 1*: A villain has kidnapped A and B. He will either (1) inflict a migraine on A, or (2) inflict a one-in-a-zillion chance of death on B. You must choose which.

If you choose (1), A will have a complaint about a migraine. This complaint is relevant, for it is sufficiently strong relative to the expected strongest complaint associated with (2), which is equal to the strength of a complaint about death multiplied by one in a zillion. So, the expected sum of strength-weighted, relevant complaints associated with (1) is equal to the strength of a complaint about a migraine. If you choose (2), there is a tiny chance that B will have a complaint about death. This complaint is also relevant, for it is sufficiently strong relative to the expected strongest complaint associated with (1), which is equal to the strength of a complaint about a migraine. So, the expected sum of strength-weighted,

relevant complaints associated with (2) is equal to the strength of a complaint about death multiplied by one in a zillion. Since a complaint about a migraine is stronger than a complaint about death multiplied by one in a zillion, MRC implies that you should choose (2).

MRC also avoids the second horn of my reductio trilemma. Recall Villain 2.

> *Villain 2*: A villain has kidnapped ten zillion X people and ten zillion Y people. He will either (1) inflict a migraine on each X person, or (2) randomly select and kill ten Y people. You must choose which.

If you choose (1), ten zillion people will have a complaint about a migraine. These complaints are not relevant, for they are not sufficiently strong relative to the expected strongest complaint associated with (2), which is equal to the strength of a complaint about death. So, the expected sum of strength-weighted, relevant complaints associated with (1) is zero. If you choose (2), ten people will have a complaint about death. These complaints are relevant, for they are sufficiently strong relative to the expected strongest complaint associated with (1), which is equal to the strength of a complaint about a migraine. So, the expected sum of strength-weighted, relevant complaints associated with (2) is equal to the combined strength of ten complaints about death. Since the combined strength of ten complaints about death is greater than zero, MRC implies that you should choose (1).

However, as we should now expect, MRC falls on the third horn of my reductio trilemma. The reasoning showing that MRC implies that you should choose (2) in Villain 1 also shows that MRC implies that you should choose (2) for each pair in Villain 3. Since MRC implies that you should choose (1) in Villain 2, MRC has inconsistent implications across Villain 2 and Villain 3.[22]

To avoid this inconsistency, MRC could be applied not to individual acts, but rather to sequences of acts.[23] But, as we have already seen, this causes other problems. If MRC applies to sequences of acts, it has implausible implications in cases like Long Life.

C. Partial Partial Aggregation

Seth Lazar has recently developed a partially aggregative view that he argues avoids problems with risk.[24] His view is, roughly, a hybrid of ARC,

22. I also present this objection to MRC in Horton, "Aggregation, Complaints, and Risk," 68–70.

23. Lazar, "Limited Aggregation and Risk," 141–42, 158; Tadros, "Localized Restricted Aggregation," 197–202; and Lazar and Lee-Stronach, "Axiological Absolutism and Risk," 104–8.

24. Lazar, "Limited Aggregation and Risk."

ARC 2.0, and MRC. It is correspondingly complicated, consisting of fourteen individually complex clauses. However, we need not consider the details here.

Lazar's view implies that you should choose (2) in Villain 2 and thereby falls on the second horn of my reductio trilemma. Surprisingly, Lazar is aware that his view has implications of this kind, and he suggests that proponents of partially aggregative views should embrace them.[25]

If proponents of partially aggregative views embrace these implications, they draw a sharp distinction between cases involving certainty and cases involving risk. They continue to hold that, in cases involving certainty, there is no number of people you should save from a minor burden rather than saving one person from a severe burden. But they also hold that, in cases with the structure of Villain 2, you should save *N* zillion people from a minor burden rather than saving *N* people from a severe burden, regardless of the size of *N*.

We can call this view *Partial Partial Aggregation* (PPA). It seems bizarre to draw such a sharp distinction between cases involving certainty and cases involving risk. But there is another, stronger reason to reject PPA. This view faces the same objections as fully aggregative views while lacking their theoretical plausibility.

Suppose that in

> *Lollipops for Lives*: If you grant the Devil permission to kill a few people in one half of the universe, he will give everyone in the other half a lollipop.[26]

Cases like this present the strongest objection to fully aggregative views. Intuitively, you should not trade lives for lollipops, no matter how many lollipops are on offer. But fully aggregative views imply that, if enough lollipops are on offer, you should make the trade, for the small benefits to the many will together outweigh the severe burdens to the few.[27]

If PPA avoided the problems with standard partially aggregative views while capturing our intuitions about cases like Lollipops for Lives, that might justify accepting it. But PPA does not capture these intuitions. Since the Devil has not told you who will be killed, Lollipops for Lives is

25. Lazar, "Limited Aggregation and Risk," 132, 157.

26. This case is a variation of a case presented in Larry Temkin, "Aggregation within Lives," *Social Philosophy and Policy* 26 (2009): 1–29.

27. Some proponents of fully aggregative views have responded by arguing that our antiaggregative intuitions are systematically unreliable. I am not sure how much weight to place on these arguments. For discussion, see Norcross, "Comparing Harms," 146–47; Parfit, "Justifiability to Each Person," 385–86; John Broome, *Weighing Lives* (Oxford: Oxford University Press, 2004), 56–57; Michael Huemer, "In Defence of Repugnance," *Mind* 117 (2008): 899–933, 907–10; Halstead, "Numbers Always Count," 797; Horton, "Aggregation, Complaints, and Risk," 72–73; and Tadros, "Localized Restricted Aggregation," 175–76.

relevantly like Villain 2. So, PPA implies that, if enough lollipops are on offer, you should make the trade. And, unlike proponents of fully aggregative views, proponents of PPA cannot offer the most natural explanation of this implication.

IV. SUMMARY

In Section II, I developed a risk-based reductio argument against partially aggregative views. I then argued that the only plausible response to this reductio is to accept a fully aggregative view. In Section III, I applied this reductio to a range of partially aggregative views from the literature, including some that were specifically designed with risk in mind.

I conclude that we should accept a fully aggregative view. As I said earlier, this does not mean accepting act utilitarianism or any form of act consequentialism, nor does it mean rejecting moral options or moral constraints. It does mean, however, that there is some number of people you should save from a minor burden, such as a migraine, rather than saving one person from a severe burden, such as death.