Aggregating Small Risks of Serious Harms

Tomi Francis (Global Priorities Institute, University of Oxford)

Global Priorities Institute | October 2024

GPI Working Paper No . 21-2024

Please cite this working paper as: Francis, T. Aggregating Small Risks of Serious Harms. *Global Priorities Institute Working Paper Series, No. 21-2024.* Available at https://globalprioritiesinstitute.org/aggregating-small-risks-of-serious-harms-tomi-francis

GLOBAL PRIORITIES INSTITUTE

UNIVERSITY OF OXFORD

Aggregating Small Risks of Serious Harms

Tomi Francis*

Abstract

According to Partial Aggregation, a serious harm can be outweighed by a large number of somewhat less serious harms, but can outweigh any number of trivial harms. In this paper, I address the question of how we should extend Partial Aggregation to cases of risk, and especially to cases involving small risks of serious harms. I argue that, contrary to the most popular versions of the ex ante and ex post views, we should sometimes prevent a small risk that a large number of people will suffer serious harms rather than prevent a small number of people from certainly suffering the same harms. Along the way, I object to the ex ante view on the grounds that it gives an implausible degree of priority to preventing identified over statistical harms, and to the ex post view on the grounds that it fails to respect the separateness of persons. An insight about the nature of claims emerges from these arguments: there are three conceptually distinct senses in which a person's claim can be said to have a certain degree of strength. I make use of the distinction between these three senses in which a claim can be said to have strength in order to set out a new, more plausible, view about the aggregation of people's claims under risk.

1 Introduction

Consider the following stylised decision situation.

> **Deaths vs Risks of Death**: There is a one-in-ten-thousand chance that eight million New Yorkers will be killed. One hundred Texans also face certain death. All of these individuals are known to you by name. You can save the Texans or eliminate the risk to the New Yorkers, but not both.[1]

Nobody will ever actually face this choice, not least because nobody is *that* good with names. But we do sometimes need to trade off small probabilities of harms to the many against larger probabilities of harms to the few. A government might have to choose between deploying part of its limited budget on the provision of life-saving medical care, or on preparedness for an unlikely pandemic. A philanthropist might wonder whether their money is better spent on relief for an ongoing disaster or on preparedness for future disasters. Voters might ask themselves whether it is worth diverting money that could be used to save lives today towards civil defence measures, given that these could save many lives in the (hopefully) very unlikely event of nuclear war. And so on. The actual decision situations are the ones

*Global Priorities Institute, Faculty of Philosophy, University of Oxford. I would like to thank Avram Hiller, Todd Karhu, Kacper Kowalczyk, Jakob Lohmar, Andreas Mogensen, Michael Otsuka, Theron Pummer and Timothy L. Williamson for helpful comments and discussion. I would also like to thank audiences at the St Andrews–GPI Joint Workshop on the Long-Term Future, the fourteenth Oxford Workshop on Global Priorities Research, and the 2024 Rocky Mountains Ethics Congress.

[1] A structurally similar case has been discussed by Steuwer (2021: 119–120).

of practical interest. But cleaned-up cases like *Deaths vs Risks of Death* can help us identify the morally salient factors at issue.

What, then, ought we to do in cases like *Deaths vs Risks of Death*? Proponents of "Full Aggregation"—the view that benefits and harms can always be aggregated—have a ready answer available: simply compare the expected numbers of lives saved.[2] Rescuing the New Yorkers would save eight hundred people in expectation, while rescuing the Texans would save one hundred. So, on this view, we should save the New Yorkers. At the opposite end of the spectrum, proponents of "No Aggregation" think that the numbers never matter, so that we should prevent a single serious harm rather than any number of less serious harms.[3] On this view, presumably we ought to save the Texans.[4]

However, both Full and No Aggregation have implications that many people find implausible. Full Aggregation implies that we should prevent a large enough number of trivial harms rather than save a life, since doing so might prevent greater aggregate harm; and this strikes many people as counter-intuitive.[5] No Aggregation goes too far in the opposite direction: it prevents us from saying, as we might like, that preventing one million people from having their left legs broken is more important than preventing one person from having both legs broken.

There is a third view which sits between these two extremes. According to Partial Aggregation (also known as Limited Aggregation), large enough numbers of lesser harms can outweigh smaller numbers of *somewhat* graver harms, but no number of trivial harms can outweigh any number of serious harms. Or, put more simply, losses of limbs can outweigh death, but headaches cannot. Since the verdicts of Partial Aggregation strike many people as extremely plausible, I shall assume this view without further argument. My project shall be to investigate what we should say about *Deaths vs Risks of Death*, and about risky cases more broadly, if we take Partial Aggregation for granted.

The most popular version of Partial Aggregation is Alex Voorhoeve's "Aggregate Relevant Claims View".[6] This works as follows:

(i) Individuals have claims against being harmed, and the strengths of these claims correspond to the extent to which they would be harmed.[7]

[2]Not every fully aggregative moral theory endorses this answer. For example, Ex Ante Prioritarianism, which has been defended by McCarthy (2006, 2008), applies a priority weighting to each person's *expected* wellbeing: our reasons to benefit people are stronger, the worse-off these people would otherwise be *in expectation*. Since the Texans are worse off in expectation if they are not helped, Ex Ante Prioritarians will somewhat prioritise increasing their expected wellbeing, and thus (depending on how heavily the worse-off are to be prioritised) may decide to save the Texans. One could, in principle, also combine Alec Walen (2020)'s "Weak Aggregation" with the view that while claims lose strength at a faster-than-linear rate as they decline in size, they always remain relevant even to much stronger claims.

[3]See Taurek 1977.

[4][Redacted] has pointed out to me that proponents of non-aggregation might be able to justify the social policy of saving the greater expected number on the basis that this policy would be in each person's interests *ex ante*. Thus, if the decision in *Deaths vs Risks of Death* falls to a government, the choice to safeguard New York might be justifiable even if interpersonal aggregation is not, despite the fact that safeguarding New York is unlikely to actually save any New Yorkers. See Taurek 1977: 312–313. Reibetanz (1998: 301) suggests something similar.

[5]See Scanlon 1998: 235.

[6]See Voorhoeve 2014.

[7]One might take the view that claims should also be weighted for additional morally relevant factors such as the importance of prioritising the worst off. Such complications will not matter for my purposes, so I shall set them aside going forward.

(ii) Claims compete when they cannot be jointly satisfied.

(iii) A claim is *relevant* if it is strong enough, compared to the strongest individual competing claim. Otherwise it is irrelevant.

(iv) We should choose whichever action maximises the satisfaction of aggregate strength-weighted relevant claims.

Unfortunately, the Aggregate Relevant Claims View, as specified above, does not yield a useful verdict on whether we should rescue the Texans or the New Yorkers. The problem is that the strengths of claims are taken to be given by the extent to which individuals *would* be harmed, unless you were to prevent the harm in question.[8] But the point of *Deaths vs Risks of Death* is that we don't know whether the New Yorkers will die if this is not prevented. In order to get a useful recommendation as to what we should do, we must therefore find a way to extend the Aggregate Relevant Claims View to cases involving risk.

At a first pass, the main thing we need to decide is whether the New Yorkers' claims are relevant. Two ways of thinking about this immediately present themselves. On the first, each individual New Yorker can have only a tiny claim to be rescued, because the strength of their claim against death should be discounted by the improbability that their death would actually eventuate if one were to save the Texans instead. These claims will then be treated the same way as claims against the certainty of a harm one-ten-thousandth the size of death. On the plausible assumption that such harms would be irrelevant to certain death, the claims of the New Yorkers will also be irrelevant, leading us to side with the unopposed claims possessed by the Texans.

The second way of thinking about the relevance of the New Yorkers' claims is to attend to the fact that, if the one-in-ten-thousand probability does eventuate, the New Yorkers stand to *die*. Risks of death, one might think, should always count as relevant when they compete with certain death. On this approach, the New Yorkers will have relevant claims to be saved, and the aggregate strength of these claims will presumably exceed the aggregate strength of the claims of the Texans.

I shall argue in this paper that we should take the second view, not the first. I'll begin by sketching the two most popular approaches to risk in the literature on Partial Aggregation, which are the ex ante and ex post views. The most common versions of these views agree that the claims of the New Yorkers should be considered irrelevant. However, I shall argue that both are unsatisfactory. The main problem with the ex ante view is well-known: it gives implausibly extreme priority to preventing identified over statistical harms. I shall flesh out this objection in §2. The ex post view can be objected to on various grounds, but the objection I shall press in this paper is that it fails to respect the separateness of persons. I shall make this objection precise in §3.

If successful, these two objections to the two standard views do not only show that these views are mistaken. As I shall show in §4, it turns out that avoiding these two objections requires that claims against small risks of serious harms must be relevant to claims against being harmed to the same extent with certainty. Since we should indeed avoid both objections, this answers the question we started with. We should rescue the many rather than

[8] That is, on a straightforward reading, the Aggregate Relevant Claims View might be interpreted as what Steuwer (2022: 72) calls "actualist ex post limited aggregation".

the few in cases like *Deaths vs Risks of Death*, because the claims of the many are indeed relevant.

To help explain how we can reconcile respecting the separateness of persons with taking statistical harms seriously, I shall distinguish between three senses in which a claim can be said to have a certain degree of "strength". I will use this distinction to present what I shall call the "Ex Ante Claims, Ex Post Relevance" view. I believe that this view brings us closer to getting things right about the aggregation of people's claims in cases of risk. I shall respond to two important objections to it in §5, and then conclude in §6.

2 Against the Ex Ante View

2.1 Two Standard Views

There are two standard approaches to the problem of partially aggregating people's claims in cases of risk. On what is usually called the "ex post" approach, whenever we can be certain that *someone* will be gravely harmed if we do X, though we are not sure who, we should treat this situation in the same way as we would if we knew who would suffer the grave harm in question.[9] Many ex post views also say that when there is a small chance that *anyone* will be harmed, the claims associated with these harms must all be small, and will in particular be rendered irrelevant by strong claims, such as claims against certain death. These views therefore imply that we ought to save the Texans rather than the New Yorkers, since there is then only a small probability that anyone will die.[10]

The second standard approach to partially aggregating people's claims in cases of risk is the *ex ante* approach. On the ex ante view, we attend to the expected losses for each particular individual, rather than to the matter of whether *someone* will be gravely harmed.[11] When it is known that someone will be gravely harmed if we do X, but it is unknown who this will be, the ex ante approach says that many people have weak claims against our doing X, which correspond to the small expected losses suffered by each person. The ex ante view says that in such cases there is no strong claim against our doing X. The ex ante view also licenses saving the Texans in *Deaths vs Risks of Death*, since the probability-discounted claims of the New Yorkers are too small to be relevant when they compete with the claims against certain death possessed by the Texans.

On the face of it, then, the judgement that we should save the Texans in *Deaths vs Risks of Death* looks to be fairly secure. There are two standard approaches to partially aggregating people's claims in cases of risks, and both seem to agree that we should save the Texans.[12] I shall argue, however, that both of these standard approaches should be rejected.

[9] See Scanlon 1998: 208–209, Reibetanz 1998 and Otsuka 2015.

[10] Otsuka (2015: 91) might favour saving the New Yorkers, given his claim that "what becomes relevant [in the cases he discusses] is simply the expected number who would suffer [the] harm under a given course of action". I say "might" because the cases Otsuka discusses do not involve only a small probability that *anyone* will be harmed.

[11] See for instance Frick 2015a,b or Kumar 2015.

[12] Emma Curran (2022) presses this point with regard to a case involving small chances of benefiting future generations.

2.2 The Statistical Harms Problem

I begin by recapping a familiar objection to the ex ante view, which seems to me decisive. This is that the ex ante view has false implications in the following kind of case:

> **Identified vs Statistical Lives**: Eight million New Yorkers will be entered into an indeterministic lottery which will randomly select eight hundred of them. We can either:
>
> (IL) Prevent a Californian, Ted, from being prematurely killed, or
>
> (SL) Prevent the eight hundred randomly selected New Yorkers from being prematurely killed.

If we accept the ex ante view, we might reason as follows. The New Yorkers only face small risks of death if we choose (IL), because each person has only a one-in-ten-thousand chance of drawing a losing ticket. Ted, on the other hand, faces certain death if we choose (SL). Ted's claim is very strong, while the individual claims of the New Yorkers are weak enough to be irrelevant when compared to certain death. We therefore ought to choose (IL) rather than (SL).

But this conclusion is surely false.[13] Admittedly, Ted is doomed if we choose (SL), whereas each New Yorker would have only a one-in-ten-thousand chance of death if we choose (IL). These facts may hold some moral significance. They cannot, however, be important enough to justify our saving one life when we could instead save eight hundred. Furthermore, the disparity between the identified and statistical harms can be ramped up if we make use of a more outlandish example:

> **Identified Hand vs Statistical Lives**: One quintillion Andromedans, all of whom are persons with the same moral status as humans, will be entered into an indeterministic lottery which will randomly select one trillion of them. We can either:
>
> (IH) Prevent Ted, the Californian, from losing his left hand, or
>
> (SL′) Prevent the one trillion randomly selected Andromedans from being killed.

Following the same reasoning as before, it seems that if we accept the ex ante view, we shall have to say that we should choose (IH) rather than (SL′), since a tiny risk of death is presumably irrelevant compared to Ted's certainty of losing his left hand. If this is what the ex ante view implies, it strikes me as a *reductio*.

[13] Steuwer (2021: 139–142) disagrees, and seems to think that the verdict of the ex ante view is defensible here, given that objective rather than epistemic probabilities are at issue. He gives two arguments. First: if we combine the claim that we should save many statistical lives rather than fewer identified lives with certain other plausible claims, we are left with cyclic moral preferences. While this is true, Partial Aggregation already gives rise to cyclic moral preferences even if we do not consider risky cases; see Parfit 2003: 384, Gustafsson 2015, Halstead 2016: 797–799 and Voorhoeve 2017: 153. Second: Steuwer holds that, rather than viewing (IL) as a choice which results in eight hundred New Yorkers dying, we should view it as a choice in which eight million New Yorkers have a small chance of dying. To look at things in the first way begs the question in favour of ex post views. In reply: both ways of viewing the consequences of (IL) are accurate. If (IL) is chosen, it is a fact that eight million New Yorkers face small risks of dying. But it is also a fact that eight hundred New Yorkers will die. Appeals to the apparent moral significance of the second fact do not beg the question.

2.3 Pluralism to the Rescue?

Proponents of the ex ante view are well aware of the statistical harms problem. Johann Frick (2015a: 219–223) takes cases like *Identified vs Statistical Lives* to indicate that the ex ante approach merely describes one part of morality, but not the whole of it. In addition to our moral reasons stemming from the complaints of individuals (our "reasons of equity", as Frick puts it), which generally favour saving identified lives, we are supposed to have other moral reasons which can tip the balance in favour of saving statistical lives. Frick suggests that these moral reasons might stem from the effects of our actions on people's wellbeing, rather than on the individual claims or complaints that can be pressed by each person.[14] On this pluralist view, we ought to choose (SL) in cases like *Identified vs Statistical Lives*, not because we thereby satisfy the greatest aggregate of strength-weighted relevant claims, but because we thereby greatly benefit the eight hundred New Yorkers we happen to save, whoever those people turn out to be.

This sort of reply leaves open the following question: how exactly should our reasons to make people better off be taken into account? We cannot straightforwardly aggregate these reasons, because as Seth Lazar has pointed out, this would imply that we sometimes ought to prevent headaches rather than deaths.[15] We therefore need to find an account of our reasons to benefit people which allows for aggregation in cases like *Identified vs Statistical Lives*, where deaths or similarly serious harms are at stake either way, but which disallows aggregation when headaches are being traded off against deaths. Indeed, Frick recognises that the wellbeing-based component of a plausible pluralist view must meet these requirements.[16]

The problem is that this brings us back to square one. For suppose we succeed in providing an account of our reasons of wellbeing which yields the apparently correct results in cases like *Identified vs Statistical Lives*, and in other cases of aggregation under risk, without implying that we should prevent headaches rather than deaths. In that case, what we've got looks an awful lot like an adequate account of Partial Aggregation in cases of risk. And, by stipulation, this is an account which disagrees with the monistic ex ante view when it comes to the statistical harms problem.

It therefore seems to me that pluralism of this sort is unpromising, because it is needlessly unparsimonious. I do not see why we need to appeal to reasons of equity in cases of aggregation if our reasons of wellbeing are enough to get the job done. But even if we do want to be pluralists, we are still left with the problem we started with, which is to find a version of Partial Aggregation which robustly delivers the right answers in cases involving risk. As we have seen, this view (or component) will need to rule that statistical harms can sometimes outweigh identified harms. It must therefore be different in this respect from the standard ex ante view.

[14] Frick 2015a: 222
[15] Lazar 2018: 138–139
[16] See Frick 2015a: 223, fn 48.

3 Against the Ex Post View

3.1 Social Risk

If the ex ante view has no legs, or if it needs to be supplemented, the natural response is to turn to some version of the ex post view. The ex post approach, however, also has its problems. One prominent objection to it is that it seems to imply that we should not confer packages of benefits and risks on large numbers of people, even when doing so could be justified to each person as being in their ex ante interests. A standard case goes as follows.[17] We could vaccinate one hundred million children against a serious, but non-fatal, disease. As a side effect, the vaccine would impose on each child an independent one-in-a-million risk of death. If the disease is serious enough, it seems that we ought to vaccinate all of the children, since doing so would be in each child's ex ante interests despite the tiny risk of death involved. Yet the ex post view approach appears to disagree because if we do vaccinate everyone, it is all but certain that at least one child will die.

I find this objection persuasive, but not decisive. It might seem to be in each child's interests to be vaccinated, *ex ante*. But once the vaccinations are carried out, predictably we will see that being vaccinated was not *actually* in the interests of those unlucky children who end up dying. We might reasonably hold that the facts about the actual interests of the children are more important, morally speaking, than the facts about the ex ante interests of the children, as observed from our position of ignorance prior to actually carrying out the vaccinations.

One might alternatively respond to this problem by trying to avoid it. It might be said, for instance, that those who stand to benefit from a package of benefits and risks would withdraw their claims against having those risks inflicted on them, so that there are never any live objections to courses of action which are in each person's ex ante interests.[18]

I will therefore provide a different sort of objection to the ex post approach. I shall argue that there is a deeper problem with ex post views which cannot be avoided by modifying them to deal with a few edge cases. The problem is that these views fail to respect the separateness of persons, in a sense that I shall now make precise.

3.2 The Separateness of Persons and the Individualist Restriction

According to Full Aggregation, we may always aggregate benefits and harms. What, if anything, is wrong with this view? One answer is simply that it implies that we should prevent a large number of headaches rather than one death, and that this is implausible. But this answer, by itself, is not very satisfying. *Why* is it wrong to prevent the headaches, given that one would thereby prevent the greater aggregate harm? Opponents of Full Aggregation typically answer that in attending to the aggregate strength of this harm, we miss the fact that there is no person who stands to suffer a harm of that aggregate strength.[19] If we fail to prevent the headaches, many individuals will have a headache, and *that is all*. There is nobody who will suffer anything worse than a headache. In contrast, if we fail to prevent the death, *that individual* stands to die. To insist that we should prevent the headaches

[17] See Frick 2015a: 181–182.
[18] See Walen 2020: 72–73.
[19] See Taurek 1977: 307–308,

simply because the aggregate harms involved are greater is to fail to appreciate the moral significance of these facts about the harms faced by each individual.

This explanation suggests that we must have some place in our moral theory for individualistic comparisons of harms and benefits: comparisons on which what one individual stands to gain or lose must be compared to what another individual stands to gain or lose. One way of taking individualistic comparisons to matter is to take them to be *all* that matters. If we take this approach, we will accept No Aggregation. But, as I noted in §1, most people find this view quite implausible.

The Aggregate Relevant Claims View provides another way of respecting the separateness of persons. It can be thought of as involving two components: a non-aggregative and an aggregative component. The aggregative component directs us to maximise the satisfaction of strength-weighted claims. The non-aggregative component places a constraint on this demand: the claims in question must be relevant, where relevance is determined by comparing the individual claims to other individual claims—and never to aggregates of individual claims. It is by determining relevance in an individualistic way that we respect the separateness of persons. As Alex Voorhoeve puts it:

> On [the nonaggregative approach], for each individual taken separately, one takes in what she would have to give up if another person's competing claim were satisfied. [...] Its justification is that this form of concern for each person taken alone is a natural expression of our appreciation of the separateness of persons. [The Aggregate Relevant Claims View] takes the form of maximization under a constraint: it maximizes the sum of strength-weighted claims that are satisfied under the constraint that these claims are relevant. [...] The constraint stems from the nonaggregative approach.[20]
>
> Voorhoeve 2014: 68–70

These considerations support the following two requirements on claims. The first requirement is that claims must be possessed by particular individuals and not by aggregates of disparate individuals, for otherwise we mistakenly attach moral significance to claims (or complaints) levied by entities who lack a first-personal perspective from which claims can legitimately be pressed. Call this the principle that Only Individuals Have Claims. The second requirement is:

> **The Individualist Restriction on the Relevance and Strength of Claims**: The relevance and strength of an individual's claim against the imposition of a risk of harm cannot depend on the number of equally strong claims with which the individual's claim competes, provided it competes with at least one such claim; nor can they depend on whether the individual's claim can be jointly satisfied along with other claims.

I shall henceforth refer to this condition merely as the "Individualist Restriction on Claims". But it is worth clarifying that this condition is much weaker than that feature of Scanlonian Contractualism commonly known as the "Individualist Restriction",[21] according to which

[20] See also Voorhoeve 2017: 150.
[21] See Parfit 2003: 372.

> the justifiability of a moral principle depends only on various *individuals'* reasons for objecting to that principle and alternatives to it.
>
> Scanlon 1998: 229

Scanlon's stronger Individualist Restriction can be taken to require not only that the relevance and strengths of claims are assessed in an individualistic way, but also that competition between groups of claims is resolved in an individualistic manner. The weaker Individualist Restriction on Claims requires individualism of the former sort, but not the latter.[22] The Individualist Restriction on Claims therefore avoids the objection that it cannot account for our apparent duty to save the greater rather than the lesser number from harms of the same magnitude.[23]

We shall next see that the ex post view must violate either the Individualist Restriction on Claims or the condition that Only Individuals Have Claims. Either way, it fails to respect the separateness of persons.

3.3 Fair Lotteries and Unfair Lotteries

Let us assume that a broken finger would be irrelevant to near-certain death, but that ninety-nine million broken fingers would constitute a much greater aggregate harm than one death. Consider the following pair of cases.

> **Fair Lottery**: Ninety-nine million people have been forced to enter an indeterministic fair lottery with one hundred million numbered tickets. If a person's number comes up, their number is up (that is, they die). Additionally, ninety-nine million *other* people stand to suffer broken fingers. You can prevent the lottery from going ahead or you can prevent the broken fingers, but not both.
>
> **Unfair Lottery**: The same lottery will take place, but Ted alone has been forced to enter, and has received ninety-nine million (99%) of the tickets. If one of Ted's numbers comes up, Ted's number is up. You can prevent the lottery from going ahead or you can prevent ninety-nine million other people from suffering broken fingers, but not both.

We can characterise ex post views as saying that we must say the same thing about these two cases.[24] That is, ex post views imply one of the following two propositions.

(i) It is permissible to prevent the broken fingers in both *Fair Lottery* and *Unfair Lottery*.

(ii) It is impermissible to prevent the broken fingers in both *Fair Lottery* and *Unfair Lottery*.

[22] One way of seeing that my version of the Individualist Restriction really is weaker is to note that it is actually compatible with Full Aggregation. Utilitarianism, for instance, could be understood as saying that the claims of individuals are *always* relevant, and that the strength of a claim corresponds exactly to the associated expected loss of wellbeing.

[23] See Otsuka 2006 for this objection.

[24] Reibetanz (1998: 304) suggests that it is only when we *know* that someone will be harmed that we should depart from the ex ante view. On this view, it matters a great deal that it is not quite *guaranteed* that someone will die if we do not prevent the lottery in *Fair Lottery*: we should attend to people's ex ante claims, and therefore we should treat *Fair Lottery* and *Unfair Lottery* differently. It seems implausible, however, that whether a death is guaranteed rather than merely extremely likely should make such a great difference.

Option (i) is ruled out by our assumption that a claim against a broken finger would be irrelevant next to a claim against almost-certain death, and is in any case implausible. This leaves us with (ii). The conjunct of (ii) that is relevant for our purposes is that it would be impermissible to prevent the broken fingers in *Fair Lottery*. The claims of the ninety-nine million people who stand to have their fingers broken in *Fair Lottery* must therefore be irrelevant, because otherwise they would outweigh the claims of those who are made to enter the lottery.

Consider next the following variant of *Fair Lottery*:

> **Single-Person Fair Lottery**: The same lottery from *Fair Lottery* will take place, but only Bill is made to enter it. He will be assigned one numbered ticket. If this number comes up, Bill's number is up. Ninety-nine million other people stand to suffer broken fingers. You can prevent Bill's possible death from the lottery or the broken fingers, but not both.

It would obviously be permissible, in Single-Person Fair Lottery, to prevent the broken fingers rather than the possible death. Doing this means that Bill will face a one-in-a-hundred-million risk of death. But that's ok. Car journeys come with small risks of death too: roughly one fatal injury every ten million trips in the United States.[25] Some of these fatal injuries involve pedestrians. Yet it is not impermissible to drive somebody to the hospital for something as significant as a broken finger, just because one would thereby impose tiny risks of death on pedestrians.[26] So, it is permissible to subject Bill to a tiny risk of death in *Single-Person Fair Lottery* in order to prevent ninety-nine million people from having their fingers broken. Those ninety-nine million people's claims against having their fingers broken must therefore be relevant.

3.4 The Argument from the Individualist Restriction on Claims

To summarise what we have seen so far, all plausible ex post views must imply the following two propositions.

(i) In *Fair Lottery*, the claims of the ninety-nine million people against having their fingers broken are irrelevant.

(ii) In *Single-Person Fair Lottery*, the claims of the ninety-nine million people against having their fingers broken are relevant.

This brings me to my main objection to the ex post view, which is that propositions (i) and (ii) jointly violate the Individualist Restriction on Claims. To see this, first note that the Individualist Restriction on Claims implies that the strengths of the claims against risks of death in *Fair Lottery* and *Single-Person Fair Lottery* cannot depend on the number of equally strong claims with which they can be jointly satisfied. Thus, these claims are equally strong; and since Only Individuals Have Claims, there are no other claims against risks of death at issue.

The Individualist Restriction on Claims also implies that whether a claim counts as relevant cannot depend on the *number* of equally strong claims with which it competes, provided it

[25] See Beck, Dellinger, and O'Neil 2007: 214.

[26] Seth Lazar (2018: 118) makes the stronger claim that we may impose tiny risks of death even if we are merely driving to the store to buy popcorn and chocolate. This stronger claim is surely correct as well.

competes with at least one such claim. In particular, if the claims against broken fingers are relevant despite competing with Bill's claim against a tiny risk of death in *Single-Person Fair Lottery*, then they must remain relevant when they compete with the ninety-nine million claims against risks of death—which are, as we have seen, equally strong—in *Fair Lottery*.[27] Our conclusion is that if the claims against broken fingers are relevant in *Single-Person Fair Lottery*, then they are relevant in *Fair Lottery*. Put another way, if (ii) is true, then (i) is false; and since (ii) *is* true, (i) is false.

3.5 Ex Ante and Ex Post Claims

That's the gist of my central argument against the ex post view. But there is one respect in which the version of it I just gave is too quick. It tacitly assumes that people's claims always count straightforwardly in favour of or against a particular course of action, and that the strengths of these claims depend solely on the probabilities and severities of the respective potential harms. In short, it assumes that the claims we are dealing with are people's *ex ante claims*. But there are other ways of understanding people's claims.

Here's one of them. Let's say that a person's *ex post claim* corresponds to the amount of wellbeing an individual loses in a particular state of nature. (By a state of nature, I mean something like a particular lottery ticket being drawn or a particular roll of the dice: an event which is both statistically and causally independent of which action you take, and which, given your decision, fully determines all the morally relevant consequences that follow.)[28] Ex post claims do not straightforwardly tell for or against a particular action from the perspective of some individual, but only do so relative to particular states of nature. Someone who stands to receive a package of risks and benefits might have ex post claims in favour of receiving this package in those states of nature where the risks do not eventuate, but will likely have ex post claims against receiving this package in those states of nature where the risks do eventuate.

Consider now what we might call the *State-Relative Ex Post View*. This view directs us to assess ex post claims for relevance within their respective states of nature according to the version of the Aggregate Relevant Claims View applicable to cases of certainty. That is, we discard all of those ex post claims which are not sufficiently strong compared to the largest competing claims in the same states of nature. We then aggregate the probability-weighted strengths of the rest.

Now, as far as I know, nobody actually accepts this view; and that's a good thing, because it turns out to be very implausible.[29] The reason I bring it up is that the State-Relative Ex Post View is indeed an ex post view in the sense that it sees no significant moral difference between *Fair Lottery* and *Unfair Lottery*, yet it also satisfies the Individualist Restriction on

[27] Johann Frick (2015a: 196) provides a related objection to the ex post view. He rhetorically asks why a person who is killed in a case like *Single-Person Fair Lottery* should have a weaker claim than a person who is killed in a case like *Fair Lottery* just because, if the latter person had not been killed, then it is likely that someone else would have been killed instead. He goes on to suggest that the ex post verdict here involves combining the complaints of different individuals in an objectionably aggregative way. My own argument from the Individualist Restriction on Claims can be viewed as a way of making precise the sense in which this is so.

[28] The stipulation that states of nature are statistically *and* causally independent of the agent's actions is intended to restrict our attention to those cases where it does not matter whether we accept causal or evidential decision theory (on which see Nozick 1969 and Lewis 1981.)

[29] See Steuwer (2022: 74–77).

Claims if we interpret the "claims" in this principle as ex post claims. It therefore illustrates one important way in which ex post theorists might be able to escape my argument: they might offer an alternative account of the notion of a person's claim, on which their view satisfies the Individualist Restriction on Claims.

Will the notion of an "ex post claim" do the required work? I think that it will not, because it offers no plausible way of explaining why the ex post claims of the many against broken fingers in *Fair Lottery* should be irrelevant in those states of nature in which nobody stands to die. It therefore cannot explain why we should prevent a 99% risk of death rather than any number of broken fingers. But there is another way of understanding claims which might seem to do the trick.

3.6 Anonymous Claims

Rather than appealing to ex post claims, we could appeal to what I shall call *anonymous claims*. We might note that in *Fair Lottery*, there is a 99% chance that *someone* will die, and conclude that we can therefore attach a claim to this *someone* which is strong enough to render the claims against broken fingers irrelevant. Michael Otsuka can be interpreted as offering a view of this sort.[30] The important question for our purposes, however, is whether we can appeal to anonymous claims while respecting the separateness of persons. Otsuka suggests that we can. Here is a paraphrase of his argument.

> Anonymous claims attach to *someone*, but seem not to attach to any particular person. Yet this first impression is misleading. Anonymous claims do not really attach to fictional people, or to aggregates of disparate people. If the lottery goes ahead in *Fair Lottery*, it is true that each individual person is unlikely to die. But it is very likely that *someone* will die. If a death eventuates, the victim will be an actual flesh-and-blood person: one of the individuals in question.[31] While we cannot say who that individual will be, we should nevertheless assign them a strong claim against the imposition of death.

This seems quite persuasive on first reading, but the devil is in the details. Who exactly is supposed to possess the supposed anonymous claim? If we are merely epistemically uncertain about the results of the lottery, and some particular person's number will in fact come up, then an obvious answer is available: *that person*—suppose it's Bill—has the claim. Since Bill is a particular individual, saying that Bill has a strong claim does not violate the principle that Only Individuals Have Claims.

However, we run back into trouble when we return to the *Single-Person Fair Lottery* case. For we can imagine that, in that single person case, Bill's number will also come up, even though this is epistemically unlikely. (Suppose we assign Bill the very same ticket in either case.) What are we to say about Bill's claim here? There are two options. We could say that Bill retains a strong claim against death in the single-person case, or we could say that he does not.

Suppose first that Bill retains his strong claim in the single-person case. It cannot plausibly be the case that Bill's claim would still be strong even if we do *not* stipulate that his number will in fact come up. If that were true, we would always have to save Bill in *Single-Person*

[30] See Otsuka 2015. A version of this view is fleshed out, though not endorsed, by Steuwer 2021: 118.
[31] Otsuka 2015: 86.

Fair Lottery, and that is implausible for the reasons I mentioned earlier. So we must say that Bill has a strong claim if his number *will* come up, but not if it will not. But this leaves us without any useful guidance on what to do, given our actual epistemic situation. The point of a moral theory which incorporates risk is to give us useful advice given that we don't know how things will go. A theory which attributes a strong claim to Bill if his number will come up, but not if it won't, fails on this count. (It also fails to prohibit saving the fingers in *Fair Lottery* in the case where, by luck, nobody would be killed if the lottery were to go ahead.)

Suppose next that Bill does not retain a strong claim against death in the single-person case, even though his number will in fact come up. We assumed earlier that he *does* have a strong claim against death if his number will come up in the full *Fair Lottery* case. But saying both of these things at once violates the Individualist Restriction on Claims. For we are making the strength of Bill's claim depend on whether or not *other* people stand to suffer risks of death.[32] To do this is to fail to assess people's claims one-by-one.

What all this shows is that appealing to the existence of a strong anonymous claim does not block the argument from the Individualist Restriction on Claims if the anonymous claim in question is possessed by Bill.

But if Bill cannot possess a strong anonymous claim in *Fair Lottery*, then by repeating the argument we can show that Ted cannot have such a claim either; neither can Joanna, nor Elizabeth, nor any of the others who will be entered into the lottery. All of these people face a one-in-a-hundred-million risk of death, not a ninety-nine percent risk of death. And these really are all the people—if you're not sure about this, you can check it in the usual way, by lining them up and counting them.[33]

Since there are no other people, there is no individual who faces a ninety-nine percent risk of death, and therefore no individual who could possess an anonymous claim against the ninety-nine percent risk of death. The principle that Only Individuals Have Claims implies that there are no other entities which could possess the strong anonymous claim (or any claim at all). So there is no such claim. Consequently, the appeal to anonymous claims does not allow the ex post view to escape the charge that it fails to respect the separateness of persons.

3.7 The Response from Indeterminacy

It might be objected that my argument against the anonymous claims view presupposes that if a strong anonymous claim against almost-certain death is operative in *Fair Lottery*, then some particular person *determinately* has this claim.[34] (By an operative claim, I just mean a claim that takes effect in the relevant choice situation, whether or not it is relevant.) A proponent of the anonymous claims view may well reject this strong reading of the condition that Only Individuals Have Claims. They could say, with some justification, that since the lottery in *Fair Lottery* is indeterministic, it is true that *somebody* has a strong anonymous

[32] See also Frick 2015a: 196.

[33] At this point it might occur to the reader that, while there are only 99 million people, perhaps these people should be designated by some means other than their names. Perhaps, for instance, one of them should be designated as "the person who will be unlucky in the lottery, if anyone is; otherwise, Bill." Bastian Steuwer (2021) has suggested a strategy like this. I will address this worry in §5.

[34] I am indebted to [Redacted] for suggesting this objection to me.

claim, but the identity of this person is indeterminate. And this may seem to block my argument.

To be more precise, my argument against the anonymous claims view appears to assume the following, stronger reading of the principle that Only Individuals Have Claims:

> **OIHC1**: For any operative claim C, there is some particular person S such that, determinately, S possesses C.

Proponents of the anonymous claims view might justifiably reject this stronger reading, and instead endorse only the following, weaker reading:

> **OIHC2**: For any operative claim C, it is determinately the case that, for some particular person S, S possesses C.

I agree that considerations stemming from the separateness of persons do not force us to endorse OIHC1 over OIHC2. But I deny that this blocks my argument against the anonymous claims view. The reason is that we would like to say that Bill *determinately* has at most a weak claim in *Single-Person Fair Lottery*, else we will have to say that it is indeterminate whether we can drive to the hospital. And of course, what goes for Bill goes for everyone else. Thus, all of the following ninety-nine million (henceforth m) propositions are true. (The subscripts under *Single-Person Fair Lottery* denote who is entered into the lottery.)

(1) It is determinately the case that Bill does not have a strong claim in *Single-Person Fair Lottery*$_{Bill}$

(2) It is determinately the case that Ted does not have a strong claim in *Single-Person Fair Lottery*$_{Ted}$

(3) It is determinately the case that Joanna does not have a strong claim in *Single-Person Fair Lottery*$_{Joanna}$

(4) It is determinately the case that Elizabeth does not have a strong claim in *Single-Person Fair Lottery*$_{Elizabeth}$

$\vdots$

(m) It is determinately the case that Rufus does not have a strong claim in *Single-Person Fair Lottery*$_{Rufus}$

By applying the Individualist Restriction on Claims m times, we then obtain

(1′) It is determinately the case that Bill does not have a strong claim in *Fair Lottery*

(2′) It is determinately the case that Ted does not have a strong claim in *Fair Lottery*

(3′) It is determinately the case that Joanna does not have a strong claim in *Fair Lottery*

(4′) It is determinately the case that Elizabeth does not have a strong claim in *Fair Lottery*

$\vdots$

(m′) It is determinately the case that Rufus does not have a strong claim in *Fair Lottery*

Propositions (1′) to (m′) jointly imply that it is determinately false that any of the ninety-nine million particular individuals entering the lottery in *Fair Lottery* have a strong claim.

Now, on the anonymous claims view, there is a strong claim C, namely the claim against a 99% risk of death, which is operative in *Fair Lottery*. OIHC2 says that it is determinately the case that, if claim C is operative in *Fair Lottery*, at least one of these ninety-nine million particular people has a strong claim; and we have just seen that this consequent is determinately false. We can therefore conclude that it is not the case that C is operative in *Fair Lottery*, which contradicts the anonymous claims view. OIHC2 is therefore sufficient for my argument.

3.8 Taking Stock

Perhaps the ex post view might yet be defended by appealing to another notion of a person's claim, which is distinct from the ex ante, ex post and anonymous claims discussed so far. I cannot think of a plausible candidate notion, but that does not mean that there are no other possibilities. Still, antecedently there seems to be good reason to favour the ex ante understanding of claims. The function of claims is to provide prima facie moral justification to perform or to avoid particular actions, on the basis of how the actions in question affect the balance of risks and benefits faced by each individual. That being the case, it makes sense to think that claims should indeed simply tell *for* or *against* particular courses of action, that they should attach to particular individuals, and that they should be sensitive to each individual's overall assessment of the balance of risks. These are the main characteristics of ex ante claims. So, I think that we should understand claims in this way, and that we should adopt the version of the Individualist Restriction which applies to ex ante claims (henceforth the Individualist Restriction on Ex Ante Claims). If we do, this rules out the ex post approach.

We cannot, however, reasonably agree with the standard ex ante view that identified lives (or identified hands, for that matter) should be prioritised over any number of statistical lives. We should instead adopt the position that statistical harms, if there are enough of them, can always outweigh identified harms of the same severity. We will next see that this position, together with the Individualist Restriction on Claims, settles the question we started with, in that it implies that the New Yorkers' claims are relevant in *Deaths vs Risks of Death*. We can call this the *Argument from the Relevance of Statistical Harms*. Afterwards, we shall see that taking statistical harms seriously really can be reconciled with the separateness of persons. I will demonstrate this by setting out a new variant of ex ante Partial Aggregation, called the *Ex Ante Claims, Ex Post Relevance View*, which does exactly that.

4 A Variant of the Ex Ante View

4.1 The Argument from the Relevance of Statistical Harms

Let H be any harm. Consider the following generalisation of the Identified vs Statistical Lives case:

> **Identified vs Statistical Harms**: Eight million New Yorkers will be entered into an indeterministic lottery which will randomly select eight hundred of them. You can either:

(IH) Prevent Ted, the Californian, from suffering harm H, or

(SH) Prevent the eight hundred randomly selected New Yorkers from suffering harm H.

Any reasonable theory of Partial Aggregation should imply that you ought to choose (SH) here, and that this is so because you thereby satisfy the relevant claims of the New Yorkers.[35] In §2 I argued for this position in the specific case where the harm H is death, but these arguments work just as well for other sorts of harms. So let us take it for granted. We will also need the Individualist Restriction on Ex Ante Claims.

The Individualist Restriction on Ex Ante Claims implies that, since each New Yorker has a relevant claim in *Identified vs Statistical Harms*, they would retain a relevant claim of the same strength against the imposition of the same risk of harm, even if they were the only one entering the lottery (provided they face precisely the same risk of harm H). This is because the relevance and strength of an individual's claim cannot depend on whether or not their claim can be jointly satisfied together with the claims of other individuals. The Individualist Restriction on Ex Ante Claims further implies that the New Yorkers' claims against risks of harm H retain their relevance and strength when they are concentrated into a single state of nature, so that there is a one-in-ten-thousand chance that all New Yorkers will suffer H, and otherwise none of them will suffer any harm. This is exactly the sort of situation we face in *Deaths vs Risks of Death*.

What we have shown is that the relevance of statistical harms in cases like *Identified vs Statistical Harms* and the Individualist Restriction on Ex Ante Claims jointly imply that

> **Small Risks Are Relevant**: Other things being equal, a claim against even a small risk of harm H is always relevant to a claim against certainty of harm H.

It's also worth mentioning that *Small Risks Are Relevant* enjoys independent support from one commonly held criterion for treating certain claims as irrelevant. This is the idea that a claim is irrelevant only if it would be *disrespectful* to appeal to it in justifying a decision not to satisfy a stronger claim with which it competes.[36] When the small claims in question are claims against trivial harms like minor headaches or sore throats, it plausibly is disrespectful to take these claims into account. When the claims are "small" only in the sense that they involve small risks of being seriously harmed, however, matters appear quite different. Recall our original case of *Deaths vs Risks of Death*. Suppose you do prevent the one-in-ten-thousand risk that all eight million New Yorkers will die, rather than saving the one hundred Texans from certain death. Perhaps, if my arguments have been unsound, you thereby do something wrong. But it does not seem to me that what you do is *disrespectful* to the Texans.

So, small risks of serious harms are relevant to certainties of serious harms. We should therefore rescue the New Yorkers (if there are enough of them) in *Deaths vs Risks of Death*. This settles the question we started with. But two new questions arise in its place. First, can taking statistical harms seriously be reconciled with respecting the separateness of persons, or are these positions in some way inconsistent? Second, how can it be that a small risk

[35] Recall that a pluralist ex ante view will still need to include a wellbeing-based component which recommends preventing many statistical harms over preventing fewer identified harms.

[36] See Kamm 2007: 34, Voorhoeve 2017: 153.

of a serious harm is strong enough to remain relevant when it competes with certainty of the same sort of harm, yet not strong enough to render claims against comparatively trivial harms irrelevant in cases like *Single-Person Fair Lottery*? I will answer these in reverse order.

4.2 Three Concepts of the Strength of a Claim

We have seen that claims against small risks of serious harms are strong, in the sense that they are capable of remaining relevant even when they compete with claims against certainty of the same sorts of harms. Yet we have also seen that they are weak, in the sense that claims against comparatively trivial harms remain relevant when they compete with a single claim against a small risk of a serious harm, as in cases like *Single-Person Fair Lottery*. The key to reconciling these apparently contradictory statements is to notice that these two senses in which a claim can have strength really *are* two senses in which a claim can have strength. In fact, there are three distinct conceptual roles played by the notion of strength in the Aggregate Relevant Claims View. It will help to give them names:

(i) The *aggregative strength* of a claim corresponds to the extent to which a claim gives us moral reason to perform a given action, provided that claim is relevant.

(ii) The *relevance-retaining strength* of a claim corresponds to the extent to which a claim is able to remain relevant, despite competing with other claims.

(iii) The *irrelevance-making strength* of a claim corresponds to the extent to which a claim is able to render other claims irrelevant.

The aggregative strength of a claim is the notion of strength at play when we are enjoined by the Aggregate Relevant Claims View to "maximise the satisfaction of strength-weighted relevant claims". When it is said that "a claim is relevant when it is strong enough, compared to the strongest competing claim", this means (using our new vocabulary) that "a claim is relevant when its relevance-retaining strength is great enough, compared to the competing claim with greatest irrelevance-making strength".

The traditional view is that all three notions of strength coincide. I believe that this traditional view is mistaken. It seems to me that the cases we have considered provide considerable evidence that the relevance-retaining strength and the irrelevance-making strength of a claim can come apart. It is implausible that claims against small risks of serious harms should only have weak relevance-retaining strength, else, given the Individualist Restriction on Ex Ante Claims, we would have to prioritise identified harms over any number of statistical harms of the same sort. Yet it is also implausible that claims against small risks of serious harms should have great irrelevance-making strength, else we would have to satisfy single claims against small risks of serious harms rather than large numbers of claims against certainty of comparatively trivial harms (which might nevertheless correspond to greater expected losses of wellbeing). We should conclude that these two notions of claim strength *do* sometimes come apart, so that claims against small risks of serious harms can have great relevance-retaining strength, but only very limited irrelevance-making strength.

4.3 The Ex Ante Claims, Ex Post Relevance View

The position I really mean to defend in this paper is simply the view that *Small Risks Are Relevant*. However, using the different notions of claim strength I have just outlined, I

shall now provide a more complete theory of the aggregation of people's claims in cases of risk. This is partly in the spirit of providing a suitable target for potential objections, and partly in order to satisfy the reader that the Argument from the Relevance of Statistical Harms does not proceed from inconsistent premises. The theory will satisfy the condition that *Small Risks Are Relevant*, as well as our conditions stemming from the separateness of persons: the Individualist Restriction on Claims and the principle that Only Individuals Have Claims. It will also imply that we should prevent large numbers of statistical harms rather than small numbers of identified harms of the same sort.

In order to provide this theory, I will need to take a stand on whether priority should be given to identified over statistical victims. In what follows, I shall take the view we should not give priority to preventing identified over statistical harms in a straight choice between the two. More generally, I shall need to decide how the aggregative and irrelevance-making strengths of claims should be determined; I shall assume that they are proportional to the expected losses of wellbeing with which they are associated. These decisions are mostly taken with an eye to simplicity. It is possible to formulate other sorts of theories satisfying my desiderata—in particular, it is possible to give *some* limited priority to identified over statistical victims—but these theories are more complicated.

My proposal, then, is that the aggregative and irrelevance-making strengths of ex ante claims should be given by their contribution to expected wellbeing, but that the relevance-retaining strengths of claims should depend instead on the size of the harm to which the individual in question *might* be subjected, if things go badly. Call this the *Ex Ante Claims, Ex Post Relevance View*. It works as follows:

The Ex Ante Claims, Ex Post Relevance View:

(i) Individuals have ex ante claims against being subjected to risks of harm. The aggregative and irrelevance-making strengths of these claims correspond to the extent of their expected losses of wellbeing.[37] The relevance-retaining strengths of these claims correspond to the maximum extent to which they might be harmed.

(ii) Ex ante claims compete when they cannot be jointly satisfied.

(iii) An ex ante claim is *relevant* if its relevance-retaining strength is great enough, compared to the individual competing ex ante claim with greatest irrelevance-making strength. Otherwise it is irrelevant.

(iv) We should choose whichever action maximises the total satisfaction of aggregative-strength-weighted relevant ex ante claims.

To see how this view works, it may help to consider how it applies to some of our previous examples. In both *Deaths vs Risks of Death* and *Identified vs Statistical Lives*, we must trade off a smaller number of claims against certain death against a larger number of claims against small risks of death. On the Ex Ante Claims, Ex Post Relevance View, the relevance-retaining strength of a claim against a small risk of death is given by the magnitude of the harm of death, which is *not* discounted by the small probability with which it occurs. It therefore implies that in both of these examples, the claims against small risks of death

[37] If an expected harm is counterbalanced by a larger expected benefit, the full package will not lead to an expected loss of wellbeing, and therefore the individual in question will not have a claim against receiving this package. One might want to modify the Ex Ante Claims, Ex Post Relevance View to incorporate a harm-benefit asymmetry, but I leave such complications aside here.

are relevant. Since the total aggregative strength of the claims against small risks of death exceeds the total aggregative strength of the claims against certain death, the Ex Ante Claims, Ex Post Relevance View implies that we should save the many in both cases.

In *Unfair Lottery*, we must either satisfy a single claim against almost-certain death, or a large number of claims against broken fingers. On the Ex Ante Claims, Ex Post Relevance View, the irrelevance-making strength of the claim against almost-certain death is great, since this is given by the probability-discounted magnitude of the harm of death. Since the probability of death is high, the discounting is only slight. The irrelevance-making strength of this claim is enough to overpower the relevance-retaining strengths of the claims against broken fingers (which are given by the undiscounted magnitudes of these harms). Thus, the broken fingers are counted as irrelevant and we should prevent the almost-certain death.

In *Fair Lottery*, however, things are different. The claims against broken fingers now compete with many claims against small risks of death. While the small risks of death have great relevance-retaining strength, their irrelevance-making strength is only slight, since this is given by the probability-discounted magnitude of the harm of death. The Ex Ante Claims, Ex Post Relevance View therefore implies that the claims against broken fingers are relevant, and that we should therefore prevent them if there are enough of them.

Hopefully that is enough to get a feel for how the Ex Ante Claims, Ex Post Relevance View works. I will soon set out and respond to two of the most serious objections to the view. But first I think it's worth mentioning some of its advantages. Its most important advantage, of course, is the one it was designed for: it takes statistical harms seriously while still respecting the separateness of persons. Unlike some ex post views, it also has the welcome implication that (at least in two-option cases) we may do what is in each person's ex ante interests even if this will foreseeably result in some people being seriously harmed, for in such cases nobody has an overall ex ante claim against receiving the whole package of benefits and harms. A final advantage is that if, as some philosophers have argued, considerations of chaos mean that any action we take is bound to eventually lead to *someone* being seriously harmed, the Ex Ante Claims, Ex Post Relevance View does not then absurdly imply that all moderate benefits should count as irrelevant even when no serious harms are immediately at issue.[38] This is because in such cases we make no dramatic difference to any particular person's ex ante prospects.

5 Objections

Let me now turn to two objections to the Ex Ante Claims, Ex Post Relevance View. The first objection concerns an implication of the view which I expect many people will find objectionable: it recommends preventing a large number of trivial harms rather than a small number of serious *statistical* harms. The second objection, which I call the "Problem of Designation", concerns the way in which we should designate people for the purposes of assigning ex ante claims.

[38]See Mogensen and MacAskill 2021 on the paralysis problem. Charlotte Unruh (2023) has shown that the paralysis problem for the proponent of the doctrine of doing and allowing harm can be solved by imposing constraints against doing harm only when we make thereby make people significantly worse off in expectation. The way in which the Ex Ante Claims, Ex Post Relevance View solves the parallel problem for Partial Aggregation is roughly analogous.

5.1 Trivial Harms vs Statistical Deaths

Consider the following variant of *Fair Lottery*.

> **Fair Lottery (trivial harms)**: If nothing is done, one Californian will be chosen at random from the entire population and killed. Meanwhile, everyone from Texas will suffer a comparatively trivial harm, such as a migraine. You can prevent the death or the trivial harms, but not both.

Assume that the population of Texas is large enough that the trivial harms to the Texans, in aggregate, would outweigh the death of one Californian (or replace the population of Texas with some larger population). The Ex Ante Claims, Ex Post Relevance View, along with the standard ex ante view, then implies that we should prevent the trivial harms in *Fair Lottery (trivial harms)*. This is because the claims of the Californians have only weak irrelevance-making strength in light of the small probability of death that each individual faces. The claims of the Californians are therefore incapable of rendering the weak claims of the Texans irrelevant. Since the claims of the Texans have, by stipulation, greater total aggregative strength, we then maximise the satisfaction of relevant claims by preventing the trivial harms.

Is it a problem that the Ex Ante Claims, Ex Post Relevance View says this? Joe Horton thinks that it is: any theory of Partial Aggregation which implies that we should prevent the trivial harms in such cases, he says, must be inadequate.[39] He gives two closely related arguments to this effect. The first is that it would be bizarre to hold that one should give absolute priority to saving identified lives over preventing trivial harms, and yet at the same time endorse preventing trivial harms rather than statistical deaths.[40] The second is that these two positions are in tension with each other in the sense that we cannot robustly capture the intuition that we should prevent deaths rather than any number of trivial harms—which is shared by all proponents of Partial Aggregation—if we accept that we should sometimes prevent trivial harms rather than *statistical* deaths.[41] I think that both arguments can be answered.

The first argument can be addressed by providing a moral justification for favouring identified over statistical lives in the sorts of cases at issue. In general, it is not bizarre to treat two decision situations differently if there is a morally relevant difference between them. And there is indeed such a difference in the present case: when a particular person is certain to die unless they are helped, *that person* has a powerful complaint against being left high and dry. In contrast, when it is certain that somebody will die, but it is unclear who that will be, no particular individual has this sort of powerful complaint. It is precisely this difference which the Ex Ante Claims, Ex Post Relevance View (and the standard ex ante view) points to.

Moreover, the moral significance of this difference is underwritten by the Individualist Restriction on Ex Ante Claims. In the single-person variant of *Fair Lottery (trivial harms)*, the small risk of death to one person presumably does not render the trivial harms irrelevant, else we would not be permitted to drive to the pharmacy to buy aspirin.[42] The Individualist

[39] See also Rüger 2018: 248–250.
[40] Horton 2020: 517, 528.
[41] Horton 2020: 516, 528–529.
[42] See Walen 2020: 82–83 or Lazar 2018: 118 on this sort of case.

Restriction on Ex Ante Claims then implies that the trivial harms must remain relevant in the original, many-person variant of the case.

Horton's second argument can be answered in the following way. We can argue that there is indeed an intuitive difference between, on the one hand, cases where identified deaths are traded off against trivial harms or benefits, and, on the other, cases where statistical deaths are traded off against trivial harms or benefits. If that is so, the intuition that we should prevent deaths in the former sort of case is not undermined by the claim that we should sometimes prevent trivial harms in the latter sort of case.

To see that there is an intuitive difference between cases involving identified deaths and cases involving statistical deaths, consider the manufacture of in-flight entertainment systems. The manufacture of these systems, like all manufacturing activity, involves direct and indirect increases in air pollution. Air pollution is deadly: on one estimate, it caused around 3.3 million deaths globally in the year 2015.[43] Ongoing annual deaths, while they probably vary somewhat from year to year, presumably remain roughly in that ballpark.

How many of these deaths might be attributed to the manufacture of such systems? It is not easy to say with any precision. But the current total size of the global market for in-flight entertainment systems is generally estimated as being between three and ten billion US dollars per year, or around 0.003% to 0.01% of the global economy. If air pollution is proportional to economic activity, these figures would imply that the production of in-flight entertainment systems is responsible for somewhere between 100 and 330 deaths per year. The real numbers could, of course, be substantially different: it may very well be, for instance, that the high-tech components which go into in-flight entertainment systems involve much less pollution per dollar than, say, a car. But the point is this: it is extremely plausible that *at least one* statistical life has been, and will be, lost as a result of air pollution associated with the manufacture of these systems.

Granting that this is indeed a fact, what should we make of it? If there is no morally relevant difference between identified and statistical deaths, and if we accept Partial Aggregation, it seems that we ought to stop manufacturing in-flight entertainment systems. After all, the benefits of these systems are clearly trivial compared to certain death—even for frequent flyers. But even if this conclusion is correct, I take it that it is far from obvious.[44]

Imagine, however, that the facts were somewhat different. Suppose that the manufacture of in-flight entertainment systems involved not air pollution, but the leakage of a certain chemical into the water supply. Suppose that this leakage would be impossible to avoid so long as we continue to manufacture these systems, and that it would be sure to kill several known individuals, who have a particular rare allergy to that chemical, per year. In *this* case, I submit, the claim that we ought to stop manufacturing in-flight entertainment systems seems significantly more plausible. My point is that whether or not there is a genuine moral difference between these two cases, at the very least there is an intuitive difference. So, *pace* Horton, preventing trivial harms rather than statistical deaths does not preclude capturing the key anti-aggregative intuition that we should prevent identified deaths rather than any number of trivial harms.

[43] See Lelieveld et al. 2015

[44] Norcross (1998: 159–167) uses an example like this to argue that the non-aggregative intuition that trivial harms can never outweigh deaths is not intuitive after all. If there is a morally relevant difference between cases involving identified and statistical deaths, as the Ex Ante Claims, Ex Post Relevance View and the standard ex ante view predict, this also serves to defuse Norcross's objection.

There is one further objection that could be levied against the implication that we should prevent the trivial harms in *Fair Lottery (trivial harms)*, which is that this implication is simply implausible. I believe that the case of in-flight entertainment systems, and other similar cases, should make us think twice about this sort of intuition. But if you remain gripped by the thought that we must prevent statistical deaths rather than any number of trivial harms, you might take the lesson here to be that we had better reject the Individualist Restriction on Ex Ante Claims. In that case, the traditional view that the rejection of Full Aggregation can be justified by an appeal to the separateness of persons may need to be revised.

5.2 The Problem of Designation

The problem of designation can be illustrated by the following case.

> **Death by Designation**: A villain has released a virus which has infected every inhabitant of a large city, all of whom are known to you by legal name. One of those people is Superman. Superman stands to die from the virus. The other victims of the virus will be substantially harmed, but in a way that falls slightly below the threshold of relevance when compared to certain death. You are certain that Superman is either Clark Kent or Bruce Wayne (both of whom live in the city), but as far as you know, these possibilities are equally likely. You can cure Superman or the other inhabitants of the city, but not both.

In order to decide what to do on the Ex Ante Claims, Ex Post Relevance View (and also on the standard ex ante view), we need to assign ex ante claims to the moral patients at issue. The problem is that there seem to be two ways of doing so. We could assign a claim against a 50% probability of death to Clark Kent, assign a similar claim to Bruce Wayne, and assign claims against the substantial harm to the remaining people. Having assigned claims in this way, the Ex Ante Claims, Ex Post Relevance View will imply that the substantial harms of the many are relevant, since they are strong enough to compete with one-in-two risks of death. We will then conclude that we should cure the many, rather than prevent Superman's death.

An alternative way of looking at it is to note that *Superman* is certain to die. We may then assign a claim against certain death to Superman, and assign claims against the substantial harms to all of the remaining people. Since these harms are, by stipulation, irrelevant to certain death, we will then conclude that we must rescue Superman.

This raises two related problems for the Ex Ante Claims, Ex Post Relevance View (and, indeed, for any ex ante view). The first problem is that these moral judgements can't both be correct. If what we ought to do depends on the assignment of ex ante claims, there must therefore be a principled way of designating individuals for the purposes of assigning these claims. But it is unclear at the outset what this principled method of designation might be.[45]

The second problem is that even if there is a principled way of choosing designators, perhaps the right way to do it undermines the arguments against the ex post view—in particular,

[45] Anna Mahtani (2017) considers a similar problem for the ex ante pareto principle. She argues that there is no principled way of assigning designators, and that therefore the ex ante pareto principle, as usually stated, is ill-defined.

the argument against the anonymous claims view—which I pressed in §3. Bastian Steuwer has proposed an account of designation like this, which we can call the Ranking Account.[46] According to the Ranking Account, we should designate individuals as "the person who loses most if X is done", "the person who loses second-most if X is done", and so on. We then assign ex ante claims using these designators. In *Death by Designation*, for instance, we would assign a strong claim against certain death to *the person who stands to die from the virus* (i.e., Superman). The Ranking Account underwrites the anonymous claims view.

I don't have a general solution to these two problems of designation. But I do think that they can be answered if we restrict our attention to an important class of cases, namely those in which the moral agent and moral patients are in the same epistemic position.[47] (Note that this condition is not satisfied in *Death by Designation*, since Superman knows more than you do.) My proposal is that when the moral agent and all moral patients have the same information, one's choice of designators should satisfy the following coherence requirement, which ensures that one's choice of designators matches up with each person's assessments of the risks they face from *their* perspective:

> **Coherence Requirement on Designator Choice**: It is admissible to use a system of designators $D_1, \ldots, D_n$ for the purposes of assigning ex ante claims only if
>
> (i) each person S_i who assesses that they themselves face a risk is picked out by exactly one designator D_i, and
>
> (ii) for each designator D_i picking out person S_i, the ex ante claim assigned to D_i agrees with S_i's assessment of the risks they face.

Provided everyone has to hand an assessment of the risks that they face in a given choice, the Coherence Requirement uniquely determines the assignment of ex ante claims. For suppose we have two systems of designators $D_1, \ldots D_n$ and $D'_1, \ldots D'_m$, both satisfying the Coherence Requirement. Condition (i) implies that the number of designators must match the number of people who assess that they face risks; thus $n = m$. Condition (ii) requires that, if D_i and D'_j pick out the same person S, they must assign exactly the same ex ante claim, since both ex ante claims must agree with S's first-personal assessment of the risks they face. Hence, the two systems of designators assign the same pattern of ex ante claims.

Adopting the Coherence Requirement therefore solves the first problem of designation. It also solves the second problem of designation, because it requires us to assign the sorts of ex ante claims we would intuitively expect, rather than anonymous claims of the sort that fall out of the Ranking Account.

The question remains, of course, whether there are good reasons to endorse the Coherence Requirement over its competitors, such as the Ranking Account. I believe that there are. One of the key motivating thoughts behind the partially aggregative approach is that, in order to respect the separateness of persons, a moral theory must attach special significance to the stakes of a given situation from the first-personal perspective of individuals, since it is only individuals (and not aggregates of individuals) that *have* a first-personal perspective. As Alex Voorhoeve puts it, the non-aggregative component of the Aggregate

[46] See Steuwer 2021.

[47] The importance of distinguishing cases in which no one is able to determine who stands to be harmed from cases in which the moral agent is unable to determine this, but the moral patients are, has also been highlighted by Alec Walen (2020: 65).

Relevant Claims View "involves imaginatively placing oneself, one person at a time, in the position of each person who has a claim *and viewing the situation through her eyes* [emphasis added]".[48] This process of sympathetic identification supports the Coherence Requirement. A Contractualist approach to Partial Aggregation likewise supports the Coherence Requirement. This is because Contractualism requires that we attend to the reasonable objections which individuals can raise against a principle licensing a particular course of action.[49] Such objections must surely be raised from the standpoint of the individual concerned.[50]

I conclude that proponents of Partial Aggregation have good theoretical reasons to endorse the Coherence Requirement. Doing so solves both problems of designation, at least when it comes to those decision situations in which the moral agent and all moral patients share the same information. I leave open the question of how we should deal with more complicated cases, such as *Death by Designation*.

I also note in passing that both the Problem of Designation and the problem of trivial harms outweighing statistical deaths apply to the standard ex ante view as well as the Ex Ante Claims, Ex Post Relevance View. If these problems show that the Ex Ante Claims, Ex Post Relevance View is wrong, they show that ex ante views in general are wrong. It therefore seems to me that the Ex Ante Claims, Ex Post Relevance View is, at least, a more plausible variant of the ex ante view than the standard version.

6 Concluding Remarks

I have argued that the standard ex ante and ex post approaches to Partial Aggregation should be rejected. The standard ex ante view should be rejected because it fails to take statistical harms seriously. The ex post approach should be rejected because it fails to respect the separateness of persons. These two desiderata for a theory of Partial Aggregation turn out to imply that small risks of serious harms are always relevant to the certainty of harms of the same sort. That this is so was the main thing I intended to show in this paper.

The significance of this result depends to a great extent on whether we should indeed accept the Individualist Restriction on Claims, since this principle is really the crux of my main arguments. So, do we sacrifice anything of importance if we reject this Individualist Restriction? I think that we do.

Departures from full aggregation are usually defended on the grounds that a proper appreciation of the separateness of persons requires us to pay special attention to the significance of losses for the particular individuals who experience them. In particular, one needs to have a response to a proponent of Full Aggregation who says, in response to a case where we can either prevent one death or many headaches, that "the claims of those who stand to suffer headaches should be counted because, taken together, they are far stronger than the claim of the one person who stands to die". I do not see how one could respond to this in any other way than to insist that whether a claim should be counted depends only on its individual strength in comparison to the individual claims with which it competes, and

[48] Voorhoeve 2014: 68.

[49] See Scanlon 1982: 110; 1998: 153.

[50] My argument here does not by itself show that Contractualists must reject the ex post approach to Partial Aggregation, for it says nothing regarding the time at which an individual's complaint should be assessed and taken into account.

not on the aggregate strength of the claims with which it can be jointly satisfied. But to respond in this way is to appeal to the Individualist Restriction on Claims.

The cost of rejecting the Individualist Restriction on Claims, then, is that this principle appears to get to the heart of why Full Aggregation seems to many people objectionable. That is not to say that one *cannot* reject both Full Aggregation and the Individualist Restriction on Claims: there is nothing logically inconsistent about this combination of positions. But if we cannot justify the rejection of Full Aggregation via an appeal to the Individualist Restriction on Claims, we need an alternative justification. I leave the problem of providing such a justification as a friendly challenge to ex post theorists.

References

Beck, Laurie F., Ann M. Dellinger, and Mary E. O'Neil (2007). "Motor Vehicle Crash Injury Rates by Mode of Travel, United States: Using Exposure-Based Methods to Quantify Differences". In: *American Journal of Epidemiology* 166.2, pp. 212–218.

Curran, Emma J. (2022). "Longtermism, Aggregation, and Catastrophic Risk". GPI Working Paper No. 18-2022.

Frick, Johann (2015a). "Contractualism and Social Risk". In: *Philosophy & Public Affairs* 43.3, pp. 175–223.

– (2015b). "Treatment versus Prevention in the Fight against HIV/AIDS and the Problem of Identified versus Statistical Lives". In: *Identified versus Statistical Lives: An Interdisciplinary Perspective*. Ed. by I. Glenn Cohen, Norman Daniels, and Nir Eyal. Oxford: Oxford University Press. Chap. 13, pp. 182–202.

Gustafsson, Johan E. (2015). "Sequential Dominance and the Anti-Aggregation Principle". In: *Philosophical Studies* 172.6, pp. 1593–1601.

Halstead, John (2016). "The Numbers Always Count". In: *Ethics* 126.3, pp. 789–802.

Horton, Joe (2020). "Aggregation, Risk, and Reductio". In: *Ethics* 130, pp. 514–529.

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

Kumar, Rahul (2015). "Risking and Wronging". In: *Philosophy & Public Affairs* 43.1, pp. 27–51.

Lazar, Seth (2018). "Limited Aggregation and Risk". In: *Philosophy & Public Affairs* 46.2, pp. 117–159.

Lelieveld, J. et al. (2015). "The Contribution of Outdoor Air Pollution Sources to Premature Mortality on a Global Scale". In: *Nature* 525, pp. 367–371.

Lewis, David (1981). "Causal Decision Theory". In: *Australasian Journal of Philosophy* 59.1, pp. 5–30.

Mahtani, Annals (2017). "The Ex Ante Pareto Principle". In: *The Journal of Philosophy* 114.6, pp. 303–323.

McCarthy, David (2006). "Utilitarianism and Prioritarianism I". In: *Economics and Philosophy* 22.3, pp. 335–363.

– (2008). "Utilitarianism and Prioritarianism II". In: *Economics and Philosophy* 24.1, pp. 1–33.

Mogensen, Andreas and William MacAskill (2021). "The Paralysis Argument". In: *Philosophers' Imprint* 21.15, pp. 1–17.

Norcross, Alastair (1998). "Comparing Harms: Headaches and Human Lives". In: *Philosophy & Public Affairs* 26.2, pp. 135–167.

Nozick, Robert (1969). "Newcomb's Problem and Two Principles of Choice". In: *Essays in Honor of Carl G. Hempel: A Tribute on the Occasion of his Sixty-Fifth Birthday*. Ed. by Donald Davidson, Jaakko Hintikka, and Gabriël Nuchelmans. Dordrecht: Springer Netherlands, pp. 114–146.

Otsuka, Michael (2006). "Saving Lives, Moral Theory, and the Claims of Individuals". In: *Philosophy & Public Affairs* 34.2, pp. 109–135.

– (2015). "Risking Life and Limb: How to Discount Harms by Their Improbability". In: *Identified versus Statistical Lives: An Interdisciplinary Perspective*. Ed. by I. Glenn Cohen, Norman Daniels, and Nir Eyal. Oxford: Oxford University Press. Chap. 5, pp. 77–93.

Parfit, Derek (2003). "Justifiability to Each Person". In: *Ratio* 16.4, pp. 368–390.

Reibetanz, Sophia (1998). "Contractualism and Aggregation". In: *Ethics* 108.2, pp. 296–311.

Rüger, Korbinian (2018). "On Ex Ante Contractualism". In: *Journal of Ethics and Social Philosophy* 13.3, pp. 240–258.

Scanlon, Thomas M. (1982). "Contractualism and Utilitarianism". In: *Utilitarianism and Beyond*. Ed. by Amartya Sen and Bernard Williams. Cambridge: Cambridge University Press. Chap. 5, pp. 103–128.

– (1998). *What We Owe to Each Other*. Cambridge, Massachusetts: Harvard University Press.

Steuwer, Bastian (2021). "Contractualism, Complaints, and Risk". In: *Journal of Ethics and Social Philosophy* 19.2, pp. 111–147.

– (2022). "Limits to Aggregation and Uncertain Rescues". In: *Utilitas* 34.1, pp. 70–83.

Taurek, John M. (1977). "Should the Numbers Count?" In: *Philosophy & Public Affairs* 6.4, pp. 293–316.

Unruh, Charlotte F. (2023). "The Constraint Against Doing Harm and Long-Term Consequences". In: *Journal of Moral Philosophy* 20.3–4, pp. 290–310.

Voorhoeve, Alex (2014). "How Should We Aggregate Competing Claims?" In: *Ethics* 125.1, pp. 64–87.

– (2017). "Why One Should Count Only Claims with which One Can Sympathise". In: *Public Health Ethics* 10.2, pp. 148–156.

Walen, Alec (2020). "Risks and Weak Aggregation: Why Different Models of Risk Suit Different Types of Cases". In: *Ethics* 131, pp. 62–86.