


ARTICLE

Limits to Aggregation and Uncertain Rescues

Bastian Steuwer 

Center for Population-Level Bioethics, Rutgers University, New Brunswick, NJ, USA
Email: steuwer@cplb.rutgers.edu

Abstract

Limited aggregation holds that we are only sometimes, not always, permitted to aggregate. Aggregation is permissible only when the harms and benefits are relevant to one another. But how should limited aggregation be extended to cases in which we are uncertain about what will happen? In this article, I provide a challenge to *ex post* limited aggregation. I reconstruct a precise version of *ex post* limited aggregation that relies on the notion of *ex post claims*. However, building a theory of limited aggregation based on *ex post* claims leads to a dilemma. This shows that *ex post* limited aggregation is currently far away from being a well-defined alternative, strengthening the case for *ex ante* limited aggregation.

Consider the following (*Anne's Rescue*): Anne is a miner who is trapped in a mineshaft. We can launch a rescue mission that will, with certainty, bring Anne to daylight. If we fail to launch the rescue mission, then Anne will surely die in the mineshaft. However, undertaking the rescue mission has an opportunity cost. Instead of paying for the rescue mission we could use the resources to cure the sore throats of a very large number of people. What should we do? To many it seems that we should save Anne's life. The sore throats are not the right kind of consideration that can outweigh what is at stake for Anne. Regardless of how many sore throats we can cure, we should always save a single life rather than the sore throats. The sore throats do not add up to anything that is of greater moral significance than Anne's life.

However, few actual cases are of this sort. In *Anne's Rescue* we know with certainty what will happen if we launch the rescue mission and what will happen if we provide the pain relief. In the real world, we very often face situations in which we are unsure about the results of our action. How should we think about cases like *Anne's Rescue* in circumstances of uncertainty?

The sentiment that sore throats do not add up to the moral significance of a single life expresses skepticism about the permissibility of aggregating harms across different individuals. Such skepticism can take different forms. A more radical form, which I call *no aggregation*, holds that we should engage in pairwise comparisons between different individuals and never save a person who has a less strong claim to our aid.¹ According

¹No aggregation is easily confused with “numbers skepticism”, the view that we have no duty to save the greater number. However, it is possible to justify a duty to save the greater number without aggregation, for © The Author(s), 2021. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted re-use, distribution, and reproduction in any medium, provided the original work is properly cited.

to “no aggregation”, we should, for example, not save a very large number of people from paraplegia over a single person from death. Paraplegia, we can assume, is a substantial harm even if it is much less bad than death for the individual. A less radical form of aggregation skepticism, which I call *limited aggregation*, holds that while it is permissible that the numbers count in deciding whom to save in some trade-offs, in other trade-offs the relative numbers should not count.² For example, the numbers can count only in trade-offs between harms that are relevant to one another. Limited aggregation can then hold that the numbers count in the trade-off between life and paraplegia, but that the numbers do not count in the trade-off between life and sore throats. I will focus on limited aggregation.

How should limited aggregation be extended to cases in which we are uncertain about what will happen? This is the question I want to address. In Section I, I begin by drawing the distinction between *ex ante* and *ex post* limited aggregation. *Ex post* limited aggregation is seldom well defined in the literature and I make the attempt to provide three different ways of understanding *ex post* views. For the rest of the article, I focus on one initially promising approach that makes use of the notion of *ex post claims*. However, building a theory of limited aggregation based on *ex post* claims leads to a dilemma. I explain both horns of the dilemma in Sections II and III. This result puts pressure on the critics of *ex ante* limited aggregation. It shows that it is more difficult than one may initially think to develop a *ex post* limited aggregation as a coherent and minimally plausible alternative. The absence of such an alternative strengthens the case for interpreting limited aggregation in an *ex ante* manner.

A word on terminology: I will be talking about claims to rescue instead of using the language of complaints. But my arguments can easily be translated into the language of complaints. A complaint is the negative flipside of a claim. A strong claim to be rescued means that there would be a strong complaint if rescue is denied.

I. *Ex ante* and *ex post* limited aggregation

Let us return to the case in which there is uncertainty about whether we can save Anne. One way to resolve this case is to point out that if we try to save Anne, then there is a given chance that we will indeed end up saving Anne. But likewise, there is also the chance that our help will be futile. To account for this, we have to discount Anne’s claim to be saved by the probability that we will not be able to help Anne. Anne’s

example, by adopting a Leximin decision procedure. For a discussion on non-aggregative arguments to save the greater number see Michael Otsuka, *Saving Lives, Moral Theory, and the Claims of Individuals*, *Philosophy & Public Affairs* 34 (2006), pp. 109–35 (pp. 118–26). Furthermore, some proponents of no aggregation believe that while aggregating claims is not required, aggregative considerations are an intelligible reason that an agent may act upon. See Véronique Munoz-Dardé, *The Distribution of Numbers and the Comprehensiveness of Reason*, *Proceedings of the Aristotelian Society* 105 (2005), pp. 191–217. Other proponents of no aggregation believe that in cases of equally strong claims we should give each person an equal chance. See John M. Taurek, *Should the Numbers Count?*, *Philosophy & Public Affairs* 6 (1977), pp. 293–316.

²F. M. Kamm, *Morality, Mortality*, vol. 1 (New York: Oxford University Press, 1993), pp. 156–61, and *Intricate Ethics* (Oxford: Oxford University Press, 2007), pp. 31–77; T. M. Scanlon, *What We Owe to Each Other* (Cambridge, MA: Belknap Press, 1998), pp. 238–41; David Lefkowitz, *On the Concept of a Morally Relevant Harm*, *Utilitas* 20 (2008), pp. 409–23; Larry Temkin, *Rethinking the Good* (Oxford: Oxford University Press, 2012), ch. 3; Alex Voorhoeve, *How Should We Aggregate Competing Claims?*, *Ethics* 125 (2014), pp. 64–87; and Bastian Steuwer, *Aggregation, Balancing and Respect for the Claims of Individuals*, *Utilitas* 33 (2021), pp. 17–34.

claim would then be determined not by the harm she may eventually suffer, but rather by the prospect of harm. This approach can be called *ex ante limited aggregation*.³ It is *ex ante* because we are looking at the claims individuals can make before any uncertainty is resolved.

A competing approach is the *ex post* approach. The *ex post* approach is mainly defined negatively as a rejection of the *ex ante* perspective. Here is an example that motivates the appeal of the *ex post* approach. Imagine that we do not know whether it is Anne or someone else out of a large group of miners who is trapped in the mineshaft, but we do know that we can help with certainty whoever is trapped. If so, then it appears that we need to heavily discount Anne's claim to be rescued since the chance that we will be saving Anne *in particular* is rather low. (The same holds for everyone else out of the large group of miners.) However, if we look at the eventual outcome rather than the individual prospects, then we see that *ex post* the rescue mission one person will be saved. This suggests a first interpretation of *ex post* limited aggregation.

Actualist ex post limited aggregation. We should choose the action that satisfies the claims of the group with the strongest aggregate relevant claims. Claims are determined in accordance with the actual result of the risky action.

The actualist *ex post* view takes seriously the idea that we should be guided by the outcome rather than by the prospects individuals are facing.⁴ However, it has a rather obvious drawback. It makes the permissibility of risky actions dependent on how the risk turns out.⁵ As such it is useless in helping us to decide what we should do if we do not know how the risk will turn out. We need an *ex post* view that we can apply from the epistemic position we have before acting.

Reconsider then the case in which we do not know whether Anne or another miner is trapped. A natural answer in this case would be that what should matter to us is that *some* miner is trapped. Someone has, or will have, a claim that we should rescue him or her.⁶ We can deduce this from the fact that we have knowledge of the outcome

³Most of the discussion of *ex ante* views that are skeptical of aggregation has taken place with respect to contractualist moral theory. See in particular Aaron James, Contractualism's (Not So) Slippery Slope, *Legal Theory* 18 (2012), pp. 263–92; S. D. John, Risk, Contractualism, and Rose's "Prevention Paradox", *Social Theory and Practice* 40 (2014), pp. 28–50; Rahul Kumar, Risking and Wronging, *Philosophy & Public Affairs* 43 (2015), pp. 27–51; Johann Frick, Contractualism and Social Risk, *Philosophy & Public Affairs* 43 (2015), pp. 175–223; and Bastian Steuwer, Contractualism, Complaints, and Risk, *Journal of Ethics & Social Philosophy* 19 (2021), pp. 111–47.

⁴This is embraced most clearly by Sune Holm, The Luckless and the Doomed: Contractualism on Justified Risk-Imposition, *Ethical Theory and Moral Practice* 21 (2018), pp. 231–44. A number of other proposals come close to this by invoking to an "*ex post* epistemic perspective"; see James Lenman, Contractualism and Risk Imposition, *Politics, Philosophy & Economics* 7 (2008), pp. 99–112; Jussi Suikkanen, *Ex Ante* and *Ex Post* Contractualism: A Synthesis, *Journal of Ethics* 23 (2019), pp. 77–98; and Barbara H. Fried, *Facing Up to Scarcity* (Oxford: Oxford University Press, 2020), pp. 33–37, 59–82.

⁵Some, notably Judith Jarvis Thomson, nevertheless argue that permissibility should be understood in such an objective fashion. See Judith Jarvis Thomson, *The Realm of Rights* (Cambridge, MA: Harvard University Press, 1990), pp. 227–34.

⁶Michael Otsuka, Risking Life and Limb, *Identified versus Statistical Lives* ed. I. Glenn Cohen, Norman Daniels, and Nir Eyal (Oxford: Oxford University Press, 2015), pp. 77–93; Sophia Reibetanz Moreau, Contractualism and Aggregation, *Ethics* 108 (1998), pp. 296–311 (pp. 302–4); and Alec Walen, Risks and Weak Aggregation: Why Different Models of Risk Suit Different Types of Cases, *Ethics* 131 (2020), pp. 62–86 (p. 64).

distribution in which one miner will be dead unless we initiate the rescue mission. This someone could be described as “the person with the strongest claim to the rescue”.

Of course, in many cases we are not certain about what the outcome distribution will be. Just like in the case in which we do not know whether we can save Anne, the more realistic case is one in which we do not know whether we can save someone. What we can do then is discount the claim to be saved by the probability that no one will be saved. We can still speak of the claim that is had by “the person with the strongest claim to the rescue”. This description incorporates whoever it is who has the strongest claim to the rescue in each possible outcome.⁷

Someone’s claim ex post limited aggregation. We should choose the action that satisfies the claims of the group with the strongest aggregate relevant claims. We are concerned with the claims that someone, that is, the person with the strongest claim in each outcome, can raise.

Unlike the actualist *ex post* view, this is an interpretation of *ex post* limited aggregation that can be applied by a decision-maker. It explains nicely one core motivation of the rejection of *ex ante* limited aggregation, namely the idea that the specific identity of a beneficiary should not matter. But this interpretation struggles with a different set of cases. Imagine a risk imposition in which many people are subjected to a small risk of great harm. However, the risks are positively correlated such that either everyone gets hurt, or no one does. If we are concerned with the claim that someone can raise against this risk imposition, then we have to conclude that there is only a heavily discounted claim against the risk imposition. The likelihood that someone will be harmed is not very great. If the claim against the risk imposition is sufficiently small such that it is not relevant to a given opportunity cost like curing sore throats, then we should save the sore throats rather than avert the risk, no matter how many people would be harmed should the harm materialize. A proponent of *ex post* limited aggregation might bite the bullet here, but it seems worthwhile looking for a different interpretation of the *ex post* view that does not have these troubling implications.⁸

The core defect in failing to acknowledge the problem of correlated risks is that focusing only on the claim that someone will be harmed overlooks just how objectionable the worst-case scenario is, not just to one person who might be harmed but to everyone who might be harmed. In the case of the positively correlated risks many people would be severely harmed in this scenario, and this should give this scenario much greater weight than the scenario in which a single person would be harmed. It would be preferable if we had an interpretation of the *ex post* view that considers all these claims. We then need a view that considers all the claims individuals would bring forward in the worst-case scenario. This is the kind of *ex post* view that I want to discuss here.

⁷Otsuka, *Risking Life and Limb*. The view is more precisely reconstructed by Joe Horton, *Aggregation, Complaints, and Risk*, *Philosophy & Public Affairs* 45 (2017), pp. 54–81 (pp. 60–66) and Steuwer, *Contractualism, Complaints, and Risk*, pp. 112–19. An analogy may help here. “The person with the strongest claim to the rescue” is not unlike “the worst-off group” in the application of Rawls’s difference principle. In our rescue cases it may be that different people have the strongest claim to rescue depending on who is in need. Similarly, in the case of the difference principle it may be that different sets of people are the worst-off depending on different configurations of the basic structure of society.

⁸For further discussion on the downsides of this approach to *ex post* limited aggregation, see Horton, pp. 66–70; and Steuwer, *Contractualism, Complaints, and Risk*, pp. 119–29.

More generally, the key idea is that we should think of the claims individuals can bring forward in different possible states of the world. Suppose we are in the state of the world in which the worst-case scenario will happen, then many people have very weighty claims against the risk imposition. But then suppose we are in a different state of the world in which the risk will not materialize, then no one has a claim against the risk imposition. We can call these claims *ex post* claims. Importantly, an *ex post* claim is tied to one state of the world, that is, tied to one way things may turn out to be. *Ex post* claims embody a form of conditional reasoning. They are always claims on the condition that we are in a given state of the world. The idea of the third interpretation of *ex post* limited aggregation is then to use these *ex post* claims as the building blocks of a theory of limited aggregation.

Conditional claims ex post limited aggregation. We should choose the action that satisfies the claims of the group with the strongest aggregate relevant claims. We are concerned with the *ex post* claims individuals have in different states of the world.

I want to focus on conditional claims *ex post* limited aggregation.⁹ Because the idea of *ex post* claims is central to my argument, it may help to consider how *ex post* claims are invoked in Anne's cases. Consider first the example in which we do not know whether Anne or another miner is trapped in the mineshaft. Instead of focusing on the fact that someone, namely the miner in greatest danger, would be left dead by our refusal to start the rescue mission, we can reason as follows. There are only a fixed number of possibilities. Either the trapped miner is Anne, or she is Beth, or she is Carla, and so on. If the trapped miner is Anne, then Anne has a strong claim to the rescue mission, a claim stronger than any appeal to the sore throats we could also cure. If the trapped miner is Beth, then Beth has a strong claim to the rescue mission. If the trapped miner is Carla, then Carla has a strong claim. We should then start the rescue mission not because of the claim of "the miner in greatest danger" but rather because whichever way things are, or turn out, we know that there are strong claims to the rescue mission.¹⁰

While in this particular case the *ex post* claims unanimously favor one course of action, we know that this is not always the case. For example, what if we are unsure whether we are able to help Anne with the rescue mission? Supposing that we are in the state of the world in which Anne can be helped, Anne has a strong claim to the rescue mission. But supposing that Anne cannot be helped, the *ex post* claims of the people with sore throats favor the sore throat relief. Aggregation skeptics need a theory for how to decide cases like these. We need an answer to the question how to trade-off these different *ex post* claims.

If the chance that we cannot help Anne is very small, then we should nevertheless favor the rescue mission even given the off-chance that it turns out that it would have been better to give sore throat relief had we known what will happen. If, however, the chance that the rescue mission will be futile is very great, then we should at some

⁹For arguments that invoke conditional claims of this kind see Marc Fleurbaey and Alex Voorhoeve, Decide as You Would with Full Information!, in *Inequalities in Health*, ed. Nir Eyal, Samia A. Hurst, Ole F. Norheim, and Dan Wikler (Oxford: Oxford University Press, 2013), pp. 113–28; and Seth Lazar, Limited Aggregation and Risk, *Philosophy & Public Affairs* 46 (2018), pp. 117–59.

¹⁰For such an argument compare Fleurbaey and Voorhoeve.

point favor the sore relief even given the off-chance that the mission would have been successful after all.

A natural thought is to adopt limited aggregation and to aggregate *ex post* claims according to one's favored theory of aggregation.¹¹ By aggregating different *ex post* claims from different states of the world, we can decide when the possibility that we should have done A rather than B is more important than the possibility that we should have done B rather than A. We then need a principle that tells us which *ex post* claims we can aggregate. There are two possibilities here. First, only the *ex post* claims within one state of the world can determine whether we can aggregate claims. Second, both *ex post* claims within one state of the world and across different states of the world determine whether we can aggregate claims. As I shall argue, neither of the options is plausible. This leaves us with a dilemma for *ex post* claims. We either need to look elsewhere for constructing an *ex post* version of limited aggregation or accept some form of *ex ante* limited aggregation.

II. First option: relevance tied to a state of the world

According to the first option, whether or not claims can be aggregated is determined solely by reference to the claims in that state of the world. A basic version of this view would tell us to determine first which claims are relevant to the strongest claims in each state of the world, second discount these relevant claims, and third aggregate all discounted relevant claims. We should then perform the action that satisfies the greatest aggregate of discounted relevant *ex post* claims.

The proposed view is a natural extension of Alex Voorhoeve's Aggregate Relevant Claims view.¹² Voorhoeve's view is developed only for cases of certainty. The proposed view supplements Voorhoeve's view with an emphasis on *ex post* claims and the idea that the relevance of claims is determined only within the same state of the world. The proposed view is also a simplified version of Seth Lazar's *Ex Post Maximize Satisfaction of Claims*.¹³ Lazar's view is intended to provide a version of *ex post* limited aggregation.

Consider now the following case (*Uncertain Rescue*): As in *Anne's Rescue* we have Anne, the miner, who is trapped in a mineshaft. Again, we know that the rescue mission

¹¹A brief comment on "no aggregation". We could in this case engage in a pairwise comparison between the strongest *ex post* claim in each state of the world. We would compare, for example, Anne's *ex post* claim to the rescue mission given that she can be helped with the *ex post* claim of someone who has a sore throat given that Anne cannot be helped. A version of this idea would first discount each *ex post* claim by the likelihood that their state of the world is the actual state of the world. In either case, one concern with this response is that it takes the idea of pairwise comparisons too far. It makes the decision whom to save dependent entirely on what happens in one state of the world. While using pairwise comparisons in cases of certainty is motivated by respecting the different standpoints of individuals, pairwise comparisons between *ex post* claims are rather indicative of avoiding a worst-case scenario. In cases in which the worst-case scenario is much worse than all other outcomes, we disregard all the other possibilities and pay attention only to the worst-case. While it may make sense to believe that we should be guided only by the fate of a single individual who has much at stake, it makes little sense to believe that we should be guided only by one possible eventuality.

¹²Voorhoeve, How Should We Aggregate Competing Claims?.

¹³Seth Lazar, Limited Aggregation and Risk, pp. 139–42. I should note that Lazar ultimately rejects the view but calls it a "real contender" (p. 141). Instead, Lazar embraces a hybrid view (pp. 149–58). This hybrid view contains the *ex post* component in it. If the *ex post* view is implausible as an *ex post* view, as I go on to show, then this sheds doubt on the plausibility of Lazar's hybrid view.

will with certainty bring Anne's body to daylight and that the opportunity cost is not being able to provide pain relief for sore throats to a very large number of people. However, *Uncertain Rescue* differs from *Anne's Rescue* insofar as there is a very small chance that all help will come too late. Anne might already be dead. Although we heard life signs from Anne only a few seconds ago, it is possible that Anne will have died by the time we reach her. If this is so, there is nothing we can do for her and the rescue mission serves no purpose.¹⁴ Should this very small chance make much of a difference? It is hard to see why. It is overwhelmingly likely that we can still save Anne and the gains we can achieve by not trying to save Anne are of much less significance than what is at stake for Anne. Importantly, proponents of limited aggregation consider their theory to be of practical relevance. In any real-world scenario there will always be a small chance that rescue will be futile. If limited aggregation fails to account for *Uncertain Rescue*, then it appears to be practically inert.

Devastatingly, the present option to extend limited aggregation to uncertainty fails to account for this judgment. Here is why. The first step is to identify which are the different *ex post* claims in each state of the world and which claims are relevant. There are only two states of the world, S_1 in which Anne is alive and S_2 in which Anne is dead. Suppose that we are in S_1 , the state of the world where Anne is still alive. In this case, Anne has a strong claim to the rescue option. Her life is at stake. Everyone in the large group has only a weak claim to the dispersal of the small benefit. But given that Anne's claim is much stronger than the claim of everyone in the group, her claim is the only relevant one. By contrast, suppose we are in S_2 , the state of the world where Anne is already dead. Then Anne has no claim. There is nothing we could do for her. Here as well, everyone in the large group has a weak claim to the dispersal of the small benefit. These claims are relevant by default given that there is no claim at all to the rescue mission.

In the second step, we have to discount everyone's claim by the likelihood that their associated state of the world obtains. Anne's weighty claim is discounted by the likelihood of Anne already being dead, which is very low. Her claim remains very strong. Discounting the weak claims of the many group members will further weaken them since the probability of Anne being still alive is very high. Nonetheless, the third step allows us to aggregate all relevant claims in the end. If there are enough members of the group, then they together will outweigh Anne's claim.¹⁵

¹⁴We can assume that Anne lives in a society which attaches no special meaning to burial rites. This explains why, if Anne is dead, the rescue mission would have provided not even small benefits to Anne's loved ones in terms of coming to terms with their loss.

¹⁵I mentioned earlier that this is a simplified version of Lazar's view. For interested readers, here is how the view is simplified and why this does not affect my argument. (1) Instead of talking about claims and relevant claims, Lazar talks about interests and claims. Lazar believes that we can aggregate all claims, not only relevant ones, but that only some interests are protected by a claim. (2) An interest is protected by a claim if and only if the person whose interest it is would be permitted to save themselves rather than everyone with a relevant competing interest combined. (3) An interest is relevant in turn if that person would be permitted to save themselves rather than the initial person. (4) The differences in Lazar's approach do not affect my argument. In the above argument one can replace "claim" with "interest" and "relevant claim" with "claim" and we have translated the argument into Lazar's approach. The interests of the members of the large group would still be protected by claims because they are unopposed. Anne's interest would similarly survive the more complicated test of Lazar's for being protected by a claim. (5) One further difference is the following: Lazar determines *ex post* interests counterfactually by comparing the well-being of the person given the chosen action with the counterfactual well-being given the alternative action. If we

This does not correspond to what we intuitively thought about the case. The introduction of even a minimal chance that Anne cannot be helped tips the balance against Anne. Anne is almost certain to be saved from death if we intervene, but the tiny possibility that she may not be makes all the difference here. Since in reality there is always an off-chance that our rescue efforts will be futile, this would render limited aggregation practically inert. The problem is that this method allows relevant claims to arise in a given state of the world too easily. Even a fairly small gain can become a relevant claim if it is sufficiently larger than its competitors in that state alone. Provided that there are sufficiently many of these small gains, they can then, in the end, outweigh the relevant claims of other states of the world.

III. Second option: relevance inside and across states of the world

To avoid the problem that claims can easily arise in one state of the world and outweigh claims in other states of the world, we can opt for the second option that I distinguished. Following this option, it is both the claims within the same state of the world and claims in different states of the world that determine whether aggregation is permissible. This option is not an ad hoc adjustment of our view to avoid the problem I just outlined. It can also be justified by appealing to a core idea of limited aggregation.

In her justification of a limited aggregationist view, Frances Kamm introduces the idea of irrelevant utilities.¹⁶ The idea is that certain utilities or claims are not important in the face of other more significant claims. To take Kamm's example, it would be inappropriate and disrespectful to consider the claim to be cured from a sore throat when deciding between whom to save from death. A similar idea can be applied to risky cases. It can seem inappropriate and disrespectful to consider the claim to be cured of a sore throat in one possible state of the world when it might be the case that we are facing a life-and-death decision. This could explain what has gone wrong with *Uncertain Rescue*. We have allowed the possibility that people will lose out on a sore throat treatment to overpower the consideration that Anne may lose her life due to our refusal to rescue her.

Kamm's view has one relevance test in the case of certainty. Claims are relevant, and thus allowed to be aggregated, only if it would not be disrespectful to consider the weaker claim in light of the gravity of the stronger claim. In uncertain cases we could use a relevance test that looks at relevance not just limited to the particular state of the world. Instead claims have to be relevant to the strongest competing claim across all states of the world.

But this extended relevance test for claims fails. Consider *Desperate Rescue*: we are again uncertain about whether Anne, the miner, is still alive in the mineshaft. We have not heard life signs for a long time and the rescue team is losing hope. There is only a very small chance that Anne is still alive, saving her now would be a miracle. The rescue mission is costly, and the recourses could be used to save a large group of people from moderate chronic pain. Moderate chronic pain, we can suppose, is

compare the action of providing pain relief for sore throats, we cannot observe, however, whether Anne is dead or alive. So how should we assess Anne's interest here? For Lazar we should take the expected value for Anne, whereas I advocate distinguishing between different states of the world, even if they are epistemically indistinguishable. Given the near certainty of Anne still being alive, this small difference has no bearing on my argument either.

¹⁶Kamm, *Morality, Mortality*, vol. 1, pp. 144–63.

not relevant to death in cases of certainty. Supposing that Anne is still alive, Anne's *ex post* claim is the only relevant one. But in the case that Anne is already dead, the group members have claims to be relieved of the chronic pain. However, none of their claims are relevant to Anne's claim in the eventuality that Anne is still alive. We should try saving Anne, regardless of how unlikely it is that our intervention will be successful. A very small chance here would make all the difference. Not only is this absurd in this case, it also cannot be plausibly justified. The idea that the mere possibility of death should make it disrespectful to consider lesser claims is not convincing.

A more plausible relevance test is one in which relevance is determined only after discounting the claims by their likelihood. The idea that considering small claims in the presence of a substantial claim is disrespectful is certainly more plausible when the claims were discounted by their likelihood. This revision also explains what is wrong with the answer that the previous proposal gave in *Desperate Rescue*. Anne's *ex post* claim does not appear weighty enough to us given the small likelihood that she is still alive. Given this only small chance, Anne's *ex post* claim should not render the claims against chronic pain relief irrelevant. We can revise the test and adopt the following two-stage relevance test. We first discount all claims by the likelihood that their state of the world is the actual world. In the second stage we determine which claims are relevant to the strongest claim overall.

The two-stage relevance test struggles to account for cases where risk is dispersed among various states of the world. Consider *Anonymous Rescue*. A large group of people is trapped on a sinking ship. We are able to communicate with the ship and know that at most one person is still alive. We do not know who among the 10,000 crew and passengers is the person who might still be alive. There is also an about 50 per cent chance that none of the 10,000 is still alive. We have the choice between a rescue mission or providing a small and certain benefit to a very large group of people. In this scenario there are 10,001 states of the world. In each of the 10,000 states of the world where one person is still alive, that person's *ex post* claim is relevant and outweighs all other claims. In $S_{10,001}$, the state of the world where there is nothing we can do for the people on the ship, the claims of the large group members are relevant. The *ex post* claims of each passenger must be discounted by 1 in 20,000. It is quite likely that the *ex post* claims of the passengers will then not be relevant to the *ex post* claims of the group members receiving a small benefit. Because the risk is dispersed across states of the world, we decide not to try to save the person on the ship. However, if there had been a single, identified person on the ship, her claim would not have been discounted heavily enough to be rendered irrelevant. On the contrary, her claim would have rendered the claims to the small benefit irrelevant. Such an identified victim bias is a motivation for *ex post* views and should not be a component of them.

One way to resist this implication is to protest that my way of setting up the problem was erroneous. It was false to distinguish between the first 10,000 states of the world. A state of the world is a set of possible worlds, or a model of possible worlds, that leaves no relevant aspect of the world undescribed. A state of the world is not a full description of a possible world. One might protest that I *overdescribed* the states of the world. If we should treat the expectation of a 50 per cent chance of saving someone to be equivalent to a 50 per cent chance of saving a particular person, then this is because the identity of the person to be saved does not matter. If the identity of the person does not matter, the protest goes, it is because the identity of the person is not a relevant feature in this case. Consequently, the states of the world do not differ in any relevant aspect. What should matter is that some person on the ship might die rather than who exactly has the claim

to be rescued. Anonymizing for the victim, the different possible worlds do not differ in any relevant respect.

When we frame the decision problem for risky cases, we inevitably have to group possibilities together. Often there will be small differences in possible outcomes that do not have any moral relevance. Our criterion for how to group possibilities will depend on what we think is morally relevant in this case. The *ex post* proponent can now argue that since we should not be biased in favor of identified lives, the specific identity of a victim does not matter morally either. Hence, we should not divide outcomes in which only the identity of the victim differs into different states of the world.

I should note that this is a radical way of looking at cases like this, one that not even proponents of *ex post* views rely upon. Consider, for example, the conflict between *ex post* egalitarianism and the *ex ante* Pareto principle.¹⁷ As an illustration imagine the following case. There are two boys, Adam and Bill, who will go blind in the absence of any intervention by us. We can intervene in one of two ways. The egalitarian treatment assures both an impaired vision that is better than blindness but worse than full sight. The risky treatment can save one of the two from any impairment whatsoever, but the other boy will suffer an impairment worse than the egalitarian treatment would have given him. The risky treatment is, in addition, in the rational self-interest of both boys. The utilities of the different options are reproduced in Table 1.

Ex post egalitarianism favors the egalitarian treatment even though the risky treatment is in the *ex ante* interest of both Adam and Bill. Proponents of *ex post* egalitarianism agree with this framing of the debate.¹⁸ Fleurbaey and Voorhoeve, for example, dub their argument an argument against *ex ante* Pareto. Yet the redescription strategy that I invoked earlier disagrees. Since the specific identity of who is benefited and who is harmed should not matter, we should not distinguish here either between S_1 and S_2 . This is not a case of uncertainty at all and hence no question of *ex ante* Pareto arises! The decision problem would look like Table 2.

The proposal to reframe the decision problem is radical then because instead of arguing for the idea that the specific identity of persons does not matter as a conclusion, it invokes it as a premise of an argument. We start out with the idea that the specific identity of persons should not matter. The case of Adam and Bill is, therefore, transformed into the simpler Table 2. In Table 2 we are facing no question of uncertainty and hence no debate between *ex ante* and *ex post* views. We are only facing a question on whether inequality matters. The set-up of the decision problem makes *ex ante* views unintelligible in this case. But they are a genuine alternative to the *ex post* approach here. One can choose the risky treatment because one believes in *ex ante* Pareto and not because one values total utility over equality in this case.

Ex post egalitarians do not subscribe to this argument. They want to argue, not assume, that identities do not matter. Their typical argument proceeds from Table 1 and reasons that there are only two options and that in each of the two options there is a strong claim in favor of the egalitarian treatment. This argument is dialectically more powerful since it starts with a commonly agreed upon framing of the decision problem and then argues that the specific identity of the recipient of the benefit should not matter.

¹⁷See Matthew D. Adler, *Well-Being and Fair Distribution* (Oxford: Oxford University Press, 2012), pp. 496–530; Fleurbaey and Voorhoeve; and Ittay Nissan-Rozen, How to be an Ex-Post Egalitarian and an Ex-Ante Paretian, *Analysis* 77 (2017), pp. 550–58.

¹⁸Adler, *Well-Being and Fair Distribution*, pp. 496–530; Fleurbaey and Voorhoeve.

Table 1. Adam and Bill

Alternative	Person	S ₁ (Adam could benefit)	S ₂ (Bill could benefit)
Egalitarian Treatment	Adam	0.8	0.8
	Bill	0.8	0.8
Risky Treatment	Adam	1	0.65
	Bill	0.65	1

Table 2. Adam and Bill (Redescribed)

Alternative	Person	S ₁ (Someone could benefit)
Egalitarian Treatment	Adam	0.8
	Bill	0.8
Risky Treatment	Winner	1
	Loser	0.65

Even though I think these are good reasons to reject the reframing of the decision problem, I now want to set aside this problem. Because even if we accept this radical rethinking of cases of uncertainty, it does not help the case of conditional claims *ex post* limited aggregation. Reframing the decision problem helps this view with *Anonymous Rescue*, but it does not help with a related case. In *Anonymous Rescue* all 10,000 people faced the same fate, death. This is why the alternative way of framing the decision problem would only speak of two states of the world, one where someone is alive and another one where no one is alive, both of which are equally likely. If we frame the problem this way, we should try to rescue the person rather than giving the small benefit to any number of persons. But plausibly we should also try to rescue one person from, for example, the loss of a limb, rather than giving the small benefit to any number of persons. So, if all 10,000 people are facing the loss of a limb, we can again redescribe this as one state of the world where someone is facing the loss of a limb. Suppose, however, that one of the 10,000 is facing the loss of a limb, another person is facing permanent paraplegia, a third person is facing chronic pain worse than paraplegia, and so on. All 10,000 persons are facing a different harm that is different in morally relevant respects. All 10,000 persons are facing a harm between the loss of a limb and death. In this variation the redescription strategy is no longer possible. These are genuinely different states of the world. The problem of *Anonymous Rescue* reappears here. The risk is spread out across various states of the world in which different people face the harm. This makes all the difference for the redescription strategy, even though it should not for someone committed to the *ex post* view. Even worse, if every person were to face the loss of a limb, we should try preventing this loss. But if some people are facing a more serious loss, then we should no longer try to prevent this loss. Because these people are facing a more severe hardship, the strategy of reframing the decision problem no longer works. Their *ex post* claims have to be counted as belonging to separate states of the world and discounted separately. This way they become irrelevant.

Could the redescription strategy work here once more? The suggestion would be that instead of saying that we are facing 10,000 different states of the world in which different harms occur, we are facing just a single state of the world in which we could save a person from *at least as much* harm as a loss of a limb. Because there is only one state of the world that corresponds to the possibility that we can help someone on the ship at least this much, this *ex post* claim is not heavily discounted. While this suggestion would yield the correct result in the varied scenario of *Anonymous Rescue*, it loses what was appealing and plausible in the earlier redescription strategy. The protest against the initial set-up of *Anonymous Rescue* was that it distinguished between scenarios that are morally alike. If we cannot point out a morally relevant difference between different scenarios, we should treat them as one state of the world. But in the varied scenario the different scenarios *are* morally different. Saving a person from death is more important than saving a person from the loss of a limb.

At any rate, even if we believe in a contextual answer to the question of what counts as a morally relevant difference, another variation of *Anonymous Rescue* shows it to be implausible. The current proposal holds that we can redescribe the case by focusing on the smallest common denominator. To be able to describe the case as involving only one state of the world, we need to say that we are able to save someone from “at least” as much harm. But sometimes we need to invoke the additional benefits which this redescription strategy ignores. A numerical example is presented in Table 3. The numbers represent the strength of an individual’s *ex post* claim. The second column represents the relevant claims to the rescue mission, the third column represents the relevant claims against the rescue mission. For simplicity we can assume that the strength of these claims tracks health benefits that fall between 1 and 0. Let us assume that a claim becomes irrelevant if the stronger claim is at least six times as strong as the weaker claim. This explains why except for S_{11} the claims against the rescue mission are irrelevant once we are able to save someone on the ship.

If we discount the *ex post* claims in S_2 to S_{11} individually in each state of the world then the claims of the persons with the smaller benefit will always be relevant. In S_2 , for example, the discounted claim in favor of the rescue mission has a strength of 0.05 which is as strong as the discounted claim against the rescue mission in S_1 . This is where the redescription strategy should come into play. But redescribing S_2 to S_{11} as a harm of at least as much as 0.55 will not help here. Discounted by the 50 per cent likelihood that the rescue mission will not achieve anything, the strength of the *ex post* claim will be 0.275. This is not six times greater than the *ex post* claim against the rescue mission that delivers the smaller benefit to many. Yet, if we contrast this to a case in which an identified person on the ship had a claim of the strength of 0.6, we should prioritize that person. Why? Because in this scenario the *ex post* claim would be of a strength of 0.6 and thereby six times as large as the claim of the large group. This strikes me as a weak point for the proposal. In almost all possible outcomes, save for one fairly unlikely one, we are averting a less severe harm on the ship. Yet nevertheless in the case of an identified person with a claim of the strength of 0.6 the rescue mission would be required, unlike the variation shown in Table 3. Note that the probability of S_{11} being the actual world can be made infinitesimally small and the argument goes through, nevertheless. The small chance that we might only help the person less than 0.6 lowers the threshold for our redescription strategy and thereby renders the discounted *ex post* claim relevant to the *ex post* claims against the rescue mission.

Table 3. Variation on *Anonymous Rescue*

State of the world	<i>Ex post</i> claim to rescue mission	<i>Ex post</i> claim against rescue mission
S ₁ (p = 0.5)	–	0.1 by a large group
S ₂ (p = 0.05)	1	–
S ₃ (p = 0.05)	0.95	–
S ₄ (p = 0.05)	0.9	–
S ₅ (p = 0.05)	0.85	–
S ₆ (p = 0.05)	0.8	–
S ₇ (p = 0.05)	0.75	–
S ₈ (p = 0.05)	0.7	–
S ₉ (p = 0.05)	0.65	–
S ₁₀ (p = 0.05)	0.6	–
S ₁₁ (p = 0.05)	0.55	0.1 by a large group

IV. Conclusion

Thinking about *ex post* claims leaves us in a dead end. None of the principles that tell us when claims can be aggregated are plausible. Determining when claims can be aggregated only by looking at one state of the world allows aggregation too easily. As demonstrated by *Uncertain Rescue*, a small chance that no aid can be provided already opens the door to aggregation. Determining when claims can be aggregated by also looking at other states of the world makes aggregation either too difficult or too dependent on how the risk is distributed across different states of the world. As demonstrated by *Desperate Rescue*, the mere possibility of being able to save someone who has a lot at stake may rule out aggregation. Alternatively, as demonstrated by *Anonymous Rescue*, it matters whether the risk is faced by a single identified person or by different people in different states of the world.

A common feature of the failure of both approaches is that both give special emphasis to uncertainty. Both treat near-certainty radically differently from certainty. The first option radically changed its verdict once we introduced the small probability of all help coming too late. The second option radically changed its verdict once we introduced the small probability of being able to help at all or being able to help as much.

What does this mean for *ex post* limited aggregation? Perhaps the appeal to *ex post* claims is simply limited in scope. In explaining how to think about *ex post* claims I started with the case in which we know one person will be harmed but do not know their specific identity. In such a case, it is true that in whichever state of the world we are, we should prefer the many small harms to the one substantial harm. This type of reasoning embraces a dominance condition. If in all states of the world the *ex post* claims weakly prefer one action and in at least one state of the world strongly prefer this action, then we ought to perform the action. If the *ex post* claims in all states of the world are indifferent, then we ought to be indifferent. This dominance reasoning does not tell us what to do when different options are preferred by the *ex post* claims in different states of the world. Restricting the scope of *ex post* limited aggregation in such

a way would mean that it is silent on the cases I discussed in this article. It would be striking if *ex post* reasoning would not apply to any of these. Furthermore, it seems concerning that we cannot expand its core reasoning, the idea of *ex post* claims, in a plausible manner. This may indicate that the idea of *ex post* claims is inherently flawed.

A proponent of the *ex post* view may perhaps want to revert to someone's claim *ex post* limited aggregation and accept the problems that it brings with it. Whichever way we go, my argument provides a challenge for critics of the *ex ante* view. *Ex post* limited aggregation is far away from a well-developed alternative. The conclusion which I draw from this is that my argument strengthens the case for adopting an *ex ante* approach to limited aggregation in cases of uncertainty.¹⁹

¹⁹An earlier version of this article was presented at the Oxford Graduate Conference for Philosophy. On this occasion, I benefited from excellent comments from my commentator Roger Crisp, as well as from discussion with the audience. I also thank Tomi Francis, Dan Hausman, Joe Horton, Michael Otsuka, Patrick Tomlin, and Alex Voorhoeve for very helpful feedback and discussion. Special thanks go to Alex Voorhoeve, who nudged me to write the article in the first place. Lastly, I want to thank two outstanding anonymous reviewers whose reports have greatly improved the article.