

LIMITED AGGREGATION AND RISK¹

Seth Lazar, ANU

1. Introduction

Deontological ethical theory and unqualified absolutism were once near synonyms: 'do justice though the heavens fall'. Few deontologists today are so hard-line. But many still believe that *some* trade-offs that would yield unambiguously better outcomes are nonetheless wrong. Here is one such scenario—call it *Life for Headaches*: We must choose either to avert a minor, temporary headache for each member of a multitude or to save one person's life.² No matter how numerous the multitude, we ought always to save the one.

We can call this view 'moderate absolutism'.³ It seems to fit common sense for decision-making with perfect information. But in the real world, our information is never perfect. We must therefore extend moderate absolutism to decision-making under doubt. And here things get trickier. It would be absurd to treat a very low probability that I will kill an innocent person—for example, as I drive to the shops—as no less forceful a reason than if their death were certain. The force of our reasons, when acting under risk or uncertainty, must somehow be discounted by their probability of being actual.⁴

How should we proceed, then, if we can save either one person from a 1/2 probability of death or a multitude from a certain headache? What if the probability of death is only 1/10,000?

¹ [omitted]

² Alastair Norcross, "Comparing Harms: Headaches and Human Lives", *Philosophy & Public Affairs*, 26 (2) (1997), 135-67.

³ Absolutist theories say that some considerations cannot be outweighed by any amount of good. Moderate absolutism says that there is a consideration C and a good G such that no amount of G can outweigh any amount of C. Thanks to [omitted].

⁴ In this essay, I discuss only decision-making under risk, when probabilities are defined, rather than under uncertainty, when they are not.

At some point we must avert the headaches. Likewise, in a more realistic case, I can justifiably run *some* risk of killing an innocent person in a road accident, even if I'm just driving to the store to buy popcorn and chocolate. And yet, if enough of us face enough of these choices over a long enough time, then eventually we are sure to either allow or cause a death, only for the sake of comparatively trivial benefits. So, does weakening our absolutism in response to risk mean abandoning it over the long run?

This paper explores principles that can preserve common sense in *Life for Headaches*, while addressing risky cases in an intuitively attractive way. It shows that deontologists attracted to moderate absolutism can extend their views to decision-making under risk. It is part of a broader project to show how deontologists can develop a compelling account of decision-making with imperfect information. In a slogan, it is one step towards a deontological decision theory. In a world where advances in artificial intelligence have generated a compelling demand for formal versions of ethical theories adapted to decision-making under risk, it is more crucial than ever that deontologists explain how to systematically deal with duty under doubt.⁵

Here's how this will go. First, I will introduce a moderate absolutist principle—'Maximise Satisfaction of Claims' (MSC)—which does not take risk into account.⁶ Next I offer some desiderata for a successful extension of MSC to risk. My aim here is to answer the many critics of moderate absolutism, who say that this cannot work. The rest of the paper offers five different extensions of MSC, of which the last is, I think, the most promising.

2. A Principle for Limited Aggregation

⁵ There's an interesting question, raised by an associate editor, as to whether machine learning might actually vitiate the need to formalise our moral theories—feed the algorithm enough intuitive judgements, and it will articulate the underlying principles. This might be possible, though of course part of doing moral theory is figuring out when we need to revise those intuitive judgements, so the algorithm would have to be highly sophisticated.

⁶ Alex Voorhoeve, "How Should We Aggregate Competing Claims?", *Ethics*, 125 (1) (2014), 64-87: 74-75.

2.1 Maximise Satisfaction of Claims

Some—'anti-aggregationists'—justify moderate absolutism by arguing that sometimes aggregation of lesser interests is impermissible. Anti-aggregationists can either oppose *all* aggregation or—more plausibly—favour Limited Aggregation.⁷

Consider *Life for Lives*, in which you can save either one person's life—call him Xavier—or the lives of 100 others.⁸ Surely aggregation is allowed: you should save the lives of the 100.⁹ And numbers are more than tie-breakers. Consider *Life for Legs*, in which you can save Xavier's life or the legs of some number of others. For a high enough number, you should save the legs. A theory of Limited Aggregation must explain when aggregation is prohibited, when it is allowed, and what makes the difference.

Several philosophers have argued that cases like these involve competing claims, some of which are 'relevant', others 'irrelevant'.¹⁰ The headache-sufferers' claims are irrelevant to the one's claim to keep his life, so may not be aggregated, and we must save the life no matter how many headaches we could avert. But the claims of those who will die or lose their legs are relevant to the one's claim to survival, so can be aggregated. So: what explains and grounds relevance?

The simplest explanation is that one person's lesser claim is irrelevant to another's greater claim when the latter exceeds the former by a great enough amount. I will focus on a gloss on this approach offered by Alex Voorhoeve.¹¹ He argues that a person's lesser claim is irrelevant to Xavier's competing claim when that person—call her Yolanda—has so little at stake relative to

⁷ I take this term from Patrick Tomlin, "On Limited Aggregation", *Philosophy & Public Affairs*, 45 (3) (2017), 232-60.

⁸ In all of my cases throughout this paper, everyone is innocent, starts out at the same level of well-being, has as much to live for as the rest, and besides what I specify all else is equal.

⁹ Obviously some philosophers deny that we may ever aggregate in this way, preferring a simple or weighted lottery. Thanks to [omitted] for raising this point.

¹⁰ David Brink, "The Separateness of Persons, Distributive Norms, and Moral Theory", in R. G. Frey and Christopher Morris (eds.), *Value, Welfare, and Morality* (Cambridge University Press, 1993), 252-89; T. M. Scanlon, *What We Owe to Each Other* (London: Belknap Press, 1998): 238-41; Larry S. Temkin, *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning* (Oxford: Oxford University Press, 2012); Thomas Nagel, *Equality and Partiality* (Oxford: Oxford University Press, 1991); F. M. Kamm, *Morality, Mortality* (New York: Oxford University Press, 1993): ch 8-10; Michael Otsuka, "Saving Lives, Moral Theory, and the Claims of Individuals", *Philosophy & Public Affairs*, 34 (2) (2006), 109-35

¹¹ Voorhoeve, "How Should We Aggregate": 74-75 (which explicitly focuses on objective permissibility).

Xavier that, faced with the choice of whether to serve Xavier's claim or bear that much cost, she would be morally required to bear that cost.¹²

Yolanda would have a duty to bear a headache to ensure Xavier's life is saved, so Yolanda's claim not to suffer a headache is irrelevant to Xavier's claim to survival. The same is true for anyone else with a headache at stake, so their claims may not be aggregated to outweigh Xavier's claim. One cannot complain about one's claim not being served when, given the choice, one would have had to forego one's claim for the sake of the greater claim with which it competes.

Conversely, one would not be required to sacrifice one's legs to save a stranger's life. So in *Life for Legs*, Yolanda's claim to keep her legs is relevant to Xavier's claim to keep his life. The same is true for any others with legs at stake, so their claims may be aggregated, and if they are numerous enough will outweigh Xavier's claim. The same is also true, of course, in *Life for Lives*.

This is, I think, Voorhoeve's most interesting insight. Our duties of rescue and relevance are closely connected—and this connection allows us to resist aggregation without endorsing the contractualist moral theory that usually underpins it. We don't need a story about a hypothetical bargaining position to explain the individualist nature of Limited Aggregation. We just need two simple ideas. First, claims belong to individuals, and are grounded in an individual's interests alone. Claims entail normative powers, including the power to unilaterally waive one's claim. But one could not do so if the claim was grounded in others' interests besides your own. Second, you can't have a claim to receive a benefit at another's expense when, given the opportunity, you would have a duty to refrain from saving yourself if saving yourself would prevent the stronger claim being served. If this is true for you, then it is the same for everyone else whose identical claim is aligned with yours, so none of those claims should be considered.

If this is right, then we have more to go on here than just intuitions about cases. We have a theoretical justification for Limited Aggregation, which does not rely on a background

¹² F. M. Kamm, *Intricate Ethics: Rights, Responsibilities, and Permissible Harm* (New York: Oxford University Press, 2007): 297-98.

commitment to contractualism.

However, if we take seriously this connection between relevance and our duties of rescue, then we should adjust Voorhoeve and other Limited Aggregationists' individualist assumptions. They all stress that, to determine whether Xavier has a claim, we must consider only bilateral comparisons between his interest and that of every individual with whom his interest competes. This 'individualist restriction' is borne of their underlying commitment to contractualism—a commitment that I don't share. I think duties of rescue work differently from this, so it is plausible that relevance should too.

I can sometimes have a duty to bear a cost to spare a large enough group of people from each suffering a cost that is, for each of those people, smaller than the cost to me. This is true only when the lesser cost is relevant to the greater cost. I could never be required to forego saving my own life, in order to spare a multitude from suffering a headache. But for some number of people, I *can* be required to forego saving both of my legs, if that spares each of them from losing a single leg. As for duties of rescue, so for relevance: sometimes whether one has a claim is determined by how one's interest compares with the *set* of relevant competing interests.

I think this is the default for duties of rescue under certainty. And it looks even more plausible for duties of rescue under risk, as well as for extending a principle of Limited Aggregation to decision-making under risk. I take this up in 6.3 below.

These ideas are enough to articulate a principle of Limited Aggregation. For the sake of clarity, I will introduce somewhat different terminology from Voorhoeve and others.

Our topic is whom to help when people's interests compete. As I will use those terms, one's interests are the constituents of one's well-being: if one's interest is satisfied, one's well-being is

advanced; if one's interest is thwarted, one's well-being is set back.¹³ I am neutral between theories of well-being. Interests *compete* or (equivalently) *conflict* when they cannot be jointly satisfied.

When one's interest is the object of a *claim*, it enjoys a special kind of moral protection: to fail to advance that interest, *even through action that is all-things-considered permissible*, is to *wrong* the claim-bearer, *pro tanto*.¹⁴ I will write that an interest is the *object* of a claim, and that it is *protected* by a claim, interchangeably.

Some interests which one might think protected by a claim do not in fact enjoy that protection: Voorhoeve calls these 'irrelevant claims'. I will call them simply interests.

Let's dig deeper into what it means for an interest to be protected by a claim. A candidate interest—say, Xavier's interest in avoiding an untimely death—is protected by a claim if and only if it is not sufficiently outweighed by the relevant competing interests of others. What makes a competing interest relevant? Think back to *Life for Headaches*. Let's suppose that if we save Xavier, we will not be able to spare the Ys from suffering a temporary minor headache. The first question, then, is whether the competing interests of the Ys are relevant to Xavier's candidate interest. A competing interest is relevant to the candidate interest if and only if the competing interest does not sufficiently outweigh the candidate interest. So, consider one of the Ys—Yolanda again. Is her interest in avoiding a headache sufficiently outweighed by Xavier's interest in survival? Yes it is. And the same is true for each of the other Ys. Each Y's interest is sufficiently outweighed by Xavier's interest in survival. So Xavier's interest faces no relevant competing interests, and *a fortiori* is not sufficiently outweighed by any, so it is protected by a claim. By contrast, in *Life for Legs*, Xavier's interest in survival competes with the Ys' interests in keeping their legs, then since one's interest in keeping one's legs *is not* sufficiently outweighed by

¹³ Perhaps some interests are constitutively unsuited to being objects of a claim. If so, then when I use 'interests', I mean only those that *can* be the object of a claim. Note that some people reserve the term 'interests' to mean something orthogonal to individual well-being. That's not how I am using it in this paper.

¹⁴ I suspect that when an act harms you, and is *not* all-things-considered permissible, then it wrongs you just in case the harm to you is unjustified, regardless of whether you have a claim.

another's interest in survival, the Ys' interests *would* be relevant, and so if there are enough Ys for their relevant interests in keeping their legs to sufficiently outweigh Xavier's interest in survival, then the latter would not be protected by a claim.

What does it mean for an interest to be sufficiently outweighed? To answer, we must first measure interests. I focus mostly on easy cases: you have two options, so the size of an individual's interest is the difference between how she fares under each option. Matters are complicated enough without adding more options into the picture; what's more, if Limited Aggregation under risk doesn't work with simple two-way choices, then the view is hopeless. Nonetheless, at the end of this section, I show how my principle would cater for more complicated cases.

The moral weight of an interest is also affected by other factors. Perhaps the interests of the worst-off get more weight; we might also discount those of people who are responsible for their own plight or for the plight of the people with whom their interests compete. To keep things simple, I will hold these other factors constant. But it is crucial to remember that we can factor in many kinds of considerations by thinking carefully about the moral weight of each interest.

I will use the duties of rescue test to tell when an interest is sufficiently outweighed. In *Life for Legs*, Yolanda's interest is sufficiently outweighed by Xavier's if she would be required to forego saving her legs to ensure that Xavier's life is spared. In my view, one is not required to forego saving one's legs in order to spare another person's life. If the legs of each of the other Ys are also at stake, then they also have relevant competing interests, and Xavier's interest is protected by a claim if and only if it is not sufficiently outweighed by the aggregated interests of Yolanda and the other Ys. Would Xavier be required to forego saving his life to spare the Ys the loss of their legs? Views on this question will differ. If you think that one could not have a duty to forego saving one's life here, then you might think that Xavier's interest is always protected by a claim—though it might be outweighed. Alternatively you might think that, for some number of

Ys, Xavier would have to forego saving himself in order to ensure their legs were spared. In that case his interest would not be protected by a claim.

The duties of rescue test is at least a valuable heuristic, affording another lever with which to pump intuitions, in a field where cases are highly abstract. More than this, I suspect that our duties of rescue *explain* why interests are irrelevant, when they are, and why one should not aggregate irrelevant interests. However, defending this claim would distract from my goal here—which is to present an ecumenical principle that most anti-aggregationists can endorse, and then extend it to decision-making under risk. So, when presenting my principles, I define relevance in terms of whether an interest is sufficiently outweighed, and appeal to duties of rescue only for guidance and inspiration, without insisting that our duties of rescue ground Limited Aggregation. If you're unpersuaded that relevance and our duties of rescue line up, remember that the question is whether I would be required to forego conferring a benefit on myself, in order to ensure that the weightier competing interests of others are served. It is not a matter of whether I am required to impose a given cost on myself, to ensure those other interests are served.

Anyone who has an interest at stake, which is not sufficiently outweighed by the relevant competing interests of others, has a claim on your aid. What, then, should you do? When claims are at stake, only other claims can outweigh them. No matter how collectively weighty a sum of mere interests may be, if they are not protected by claims then they cannot outweigh even a single claim. Claims, however, aggregate: when you face competing claims, you should satisfy the stronger overall set, taking into account their moral weights. Hence the principle is called 'Maximise Satisfaction of Claims' (MSC).

So: serve claims first, and only if the options are on a par with respect to claims, serve interests that are not protected by claims. I will assume that once claims have been optimally satisfied, one is permitted to maximise overall morally weighted well-being.

We now have all the materials in place to state our first principle, MSC, more precisely. To

do this we need some notation. Let's call Xavier X_1 . I'll annotate X_1 's interest S as S_1^x . Yolanda is the first of the people whose interests compete with Xavier's, so I'll annotate her as Y_1 , and her interest as S_1^y . I'll call the other people whose interests are aligned with Xavier's the X s, or $X_{1...n}$ and those whose interests are aligned with Yolanda's the Y s, or $Y_{1...n}$. I'll write the aggregate of the interests of the Y s as $S_{1...n}^y$, and of the X s as $S_{1...n}^x$. When I want to refer to an arbitrary member of $X_{1...n}$ or $Y_{1...n}$, I will write X_n or Y_n , whose interest is S_n^x or S_n^y respectively.

MAXIMISE SATISFACTION OF CLAIMS (MSC)

In a choice whether to ϕ or ψ , in which $X_{1...n}$'s interests compete with $Y_{1...n}$'s interests:

1. *The moral weight of an individual's interest is determined by the difference between their well-being if you ϕ and their well-being if you ψ , subject to other weightings such as responsibility and priority.*
2. *X_n 's interest S_n^x is the object of a claim if and only if it is not sufficiently outweighed by the relevant competing interests of $Y_{1...n}$.*
3. *The interest of each member of $Y_{1...n}$ which competes with S_n^x , is relevant to S_n^x if and only if it is not sufficiently outweighed by S_n^x .*
4. *The moral weight of a claim is determined by the moral weight of the interest that is its object.*
5. *Maximise the sum of the moral weights of satisfied claims.*
6. *If more than one option maximizes the sum of the moral weights of satisfied claims (perhaps because there are no claims to satisfy), then you may choose the one that maximizes overall morally weighted well-being.*

MSC can deliver plausible verdicts in our test cases. In *Life for Headaches*, the interest of each of the headache-sufferers is sufficiently outweighed by the interest of Xavier, who will die if you avert the headaches. So they are not relevant to Xavier's interest, and Xavier has a claim on your

aid. Since Xavier's interest in survival is relevant to each of the headache-sufferers' interest in avoiding a headache, and it sufficiently outweighs each of those interests, each *Y* lacks a claim to aid. So the only claim in play is Xavier's, and you must help him, no matter how many headache-sufferers there are.

By contrast, in *Life for Legs*, Xavier's interest does not sufficiently outweigh each of the *Ys*' interests. Yolanda is not required to forego saving her legs to ensure that Xavier's life is saved. The same is true for the other *Ys*. So they all have claims. Whether Xavier has a claim depends on how many *Ys* there are—if there are enough relevant claims competing with his own that he would be required to forego saving himself to ensure that those claims were met, then he would not have a claim to aid. If Xavier does have a claim, then it is an open question whether we maximise the sum of the moral weights of satisfied claims by saving him or the *Ys*—it depends on how many *Ys* there are.

2.2 Objections to Limited Aggregation

Limited Aggregation faces two main lines of attack. The first argues that, once we address more complex cases, involving multiple options, outcomes, or beneficiaries, we get counterintuitive results. The second argues that we cannot plausibly extend Limited Aggregation to decision-making under risk. While I focus on the second objection here (because of my overarching goal to develop a deontological decision theory), it is important to pause and consider some of the most pressing objections in the first group as well. If MSC fails at the level of objective moral theory, then there may be little point extending it to decision-making under risk.

Theories of Limited Aggregation employ relevance conditions that make the value of an option depend essentially on what other options are available—even options that we know we won't choose. This can lead to counterintuitive results. Consider a case.¹⁵ You can help only one

¹⁵ Thanks to an associate editor here.

of four groups: *Life*, *Legs*, *Fingers*, *Headaches*. Suppose only one person's life is at risk, while an increasing number is at risk of each of the subsequent injuries. Let's stipulate that, if there were only one person in each group, each harm would be relevant to its neighbours, but irrelevant to the next-but-one condition. How do we work out whom to save?

One doesn't have a claim to have one's finger saved, or one's headache averted, at the cost of someone else dying. So theories of Limited Aggregation prohibit helping *Headaches* or *Fingers*. *Legs'* interests, however, are relevant to *Life's* interest: one is not required not to save one's leg to ensure another's life is saved. So if the number in *Legs* is high enough, then their claims can outweigh *Life's* claim to survival, and you ought to save *Legs*.

On standard Limited Aggregation theories, if we remove the option of saving the life from the choice set, then if the number of finger victims is high enough, we would have to save *Fingers* rather than *Legs*. If we also removed the option of helping *Legs*, then if the number in *Headaches* is high enough, we would have to avert the headaches.¹⁶ John Halstead put a sharp point on this objection, saying that, in a situation where we know that we can help *Legs*, *Fingers*, or *Headaches*, we should pay to find out whether we can also help *Life*, even though we know in advance that we would not then do so.¹⁷

MSC avoids this counterexample, because of its additional element of aggregation relative to views like Voorhoeve's. By stipulation, the number in *Legs* is high enough that, if their interests are the objects of claims, they can together outweigh *Life's* claim. If given the choice whether to save *Life* or *Legs*, for that number we should save *Legs*. Meanwhile, each member of *Fingers'* interest in keeping her finger is sufficiently outweighed by *Life's* interest in survival—by stipulation, one is required to forego saving one's finger if that would save another person's life. So, if one's interest in keeping a finger is *sufficiently outweighed* by another's interest in survival, and that interest in survival is *outweighed* by some other folks' aggregated relevant interests in keeping

¹⁶ For discussion, see Tomlin, "On Limited Aggregation": 242; Voorhoeve, "How Should We Aggregate": 78.

¹⁷ John Halstead, "The Numbers Always Count", *Ethics*, 126 (3) (2016), 789-802: 799.

their legs, then one's interest in keeping one's finger is sufficiently outweighed by that group's aggregated interests in retaining their legs. According to my version of MSC, if A is sufficiently outweighed by B, and B is outweighed by C, then A is sufficiently outweighed by C. As a result, we should save *Legs* whether or not we have the alternative of saving *Life*, because each member of *Fingers* (and a fortiori each member of *Headaches*) lacks a claim to be helped, since one's interest in keeping one's finger is sufficiently outweighed by that number of competing interests in keeping a leg.

What if we remove the option of saving a life, and stipulate that the number in *Legs* is low enough that the members of *Fingers* don't lose the protection of a claim by competing with *Legs*? Let's also stipulate that there are enough members of *Fingers* that the interests of each member of *Legs* is sufficiently outweighed by the relevant competing interests in *Fingers*. And finally, let's stipulate that there are enough people in *Headaches* that, if your only options were to save *Fingers* or *Headaches*, you would have to save the latter.

On this set-up, nobody has a claim to aid. There are enough members of *Fingers* that each member of *Legs* lacks a claim to his leg being saved. And if the aggregate weight of *Headaches* is greater than that of *Fingers*, and one's interest in keeping a leg can be sufficiently outweighed by the aggregate interests in *Fingers*, then one's interest in keeping one's finger can be sufficiently outweighed by the greater set of aggregate interests in *Headaches*.¹⁸ Finally, those in *Headaches* also lack claims, since if the aggregate of *Fingers* is enough to sufficiently outweigh one person's interest in keeping his leg, it is also enough to outweigh another's interest in avoiding a headache. Since nobody has a claim, you should simply move to the next clause (6 in the formulation above), and distribute the well-being at stake according to that. By assumption, in these cases that means serving the greater aggregate interest: that is, saving *Headaches*. And suppose we now remove *Legs* from the case. Again, nothing about doing so changes the relative magnitudes at

¹⁸ Remember that an interest in avoiding a headache is relevant to an interest in keeping a finger.

stake, so the ordering remains the same: *Headaches* take priority over *Fingers*.

In these cases, MSC fares better than other theories of Limited Aggregation. But it still faces problems. Consider again that last case, where it's *Legs* v. *Fingers* v. *Headaches*. In this choice, nobody has a claim to aid, so we should maximise morally weighted well-being, which means that we should prefer saving *Headaches* to saving *Fingers*, and saving *Fingers* to saving *Legs*.

But if we now remove the option of saving *Fingers*, then while the members of *Headaches* still lack claims (by assumption, the harm of avoiding a headache is not relevant to the harm of losing a leg), the members of *Legs* now *have* claims, since one cannot be required to forego saving one's leg in order to spare some number of others from suffering a headache. So the ordering is reversed: we should prioritise saving the legs over averting the headaches.

However, is this really a problem? We have independent reason to think that the members of *Legs* lack claims when the members of *Fingers* are present, because one would have a duty to not to save one's leg, if that is necessary to spare *n* people the loss of their fingers. There would have to be a lot of fingers at stake for this to be true. Is it so odd that we should care about the attitude that our decision expresses to a (presumably very large) group of people, even if we would not ultimately help them?

We should evaluate our choices not only by how they affect the well-being of their potential beneficiaries, but also by the attitudes of respect or disrespect that they display to those people.¹⁹ The presence of the members of *Fingers* really does change your choice situation: acknowledging claims on the part of the members of *Legs* would amount to disregarding the danger to the members of *Fingers*.

Critics have argued that Limited Aggregation is susceptible to deontic cycling, generating either individual cases, or sequences of cases, in which, for example, one morally prefers A to B, B to C, and C to A. Here, again, MSC has resources that other theories lack.

¹⁹ Voorhoeve, "How Should We Aggregate": 78; Tomlin, "On Limited Aggregation": 242; Marc Fleurbaey, Bertil Tungodden, and Peter Vallentyne, "On the Possibility of Nonaggregative Priority for the Worst Off", *Social Philosophy and Policy*, 26 (1) (2009), 258-85: 284.

Suppose that Alf has unjustifiably poisoned Billie, who has unjustifiably poisoned Celeste, who has unjustifiably poisoned Alf.²⁰ All will die in an hour if they do nothing. But Alf has the antidote, which would cure any one of them (but only one). Presumably, he has a duty to give the antidote to Billie. But Billie has a duty to give the antidote to Celeste. And Celeste has a duty to pass it back to Alf. If we attended only to their duties to one another, they would pass the bottle around until they were all dead. They should instead recognise that their situation is tied with respect to their duties to one another, and draw lots for who gets the antidote. When duties fail to deliver a verdict, move to the next principle.

The same would be true for MSC if it entailed cyclical orderings. If, with respect to claims, option A is better than option B, B than C, and C than A, then the three options are tied, and we move to our principles for the distribution of morally weighted well-being (claims aside). Notice the difference between this move and the 'pluralist' move often made by anti-aggregationists.²¹ MSC does not say that anti-aggregationism is just one set of reasons among others. It says that if your options are tied with respect to claims, then we can apply other principles.

Thinking about duties of rescue can help answer another objection to theories of Limited Aggregation, raised recently by Patrick Tomlin—though presented here in a simpler form.²² Suppose your options are: spare Xavier's life, or spare Yolanda from quadriplegia *and* spare a large number of others from suffering a minor headache. Intuitively, sparing someone from quadriplegia is *almost* as important as sparing someone's life. So shouldn't the headaches be able to tip the balance, if there are enough of them?

One possible response: bite the bullet and insist that trivial interests should be disregarded in such cases. Alternatively, we can again take inspiration from our duties of rescue. I am not required to forego averting a headache to spare someone else's life when I have as an alternative:

²⁰ This case is inspired by a similar one of self-defence, in Russell Christopher, "Self-Defense and Defense of Others", *Philosophy & Public Affairs*, 27 (2) (1998), 123-41.

²¹ Johann Frick, "Contractualism and Social Risk", *Philosophy and Public Affairs*, 43 (3) (2015), 175-223: 219.

²² Tomlin, "On Limited Aggregation": 242. I owe this version of the objection to an associate editor.

don't suffer the headache and spare someone else from quadriplegia. I don't have to bear a great cost for a great benefit, if I can realise a benefit that is almost as good, at much less cost. Similarly, suppose that I can save 1,000,000 lives at the cost of my life, or else save 999,999 lives at the cost of my hand. Even if I would be required to forego saving myself if that were necessary to save a million, I can't be required to forego saving my life just for the purpose of saving *one more* life.²³ The same is true in our aggregation case. Plausibly, the headache sufferers are not required to forego saving themselves from a headache to realise *the difference between one person avoiding quadriplegia and another avoiding death*.²⁴ So they *do* have claims in this case.

To decisively accommodate these responses in MSC, I might have to be more decisive, and argue that facts about our duties of rescue aren't merely a useful guide to relevance and claims, but actually underpin and explain them. I would also have to state the principle in a form that generalised to multiple-choice cases. However, I want MSC to accommodate a range of views about Limited Aggregation, rather than be tied only to my own. And I want to keep it simple enough to allow extending it to the much more complex circumstances of risk. I have shown that MSC is a robust contender for a theory of objective Limited Aggregation, and have offered a blueprint for how to think through further objections.

3. Desiderata for Limited Aggregation under Risk

Critics have long argued that Limited Aggregation has counterintuitive results when extended to risky cases.²⁵ In formulating my response, I extract four intuitive desiderata from these

²³ This allows us to answer a similar objection raised by Daniel Ramöller.

²⁴ As an associate editor has commented, it may make an important difference here whether there is a particular person who is spared the difference between dying and a broken leg; in another case where the beneficiaries of the two harms are different, we might get different results. That is a complication, however, that I must explore in detail elsewhere.

²⁵ Elizabeth Ashford, "The Demandingness of Scanlon's Contractualism", *Ethics*, 113 (2) (2003), 273-302; Barbara H. Fried, "Can Contractualism Save Us from Aggregation?", *The Journal of Ethics*, 16 (1) (2012), 39-66; Barbara H. Fried, "What Does Matter? The Case for Killing the Trolley Problem (or Letting It Die)", *Philosophical Quarterly*, 62 (248) (2012), 505-29; Joe Horton, "Aggregation, Complaints, and Risk", *Philosophy & Public Affairs*, 45 (1) (2017), 54-

criticisms, which any Limited Aggregationist decision theory must satisfy in order to answer those critics.

1. *Aggregate Against Risks of Relevant Harms*

When the risked harms are relevant to one another, lower individual risks of harm may be aggregated to outweigh higher individual risks of harm.

2. *Don't Aggregate Against High Risk of Greater Harm*

When the risked harms are not relevant to one another, and the probability of the greater harm is high enough, individual risks of lesser harms may not be aggregated to outweigh the individual risk of the greater harm.

3. *Aggregate Against Low Risk of Greater Harm*

When the risked harms are not relevant to one another, and the probability of the greater harm is low enough, individual risks of lesser harms may outweigh a small enough individual risk of the greater harm.

4. *Aggregate Against Accumulated Low Risks of Greater Harms*

When facing a series²⁶ of choices to spare people from a risk of a lesser harm, each of which imposes a very low individual risk of a greater harm, but where the series has a very high risk of the greater harm befalling someone, the individual risks of lesser harms may at least sometimes be aggregated to outweigh the cumulative risk of the greater harm, so that it can be permissible to initiate the series.²⁷

81; Norcross, "Comparing Harms"; Alastair Norcross, "Speed Limits, Human Lives, and Convenience: A Reply to Ridge", *Philosophy and Public Affairs*, 27 (1) (1998), 59-64; Sophia Reibetanz, "Contractualism and Aggregation", *Ethics*, 108 (2) (1998), 296-311; David Sobel, "Backing Away from Libertarian Self-Ownership", *Ethics*, 123 (1) (2012), 32-60; James Lenman, "Contractualism and Risk Imposition", *Politics, Philosophy and Economics*, 7 (1) (2008), 99-122; S D John, "Risk, Contractualism, and Rose's 'Prevention Paradox'", *Social Theory and Practice*, 40 (1) (2014), 28-50. For a very clear formulation of the problem, see Tom Dougherty, "Aggregation, Beneficence and Chance", *Journal of Ethics and Social Philosophy*, 7 (2) (2013), 1-19.

²⁶ It's an open question whether the 'series' of acts must necessarily be understood as being extended in time. Whether that makes a difference or not is likely to be particularly important for theorists who adopt ex post MSC, and so want to draw fine distinctions between different cases that would be assimilated by, e.g., Hybrid MSC II. Thanks to an associate editor here.

²⁷ As an associate editor has pointed out, if we don't already think that there is some difference between series of choices and individual choices, then Desideratum 4 might be unnecessary. However, the point of Desideratum 4 is that cases of this kind—where we initiate a series of choices each of which imposes a small risk of serious harm for

Desideratum 1 (*Aggregate Against Risks of Relevant Harms*) enjoins that our principle not be hysterical. If the harms at stake are relevant to one another, then we should allow aggregation, regardless of any disparity between the levels of risk to which different people are subjected. This is especially clear when the harms are identical. Suppose you could choose between sparing each of the Ys from a 1/10,000,000 probability of death, or instead sparing Xavier from a 1/2 probability of death. The critics of Limited Aggregation would argue that, as long as there are more than 5,000,000 Ys, you should save them, even though each individual faces a much lower risk than Xavier.²⁸

Desideratum 2 (*Don't Aggregate Against High Risk of Greater Harm*) balances the first, by insisting that Limited Aggregation should still be practically relevant. We cannot defensibly oppose aggregation only for decision-making under full information (when all probabilities are either 1 or 0, for example). Instead, for some realistic probability between 0 and 1 that an individual will suffer the greater harm, no amount of risks of the lesser harm should be enough to outweigh it.

Some anti-aggregationists would go further, resisting aggregation when the probability that *someone* will suffer the greater harm is above some threshold. However, this would be in direct tension with Desideratum 4, and is more contestable than Desideratum 2 as currently stated. If Desideratum 2 is not satisfied, then one's theory of Limited Aggregation will never in practice apply. If this strengthened desideratum is not satisfied, then one's theory will give results that some Limited Aggregationists find counterintuitive. However, if the theoretical case for those results is strong enough, we should be prepared to revise our intuitions. Moreover, I aim only to give a proof of concept, showing that anti-aggregationists can extend their theory to risk without

the sake of a small benefit, but where the series as a whole imposes a high risk of the serious harm, for the sake only of many small benefits—are the ones in which anti-aggregationism is least intuitively plausible. Since my task in this paper is to meet the opponents of anti-aggregationism head on, it is appropriate to address their concerns in this form.

²⁸ One might think that we should give some priority to dispersing risks more widely, rather than concentrating them. More on this below.

falling into the errors attributed to them by their opponents. Satisfying Desiderata 1–4 suffices for that purpose.

Desideratum 3 (*Aggregate Against Low Risk of Greater Harm*) ensures that Limited Aggregation is not paralysing in practice. Sometimes we have to tolerate risks of more serious harms to avert lesser harms. The most common criticism of all forms of moderate absolutism is its apparent failure to do this.

For example, among the possible side-effects of paracetamol are unusual tiredness and weakness, fever, and sudden pain; ibuprofen can cause intestinal haemorrhaging, dizziness, vomiting, and hypertension. We regularly take both drugs to ease minor headaches (I am doing so as I write); neither carries advice not to operate heavy machinery while under its influence. So there is some probability, when you take paracetamol or ibuprofen before driving, say, that you will suffer those rare side-effects, and lose control of the vehicle, causing serious harm to innocent bystanders. And yet it is permissible to take these medicines to treat a minor headache; we are allowed to impose small risks on others for the sake of trivial benefits to ourselves.

However, as many have noticed, in a long-enough series of choices like that covered by Desideratum 3, it eventually becomes very likely that *someone* will suffer the risked harm, while still realising only trivial goods (albeit for many). Over the course of *my* life, it is very unlikely that I will cause a serious accident because of taking paracetamol. But over a large enough population, it becomes almost certain that some such accident will happen. When we plant trees by highways, merely for the aesthetic benefit, we know that in the long run one of those trees will likely contribute to a fatal accident, which could have been avoided by not planting the tree, and foregoing the benefit. The same is true when we permit driving for frivolous purposes, like to pick up popcorn and chocolate, and when we tolerate anything other than zero blood alcohol. In all of these cases, commonsense morality says that aggregation *is* permissible, and that the trivial benefits to the many justify the very high probability that *someone* will suffer a very severe

cost. If we did not believe this, the problem of paralysis would recur.

In response to this, the many critics of anti-aggregationism have argued, in effect, that we need to endorse a further desideratum, which suggests that it is permissible to initiate and sustain series of choices of this kind, even though we know that in the long run the individually small risks will eventuate in harm. This leads to Desideratum 4 (*Aggregate Against Accumulated Low Risks of Greater Harms*), which states that, even if a long series of choices is likely to result in a greater harm at some point, while realising only trivial benefits, it can still be permissible to initiate the series, provided that, in each of those choices, the risk of the greater harm is low enough, and the overall expected benefit is high enough. We can aggregate the high probabilities of avoiding lesser harms for the many so that they outweigh the high probability that *someone* will suffer a higher-order harm (itself an aggregate of many low risks of serious harm imposed on many people).

Desideratum 4 would apparently be inconsistent with the strengthened version of Desideratum 2 considered above. Over the series of choices, we realise only trivial benefits for many, and the probability that *someone* will be harmed can be sufficiently high for (strengthened) Desideratum 2 to be triggered. Some Limited Aggregationists, then, will endorse the strengthened version of Desideratum 2, and either reject Desideratum 4 or else find some other way to accommodate it.²⁹ I present a principle that works for them in Section 5 below.

We can bring all of these desiderata together. Anti-aggregationists can answer their critics if they could defend an extension of their views to risk, which would oppose only aggregating risks of lesser harms (such as headaches) against high individual risks of greater harms (such as death). But can such a principle be formulated, generalised, and justified, on grounds other than simply that it sorts cases in the right way? Answering that question is the task of the rest of this paper.

²⁹ Thanks in particular to [omitted], for pressing me to see the different ways in which anti-aggregationists might find ways to accommodate Desideratum 4.

4. Limited Aggregation and Orthodox Decision Theory

The first step in developing a deontological decision theory is to seek guidance from orthodox decision theory. I argue elsewhere that deontologists can use its formal framework more than they might think.³⁰ We need only represent our objective moral theory with a cardinal ranking of one's available acts, given the different states the world might be in.³¹ This 'deontic value function'³² need not imply anything fundamental about the nature of our moral reasons—it is only a representation.³³ We then multiply the values of those act–state pairs by the probability of that state, and choose the option for which the sum of those products is greatest. Of course, we must make *some* changes to the decision rule—orthodox decision theory enjoins us to *maximise* expected deontic value, but deontological decision theory will afford options to act suboptimally.³⁴ But it normally seems at least subjectively *permissible* to maximise expected choiceworthiness. Limited Aggregation, however, poses a bigger problem.

MSC distinguishes between claims and mere interests, and gives the former priority over the latter. When interests compete with claims, they have no weight. We can represent this by giving mere interests zero weight whenever claims are in play. Absent claims, the interests get their regular weight. A claim's weight is determined by the weight of the interest that it protects.

This quickly yields counterintuitive results. Suppose you can save Xavier from a 99/100 risk of death or prevent a 1/10 risk of a headache for a billion Ys. Let's say that if you fail to save someone who goes on to have an illness, that is worth 0. If Xavier has a claim, then saving his life is worth 1,000,000. If the Ys have a claim to a cure for their headaches, then serving each person's claim is worth 1. Treating a group that is not in fact sick is worth 0, as is treating them when their interests are at stake, but not protected by a claim. Here is your decision table

³⁰ E.g. [omitted]. See also Mark Colyvan, Damian Cox, and Katie Steele, "Modelling the Moral Dimension of Decisions", *Noûs*, 44 (3) (2010), 503-29; Graham Oddie and Peter Milne, "Act and Value: Expectation and the Representability of Moral Theories", *Theoria*, 57 (1-2) (1991), 42-76.

³¹ For an overview, see Douglas W. Portmore, "Consequentializing", *Philosophy Compass*, 4 (2) (2009), 329-47.

³² Holly M. Smith, "The Subjective Moral Duty to Inform Oneself before Acting", *Ethics*, 125 (1) (2014), 11-38: 20.

³³ [omitted].

³⁴ [omitted].

	Xavier ill Ys ill	Xavier ill Ys not ill	Xavier not ill Ys ill	Xavier not ill Ys not ill	Expected choiceworthiness
p	0.099	0.891	0.001	0.009	
Treat Xavier	1,000,000+0= 1,000,000	1,000,000+0= 1,000,000	0+0= 0	0+0= 0	999,000
Treat Ys	0+0= 0	0+0= 0	0+1,000,000,000= 1,000,000,000	0+0= 0	1,000,000

Desideratum 2 (*Don't Aggregate Against High Risk of Greater Harm*) says that our principle should apply in some realistic cases. And yet here you are more confident of Xavier's having a claim than you are of almost anything in life, while the risk to the Ys is of only a 1/10 probability of a minor harm; yet their claims win out. Indeed, we could make the probability that Xavier has a claim arbitrarily close to 1, and the probability that the Ys have a claim arbitrarily close to 0, and yet still, there would be some number of Ys such that one should treat them rather than Xavier. If Limited Aggregation is silenced in cases like these, then it is practically irrelevant. This seems a decisive objection to incorporating MSC into orthodox decision theory.

Limited Aggregation tells us that when claims are in play, we must ignore mere interests. The decision-theoretic approach respects that condition within a given outcome. But by basing decisions on a probability-weighted average of the possible outcomes, it allows interests in one outcome to counterbalance claims in another. This explains its decisive failure.

Perhaps the Limited Aggregationist could develop a more creative deontic value function using bounded values.³⁵ Alternatively, one could ransack heterodox decision theory for a more promising decision rule (lexicographic decision theory might be one option).³⁶ Or we could return to the underlying ideas that ground our moral theory, and try to build in uncertainty at the

³⁵ [omitted].

³⁶ Chad Lee-Stronach, "Moral Priorities under Risk", *Canadian Journal of Philosophy*, Online First (2017).

ground level. The latter is my task in the next three sections.

5. Ex Ante and Ex Post Approaches

5.1 Introducing ex ante and ex post interests

Any extension of Limited Aggregation to decision-making under risk has to offer an account of two ways in which we can weigh interests in risky choices: ex ante and ex post.

In a risky decision between two alternatives—to ϕ or to ψ —everyone who has a non-zero probability of being affected has an ex ante interest. The magnitude of a person's ex ante interest is the difference between her expected well-being given that you ϕ or ψ . One's expected well-being is the probability-weighted average of one's actual well-being in the different possible outcomes of that option. For example, if Xavier is sure to live if you ϕ , but has a 1/2 probability of death if you ψ , then his ex ante interest is the difference between life for sure and 1/2 (life)+1/2 (death). The moral weight of an ex ante interest depends on its magnitude, as well as other factors like priority and responsibility—the same weights appealed to above.

Ex post interests are intuitively easy to grasp, but not always easy to measure and weigh. Instead of considering the antecedent level of risk to each person whom your action might affect, ex post interests focus on the possible outcomes that one's action might realise, and considers how people fare in those outcomes, relative to how they would have fared had you done otherwise. In other words, where ex ante interests focus on the risk of harm, ex post interests focus on what the risked harm is.

To see how things can get more complicated, we will need some more notation. Let's consider how Xavier fares in one of the outcomes if you ϕ . I'll call it outcome O_i of ϕ ing, and annotate it as O_1^\square . In a simple case, had you ψ d instead, Xavier would have lived. So his ex post interest in O_1^\square is the difference between life and death. However, suppose that, had you ψ d

instead, Xavier would have faced some non-trivial probability of dying. Then his ex post interest at O_1^\square clearly has less weight than if ψ ing would have ensured his survival. Suppose that ψ ing could have realised four equiprobable outcomes, which we can write as $O_{1\dots 4}^\square$. In O_1^\square he would have died, in the others he would have been unharmed. Had O_1^\square been the case, whether you ϕ d or ψ d would have made no difference. Had $O_{2\dots 4}^\square$ been the case, he would have lived. So his ex post interest at O_1^\square is the probability-weighted average of these differences. Since each is equally likely, this is just the average: 3/4 (life) compared with certain death.

To weigh Xavier's ex post interest, then, we must compare how he fares in a given outcome of ϕ ing with something similar to his ex ante expectation had you ψ d instead. Similar, but not identical. First, the magnitude of one's ex post interest in an outcome of ϕ ing is the difference between one's actual well-being in that outcome and one's expected well-being had one ψ d instead. This is different from one's ex ante interest, which is the difference between one's expected well-being had one ϕ d or ψ d.

Second, the probabilities of each outcome are conditional on the state of the world that led to O_1^\square being the case—if anything that is true at that outcome is causally relevant to what would have happened if you had ψ d, then it must be held constant when working out the probabilities of each outcome coming about from ψ ing. Suppose that outcome results from ϕ ing only because Xavier has a genetic condition which would affect what happens to him if you ψ instead. Then, when calculating Xavier's ex post interest at O_1^\square , we must consider only counterfactual scenarios in which Xavier has the same genetic condition.³⁷ By contrast, when calculating Xavier's ex ante interest, given that we don't know whether he has this genetic condition, we must also consider outcomes of ϕ ing and ψ ing in which he does not have the

³⁷ To keep things simple, however, all my examples will have causally independent outcomes.

condition.

In sum, the magnitude of Xavier's ex post interest S_1^x in O_1^{\square} given that you ϕ is the sum of the differences between how he fares in O_1^{\square} , and how he would have fared in each of the possible outcomes had you ψ d instead, with each difference weighted for its associated outcome's probability of coming about had you ψ d.

5.2 Ex Ante MSC

The problems with ex ante versions of anti-aggregationism are now well known.³⁸ In particular, they would breach Desideratum 1 (*Aggregate Against Risks of Relevant Harms*). We can put the case very strongly. You must choose between ϕ ing, which gives Xavier a 1/2 probability of dying, and ψ ing, so that $Y_{1,\dots,n}$ each face a 1/10,000,000 probability of death. None of the Y s have an ex ante claim: one could be required to face a 1/10,000,000 probability of death to spare someone else a 1/2 probability of death. So Xavier must have an ex ante claim. So we should help Xavier, no matter how high n is, sparing one person from a 1/2 risk of death rather than averting an arbitrarily high number of expected deaths. That's obviously wrong.

Some resolve problems like these by treating anti-aggregationism as one consideration among others, which can be outweighed by the importance of maximising overall well-being. This, too, has been shown to be a dead end.³⁹ If well-being considerations can ultimately outweigh considerations of equity (given by our anti-aggregationist theory), then, for some number of headaches, it must be all-things-considered morally better to avert those headaches than to save one person's life. This sacrifices the heart of Limited Aggregation.

³⁸ See especially: John, "Risk, Contractualism"; Horton, "Aggregation"; Frick, "Contractualism and Social Risk"; Aaron James, "Contractualism's (Not So) Slippery Slope", *Legal Theory*, 18 (03), 263-92.

³⁹ Horton, "Aggregation": 59.

5.3 Ex Post MSC

Ex post MSC is more promising. There are no fully-developed ex post theories of Limited Aggregation. So it is worth briefly pausing to explain how ex post MSC works.⁴⁰

Suppose that if you ϕ , Xavier faces a 1/100,000 probability of death, while the Y s are fine. If you ψ , the Y s each face a 9/10 probability of a minor headache, and Xavier is fine. To work out whether Xavier has an ex post claim, and its magnitude, we have to consider all the possible outcomes of ϕ ing in which Xavier dies, and work out whether his ex post interest in each scenario is the object of a claim. Ex post MSC cares about the actual harms suffered, not only about risks. So we determine relevance by considering actual harms. This means comparing death for Xavier with the probability-weighted difference between how each Y fares in O_1^\square and how she fares in each of the possible outcomes of ψ ing (conditional on the state of the world remaining as it was to lead to O_1^\square).

Xavier dies in O_1^\square , Y_r —Yolanda—is fine. In each outcome of ψ ing where Yolanda fares badly, she suffers only a headache. Xavier, by contrast, survives in every outcome of ψ ing. So the difference for Yolanda, in all those outcomes in which ψ is worse for her than ϕ , is between being fine and suffering a headache. And for Xavier, it is uniformly the difference between dying and living. Yolanda's ex post interest in avoiding a headache is not relevant to Xavier's ex post interest in his life being saved. The same is true for all the other Y s. And the same is true across all the other possible outcomes of ψ ing. So there is no question of Xavier's interest being sufficiently outweighed by relevant competing interests. Xavier has an ex post claim at O_1^\square , and indeed all the outcomes of ϕ ing in which he dies.

⁴⁰ Horton's version of ex post anti-aggregationism is on the right track, but (given that Horton does not focus on limited aggregation) obviously does not consider how to extend ex post anti-aggregationism to MSC, so it's worth briefly describing how that should be done. Horton, "Aggregation": 65. Otsuka's view is a form of ex post anti-aggregationism, but is not worked out in detail. Michael Otsuka, "Risking Life and Limb: How to Discount Harms by Their Probability", in Eyal Et Al (ed.), *Identified Versus Statistical Lives* (Oxford: Oxford University Press, 2015), 77-93.

The Y s also have ex post claims. Consider Yolanda, at O_1^{\square} , suffering a headache; Xavier is fine. To see if Xavier's ex post interest is relevant to Yolanda's ex post interest, we compare O_1^{\square} with each of the possible outcomes had you ϕ d instead. In all of those Yolanda is fine, so the difference, for Yolanda, is between a headache and being fine. But in some of those possible outcomes, Xavier dies. The difference between life and death is relevant to the difference between suffering and not suffering a minor headache. So at O_1^{\square} , Yolanda's ex post interest competes with a relevant ex post interest.

The question, though, is whether it is *sufficiently outweighed* by Xavier's relevant ex post interest. In the vast majority of the possible outcomes of ϕ ing, Xavier would have lived— ϕ ing gave him only a 1/100,000 probability of death. We cannot, therefore, give his ex post interest at O_1^{\square} the same weight as if he were sure to die if you ϕ d. We must instead discount his ex post interest for the probability that he would have been fine. Instead of comparing a certain headache for Yolanda with certain death for Xavier, we must compare Yolanda's headache with 1/100,000(death) for Xavier. I would not be required to suffer a headache in order to avert 1/100,000 expected deaths. So, Yolanda has an ex post claim, as do the other Y s.

Just as with ex post interests, ex post claims are indexed to outcomes, and must be discounted for the probability of the associated outcome coming about. So, we aggregate ex post claims on either side, weighting them all for the probability of their associated outcome being realised. The result: we are comparing 1/100,000 expected deaths with $9/10(n)$ expected headaches. For some value of n , we should avert the headaches. This satisfies Desideratum 3 (*Aggregate Against Low Risk of Greater Harm*), because we can justify taking small risks in individual cases.

On this approach, assuming that we apply a linear probability discount (which I will do for simplicity), then if two options have the same expected outcome, the moral weight of the

aggregate ex post claims in that option will be the same. Whether $Y_{1...100}$ each face a 1/10 risk of death or $Y_{1...1000}$ each face a 1/100 risk of death, their ex post claims will weigh the same in the aggregate. So Desideratum 1 (*Aggregate Against Risks of Relevant Harms*) will be satisfied. If the risked harms are of the same order, then we will maximally satisfy ex post claims by choosing the option with the better expected outcome.

Ex post MSC is a real contender. But it falls foul of Desideratum 4 (*Aggregate Against Accumulated Low Risks of Greater Harms*), as many critics of anti-aggregationism have shown for related principles. Imagine that we could repeat the case just discussed many times, and you could decide, now, whether to permit m iterations of that choice or to prevent them. You can either ϕ , initiating a series, the net effect of which is that $X_{1...m}$ each face a 1/100,000 risk of death and $Y_{1...m}$ are fine, or ψ , in which case $Y_{1...m}$ face a 9/10 probability of suffering a headache, but $X_{1...m}$ are fine. Suppose that m is 100,000, so that if you ϕ there is one expected death. Then each of the X s has an ex post claim, but none of the Y s do (one is required to suffer a headache to avert one expected death). So you should save the X s.

This is a standard complaint against Limited Aggregation: it would be paralysing, ruling out policies permitting risky activity that we typically think unproblematic. Governments and health insurance providers, for example, could not devote resources to treating minor ailments if those resources could instead save even one life. We would ban construction or infrastructure projects that, in the long run, risk causing serious injury to a very small number of people. Certain kinds of air travel would be impermissible, because they realise trivial benefits for some at the cost of a small risk of a very serious harm to others. No driving to the shop for popcorn and chocolate.

Limited Aggregationists sometimes respond by arguing that in fact there are weighty enough

interests on both sides of these cases for aggregation to be permissible.⁴¹ Perhaps more promisingly, they could distinguish between the evaluation of individual choices and the evaluation of policies or series of actions. They could argue that the permissibility of initiating a series of actions that involves high general risks of serious harm depends on whether the individual acts that compose the series involve permissibly imposing risks. So it might be permissible to initiate a series of choices like the one described in the previous paragraph, just because each risk is justified when taken in isolation, even though one knows the high general risk of serious harm that the series as a whole poses. At the same time, if one faced a single choice involving a single act, which posed the same high general risk of the great harm as in this case, but only realised trivial benefits, then that would be impermissible. I have argued elsewhere that philosophers need a theory of dynamic choice, which can explain how the permissibility of individual acts and of sequences of acts interacts (I've argued that sometimes the justification of an act depends on the justification of the series of acts of which it is part, while at other times the reverse is true).⁴² So there might be fertile ground here—anti-aggregationists might be able to develop a theory of dynamic choice that renders their view consistent with Desideratum 4, while also endorsing the strengthened version of Desideratum 2—that it is impermissible to, in a single act, impose a high general risk of serious harm (i.e. a high risk that *someone* will suffer serious harm) for the sake of only trivial benefits. I think this is drawing a long bow, but it is definitely an option, and fans of this approach will need a theory of Limited Aggregation to plug in alongside their theory of dynamic choice. Ex post MSC will serve their purposes well.

Notice, though, that for ex post MSC to work, we need to aggregate individual ex post interests, to establish whether a competing ex post interest is protected by a claim. We need, in other words, MSC, rather than one of its antecedents, complete with the Scanlon/Voorhoeve-style individualist restriction. Suppose $X_{1\dots n}$ face a $1/n$ risk of death, while $Y_{1\dots m}$ each face a $9/10$

⁴¹ See Michael Ridge, "How to Avoid Being Driven to Consequentialism: A Comment on Norcross", *Philosophy and Public Affairs*, 27 (1) (1998), 50-58.

⁴² [omitted].

risk of a headache. Ex post MSC, *with the individualist restriction*, would compare each Y 's ex post interest (headache) with each X 's ex post interest (death). It would have to discount the latter for its probability of occurring. So for a high enough n , the Y 's would have ex post claims. And for a high enough m , we should spare the Y 's rather than the X 's—averting some number of headaches rather than sparing a life. More on this in section 6.3 below.

6. Hybrid Limited Aggregation: A First Pass

6.1 Motivating a Hybrid MSC

Neither ex ante nor ex post MSC can, on its own, satisfy our desiderata. And yet each clearly captures something plausible, which the other neglects. This alone is *prima facie* grounds to develop a hybrid alternative, which retains the successes of both views but avoids their shortcomings.

Thinking about our duties of rescue lends further support to the hybrid approach. Consider the following case:

Two terrorist attacks are going to take place, one in a subway station, one in an airport departure lounge. Blast force, radius, and so on, are the same and both buildings equally densely crowded. But there is more throughput in the subway, so that attack exposes more people to a lesser risk of harm than does the airport attack. But the expected outcome in either case is the same.

If we construe duties of rescue wholly in ex ante terms, then the higher risk faced by the potential victims in the airport lounge would make a big difference to how much risk you would have to bear for their sake. And yet it does not seem to matter. Suppose that each attack, if not prevented, will cause on average 100 deaths. The central question seems to be: how much risk

are you required to bear to avert 100 expected deaths?⁴³ Whether those risks are concentrated on a few or spread across a greater number seems much less important.

But now suppose that I could spare each of $Y_{1...n}$ from a 9/10 probability of suffering a headache—but only by enduring a 9/10 probability of losing my own life. Could I be required to bear that risk, to avert 9/10(n) expected headaches? Plausibly not, no matter how high n gets. The Y s' ex post interests are irrelevant to my ex ante interest: one *would* be required to suffer a headache in order to avert 9/10 expected deaths. So no amount of expected headaches can ground a duty to bear 9/10(death) in order to prevent them.

When thinking about duties of rescue under risk, then, we should ask: how much risk can I be required to bear, to avert a bad expected outcome? In other words, how does my ex ante interest weigh against the aggregated relevant ex post interests of those whom I can try to save?

6.2 Hybrid MSC I: A First Attempt⁴⁴

Hybrid MSC I says that we should serve ex ante claims that are outweighed by relevant ex post interests, until they are outweighed by *enough* that the ex ante interest is no longer the object of a claim. The simple idea is that, as in the terrorist case, at some point one is required to bear a given risk to avert, for example, a given number of expected deaths. Ex ante interests of that kind cannot be protected by a claim.

The first step is to identify whose ex ante interests are at stake. Then, ask which of those ex ante interests is the object of a claim, by comparing each with the relevant ex post interests with which it competes. An ex ante interest is protected by a claim if and only if it is not

⁴³ N.b. 'required to bear a risk' is equivalent, in my usage, to 'required to forego saving oneself from a risk', but not equivalent to 'required to impose a risk on oneself'.

⁴⁴ In 'Can Contractualism Save Us' Barbara Fried challenges her opponents to develop a hybrid theory that is well-motivated, clear, and resolves the kinds of worries I picked out in my desiderata. Though my principle is not contractualist, it responds to that challenge (to my knowledge, it is the first of its kind). In an excellent paper that appeared soon after I first submitted this one, Joe Horton briefly considers "going hybrid", but the approach he sketches has obvious problems which lead him to dismiss it. Hybrid MSC II avoids both of Horton's objections: it gives the right verdict in his *Villain 3* case, and does not face the objection Horton posed against his ex post anti-aggregationist principle—essentially that it breaches Desideratum 4. Horton, "Aggregation": 73.

sufficiently outweighed by the relevant competing ex post interests. We can be guided by this question: would one be required to bear that much risk to avert that expected outcome? Finally, maximise satisfaction of (morally weighted) ex ante claims.

HYBRID MSC I

In a choice whether to ϕ or ψ , in which either $X_{1\dots n}$ or $Y_{1\dots n}$, but not both groups, are at risk:

1. X_n 's ex ante interest S_n^x is the difference between her expected well-being conditional on ϕ ing and her expected well-being conditional on ψ ing.
2. The moral weight of an ex ante interest is determined by its magnitude, weighted by priority, responsibility, etc.
3. For any member Y_n of $Y_{1\dots n}$, Y_n 's ex post interest S_n^y in an outcome of ϕ ing O_1^\square is the probability-weighted average of the differences between how Y_n fares in O_1^\square and how Y_n would have fared in $O_{1\dots n}^\square$.
4. The moral weight of S_n^y at O_1^\square is determined by its magnitude, weighted by priority, responsibility, etc.
5. S_n^y competes with S_n^x if and only if one of the options is better than the other for X_n , but worse for Y_n .
6. S_n^y is **relevant** to S_n^x at O_1^\square if and only if, for some possible outcome of ψ O_n^\square , the difference between Y_n 's well-being in O_n^\square and O_1^\square is not sufficiently outweighed by S_n^x .
7. When aggregating $Y_{1\dots n}$'s relevant competing ex post interests, each is discounted by the probability of its associated outcome being actual.
8. X_n 's ex ante interest is the object of a claim if and only if it is not sufficiently outweighed by $Y_{1\dots n}$'s relevant competing ex post interests.
9. The moral weight of an ex ante claim is the moral weight of the underlying ex ante

interest.

10. *Maximise the sum of the moral weights of satisfied ex ante claims.*

11. *If more than one option maximizes the sum of the moral weights of satisfied ex ante claims (perhaps because there are no ex ante claims to satisfy), then you may choose the one that maximizes overall morally weighted expected well-being.*

To see how Hybrid MSC I works, it will help to work through the four desiderata, starting with those that it satisfies well. Consider Desideratum 2 (*Don't Aggregate Against High Risk of Greater Harm*).

Suppose that ϕ ing leaves Xavier exposed to a 9/10 risk of death, fully curing the Y s; ψ fully cures Xavier, leaving the Y s exposed to a 9/10 risk of a minor headache. Per (6) above, the ex post interests of the Y s are not relevant to Xavier's ex ante interest, because the latter sufficiently outweighs each Y 's ex post interest: one would be required to suffer an actual headache to avert 9/10 expected deaths. So, per (8), Xavier has an ex ante claim—his ex ante interest in avoiding a 9/10 risk of death is not sufficiently outweighed by the Y s' relevant competing ex post interests, because their ex post interests aren't relevant to that ex ante interest. However high n gets, one is not required to absorb a 9/10 risk of death in order to avert 9/10(n) expected headaches.

The Y s, meanwhile, lack ex ante claims. Xavier's ex post interest in outcomes where he dies is clearly relevant to each Y 's ex ante interest—per (6), 9/10(headache) clearly does not sufficiently outweigh the difference between life and death. To work out whether Xavier's ex post interest sufficiently outweighs each Y 's ex ante interest, we need to discount the former for its probability of being actual, per (7). Even once that is done, Xavier's ex post interest does not sufficiently outweigh each Y 's ex ante interest, per (8): one would be required to bear a 9/10 risk of a headache to avert 9/10 expected deaths.

So, only Xavier has an ex ante claim, and no matter how many Y s there are, we should avert the risk of death, rather than the risk of headaches. So Hybrid MSC I is potentially practically

relevant—it doesn't apply only when the probability of the greater harm is 1.

But Hybrid MSC I is not overly restrictive. Desideratum 3 (*Aggregate Against Low Risk of Greater Harm*) is also satisfied. Suppose that if you ϕ , Xavier faces a 1/10,000 risk of death, and the Ys will be fine. If you ψ , Xavier will be fine, and $Y_{1\dots n}$ will each face a 9/10 risk of a minor headache. First, you have to work out whether Xavier's ex ante interest is the object of a claim, per (8), above. Clearly the Ys' ex post interests compete with Xavier's ex ante interest: ψ ing is in Xavier's ex ante interest, whereas ϕ ing is in the Ys' ex post interests. Are their ex post interests *relevant* to Xavier's ex ante interest? I think they are. One is plausibly not required to suffer a headache in order to spare someone else a 1/10,000 risk of death. Per (6), Xavier's ex ante interest does not sufficiently outweigh the ex post interests of each Y, so the latter are relevant to the former. We then have to discount each ex post interest according to its probability of occurring, as per (7), and aggregate them. For some value of n , Xavier's ex ante interest will not be the object of a claim. One plausibly *is* required to accept a 1/10,000 risk of death in order to avert 9,000,000 expected headaches, for example.

What about the Ys? Consider Y_1 , Yolanda. Her ex ante interest is in avoiding a 9/10 risk of suffering a headache. It competes with Xavier's ex post interest in avoiding death. Clearly Xavier's ex post interest is relevant to Yolanda's ex ante interest, per (6). But it plausibly does not sufficiently outweigh it: as we saw when developing ex post MSC, ex post interests are indexed to particular outcomes, and after establishing relevance, they must be discounted for the probability of that outcome coming about (per (7)). One would not be required to suffer a 9/10 risk of a headache in order to avert 1/10,000 expected deaths. So Yolanda *does* have an ex ante claim (per (8)). The same is true for the other Ys. Since (if n is high enough) only the Ys have ex ante claims, according to Hybrid MSC I, you ought to ϕ , saving the Ys. So when the risk of the greater harm is low enough, we can aggregate risks of lesser harms to override it.

Desideratum 4 (*Aggregate Against Accumulated Low Risks of Greater Harms*) also poses no special

problems. Suppose we repeat the case just considered, so that you will face m iterations of the same choice. If you ϕ , $X_{1\dots m}$ each face a $1/10,000$ risk of death and $Y_{1\dots m}$ are fine, or you can ψ , in which case $Y_{1\dots m}$ face a $9/10$ probability of suffering a headache, but $X_{1\dots m}$ are fine. If n was high enough that in the one-shot version of the case Xavier lacks an ex ante claim, then a fortiori none of the X s has an ex ante claim in this repeat version.

$Y_{1\dots m}$ plausibly lack ex ante claims as well, for some values of m . $X_{1\dots m}$'s ex post interests (in avoiding death) clearly are relevant to an ex ante interest in avoiding a $9/10$ probability of a headache (per (6)). And though we must weight each ex post interest for the probability of its associated outcome coming about (per (7)), we must then aggregate them. So if m were 10,000, say, then the moral weight of the aggregated, relevant, ex post interests of $X_{1\dots m}$ would be equivalent to one expected death. And one clearly could be required to bear a $9/10$ risk of a headache to avert one expected death, so the Y s would lack ex ante claims (per (8)).

So, for some values of m and n , neither $X_{1\dots m}$ nor $Y_{1\dots m}$ will have ex ante claims, and we can simply minimise expected harm—presumably by treating the Y s, though the alternative result is also possible. Whether we think about these options one by one, or as a sequence, we will get the same verdicts. If it's permissible to ϕ on each separate occasion, it will be permissible to choose to perform the series of actions of ϕ ing on each occasion. The mere repetition of risk-taking doesn't threaten to decisively undermine Hybrid MSC I.

However, Desideratum 1 (*Aggregate Against Risks of Relevant Harms*) remains. Hybrid MSC I improves on ex ante MSC. It will not tolerate prioritising those exposed to the greatest risk at the cost of permitting an arbitrarily high number of deaths among others. But it still allows a counterintuitive degree of inefficiency when the risked harms are the same.

Suppose you can spare Xavier from a $1/2$ risk of death or $Y_{1\dots n}$ from a $1/10,000,000$ risk of death. Ex ante MSC would say that one should save Xavier, no matter how high n gets. That's crazy, and Hybrid MSC I need not reach that implausible conclusion. Clearly, at some point

Xavier's ex ante interest in avoiding a 1/2 risk of death is sufficiently outweighed by the expected deaths if you don't save the Ys, so Xavier lacks an ex ante claim per (8). The ex post interests of the Ys are relevant—one is not required to forego saving one's life to avert 1/2 expected deaths (per (6)). And, even when discounted for probability, these ex post interests together sufficiently outweigh Xavier's ex ante interest per (7). Where is this point? 10, 100, 1,000 expected deaths? I don't know. But it cannot rise arbitrarily high. For some n , one is required to bear a 1/2 risk of death in order to avert $1/10,000,000(n)$ expected deaths.

And yet, are we satisfied with this conclusion? Suppose one can be required to bear a 1/2 risk of death to avert 100 expected deaths, so that's where Xavier's ex ante interest is sufficiently outweighed. Are we prepared to endorse a principle that permits us to avert 1/2 expected deaths when we could have averted 99? Some nonconsequentialists will find this acceptable—we tolerate inefficiencies elsewhere, so why not here? They will therefore defend a weaker version of Desideratum I, permitting a greater amount of inefficiency than I think plausible. They can endorse Hybrid MSC I. My aim, however, is to satisfy the four desiderata as stated. So our task is not yet done. But the path to a more successful alternative is now clear.

6.3 Risk and the Individualist Restriction

Before exploring that more successful alternative, it's worth pausing to note how Hybrid MSC I makes clear the pay-off of the innovation in MSC, introduced in Section 2.1: to modify the 'individualist restriction' in existing theories of Limited Aggregation. According to those theories, to see whether an interest is protected by a claim, one may make only bilateral comparisons between individuals. MSC by contrast, allows an interest to fail to be protected by a claim because it is sufficiently outweighed by the relevant competing interests of others, *in the aggregate*.

Individualist Limited Aggregation could not ground a plausible hybrid theory. Suppose you can spare Xavier a 9/10 risk of a minor headache or spare $Y_{1...n}$ a $1/n$ risk of death. Is one

required to bear a 9/10 risk of a minor headache in order to avert one expected death? Individualist Limited Aggregation cannot ask this question. That one expected death is a composite of the n different scenarios in which one member of $Y_{1,...,n}$ dies. To reach it, we have to aggregate all the Y s' ex post interests. Individualist theories of Limited Aggregation permit aggregation only when we have already established that there are claims, and don't countenance the possibility that one's interest might lack the protection of a claim just because it is sufficiently outweighed by a *group* of competing interests. So this natural way of understanding ex ante claims is unavailable.

An individualist Hybrid MSC would have to establish whether we have ex ante claims based on how one's ex ante interest compares with each individual ex post interest. The first problem: this severs the link between MSC and our duties of rescue. If we did not sever that link, then the individualist restriction would imply that the degree of risk I can be required to bear for the sake of others is a function of the largest competing ex post interest at stake, rather than the aggregate of all relevant ex post interests. So I would be required to bear the same degree of risk to save $Y_{1,...,n}$ from a 1/2 probability of death irrespective of how high n is—whether it is one or a billion. That's not a plausible view.

Second, we would face a dilemma when comparing ex ante and ex post interests. We can weight the ex post interest in one of two ways; neither is appealing. First, we could simply treat each ex post interest as having the weight of the actual harm—for example, death. But if we do that, then Xavier (who faces, recall, a 9/10 risk of a headache) will *never* have an ex ante claim, no matter what the risk to the Y s is, and how few they are. This alone is cause for concern—suppose $X_{1,...,n}$ all face a 9/10 risk of a headache, which conflicts only with Yolanda's 1/1,000,000,000 risk of death. It would be very strange to say that none of the X s has ex ante claims.

Alternatively, we can discount each ex post interest for its probability of being actual. So let's

suppose that $Y_{1...n}$ face a $1/m$ probability of death, where $1/m$ is greater than $1/n$, so the number of expected deaths is greater than 1, and can be arbitrarily high. Xavier meanwhile faces only a 9/10 risk of a headache. There is some m such that Xavier is not required to bear a 9/10 risk of a headache to avert $1/m$ expected deaths (that is, an ex post interest of avoiding death, discounted by the $1/m$ probability of it being actual). Equally, though, there is some m such that none of the Y s have ex ante claims, since one can be required to bear a $1/m$ risk of death to avert 9/10 expected headaches. So, if m is above both of these thresholds, then Xavier has an ex ante claim, but $Y_{1...n}$ don't, even though n can be arbitrarily high, and the number of expected deaths if you help Xavier can likewise be arbitrarily high. This looks like the wrong way to understand ex ante claims. For some m and some n , one *is* required to bear a 9/10 risk of a headache to avert $1/m(n)$ expected death.

7. A More Promising Hybrid View

Hybrid MSC I tolerates excessive inefficiency when the risked harms are of the same order. To deliver on our desiderata, we need to argue that when a set of ex ante claims is outweighed by relevant ex post interests, we should serve the latter—even if they don't *sufficiently* outweigh the ex ante interests. I think we can make just this argument, drawing on the objectivist version of MSC with which we began. According to that view, if a set of correlated claims (that is, claims that can be jointly served) is outweighed by a set of relevant competing interests, then even if the former are not *sufficiently* outweighed, we should still serve the weightier set of interests. We should serve ex ante claims only when they are not outweighed by the competing relevant ex post interests.

Let's say that when a set of correlated ex ante claims is together outweighed by the set of competing relevant ex post interests, the former is *defeated* by the latter. A principle that maximises overall satisfaction of undefeated sets of ex ante claims will succeed. Here is that

principle:

HYBRID MSC II

In a choice whether to ϕ or ψ , in which either $X_{1\dots n}$ or $Y_{1\dots n}$, but not both groups, are at risk:

1. X_n 's ex ante interest S_n^x is the difference between her expected well-being conditional on ϕ ing and her expected well-being conditional on ψ ing.
2. The moral weight of an ex ante interest is determined by its magnitude, weighted by priority, responsibility, etc.
3. For any member Y_n of $Y_{1\dots n}$, Y_n 's ex post interest S_n^y in an outcome of ϕ ing O_1^\square is the probability-weighted average of the differences between how Y_n fares in O_1^\square and how Y_n would have fared in $O_{1\dots n}^\square$.
4. The moral weight of S_n^y at O_1^\square is determined by its magnitude, weighted by priority, responsibility, etc.
5. S_n^y competes with S_n^x if and only if one of the options is better than the other for X_n , but worse for Y_n .
6. S_n^y is **relevant** to S_n^x at O_1^\square if and only if, for some possible outcome of ψ O_n^\square , the difference between Y_n 's well-being in O_n^\square and O_1^\square is not sufficiently outweighed by S_n^x .
7. When aggregating $Y_{1\dots n}$'s relevant competing ex post interests, each is discounted by the probability of its associated outcome being actual.
8. X_n 's ex ante interest is the object of a claim if and only if it is not sufficiently outweighed by $Y_{1\dots n}$'s relevant competing ex post interests.
9. The moral weight of an ex ante claim is the moral weight of the underlying ex ante interest.
10. A set of ex ante claims is correlated if and only if all claims in the set can be jointly served.
11. An ex post interest is relevant to a set of ex ante claims if and only if it is not sufficiently

outweighed by that set of ex ante claims.

12. *A set of correlated ex ante claims is defeated if and only if it is outweighed by the set of relevant ex post interests with which it competes.*

13. *Maximise the sum of the moral weights of undefeated sets of satisfied ex ante claims.*

14. *If more than one option maximizes the sum of the moral weights of undefeated sets of satisfied ex ante claims (perhaps because there are no undefeated sets of ex ante claims to satisfy), then you may choose the one that maximizes overall morally weighted expected well-being.*

The key difference from Hybrid MSC I comes in (10)–(13), which describe how to aggregate correlated sets of ex ante claims and weigh them against relevant ex post interests. (10) explains what it means for a set of ex ante claims to be correlated, and (11) determines which ex post interests are relevant to the evaluation of sets of ex ante claims. An ex post interest is relevant to a set of ex ante claims if and only if it is not sufficiently outweighed by that set of ex ante claims. The idea of relevance is the same as it has been throughout: if one would be required to forego saving oneself from a given cost for the sake of satisfying the interests that compete with it, then one's interest should not be counted (until all claims have been dealt with).

According to (10)–(13), necessarily, if the risked harms are of the same order, then if the set of correlated ex ante claims is defeated by the set of competing relevant ex post interests, then not serving those ex ante claims will maximise expected morally weighted well-being. Conversely, if the set of correlated ex ante claims is not defeated by the set of competing relevant ex post interests, then, necessarily, serving the ex ante claims will maximise expected morally weighted well-being. And if there are no ex ante claims at stake, then of course we also maximise expected morally weighted well-being.

This means that Hybrid MSC II fares better against Desideratum 1 (*Aggregate Against Risks of Relevant Harms*). If you ϕ , Xavier faces a 1/2 probability of death, $Y_{1\dots n}$ are fine; if you ψ , $Y_{1\dots n}$ face a 1/1,000,000 risk of death, Xavier is fine. Suppose n is 10,000,000. Xavier plausibly has an

ex ante claim: one is not required to suffer a 1/2 probability of death in order to avert 10 expected deaths. Even though the competing ex post interests are *relevant* to Xavier's interest, they don't outweigh it by enough to prevent his interest being protected by a claim. $Y_{1...n}$, on the other hand, plausibly *do not* have ex ante claims. Their ex ante interest *is* sufficiently outweighed. One *is* required to bear a 1/1,000,000 risk of death to avert 1/2 expected deaths.

So far, this is just the same as Hybrid MSC I, which would say that one should serve Xavier's claim in this case. But Hybrid MSC II insists that if Xavier's claim is not part of a *set of undefeated* ex ante *claims*, we should not serve it. In this case, the set has a single member. The ex post interests of $Y_{1...n}$ are all relevant to Xavier's ex ante interest—per (11) and (6)—because one is not required to suffer death in order to spare someone else a 1/2 risk of death. And the aggregated ex post interests of $Y_{1...n}$ defeat the set of correlated ex ante claims—in this case just that of Xavier. 10 expected deaths outweighs 1/2. So there are no undefeated sets of ex ante claims to serve, so we should maximise morally weighted expected well-being, by saving $Y_{1...n}$.

Perhaps some will believe Desideratum 1 too strong—we should allow *some* priority for concentrated risks against widely dispersed ones. If we could choose between one person losing three fingers, and n people facing a $1/n$ risk of losing four fingers, then for a high enough n we should choose the former option, even though it doesn't maximise morally weighted expected well-being.⁴⁵ We could argue that it is, other things equal, objectively worse to suffer a given harm the higher the risk of that harm one was subjected to.⁴⁶ In other words, to lose three fingers through action that had a high risk of realising that result *might* be worse than to lose four fingers through action that had a very low risk of leading to that result. If that were right, then it would justify giving some priority to concentrated risks. Or perhaps there is some other way to motivate the idea that concentrated risks should get some additional priority—for example, on the grounds that a wider dispersal of risk is fairer. If so, it should be taken into account when

⁴⁵ Thanks to an associate editor for this objection.

⁴⁶ [omitted].

assessing the moral weight of one's ex ante or ex post interest, as per (2) and (4) in Hybrid MSC II.

I think that, where the probabilities at stake concern only probability of harm—rather than, for example, probability that the person who will be harmed is liable to bear it—our reason to avoid concentrated risks is no more than a tie-breaker.⁴⁷ However, the principles in this paper are flexible enough to accommodate different views on this point. Those who think ex ante claims should get more protection can either endorse Hybrid MSC I, if they think it gets the balance right, or else endorse Hybrid MSC II together with an explanation of why we should give additional moral importance to concentrated risks. Conversely, those who think that the preference for dispersed over concentrated risk can only be a tie-breaker, should build that into the distributive principle at (14)—that is, if the options are tied with respect to both claims and overall morally weighted well-being, then choose the option that more widely disperses the risk.⁴⁸ Hybrid MSC II insists only that, if we want to give ex ante interests some additional priority, we come up with some justification for doing so. The default assumption is that ex ante and ex post interests weigh the same—since they are just two different ways of compounding probabilities and interests, this is a reasonable starting point.

Again, it's important to remember what we're trying to do here: the critics of anti-aggregationism have, on my reading, taken Desideratum I for granted, and argued that anti-aggregationists cannot plausibly accommodate it when their views are extended to decision-making under risk. I have shown that this is wrong. Now, some anti-aggregationists may push back against Desideratum 1, and argue for greater priority for ex ante interests. I've shown how their views can be accommodated within my principles, but my primary task in this paper is not to fend off objections from avowed anti-aggregationists, but to refute the critics of anti-

⁴⁷ Otsuka, "Risking Life and Limb".

⁴⁸ For a defence of this tie-breaker view, see Norman Daniels, "Can There Be Moral Force to Favoring an Identified over a Statistical Life?", in Eyal Et Al (ed.), *Identified Versus Statistical Lives* (Oxford: Oxford University Press, 2015), 110-23.

aggregationism, by showing that anti-aggregationists really can accommodate risk.

Desiderata 2–4 are handled by Hybrid MSC II in the same way as Hybrid MSC I. Suppose that ϕ leaves Xavier exposed to a 9/10 risk of death, fully curing $Y_{1\dots n}$; ψ fully cures Xavier, leaving the Y s exposed to a 9/10 risk of a minor headache. The ex post interest of each Y is not relevant to Xavier's ex ante interest: each Y is required to suffer a headache in order to avert 9/10 expected deaths. So no matter how high n gets, (from 8) Xavier has an ex ante claim: one is not required to absorb a 9/10 risk of death to avert 9/10(n) expected headaches, for any n .

The Y s, on the other hand, do not have ex ante claims. Their ex ante interest in avoiding a 9/10 risk of a headache comes up against 9/10 expected deaths. One is required to bear a 9/10 risk of a headache to avert 9/10 expected deaths. So only Xavier has an ex ante claim, and you should help him, no matter how high n is. This satisfies Desideratum 2 (*Don't Aggregate Against High Risk of Greater Harm*). According to Hybrid MSC II, *only* ex ante interests enjoy this kind of absolute protection against being overridden by aggregated lesser ex post interests.

Desideratum 3 (*Aggregate Against Low Risk of Greater Harm*) is also satisfied. Consider our earlier case: if you ϕ , Xavier faces a 1/10,000 risk of death, and $Y_{1\dots n}$ will be fine. If you ψ , Xavier will be fine, and $Y_{1\dots n}$ each face a 9/10 risk of a minor headache. Per (6), the ex post interest in avoiding a headache *is* relevant to an ex ante interest in avoiding a 1/10,000 risk of death (one is not required to forego avoiding a headache to avert a 1/10,000 risk of death for someone else). As before, per (8), if n is high enough, then Xavier lacks an ex ante claim: one *is* required to bear a 1/10,000 risk of death to avert 9/10(n) expected headaches. The Y s do have ex ante claims—though the ex post interest in avoiding death is clearly relevant to the ex ante interest in avoiding a 9/10 risk of a headache, when working out whether it sufficiently outweighs the ex ante interest, we need to discount it for its probability of coming about (as per (7)). And, per (8), one is not required to bear a 9/10 risk of a minor headache in order to avert 1/10,000 expected deaths. So only the Y s have ex ante claims, and you should help them.

The same is true for Desideratum 4. If we repeat the previous case m times, we still find that none of the X s has an ex ante claim, since if one is required to bear a $1/10,000$ risk of death in order to avert $9/10(n)$ expected headaches, then a fortiori this is also true for $9/10(mn)$ headaches. As such, per (8), when mn is high enough, Xavier lacks an ex ante claim.

If m is below some bound, that is, if the case will only be repeated a few times, so that the risk of death to be averted remains low, then $Y_{1\dots m}$ might still have ex ante claims, as in the previous case, and you should help them. But if m is above that bound, for example if m is 9,000, so that the aggregate expected deaths are $9/10$, then $Y_{1\dots m}$ lack ex ante claims, and you should maximise morally weighted expected well-being. Each Y 's ex ante interest competes with the ex post interests of *all* of $X_{1\dots m}$. Each of those ex post interests is relevant to each Y 's ex ante interest, per (6)—one is not required to suffer death to spare someone else a $9/10$ headache risk. We then discount the ex post interests by probability, and sum them, as per (7). So if m is 9,000, the combined weight of the X s' ex post interests is $9/10$ expected deaths, and each Y lacks an ex ante claim: one is required to bear a $9/10$ risk of a headache in order to avert $9/10$ expected deaths. Since neither the X s nor the Y s have ex ante claims, we can maximise morally weighted expected well-being, which, for at least some values of m and n , will mean helping the Y s, rather than the X s. So we can aggregate against accumulated risks of the greater harm.

Is Hybrid MSC II genuinely explanatory? Yes. Here is a principle that would be roughly extensionally equivalent to Hybrid MSC II: *Always Aggregate, Except When Risks of a Lesser Harm Compete with a High Enough Individual Risk of a Greater Harm*. This principle tells us almost nothing. It cannot help us figure out what 'high enough' means, or what 'lesser' and 'greater' mean here, and it offers no explanation for its verdicts. Hybrid MSC II explains where that threshold should be, and why, which harms count as greater and lesser, and why, and how to understand and justify aggregation when it is permissible, and it defines and explains ex ante claims and ex post interests, all in a principle that is deeply tied into basic currents of deontological thought, in

particular our duties of rescue. This is not just a matter of redescribing intuitions about cases.

If claims take priority over interests, then why should we ignore ex ante claims that are merely outweighed (not *sufficiently* outweighed) by relevant competing ex post interests? Here I do have to depart from my earlier assertion that, when claims are at stake, mere interests must be disregarded. Relevant ex post interests can, in the aggregate, undermine an ex ante claim. An interest's relevance to the competing interests, not the fact that it is the object of a claim, is the necessary condition for aggregation of interests being admissible.

But if ex post interests are relevant, and can outweigh ex ante claims, then why aren't they too protected by claims? Why shouldn't we say that failure to satisfy these relevant ex post interests wrongs the victim? That would be just as consistent an extension of objectivist MSC, but it would force us towards the results of ex post MSC, and in particular lead to violating Desideratum 4. Here I think Hybrid MSC II stands in need of more argument.

Claims are fundamentally grounded in the interactional dimension of morality, and in particular the requirement to treat one another with appropriate respect. Whether or not I treat you with respect cannot simply be determined by objective facts such as whether or not your interest is satisfied. Instead, there must be subjective elements as well—in particular, the fact that I knew, or should have known, that my action would have the effect, say, of thwarting your interest.⁴⁹

In any case where the victims of your choice have relevant ex post interests, but don't have ex ante claims, this must be because each individual is subjected only to a small probability of harm (given that the harm is serious, if the probability were high, they would have an ex ante claim, not just a relevant ex post interest). On my view, that fact alone explains why, if their interest is ultimately thwarted through action that is all-things-considered justified—because it serves an undefeated set of ex ante claims, for example—*they* don't have a residual complaint,

⁴⁹ For this kind of view, see, inter alia, David McCarthy, "Rights, Explanation, and Risks", *Ethics*, 107 (2) (1997), 205-25; Michael J. Zimmerman, *Living with Uncertainty: The Moral Significance of Ignorance* (Cambridge: Cambridge University Press, 2008); Jonathan Quong, "Rights against Harm", *Aristotelian Society Supplementary Volume*, 89 (1) (2015), 249-66.

consistent with their being pro tanto wronged. An individual having a claim against being harmed, of the kind that grounds a residual complaint against all-things-considered permissible action, depends not only on whether they actually end up suffering harm, but on the risk to which they are subjected. There is an essential epistemic component to claims.

Again, we can draw inspiration from duties of rescue. These too have an essential epistemic component. Right now, if I were to pick up the phone and call the right person in Brazil, say, I could probably avert a murder. Doing so would cost me hardly anything, and would save a life. But I am not currently breaching a duty of rescue, because I have no idea where in Brazil a murder is currently taking place. I only have duties of rescue when my evidence—or, the evidence I would have if I were being appropriately diligent—indicates that I can help someone at a reasonable expected cost to myself. As with our duties of rescue, so with our claims to aid: in both cases, they are defined in epistemic terms.

We should protect only ex ante interests with claims. But ex post interests matter too. Indeed, as they are just a different way to compound probabilities and people's well-being in risky choices, unless we have some specific explanation for weighing them differently, the default assumption is that they have the same weight as ex ante interests. We reflect the interactional dimension of morality by protecting only ex ante interests with claims, and by basing our choice of an option on how it fares with respect to ex ante claims. We reflect the continued importance of ex post interests by giving them the same weight as the ex ante interests with which they compete, unless there is some specific reason to weigh them differently (for example, on grounds of fairness).

Hybrid MSC II shows that a sensible extension of Limited Aggregation to risk is possible. Of course, some adherents of Limited Aggregation will not be satisfied by it—most likely because they favour a strengthened version of Desideratum 2 (*Don't Aggregate Against High Risk of Greater Harm*), according to which what matters is not the risk of any particular person suffering the

greater harm, but the risk that *someone* will suffer the greater harm. Hybrid MSC II would say that there is a big difference between a standard case of *Life for Headaches*, where the person whose life is at stake is known, and one where we know that one out of a billion people will die if we avert the headaches, but we don't know who. In the former case, the person whose life is at stake has an undefeated ex ante claim, which cannot be outweighed by any number of headaches that aren't protected by claims. In the latter case, nobody has an ex ante claim, so it is permissible to maximise expected morally weighted well-being. Hybrid MSC II says that claims are part of the interactional dimension of morality, and so should be sensitive to the agent's epistemic position, as with duties of rescue. Those who disagree with this verdict, and think that risky cases with this structure should be assimilated to the original *Life for Headaches* case, should reject Hybrid MSC II and adopt ex post MSC instead. They then need to figure out how to accommodate cases like those raised by Desideratum 4, where we can initiate series of choices that realise only minor benefits, through acts that individually create small individual risks of greater harms, where the series as a whole creates a high general risk that someone will suffer the greater harm. They can either argue that the real life choices that we make, which seem to be accurately described by Desideratum 4, actually do not meet that description, or else they can develop a theory of dynamic choice to explain why it can be impermissible to create a high risk that someone will be harmed through an individual action, but permissible to create the same risk by initiating a series of actions.

8. Conclusion

I have offered a menu of options for deontologists tempted by the standard verdict on *Life for Headaches*, but concerned that they will not be able to sensibly extend their view to decision-making under risk. First, they can look deeper into heterodox decision theory; there may be resources there for improving on the obvious failure of its orthodox counterpart. I am not

optimistic, but nothing I have said here rules this route out.

The second choice is ex post MSC. This should suit those who think that the ultimate right- and wrong-makers are facts about the objective world, which don't take our epistemic position into account. Of course, ex post MSC breaches Desideratum 4 (*Aggregate Against Accumulated Low Risks of Greater Harms*), so these folks must concentrate their energies on either rejecting this desideratum or developing a theory of dynamic choice to accommodate it. They should aim to show that imposing or allowing the same general risk of harm through an individual act, and by initiating a series of acts, is qualitatively different, so that the first can be impermissible, the second permissible. That would allow them to answer anti-aggregationism's critics. There are some promising avenues to pursue here. But it is hard to believe that people's ex ante interests can simply be ignored. After all, in almost every decision-making situation that we actually face, they are all we have to go on.

The third option is Hybrid MSC I. This is perhaps more elegant than its replacement, and suits those who think that the ultimate right- and wrong-makers take epistemic position into account, and who are ultimately not too opposed to some inefficiency in morality, when it comes to saving lives. The obvious challenge is to defend that inefficiency. Of course, one could argue that we have more extensive duties of rescue than I have assumed, thereby reducing the degree of inefficiency that this principle would tolerate.

And finally there is Hybrid MSC II. It hits all the intuitive touchstones, and jibes well with the view that our claims on others must be sensitive to their information, while we should nonetheless give ex post interests a prominent role in our moral theory. It seems to balance ex post and ex ante interests in a well-motivated and appealing way. It shows that anti-aggregationists can extend their theory to decision-making under risk without falling into any of the errors of which their opponents have accused them. More importantly, together with these alternative principles, it shows that the prospects are good for deontological decision theory.