From the holdings of the Department of Special Collections , Princeton University Library, Princeton, NJ

These documents are either in the public domain, or can only be used for educational and research purposes (Fair use) as per U.S. Copyright law (text below). By accessing this file, all users agree that their use falls in the public domain or within fair use as defined by the copyrightlaw. We do not charge any permission or usage fees for the publication of images of material inour collections, including those provided by Princeton University Library, via our website, catalog, or directly from staff.For more information regarding copyright, credit, and citation please review the information onthe Copyright, Credit and Citation Guidelines page of our website.
https://library.princeton.edu/special-collections/policies/copyright-credit-and-citation-guidelines. If you have any questions, including how to order higher-quality images from the collection, contact Special Collections Public Services via our Ask Us! Form. https://library.princeton.edu/special-collections/ask-us

---

Doctoral Dissertations, 2011-2026 Philosophy

2021

MAKING THE NUMBERS COUNT WITHOUT COUNTING THE NUMBERS: A NONCONSEQUENTIALIST ACCOUNT OF INTERPERSONAL TRADE-OFFS

Zhang, Yuan

Princeton, NJ : Princeton University

---

http://arks.princeton.edu/ark:/88435/dsp01mk61rm08r

U.S. Copyright law: The copyright law of the United States (Title 17, §108, United States Code) governs the making of photocopies or other reproductions of copyrighted material. Under certain conditions specified in the law, libraries and archives are authorized to furnish a photocopy or other reproduction. One of these specified conditions is that the photocopy or reproduction is not to be used for any purpose other than private study, scholarship, or research. If a user makes a request for, or later uses, a photocopy or reproduction for purposes in excess of fair use that user may be liable for copyright infringement. This institution reserves the right to refuse to accept a copying order if, in its judgment, fulfillment of the order would involve a violation of copyright law. https://www.copyright.gov/title17/title17.pdf

MAKING THE NUMBERS COUNT WITHOUT COUNTING THE NUMBERS: A NONCONSEQUENTIALIST ACCOUNT OF INTERPERSONAL TRADE-OFFS

Erik Yuan Zhang

A DISSERTATION

PRESENTED TO THE FACULTY

OF PRINCETON UNIVERSITY

IN CANDIDACY FOR THE DEGREE

OF DOCTOR OF PHILOSOPHY

RECOMMENDED FOR ACCEPTANCE BY

THE DEPARTMENT OF

PHILOSOPHY

Advisers: Professor Johann Frick

Professor Michael Smith

September 2021

© Copyright by Yuan Zhang, 2021. All rights reserved.

Abstract

This dissertation develops and defends a nonconsequentialist account of interpersonal trade-offs. The account finds application in situations in which different individuals have interests in being aided that cannot be jointly satisfied.

Chapter One articulates a theoretical foundation for reasoning about interpersonal trade-offs. The framework I defend is an *individualist* one: In adjudicating between the competing interests of different individuals, our decision should be justifiable to *each* of those affected. I show that individualist moral theories, most notably Scanlonian contractualism, can be sensitive to the significance of numbers without resorting to the familiar consequentialist idea of maximizing aggregate well-being. In short, nonconsequentialists can make the numbers count without directly counting the numbers.

Chapter Two further develops one key result arrived at by the end of Chapter One, namely, individualist moral reasoning, suitably understood, supports a general approach to interpersonal trade-offs commonly known as *Limited Aggregation*. I propose a novel development of Limited Aggregation, which I call the Partial Primacy Account (PPA). I show that PPA is not only theoretically well-motivated, it is also capable of convincingly handling the purported counterexamples to Limited Aggregation that have been advanced in the recent literature.

Finally, in Chapter Three, I examine an issue that arises in the context of adjudicating between competing individual interests under conditions of risk and uncertainty. Here, I wade into the ongoing debate between the *ex ante* and the *ex post* approaches to risk, and I offer a qualified defense of the *ex post* framework. My defense is qualified in that it allows for the relevance of *ex ante* reasoning in a range of trade-off

situations. My discussion can be thus understood as advocating for a hybrid account of risk.

Table of Contents

Abstract ..... **iii**

Acknowledgement ..... **vi**

Chapter One: Moral Contractualism and the Numbers Problem **1**

1. Introduction 1

2. An Overview of Contractualism 4

3. Clearing the Ground 8

4. Argument from Normative Disabling 16

5. Argument from Equal Consideration 22

6. Conclusion 33

Chapter Two: How Should the Numbers Count? **35**

1. Introduction 35

2. The Relevance Approach and Its Discontents 38

3. The Relativity of Relevance and the Partial Primacy Account 45

4. Replies to Tomlin and Horton 50

5. Replies to van Gils and Tomlin 54

6. Multi-Option Cases 59

7. Conclusion 62

Chapter Three: Interpersonal Trade-Off and Risk **63**

1. Introduction 63

2. *Ex Ante* vs *Ex Post* Models 64

3. Rejecting the *Ex Ante* Model 69

4. Defending the *Ex Post* Model 77

5. A Further Objection 87

6. Conclusion 92

Bibliography 94

Acknowledgement

I have many, many people to thank for their love, help, and support during the past few years. Although this acknowledge is brief, the extent of my gratitude for these individuals is anything but.

My foremost gratitude goes to my advisors and teachers at Princeton: Johann Frick, Michael Smith, and Gideon Rosen. Each has presented to me a view of the world with which I will wrestle for the years to come. It is their brilliance and insight that I can only hope to emulate in my own work. It is their encouragement and affection that I will cherish and hold dear forever.

I have had the good fortunate to share the past few years of my life with many wonderful individuals whom I so gratefully call my friends. They have been an unfailing source of love and support, a constancy that kept me grounded, hopeful, and going. I will not attempt a list of names, but one person in particular deserves my deepest gratitude: Thank you, Amy, for being by my side every step of my adult life.

Finally, I want to thank my parents. I love you, mom and dad, for giving me life, and for filling it with your love, the depth of which I cannot put into words.

Chapter One: Moral Contractualism and the Numbers Problem

1. Introduction

Suppose that we can, at little or no cost to ourselves, rescue either one stranger from death or a greater number of others from the same fate (***Life vs. Lives***). What should we do? As far as commonsense morality is concerned, the answer is straightforward: We ought to save the greater number.[1] How this commonsense prescription is to be theoretically justified, however, has proven to be a surprisingly difficult issue, raising long-standing challenges for both consequentialist and nonconsequentialist theories of morality.

Standard forms of consequentialism supply us with an initially appealing rationale: We ought to save the greater number because only by so doing will we maximize the sum total of people's well-being. Yet this aggregative rationale itself leads to well-known problems in other cases of interpersonal trade-off. For example, in a choice between saving a person's life and preventing a large number of temporary headaches (***Life vs. Headaches***),[2] the rationale apparently calls for us to leave the single person to die, provided that more well-being *in the aggregate* is generated by the alleviation of headaches.

Aggregative reasoning of this kind has long been viewed by nonconsequentialists with suspicion. Most memorably, John Rawls has remarked that moral theories under which the losses of some may be straightforwardly offset by a greater aggregate gain of

[1] John Taurek famously dissents from commonsense morality on this point. John Taurek, "Should the Numbers Count?" *Philosophy & Public Affairs* 6 (1977): 293-316.

[2] I borrow the case of *Life vs. Headaches* comes from Alastair Norcross, "Comparing Harms: Headaches and Human Lives," *Philosophy & Public Affairs* 26 (1997): 135-67.

others fail to take seriously the separateness of persons.[3] In an effort to provide a clear and principled alternative to consequentialism's aggregative rationale, leading nonconsequenitalists have put forward the suggestion that interpersonal morality has a mode of justification that is fundamentally *individualist* in character.[4]

In contemporary ethics, the fullest development of an individualist approach to morality is found in the contractualist theory of T.M. Scanlon. Scanlon institutes into his moral contractualism the condition that the justifiability of a moral principle depends only on its implications for *single persons*, and not on its implications for collections of people. Under this *Individualist Restriction*, as the condition has come to be known, considerations of aggregate well-being do not in themselves carry any justificatory weight in determining the correct principles that regulate our interpersonal conduct. The Individualist Restriction occupies an important place within moral contractualism; as Scanlon says, it is "central to the guiding idea of contractualism, is also what enables it to provide a clear alternative to utilitarianism and other forms of consequentialism."[5]

The trouble, as Scanlon himself acknowledges, is that contractualism appears to go too far in the opposite direction. No doubt the theory's individualist focus handles well the case of *Life vs. Headaches*, yet it also leaves us without a ready explanation for our judgment in *Life vs. Lives*, where numbers really do seem to matter. Scanlon's own

---

[3] John Rawls, *A Theory of Justice* (Cambridge: Harvard University Press, 1971), pp. 19-24.

[4] For example, Thomas Nagel, "Equality," in Mortal Questions (Cambridge: Cambridge University Press, 1979), pp. 106–27.

[5] T.M. Scanlon, *What We Owe to Each Other* (Cambridge: Harvard University Press, 1998), p. 229.

solution to this conundrum, in the form of his *tie-breaking argument*, has persuaded few philosophers,[6] and contractualism's treatment of interpersonal aggregation has remained one of the most heavily criticized aspects of the theory till this day.[7]

In this chapter, I aim to develop a novel solution to the Numbers Problem within the individualist framework of moral contractualism. The solution I offer unfolds in several stages. I begin, in Section 2, by presenting a brief overview of the contractualist framework, with an eye towards clearly locating the Individualist Restriction within that theoretical structure. Next, in Section 3, I first disarm two lines of reasoning which threaten to undermine, from the very outset, a contractualist vindication of interpersonal aggregation, and I then consider and reject Scanlon's attempted solution to the Numbers Problem. This paves the way for the presentation of my own solution in Sections 4 and 5. In these sections, I argue that contractualist reasoning leads to the endorsement of the *Limited Aggregation* view, which instructs us to take aggregative considerations into account in some cases of

---

[6] Indeed, Scanlon himself has voiced dissatisfaction over his treatment of interpersonal aggregation. As he later remarks, "[w]hen I was finishing my book I was under no illusion that I had solved the problem of aggregation, and the part of the book devoted to this question was the part that seemed to me least satisfactory." T.M. Scanlon, "Replies," *Social Theory & Practice* 28 (2002): 337-58, p. 354.

[7] For example, Michael Otsuka, "Scanlon and the Claims of the Many versus the One," *Analysis* 60 (2000): 288-93; F.M. Kamm, "Owing, Justifying, and Rejecting," *Mind* 111 (2002): 323-54, p. 348; David Wasserman and Alan Strudler, "Can a Nonconsequentialist Count Lives?", *Philosophy & Public Affairs* 31 (2003): 71-94; Joseph Raz, "Numbers, With and Without Contractualism," *Ratio* 16 (2003): 346–67; Otsuka, "Saving Lives, Moral Theory, and the Claims of Individuals."; and Derek Parfit, *On What Matters Volume 2* (Oxford: Oxford University Press, 2011), pp. 191-212.

trade-offs but not in others, and thereby captures our divergent intuitions in *Life vs. Headaches* and *Life vs. Lives*. Along the way, I mark and make use of several distinctions whose significance for contractualist theorizing has not, in my view, been fully appreciated.

If successful, my discussion taken as a whole secures several important results. First, it rescues moral contractualism from the theory's long-standing problem with interpersonal aggregation. Second, it demonstrates how, starting from an individualist mode of justification, we can nevertheless arrive at a view that allows the numbers to count without directly counting the numbers. Third and finally, it articulates a theoretical rationale for accepting Limited Aggregation as the correct general approach to issues of interpersonal trade-offs.

2. An Overview of Contractualism

In contemporary ethics, the contractualist moral theory developed by Scanlon has distinguished itself as one of the most systematic nonconsequentialist accounts of interpersonal right and wrong, providing a powerful alternative to its consequentialist rivals. At the heart of Scanlon's theory lie three distinctive theses about the subject-matter of interpersonal morality:

1. The guiding ideal of interpersonal morality is that of *justifiability to each person*.
2. The substantive account of justifiability to each person is given in terms of the *reasonable rejectability* of principles that regulate our interpersonal conduct.

3. The determination of a principle's reasonable rejectability is subject certain constraints, chiefly among which is the *Individualist Restriction*.

Contractualism, as it is here understood, is first and foremost a *reductive* doctrine. The theory reduces normative facts of one kind, namely, facts about the deontic status of actions, to normative facts of another kind, namely, facts about the reasonable rejectability of conduct-governing principles. It thus turns out the core claim of contractualism is one about the metaphysics of normativity. Notably, the theory seeks to explain facts about right and wrong not in some non-normative terms, but in terms of the *normative reasons* each of us has for accepting or rejecting principles that serve to regulate our conduct and behaviour.[8]

It is worth observing that contractualism does not claim its account of right and wrong exhausts the whole of morality. Instead, the theory is meant to cover only the important domain of *interpersonal morality*, which centrally has to do with "our duties to other people, including such things as requirements to aid them, and prohibitions against harming, killing, coercion, and deception."[9] Scanlon dubs this part of morality "what we owe to each other," but he is adamant that some of our reasons for action fall outside of its scope.[10] For example, we have reasons not to flood the Grand Canyon or turn it into a site

[8] For Scanlon's settled view on the metaphysics of reasons, see his defense of metaethical non-naturalism in T.M. Scanlon, *Being Realistic About Reasons* (Oxford: Oxford University Press, 2014).

[9] Scanlon, *What We Owe to Each Other*, p. 6.

[10] Ibid., pp. 171-77 and p. 219.

for trash disposal, and some of these reasons obtain simply because the Grand Canyon is impersonal valuable, and not just because its destruction would negatively affect the interests of persons. That contractualism recognizes the reality of impersonal values is an important point to keep in mind, and I will return to it in Section 5 below.

At its core, contractualism tells us that an act is wrong if its performance under the circumstances would not be justifiable to each of those affected. The theory further advances a substantive account of justifiability, one that holds that an act fails to be justifiable if its performance would be disallowed by conduct-regulating principles that no one could reasonably reject. Ultimately, the theory locates the moral wrongness of an action in there being an objection of a certain kind to principles that would permit its performance. Here, three aspects of this *objection model of moral wrongness* call for further discussion.

First, objections to a principle have to be *reasonable*, in the sense of being supported by good reasons. But the mere presence of a reasonable objection does not all by itself make a principle reasonably rejectable. After all, even good reasons for objecting to a principle need to be balanced against reasons in the principle's favor. An objection renders a principle reasonably rejectable only if the objection in question is *sufficiently forceful*.

Second, the forcefulness of a reasonable objection is to be assessed in a *comparative* manner. This condition is meant to imply that, given a choice between two competing principles, even if one stands to be burdened a great deal more under Principle 1 than Principle 2, one's objection to Principle 1 need not amount to a reasonable rejection

if the alternative principle would impose an even greater burden on someone else.[11] Thus, how forceful an objection is turns on what alternative principles there are and what objections these alternative principles confront.

Third, and important for our purposes, the assessment of a principle's reasonable rejectability is crucially subject to the constraints imposed by the Individualist Restriction, according to which:

> ***Individualist Restriction***: Whether a principle is reasonably rejectable depends only on the objections raised by *single individuals* on the basis of their respective *personal reasons* for objecting to that principle and its alternatives.

An individual's personal reasons for objecting to a principle are those considerations having to do with the principle's implications for her *own* well-being, interest, or status. Impersonal reasons, such as those having to do with people's welfare in the aggregate or the grandeur of some natural wonder, do not themselves provide grounds for reasonable objection, since they go beyond the reasons an individual can cite *on her own behalf* for how she wants to be treated by others.[12]

Thus understood, the Individualist Restriction bars interpersonal aggregation on two distinct fronts. First, by limiting the grounds of reasonable objection to people's

---

[11] T.M. Scanlon, "Contractualism and Utilitarianism," in *Utilitarianism and Beyond*, ed. Amartya Sen and Bernard Williams (Cambridge: Cambridge University Press 1982), p. 111. See also Scanlon, *What We Owe to Each Other*, pp. 195-97.

[12] Ibid., 219-20.

personal reasons, the Individualist Restriction bars consideration of aggregate well-being from directly figuring in the *content* of an individual's objection. Second, by insisting that only objections from single individuals may matter for the purpose of reasonable rejection, the Individualist Restriction bars the aggregation of objections *across* different individuals.

The Individualist Restriction flows naturally from contractualism's most central and most attractive idea, that of justifiability to *each* person; it gives expression to the distinctively nonconsequentialist thought that an act is wrong only if there is someone regarding whom acting in such a way treats *her* in an unjustifiable manner. Additionally, the Individualist Restriction affords contractualism clear advantages as a general framework within which to develop an account of interpersonal aggregation. In *Life vs. Headaches*, for instance, the Individualist Restriction provides a principled basis for blocking the problematic transition from the axiological claim that preventing the many headaches maximizes aggregate welfare to the deontic claim that we ought to therefore save the greater number.

The challenge, as evidenced by the case of *Life vs. Lives*, is for contractualism to explain why we should nevertheless save the greater number in at least some cases of interpersonal trade-offs. I begin to meet this challenge in the next section.

3. Clearing the Ground

It is tempting to think that the entire project of vindicating interpersonal aggregation within the contractualist framework is doomed to fail. After all, contractualism as a moral theory embraces a mode of justification that is wholly antiaggregative across persons. As we saw in the preceding discussion, contractualism holds that only the personal reasons each

individual can cite on her own behalf may matter for determining the correct principles that govern our interpersonal conduct. Given this, any principle of interpersonal trade-offs that allows numbers to count seems to run afoul of the theory's signature commitment to the Individualist Restriction.

While this line of reasoning is persuasive at first glance, it is guilty of what we may call a *level confusion*. Contractualism is indeed antiaggregative at the level of *justification*, in the sense that the justifiability of a moral principle depends only on the personal reasons of single individuals. What the above reasoning overlooks, however, is that an individual might have strong personal reasons for favoring a principle which, as part of that principle's *content*, mandates us to sometimes take into account the relevance of numbers. We should therefore clearly distinguish between two levels at which aggregative considerations might enter within the contractualist framework: They might enter at the level of justification, as reasons for accepting or rejecting some principle of interpersonal trade-offs over its alternatives, or they might enter at the level of content, as considerations a principle tells us to take into account in adjudicating cases of interpersonal trade-offs. The Individualist Restriction bars aggregation at the level of justification and not at the level of content.

Of course, whether or not individuals have strong personal reasons in favor of adopting a principle that allows for some measure of interpersonal aggregation is a matter for substantive debate, and I will go on to argue for an affirmative answer. For now, the important thing to keep in mind is that the presence of such personal reasons is not ruled out by contractualism from the outset.

Turning next to the case of *Life vs. Lives*, consider a situation in which we can rescue either A from death or each of B and C from the same fate. It is tempting to reason

about the situation in the following way. Each of B and C stands to lose his life if left unaided. But A's well-being is equally at stake. In the light of the Individualist Restriction, the personal reasons given by the well-being of B and C may not be combined. Nor may the impersonal worseness of the death of two people be directly cited as grounds for saving the larger group. Therefore, the reasoning concludes, the personal reasons on each side are perfectly balanced, leaving us without a decisive basis for saving the greater number.

Despite its seeming plausibility, the above line of reasoning overlooks an important factor. It is a key feature of contractualism that considerations of individual well-being do not exhaust our grounds for reasonable objection.[13] This is as it should be, since the personal reasons we have for how we want to be treated by others are many and varied. No doubt an action can be unjustifiable to us because it diminishes our well-being, but an action can also demean us, disrespect us, or treat us unfairly, and these are all forms of treatment to which we can reasonably object.

Indeed, Scanlon's own solution to the Numbers Problem appeals to considerations that go beyond a simple concern for individual welfare. Regarding the case of *Life vs. Lives*, Scanlon has the following to say:

> [E]ither member of the larger group might complain that [a principle which permits us to be insensitive to numbers] did not take account of the value of saving his life, since it permits the agent to decide what to do in the very same way that it would have permitted had he not been present at all… The presence of the additional

[13] Scanlon, *What We Owe to Each Other*, p. 217. See especially Scanlon's discussion and dismissal of what he calls "welfarist contractualism."

> person [under such a principle] … makes no difference to what the agent is required to do or to how she is required to go about deciding what to do. This is unacceptable, the person might argue, since his life should be given the same moral significance as anyone else's in this situation…"[14]

Scanlon's proposal, in essence, is that giving equal moral significance to everyone's life requires that when the competing interests of A and B are evenly balanced, the additional presence of C should make a difference by breaking the tie in favor of saving the larger group.[15]

---

[14] Scanlon, *What We Owe to Each Other*, p. 232. Scanlon's argument draws on ideas from F.M. Kamm. See F.M. Kamm, *Morality, Mortality Volume 1* (New York: Oxford University Press, 1993), chapter 6.

[15] It is important to get clear on the content of the objection Scanlon is here attributing to each of B and C. The objection is *not* that a principle of interpersonal trade-off fails to properly take into account the value of saving one's life whenever one's presence or absence makes no difference to which group we ought to save under that principle. If *that* were the content of the objection, then it would find expression in the following case as well:

> ***Addition***: We can save either five people from death or one person from death. At a later time, one different person, X, is added to the latter group.

Under any principle which requires us to save the greater number in *Addition*, the inclusion of X makes no difference to which group we ought to save. But some such principle is precisely what Scanlon hopes to vindicate by his tie-breaking argument. Iwao Hirose interprets Scanlon's argument in this way. See Iwao Hirose, "Skepticism about Value Aggregation," in *The Oxford Handbook of Value Theory* eds. Iwao Hirose and Jonas Olson (New York: Oxford University Press, 2015): 300-314, at 311-12. Instead, the correct way to understand Scanlon's

Scanlon's tie-breaking argument, however, immediately confronts two serious worries. First, if our reason for rescuing the larger group lies in a concern that each individual's presence should make a difference to whom we are to save, then it is consistent with satisfying this concern that, rather than saving B and C over A outright, we should employ the device of a weighted lottery, in which the presence of each individual makes a difference by affecting people's chance of being saved.[16]

Second, it is unclear how Scanlon's proposal can account for the intuition that, as the disparity in numbers increases between the two competing groups, we commit a greater wrong if we fail to save the larger group. For whether we are confronted with a choice between saving one life and two others or between saving one life and one million others, Scanlon's tie-breaking argument generates individual objections with the same force against rescuing the smaller group.[17]

---

suggestion is roughly along the following lines: A principle of interpersonal trade-off fails to properly take into account the value of saving one's life if, in situations in which one's life is at stake, but in which apart from one's presence the competing claims are tied, the principle in question permits an agent to do what she would be permitted to do under that principle has one not been present at all. On this point, see also Norcross, "Contractualism and Aggregation," p. 305.

[16] See for example David Wasserman and Alan Strudler, "Can a Nonconsequentialist Count Lives?", *Philosophy & Public Affairs* 31 (2003): 71-94; and Otsuka, "Saving Lives, Moral Theory, and the Claims of Individuals," pp. 111-18.

[17] Of course, as the disparity in numbers increases, more individuals are in a position to raise an objection under Scanlon's tie-breaking argument. But since the force of individual objections may not be combined under the Individualist Restriction, this fact cannot account for the intuition at issue.

Worse still, Scanlon's proposal runs into trouble when we try to generalize it to cases where the harms faced by individuals are *unequal* in magnitude. Consider a case where we can save either a person's life or each of a large number of people from a slightly lesser harm, say, permanent paraplegia (***Life vs. Paraplegias***). Intuitively, we should forgo the life-saving option if sufficiently many paraplegias can be thereby avoided. In order for the tie-breaking argument to gain a foothold in this case, however, we must first grant that several individuals' interests in avoiding paraplegia may *together* create a tie with a competing interest in avoiding death. This is because a single person's interest in avoiding paraplegia, by itself, is insufficient for this purpose. At this point, a problem emerges for Scanlon. We are right to wonder why the tie-breaking argument does not carry over with equal success to the case of *Life vs. Headaches*. In other words, what is to prevent a tie from being created between several individuals' interests in avoiding a minor headache and a competing interest in avoiding death, with the result that the additional presence of people facing headaches may now break the tie in favor of saving the greater number?

While Scanlon does not address this problem head-on, his discussion does contain a suggestive proposal. Here is the key passage:

> [I]t seems that our intuitive moral thinking is best understood in terms of a relation of "relevance" between harms. If one harm, though not as serious as another, is nonetheless serious enough to be morally "relevant" to it, then it is appropriate, in deciding whether to prevent more serious harms at the cost of not being able to prevent a greater number of less serious ones, to take into account the number of harms involved on each side. But if one harm is not only less serious than, but not

> even "relevant to," some greater one, then we do not need to take the number of people who would suffer these two harms into account in deciding which to prevent, but should always prevent the more serious harm.[18]

If we incorporate the distinction between relevant and irrelevant harms into Scanlon's tie-breaking argument, then Scanlon's proposal, in effect, is that only competing harms that are relevant may aggregate to create ties, and only harms that are relevant may perform the function of tie-breaking.

I have several misgivings about Scanlon's proposal even in its improved form. First, it is plainly not enough to merely *assert* that there is a normatively significant distinction between relevant and irrelevant harms within the context of interpersonal aggregation. What we need is a *justification* for maintaining that the distinction holds the kind of normative significance that Scanlon in the quoted passage says it holds. But this is a justification that Scanlon nowhere provides.

Second, if Scanlon is right in saying that we do need to take into account the number of relevant harms in deciding which group to save, then it needs to be shown that this kind of aggregation is consistent with the strictures imposed by the Individualist Restriction. Once again, this is an explanatory burden which Scanlon fails to discharge. But until such an explanation is given, the distinction between relevant and irrelevant harms remains unavailable to contractualism as a fix to the theory's problem with numbers.

---

[18] Scanlon, *What We Owe to Each Other*, pp. 239-40. Scanlon here draws on ideas from Kamm. See Kamm, *Morality, Mortality Volume 1*, chapters 8-10.

Third and finally, if we do succeed in developing a justification for allowing only relevant claims to aggregate, and if we further succeed in showing that this justification is compatible with contractualism's Individualist Restriction, then the tie-breaking argument simply becomes *redundant*. We can then directly appeal to this justification in vindicating interpersonal aggregation within the general framework of moral contractualism, without resorting, in a roundabout way, to the reasoning of the tie-breaking argument. My goal in the next two sections is to unearth precisely such a justification.

For ease of exposition, I will call the principle which mandates us to always maximize aggregate well-being in deciding whom to save *Pure Aggregation*. The diametrically opposing principle, which mandates us to always satisfy the *individually* strongest interest in trade-off situations, irrespective of the number of competing interests that are also present, I will call *Pure Antiaggregation*.

To preview, the argument in what follows develops in two major steps. First, in Section 4, I articulate a contractualist rationale for rejecting the dictates of Pure Aggregation. The challenge, however, is to ensure that such a rationale does not overgeneralize in ways that preclude us from being sensitive to numbers in cases such as *Life vs. Paraplegias*. Next, in Section 5, I put forward what I call *the argument from equal consideration* in favor of saving the greater number in cases like *Life vs. Lives* and *Life vs. Paraplegias*. The argument seeks to establish that individuals have strong personal reasons, based on the ideal of equal consideration, for rejecting the dictates of Pure Antiaggregation. It will thus emerge from my discussion that what contractualist reasoning supports is the Limited Aggregation approach, which avoids the reasonable objections to both of its alternatives.

4. Argument from Normative Disabling

The contractualist case against Pure Aggregation might seem very straightforward. In thinking about the case of *Life vs. Headaches*, contractualism's Individualist Restriction focuses our attention on the personal reasons each individual can cite on her own behalf for wanting to be saved. Since an individual's interest in avoiding death is obviously weightier than any *one* person's competing interest in avoiding a minor headache, we are straightforwardly led to the conclusion that we should save the life of the single person, irrespective of the aggregate well-being generated by the alleviation of minor headaches.

But this simple way of diagnosing the defect of Pure Aggregation lands contractualism in a serious problem; it disables contractualism from discerning any morally relevant difference between *Life vs. Headaches* and *Life vs. Paraplegias*. After all, a person's interest in avoiding paraplegia, too, is less weighty than a competing interest in avoiding death. If contractualism is to have any hope of capturing the correct verdict in the latter case, we need a more nuanced rationale for rejecting Pure Aggregation.

To this end, consider two types of conflict between our practical reasons. In the first, a practical reason may carry less normative force than a conflicting reason, and it is in this sense outweighed. An outweighed reason, however, need not lose its normative force as a genuine reason. For instance, I have a stronger reason to save A's life rather than prevent B's paraplegia, but this does not mean B's interest in avoiding paraplegia is a normatively irrelevant consideration.

But our practical reasons can conflict in deeper ways. In certain contexts, a consideration which typically constitutes a reason for action can be robbed of its normative force altogether. The following example provides one illustration:

> ***Rare Stamp***: A rare stamp is on sale at a neighborhood shop, with a going price of $900. Being an avid stamp collector, I correctly judge that acquiring the stamp would greatly improve the quality of my collection. Money, however, is tight, so I immediately begin to look for ways to obtain the necessary fund. After exhausting all available avenues to no avail, it occurs to me that I am the beneficiary of my uncle's will, who is happy and flourishing at the age of seventy-two. Hastening the death of my uncle, if I so choose, will allow me to secure my inheritance early.[19]

Do I have a reason to bring about the death of my uncle under the prescribed circumstances? I have the clear intuition that the answer here is No, and I think we should take this intuition at face value. Even if I am in urgent need of money to further some hobby project of mine, the fact that I can obtain the necessary fund only by hastening the death of my uncle — say, by hiding his asthma medication behind a shelf in the attic — is not a reason for me to prematurely end his life.

In saying this, I do not mean to suggest that one's interest in furthering one's own projects is not typically reason-providing. Rather, my contention is that a consideration which is typically reason-providing can lose its normative force in certain contexts. When

[19] The case draws its inspiration from one presented by Scanlon. See Scanlon, *What We Owe to Each Other*, p. 156.

this happens, we may say that the consideration in question is *disabled*, which differs from its being simply outweighed.[20]

It might help our understanding of normative disabling in the practical domain by considering a seeming analogue in the epistemic domain, in the form of *undercutting defeaters*. Suppose that as I wander through a gallery, I chance upon a room whose walls appear to me as if they are painted blue. Typically, my perceptual experience that the walls are painted blue is a good reason for me to believe that the walls are indeed painted blue. But suppose that I am later informed by a gallery staff that the room I chanced upon is illuminated in such a way that ensures the walls appear blue whatever color they are painted. In view of this change in my evidential circumstance, the reason-giving force of

[20] We should further distinguish normative disabling from the phenomenon of reasons being *excluded*. The pioneering work on exclusionary reasons is of course due to Joseph Raz. Joseph Raz, *Practical Reasons and Norms* (Oxford: Oxford University Press, 1999), Chapter 1. Raz defines exclusionary reasons quite broadly, as any second-order reason not to take some first-order reason into account in one's practical deliberation. Importantly, an excluded reason is nevertheless a genuine reason, albeit one that should not be taken into account in an agent's deliberation. By contrast, when the normative force of a consideration is disabled, it is no longer a reason at all under the circumstances. To see how disabling can come apart from excluding, consider a case in which my psychology is such that I am easily and consistently misled by the reason-giving force of my own interest as compared to the interests of others. In certain contexts, I might have a second-order reason not to take my prudential concerns into account in my deliberation, so as to avoid the distorting effects of my own psychology. But this need not imply that my own interest is normatively disabled under the circumstances.

my perceptual experience is undercut; the experience is no longer a reason for me to believe that the walls are painted in the color that I perceive.[21]

The general point to which I am calling attention is that the reason-giving force of a consideration is not invariant across contexts. This is true of reasons for belief as well as reasons for action. Now, what explains the occurrence of disabling in the practical domain? Here is a proposal that I find attractive. When a consideration which is typically reason-providing is disabled, I suggest this is because the presence of some value makes it inappropriate for the consideration to weigh against the value in question.[22]

Consider, again, the case of *Rare Stamp*. There, the improvement of my stamp collection runs up against the preservation of my uncle's life. If you, like me, think that I do not have a reason to end my uncle's life under the circumstances, you are apt to wonder why. After all, such an improvement does give me a reason to adopt other means instrumental to its realization, such as taking out a short-term loan or calling up a friend for assistance.

On the proposal I am here advancing, disabling occurs because the value of human life sets certain limits on what kind of benefit may appropriately weigh against a life's preservation; it is inconsistent with the value of human life for the improvement in the quality of stamp collection to count in favor of bringing about a person's death. No such

---

[21] For a classic discussion of undercutting defeaters, see John Pollock, "Defeasible Reasoning," *Cognitive Science* 11 (1987): 481-518.

[22] Here, I draw on Victor Tadros's rejection of Pure Aggregation, to whose discussion of normative disabling I am much indebted. See Victor Tadros, "Localized Restricted Aggregation," pp. 172-74.

inappropriateness exists when the option under consideration is instead taking out a short-term loan, since this option does not similarly put anyone's life in serious jeopardy.

This account of disabling, I believe, holds important lessons for how we should think about the issue of interpersonal aggregation. In deciding between saving a person's life and curing a large number of headaches, each person's interest in avoiding a minor headache, much like an interest in improving one's stamp collection, is simply too trivial to ground a reason in favor of leaving another person to die. In other words, the reason-giving force of such a minor interest is disabled in a life-and-death context, for the value of human life makes it inappropriate for the alleviation of a minor headache to weigh against another person's survival.

A proper contractualist analysis of *Life vs. Headaches* would then run as follows. It is true that preventing the many headaches generates more well-being in the aggregate. But, under the Individualist Restriction, aggregative considerations do not themselves provide grounds for reasonable objection; only personal reasons matter. Now, Pure Aggregation instructs us to forgo the life-saving option in *Life vs. Headaches*. Is such a principle reasonably rejectable? Yes. But not for the simple reason that an interest in avoiding death is weightier than any competing interest in avoiding a minor headache. Rather, the reason is that an interest in avoiding a minor headache is one whose normative force the value of human life disables. By allowing these disabled interests to count, the person facing death may rightly object that Pure Aggregation fails to properly respect the value that her life has.

This way of rejecting Pure Aggregation enables contractualists to discern a key difference between *Life vs. Headaches* and *Life vs. Paraplegias*. In the former case, the

reason-giving force of each person's interest in avoiding a minor headache is wholly disabled. Consequently, no one facing a minor headache is in a position to raise a *reasonable* objection if we do not take her interest into account at all in our process of adjudication; her interest is in this sense normatively *irrelevant*.

By contrast, while an individual's interest in avoiding permanent paraplegia is less weighty than a competing interest in avoiding death, it is nevertheless weighty enough to be a relevant consideration, one that retains its reason-giving force even in a life-and-death context. As a result, an individual who stands to become permanently paraplegic deserves a hearing within contractualist reasoning; the possibility of its being unjustifiable to her to choose the life-saving option remains open.

Before I proceed, let me briefly note one particularly attractive feature of my contractualist rejection of Pure Aggregation. Too often, in thinking about cases like *Life vs. Headaches*, noncontractualist critics of Pure Aggregation freely move from the claim that a single headache is a normatively irrelevant harm because it is so trivial to the conclusion that a large number of headaches are too normatively irrelevant. But when this transition is made in the absence of an explicit commitment to contractualism's Individualist Restriction, it contains a highly questionable assumption, namely, if some harm considered on its own is trivial, then an aggregation of such trivial harms must be itself trivial.[23]

[23] John Halstead makes this criticism against critics of Pure Aggregation in his "The Numbers Always Count," *Ethics* 126 (2016): 789-802, at 794-95. The criticism is mainly directed at Alex Voorhoeve's defense of Limited Aggregation. The assumption is also found in Tadros's defense of Limited Aggregation. Tadros, "Localized Restricted Aggregation," p. 175.

On the present proposal, by contrast, the aggregate well-being generated by the alleviation of many headaches is irrelevant, *not* because it is trivial, but because the Individualist Restriction limits our attention to only the personal reasons each individual can cite on her own behalf for wanting to be saved. For this reason, a contractualist can happily grant the axiological claim that a non-trivial amount of aggregate well-being is produced if we prevent the many headaches, while still holding onto the deontic claim that we ought to nevertheless choose the life-saving option.

5. Argument from Equal Consideration

The central task of this section is to unearth an individualist justification for saving the greater number in cases in which we clearly should. I begin, however, by addressing two theoretical questions. First, what exactly is the theoretical ambition of contractualism in relation to the Numbers Problem?

Recall a remark earlier made that contractualism does not claim its account of right and wrong exhausts the whole of morality. Instead, the theory is only meant to capture one important part of morality, which Scanlon dubs "what we owe to each other." Now, some philosophers have put forward the suggestion, as a way for contractualism to avoid the Numbers Problem, that the theory should make better use of the kind of moral pluralism to which it is antecedently committed. The suggestion is that while saving the lesser number in a case like *Life vs. Lives* would not be wrong under the contractualist ideal of

justifiability to each person, it would still be wrong *all-things-considered*, once non-contractualist, aggregative reasons are brought into the picture.[24]

This suggestion does have some appeal, but I believe it is not one that a contractualist should accept. Although contractualism's theoretical ambition does not extend to the whole of morality, the theory is nevertheless meant to provide a complete account of *interpersonal morality*, part of whose content is our duty to aid other people. How our duty to aid operates in cases of interpersonal trade-offs should therefore be of central concern to contractualism. In short, unless contractualists are prepared to make the hugely concessive move that their theory offers only an incomplete account of interpersonal morality, they should not settle for the suggestion as it stands.[25]

In saying this, I might give the impression that contractualism should in no way appeal to aggregative considerations in resolving the issue of interpersonal aggregation. This brings me to the second question I want to raise, namely, what role, if any, may considerations of aggregate well-being play within a contractualist mode of justification?

One thing is clear — contractualism's Individualist Restriction bars considerations of aggregate well-being from directly figuring in the content of reasonable objections. Someone who objects to Pure Antiaggregation by directly calling our attention to the fact that it is a worse thing if there are more deaths will have done nothing to show that she

---

[24] For example, Derek Parfit, *On What Matters Volume 2* (Oxford: Oxford University Press, 2011), at pp. 216-17.

[25] Interestingly, Johann Frick has recently proposed that contractualism should scale back its theoretical ambition in precisely this way. See Johann Frick, "Contractualism and Social Risk," *Philosophy & Public Affairs* 43 (2015): 175-233, at 219-33.

herself is treated in an unjustifiable manner. After all, this is not an objection that she may raise on her own behalf.

To be sure, some philosophers deny the very idea of aggregate value.[26] They deny that there is any more disvalue in the deaths of a million people than the death of a single person. This is no part of the contractualist view.[27] What a contractualist denies, in keeping with the Individualist Restriction, is that facts about aggregate well-being can *themselves* supply grounds for reasonable objection. The distinctively nonconsequentialist element in contractualism lies in the denial that, in the domain of interpersonal morality at least, an action can be wrong simply *because* it fails to maximize aggregate value. To put the point differently, what a contractualist denies is *not* axiological aggregation, but the *direct* relevance of aggregative axiology to interpersonal morality.

But this leaves open the suggestion that aggregate value can play an *indirect* role in the contractualist framework. It is this suggestion that I aim to explore in what follows. Specifically, my suggestion is that facts about aggregate well-being can serve as the *background* against which individuals can derive personal reasons for rejecting principles of interpersonal trade-offs that are wholly insensitive to numbers.

The general phenomenon of impersonal values playing an indirect role in generating personal objections should be a familiar one to contractualists. Consider the following passage from Scanlon:

---

[26] For example, Taurek, "Should the Numbers Count?"

[27] See Scanlon's discussion of impersonal values in *What We Owe to Each Other*, pp. 218-23 and his discussion of the significance of impersonal pain and suffering on pp. 177-87.

> Impersonal reasons do not, themselves, provide grounds for reasonably rejecting a principle. But these reasons do play a significant role in determining other grounds for reasonable rejection… So, for example, part of what it means to say that the Grand Canyon is of value is that visiting and enjoying it is worthwhile. From the point of view of those who might engage in these activities in the future, then, there is reason to reject a principle that would allow someone to decide to flood the Grand Canyon without taking these benefits into account. These reasons for rejecting a principle are what I called above personal reasons, but their force as reasons depends in part on further judgments of impersonal value, namely the judgment that these objects are worth seeing and should be admired.[28]

The last sentence here is key. For someone who desires the preservation of the Grand Canyon, it matters, in assessing the force of her objection to the canyon's destruction, to take into account the fact that she has good reasons to want to engage with the Grand Canyon in the appropriate ways, reason that are provided by the impersonal value of the Grand Canyon itself. But it remains to be seen how the impersonal value of aggregate well-being can give rise to personal reasons in favor of saving the greater number.

To resolve this issue, I begin with an idea underlying Scanlon's attempted solution to the Numbers Problem. The idea is that we have reasons to treat everyone's interest with equal consideration, violations of which give rise to personal objections under the contractualist framework. Scanlon develops this idea in the direction of his tie-breaking

[28] Scanlon, *What We Owe to Each Other*, p. 220.

argument, claiming that giving everyone's interest equal moral significance in *Life vs. Lives* requires the presence of some individuals to play a tie-breaking role. Scanlon's move, as we have seen, proves unsuccessful in the end. I plan to develop the idea of equal consideration in a different direction.

What does it take to treat the same interests of two different individuals with equal consideration? A natural answer is that when the interests of two individuals, A and B, are in fact equally weighty, one should regard the satisfaction of their interests as equally important. Of course, when the interests of A and B cannot be jointly satisfied, the conflict may be resolved by flipping a fair coin or simply making a choice between the two parties. What equal consideration rules out, absent special justification, is the judgment that the interest of one party is more worthy of fulfillment than that of the other.

Consider a simple illustration. Suppose we decide that, if A's life is ever threatened, we will devote X amount of valuable resources to the saving of A's life. Suppose we also decide that, if B's life is ever similarly threatened, then we will devote to B only half of the amount of resources we would devote to A. Holding fixed that the lives of A and B do not differ in any morally relevant respects, our decisions reveals an inequality in the consideration we show to the interests of A and B. B may rightly object, on her own behalf, that her interest in avoiding a lethal threat is unjustifiably given less importance than A's interest in avoiding the same. Moreover, the inequality at issue is revealed by a difference in the amount of value that we are prepared to forgo for the respective interests of A and B. Call this understanding *equal consideration in the value-forgoing sense*.

I want to now argue that failing to save the greater number in *Life vs. Lives* also violates the ideal of equal consideration in the value-forgoing sense. My argument makes

indirect appeal to claims about aggregate value. Consider, once again, the situation in which we can rescue either A from death or each of B and C from the same fate. What objection can each of B and C raise on her own behalf against a decision to save A?

Let us begin by reframing the normative situation in *Life vs. Lives* in terms of two options that involve forgoing value. On the one hand, there is the option of saving A's life, which involves forgoing the aggregate value of saving two other lives. On the other hand, there is the option of saving the lives of B and C, which involves forgoing the value of saving one other life. We may now ask: If I opt to save A's life under the circumstances, does my decision violate equal consideration for the interests of all involved?

Now, in deciding to save A under the circumstances, I judge that A's life has the kind of importance that provides me sufficiently strong reason to forgo the aggregate value of saving two other lives. If I judge A's life to have this kind of importance, then equal consideration requires that I regard B's life as equally important. This means that I should regard B's life, too, as sufficiently important to be able to counterbalance the opposing disvalue of the deaths of two people.

Next, if the importance of B's life gives me sufficiently strong reason to save her even when opposed by the disvalue of two deaths, then my reason to save B is even *stronger* in a situation where the importance of her life is opposed by the disvalue of just one death. In other words, if the balance of reason tells in favor of saving a person's life at the cost of two deaths, then the balance of reason would be even more in that person's favor when the cost is not two deaths but only one.

And now we come to the realization that *Life vs. Lives* describes precisely a situation in which the controlling moral question is: Is there stronger reason to choose the option that saves A's life at the cost of two deaths or the option that saves B's life (together with C's) at the cost of one death? Equal consideration for the lives of A and B implies that the reason to choose the latter option is stronger. We may formulate the underlying principle of equal consideration as follows:

> **Equal Consideration 1** (EC1): If X's interest and Y's interest are equally weighty and V1 is more valuable than V2, equal consideration requires that if one judges that there is reason of a certain strength to forgo V1 for the sake of X's interest, then one judges that there is comparatively stronger reason to forgo V2 for the sake of Y's interest.

In effect, if one decides to save A in *Life vs. Lives*, then each of B and C may raise a reasonable objection in the following terms: If you regard A's life as being so important that it provides you with sufficiently strong reason to forgo two other lives, then since my life is just as important, giving equal consideration to the importance of my life means that you should regard yourself as having comparatively stronger reason to forgo just a single life for my sake.

The next step in my strategy is to generalize this argument from equal consideration to the case of *Life vs. Paraplegias*. To do so, we need to first ask what equal consideration in the value-forgoing sense requires when the harms faced by individuals are unequal. Consider another simple illustration. Suppose we decide that, if A's life is ever threatened,

we will devote X amount of valuable resources to the saving of A's life. Suppose we also decide that, if B is ever threatened with paraplegia, no amount of resources will be devoted to the alleviation of B's plight. These decisions reveal an inequality in the consideration we show to the interests of the two parties. B may rightly object that, while her interest in avoiding paraplegia is indeed less weighty than A's interest in avoiding death, we should nevertheless be prepared to allot to her an amount of valuable resources that is *proportional* to the relative weightiness of the respective interests of A and B. The following principle thus seems to me to have great intuitive force:

> **Equal Consideration 2** (EC2): When X's interest is weightier than Y's interest, equal consideration requires that the values one judges oneself to have sufficiently strong reason to forgo for the sake of their respective interests should not be disproportional to the relative weightiness of their interests.

Now, suppose I decide to take the life-saving option in *Life vs. Paraplegias*; I judge that the life of the person I save has the kind of importance that provides me with sufficiently strong reason to forgo the aggregate value of preventing a very large number of paraplegias. Equal consideration then requires that I not assign disproportionally less importance to the interest of each of those individuals facing paraplegia. But if the aggregate value of preventing paraplegias becomes sufficiently large, to the extent it may well dwarf the value of saving a life, then it would be disproportional to regard myself as having sufficiently strong reason forgo this very large amount of value for the sake of saving a life, yet I have no stronger reason to forgo significantly less value for the sake of

preventing a person's paraplegia. In other words, when there are sufficiently many paraplegias at stake, choosing the life-saving option stands to violate the proportionality constraint introduced by EC2. When this happens, each of those individuals facing paraplegia may reasonably object that her interest is not given equal consideration.

Let me note that the foregoing rationale for saving the greater number in *Life vs. Paraplegias* does not extend to the case *of Life vs. Headaches*. This is so for two reasons. First, as I said in Section 4, the reason-giving force of an individual's interest in avoiding a minor headache is wholly disabled when it runs up against a competing interest in avoiding death. Because of this, no one facing a minor headache has even the standing to raise a reasonable objection based on the ideal of equal consideration, for the simple reason that her own interest is normatively irrelevant in the situation. In short, a disabled interest should not be given equal consideration; it should be given no consideration at all.

Second and relatedly, since no one is in a position to object if we do not take into account interests having to do with the alleviation of minor headaches in a life-and-death situation, there is a strong contractualist basis for altogether excluding such aggregative considerations from having even an indirect bearing on our adjudication of interpersonal trade-offs. This marks yet another important difference between *Life vs. Headaches* and *Life vs. Paraplegias*. In the latter case, while people's interests in avoiding paraplegia, taken individually, are indeed less weighty than a competing interest in avoiding death, to simply ignore these interests is something to which each of those facing paraplegia can reasonable object. Although contractualism forbids aggregate value from having a direct bearing on interpersonal morality, I have suggested that these aggregative interests may

nevertheless matter in an indirect way, as part of the background for assessing whether equal consideration is shown to the interests of different individuals.

I think the rationale I provide for saving the greater number is plausible in its own right, but it might help to convince those who remain skeptical by noting several advantages of my proposal. First, my proposal generates an explanation for why, as the disparity increases between the competing groups, we commit a greater wrong if we fail to rescue the greater number. This is because my proposal allows aggregative considerations to matter, albeit in an indirect way. In *Life vs. Lives*, for example, when B is in a group of two, and we decide to save A instead, we in effect say that A's life is important enough to give us sufficiently strong reason to forgo the value of saving two other lives, but it is not the case that B's life gives us comparatively stronger reason to forgo the value of saving one other life. But when B is in a group of a million, and we once again decide to save A, the inequality in the consideration we show to the lives of A and B is even greater. In effect, we say that while A's life gives us sufficiently strong reason to forgo the value of saving one million lives, B's life does not provide us with comparative stronger reason to forgo the value of saving just one single life. B therefore has an even stronger personal objection, based on the ideal of equal consideration, when he is in a group of a million than when he is in a group of two.

Second, my rationale does not invite the reply that we should employ the device of a weighted lottery in resolving cases like *Life vs. Lives* and *Life vs. Paraplegias*. Recall that what makes Scanlon's tie-breaking argument susceptible to this reply is that weighted lottery allows each individual's presence to make a difference. But my proposal does not rely on the idea of allowing individuals to make a difference. Instead, it relies on the

comparative weight of the reasons we judge ourselves to have, and how these judgments relate to the ideal of equal consideration as specified by EC1 and EC2. In general, when one's decision runs contrary to a valid requirement of equal consideration, the correct way to resolve this violation is not to run a lottery, weighted or not, but to conform one's conduct to what equal consideration requires.

I will end this section by considering one final objection to my proposal. The objection can be put in simple and direct terms: Why go through all this trouble to allow aggregative considerations to matter in this roundabout way? Why not just admit that we should save the greater number because it produces more aggregate well-being overall? I have two things to say in reply. First and most obviously, such an admission would lead us right back to the implausible implications of unrestricted aggregation; accepting the verdict that we should prevent many minor headaches at the cost of leaving someone to die, I submit, is simply too big of a bullet to bite. This in one way explains why contractualism's Individualist Restriction provides a promising starting point for theorizing about the issue of interpersonal aggregation, for it removes the central difficulty in account for the correct verdict in cases like *Life vs. Headaches*.

Moreover, even setting aside the case of *Life vs. Headaches*, an explanation for why we should save the greater number that directly appeals to aggregative reasons would leave something significant unaccounted for. In thinking about the case of *Life vs. Lives*, Elizabeth Anscombe astutely raises the question who in particular is *wronged* if we do not save the greater number?[29] I think the correct answer to Anscombe's question is that

[29] Elizabeth Anscombe, "Who is Wronged?" *Oxford Review* 5 (1967): 16-17.

everyone in the larger group is wronged; saving the lesser number under the circumstances would be unjustifiable to each of them. But if the ultimate reason for saving the greater number lies directly in considerations of an aggregative kind, then no one in the larger is in a position to say, on her own behalf, that she has been treated unjustifiably. For this reason, my contractualist rationale for the saving the greater number, which embeds aggregative considerations within an ideal of equal consideration, violations of which allows each individual to raise a complaint on her own behalf, is well positioned to give the correct answer to Anscombe's question. My proposed solution to the Numbers Problem therefore has much to recommend it.

Here is where we stand. The argument from normative disabling indicates that Pure Aggregation yields results in cases like *Life vs. Headaches* that fail to accord with respect for the value of human life. On the other hand, as we have just seen, Pure Antiaggregation fares no better; it violates the ideal of treating people's interests with equal consideration. These are grounds to reasonably reject both principles. The correct approach to interpersonal trade-offs, then, is the Limited Aggregation view, which instructs us to count the numbers in *Life vs. Lives* and *Life vs. Paraplegias* but not in *Life vs. Headaches*, and thereby avoids the reasonable objections that plague both of its alternatives.

6. Conclusion

The Numbers Problem has long fixated the attention of moral philosophers, and in my view, deservedly so. The challenge of accounting for the moral relevance of numbers is one with which any moral theory vying for our allegiance must grapple, and the challenge is recalcitrant and vexing not least because it imposes more than just one condition of

adequacy. In addition to capturing the correct verdicts in the leading cases of *Life vs. Lives*, *Life vs. Paraplegias*, and *Life vs. Headaches*, a satisfactory solution to the Numbers Problem should give the correct answer to Anscombe's question and explain why we sometimes commit a greater wrong when there is more at stake. I have in this chapter put forward a contractualist treatment of interpersonal aggregation, one that hopes to do justice to these many aspects of the Numbers Problem in a systematic and unified way. If there is one overarching theme to my discussion so far, it is this: We should not expect the Numbers Problem to admit of an uncomplicated solution.

Chapter Two: How Should the Numbers Count?

1. Introduction

Suppose that we can, at no cost to ourselves, rescue either one stranger from death or some number of others from permanent paraplegia (***Life vs. Paraplegias***). Does the number of individuals facing paraplegia matter to our moral decision-making? Many of us would answer "yes"; if sufficiently many individuals stand to become paraplegic, then we should forgo the life-saving option. Next, suppose that we can, again at no cost to ourselves, rescue either one stranger from death or some number of others from a temporary headache (***Life vs. Headaches***). Does it matter in this case how many individuals stand to suffer a headache? Many of us would answer "no"; it would be wrong to leave the single person to die, no matter the number of headaches that we can help to alleviate.

These judgments suggest that there are limits to our acceptance of interpersonal aggregation. In *Life vs. Paraplegias*, we are inclined to accept that people's interests may aggregate in order to collectively defeat a countervailing interest. In *Life vs. Headaches*, by contrast, we are inclined to reject such aggregative moral reasoning. The challenge for those who want to defend this combination of judgments is to explain why interpersonal aggregation is an acceptable way to adjudicate some cases of interpersonal trade-off but not others.

By far the most influential approach to this challenge appeals to a condition of *relevance* in delimiting the bounds of acceptable aggregation. The approach has it that a person's interest can succeed or fail to be relevant to another person's competing interest, depending on whether the two interests are *sufficiently close in strength*. According to this

Relevance Approach, an aggregation of individual interests may collectively defeat only those interests to which they are relevant. Here is how the Relevance Approach is meant to apply to our two leading cases. In *Life vs. Headaches*, a person's interest in avoiding a headache is simply too weak to be relevant to another person's interest in avoiding death. Hence, no number of headaches could outweigh the moral importance of saving a life. By contrast, an interest in avoiding paraplegia and an interest in avoiding death are sufficiently close in strength for the two interests to be relevant to one another. For this reason, if sufficiently many individuals stand to become paraplegic, then they are the ones we should save.

In the past, this kind of approach has been favored by several prominent philosophers, at least in outline.[30] Lately, however, the approach has become the subject of sustained criticism.[31] In a series of recent papers, a number of philosophers have advanced

---

[30] For example, F.M. Kamm, *Morality, Mortality Volume 1* (New York: Oxford University Press, 1993), chaps. 8-10; Samantha Brennan, "Thresholds for Rights," *Southern Journal of Philosophy* 33 (1995): 143-68; and T.M. Scanlon, *What We Owe to Each Other* (Cambridge, Mass.: Harvard University Press, 1998), pp. 239-40;

[31] Besides the extensional worries I go on to discuss, Limited Aggregation has been criticized on the grounds that it implies an objectionable kind of intransitivity and that it breaches a purported principle of rationality commonly known as the Independence of Irrelevant Alternatives. For example, Alastair Norcross, "Contractualism and Aggregation," *Social Theory & Practice* 28 (2002): 303-14, at 306-09; Derek Parfit, "Justifiability to Each Person," *Ratio* 16 (2003) 368-90, at 384; and John Halstead, "The Number Always Count," *Ethics* 126 (2016): 789-90, at 797-99. Proponents of the view have in turn responded to these well-known worries, and their replies are, to my mind at least, satisfactory. See Alex Voorhoeve, "How Should We Aggregate Relevant Claims?" *Ethics* 125 (2014): 64-87, at 76-79; Alex Voorhoeve, "Why One Should Count Only Claims with Which One Can Sympathize," *Public Health*

seemingly devastating counterexamples to the Relevance Approach, raising serious doubts about its plausibility in cases that go beyond the structural simplicity of *Life vs. Paraplegias* and *Life vs. Headaches*.[32]

My central aim in this chapter is to offer a development of the Relevance Approach, one that is both theoretically well-motivated and immune to the purported counterexamples advanced by the recent critics. I begin, in Section 2, by introducing the basic framework of the Relevance Approach, and I present the objections raised by Patrick Tomlin and Joe Horton. Next, in Section 3, I articulate my preferred account of interpersonal aggregation, one that develops the Relevance Approach in novel directions. I then explain how my account convincingly handles the objections from Tomlin and Horton in Section 4. I proceed, in Section 5, to further highlight some of the distinctive features of my account, and I resolve two ambiguities in the Relevance Approach identified by Aart van Gils and Patrick Tomlin. Finally, in Section 6, I address an objection due to Derek Parfit by explaining how my account generalizes to cases of interpersonal trade-off where multiple options are involved.

Ethics 10 (2017): 148-56, at 152-53; and Victor Tadros, "Localized Restricted Aggregation," in *Oxford Studies in Political Philosophy Volume 5*, eds. David Sobel, Peter Vallentyne, and Steven Hall (Oxford: Oxford University Press, 2019): 171-204, at 185-88.

[32] Patrick Tomlin, "On Limited Aggregation," *Philosophy & Public Affairs* 45 (2017): 223-60; Joe Horton, "Always Aggregate," *Philosophy & Public Affairs* 46 (2018): 160-74; and Aart van Gils and Patrick Tomlin, "Relevance Rides Again? Aggregate and Local Relevance," in *Oxford Studies in Political Philosophy Volume* 6, eds. David Sobel, Peter Vallentyne, and Steven Hall (Oxford: Oxford University Press, 2020): 221-255.

2. The Relevance Approach and Its Discontents

2.1 Voorhoeve's ARC

To date, the best-known articulation of the basic framework of the Relevance Approach comes from Alex Voorhoeve, in the form of his *Aggregate Relevant Claims* (*ARC*).[33]

ARC

(1) Each individual whose well-being is at stake has a claim on you to be helped.

(2) Individuals' claims *compete* just in case they cannot be jointly satisfied.

(3) An individual's claim is *stronger*: (a) the more her well-being would be increased by being aided; and (b) the lower the level of well-being from which this increase would take place.

(4) A claim is *relevant* if and only if it is sufficiently strong relative to the strongest competing claim.[34]

---

[33] Alex Voorhoeve, "How Should We Aggregate Relevant Claims?" p. 66.

[34] Patrick Tomlin distinguishes between two interpretations of ARC's relevance condition:

> *Anchor by Strength*: The relevance of a claim C is determined by comparing C to the strongest claim in the trade-off situation overall, including claims that favor the same option as C.
>
> *Anchor by Competition*: The relevance of a claim C is determined by comparing C to the strongest claim that favors an option that conflicts with the option C favors.

My development of the Relevance Approach takes Anchor by Competition to be the correct model for determining relevance. See Tomlin, "On Limited Aggregation," pp. 238-40.

(5) We should choose an alternative that satisfies the greatest sum of strength-weighted, relevant claims.

Voorhoeve moreover proposes an interesting test for determining when two competing interests are relevant to one another.[35] He contends that A's interest is sufficiently close in strength to B's competing interest just in case, if A *herself* were confronted with the decision to either satisfy her own interest or B's interest, A would be morally permitted to favor her own. Here is how Voorhoeve's test applies to our two leading cases. Any person facing a minor headache would be morally required to prioritize the saving of a life over the alleviation of her own minor discomfort. Hence, a claim against a headache is not relevant to a competing claim against death. By contrast, an individual would not be morally required to allow herself to become permanently paraplegic in order to save the life of a stranger. For this reason, a claim against paraplegia and a claim against death are relevant to one another.[36]

However, critics have charged that ARC in particular and the Relevance Approach in general lead to unacceptable conclusions. In order to better state and appreciate the force of these objections, let us make the following stipulations:

- A claim against major harm is relevant to a claim against moderate harm.
- A claim against moderate harm is relevant to a claim against minor harm.

[35] Voorhoeve here draws on ideas from F.M. Kamm. F. M. Kamm, *Intricate Ethics: Rights, Responsibilities, and Permissible Harm* (New York: Oxford University Press, 2007): 297–98.
[36] Voorhoeve, "How Should We Aggregate Competing Claims," p. 72.

- A claim against major harm is *not* relevant to a claim against minor harm.

According to the critics, a set of harms that stand in the above relevance relations to one another is enough to spell trouble for the Relevance Approach.

2.2 Tomlin's Objection

Patrick Tomlin observes that cases like *Life vs. Headaches* and *Life vs. Paraplegias* involve competitions between *homogenous* groups. That is to say, they involve competitions between groups whose sizes remain constant and whose members stand to be benefitted or burdened in identical ways.[37] Tomlin objects that once we look to situations of interpersonal trade-off that go beyond these strictures, ARC is vulnerable to devastating counterexamples. Consider Tomlin's central case:

Case One:

Stage 1: We can save either Group A, which consists of one person facing major harm, or Group B, which consists of ten people facing moderate harm.

Stage 2: One person facing minor harm is added to Group A. One million people facing minor harm are added to Group B.[38]

[37] Tomlin, "On Limited Aggregation," p. 238

[38] This is a modified version of Tomlin's Case 5. Ibid., p. 242.

Suppose that, at Stage 1, the ten claims in Group B outcompete the one claim in Group A, but only *barely* so, such that if a single relevant claim is added to Group A, the balance would tip in Group A's favor.

ARC implies that we should save Group B at Stage 1. With the advent of Stage 2, however, ARC reverses its verdict. For the one million claims against minor harm added to Group B are each irrelevant under ARC, since they each compete with a claim against major harm. By contrast, the single claim against minor harm added to Group A *is* relevant under ARC, since the strongest individual claim it competes with is a claim against moderate harm. ARC therefore implies that, at Stage 2, we should switch our decision from saving Group B to saving Group A.

This implication of ARC is highly counterintuitive. As Tomlin points out, it is bizarre to think that, while we should save Group B rather than Group A at Stage 1, by adding claims of *identical* strength to each group*, but with the numbers heavily in favor of Group B*, we should now switch to saving Group A. To state Tomlin's objection more explicitly, ARC violates a plausible principle of interpersonal trade-off which I will call:

Tomlin's Principle: In a choice between two conflicting options, X and Y, if:

(a) X is more morally choiceworthy than Y at t1;

(b) X and Y each satisfies additional claims of individually identical strength at t2;

(c) X satisfies no fewer additional claims than does Y at t2;

(d) There are no other morally significant changes at t2;

then X is more morally choiceworthy than Y at t2.

2.3 Horton's Objection

Tomlin's objection to ARC is forceful. But the objection at least leaves open the prospect of developing another version of the Relevance Approach that does not suffer from the same shortcoming. Regarding this prospect, however, Joe Horton is not at all hopeful. Appealing likewise to cases involving non-homogenous groups, and building on the objection from Tomlin, Horton contends that *all* versions of the Relevance Approach fall prey to a fatal dilemma.[39] The thrust of Horton's dilemma is that any variant of the Relevance Approach violates either Tomlin's Principle or another plausible principle of interpersonal trade-off which I will call:

> **Horton's Principle**: In a choice between two conflicting options, X and Y, if:
>
> (a) X satisfies sets of individual claims X1 and X2;
>
> (b) Y satisfies sets of individual claims Y1 and Y2;
>
> (c) Claims in X1 and claims in Y1 are relevant to one another;
>
> (d) Claims in X2 and claims in Y2 are relevant to one another;
>
> (e) X1 outcompetes Y1 and X2 outcompetes Y2;
>
> then X is more morally choiceworthy than Y.

Now, consider ***Case Two***:

[39] Horton, "Always Aggregate," pp. 168-71.

Stage 1: We can save either Group A, which consists of one person facing major harm, or Group B, which consists of one million people facing minor harm.

Stage 2: Fifteen people facing moderate harm are added to each group.[40]

Let us make two suppositions about the case. Suppose that fifteen claims against moderate harm outcompete one claim against major harm, and suppose that one million claims against minor harm outcompete fifteen claims against moderate harm.

At Stage 1, any version of the Relevance Approach implies that we should save Group A rather than Group B. In view of this result, it seems that Tomlin's Principle straightforwardly mandates us to save Group A at Stage 2 as well. This is because, at Stage 2, claims of individually identical strength are added to both groups and because no fewer claims are added to Group A than to Group B.

However, if a version of the Relevance Approach view is to respect Horton's Principle, then it must imply the contrary conclusion. We can appreciate this result more clearly by considering the following three single-stage cases. Suppose that in:

Case Three: We can save either Group C, which consists of one person facing major harm, or Group D, which consists of fifteen people facing moderate harm.

[40] This is a modified version of Horton's Case 6. Ibid., p. 168.

Case Four: We can save either Group E, which consists of fifteen people facing moderate harm, or Group F, which consists of one million people facing minor harm.

Case Five: We can save either Groups C + E or Groups D + F.[41]

Given our earlier suppositions, we should save Group D rather than Group C in *Case Three*, and we should save Group F rather than Group E in *Case Four*. It is therefore eminently plausible that we should save Groups D + F rather than Groups C + E in *Case Five*. Indeed, that is what Horton's Principle mandates us to do.

But notice that Groups C + E and Groups D + F are equivalent to, respectively, Group A and Group B in our original *Case Two*. It is therefore likewise eminently plausible that we should save Group B rather than Group A in the original case. Once again, that is what Horton's Principle mandates us to do.

And herein lies Horton's dilemma. It appears that, in *Case Two*, any version of the Relevance Approach is forced to violate either Tomlin's Principle (if it permits us to save Group B) or Horton's Principle (if it permits us to save Group A). I share the view that both principles as stated are exceedingly plausible. If Horton's dilemma withstands scrutiny, then the Relevance Approach is in deep trouble.

In the next section, I first identify the feature of Voorhoeve's ARC which makes it susceptible to the objections from Tomlin and Horton. I then articulate a new version of

[41] The cases are adopted from Horton. Ibid., p. 169.

the Relevance Approach, which I go on to ague in Section 3 satisfactorily addresses these objections.

3. The Relativity of Relevance and the Partial Primacy Account

3.1 The Relativity of Relevance

Voorhoeve's ARC employs what we may call a uni-standard test for relevance. Under ARC, the relevance status of a claim is always determined by comparing the claim to the strongest claim with which it competes. Once a claim is deemed by ARC to be irrelevant in accordance with this single standard, the claim ceases to have any role to play in our process of adjudication.

But it is important to note that relevance is a *relative* notion. A person's claim can be irrelevant relative to one set of competing claims but not relative to another. For example, while an interest in avoiding death renders irrelevant a competing interest in avoiding a minor headache, this does not mean the latter interest is similarly rendered irrelevant when it instead competes with an interest in avoiding, say, a broken finger.

Suppose that we must choose between saving either Group A, which consists of one person facing major harm and ten people facing moderate harm, or Group B, which consists of fifteen people facing moderate harm and a large number of people facing minor harm. It is true that, relative to the claim against major harm, the large number of claims against minor harm are irrelevant. But it would be a mistake to conclude from this that the claims against minor harm have no role to play in our process of adjudication. After all, these claims may be used to collectively defeat the ten competing claims against moderate harm, relative which they *are* sufficiently close in strength.

Indeed, not allowing the claims against minor harm to play this more restricted role elicits a serious objection. We may imagine each of those facing minor harm say, "Look, I understand that my interest is irrelevant relative to the competing interest in avoiding major harm, so that I have no reasonable complaint if my interest is not taken into account in a choice between preventing major harms and minor harms. But that is not what I am asking. I am asking that my interest be able to defeat the competing claims against moderate harm, relative to which my interest *is* relevant. Not allowing my claim to play this more restricted role fails to give my interest the significance that it has."

In developing an adequate version of the Relevance Approach, then, we should keep in mind that relevance is a relative notion, which means that multiple standards of relevance may be operative in a trade-off situation. I am therefore in agreement with a suggestion recently made by Victor Tadros, the basic idea behind which is that if one claim is not relevant to a second, competing claim, then it may not play a role in outweighing the second claim. But it may still play a role in outweighing other competing claims in relation to which it is relevant. Tadros calls this idea *Local Relevance*.[42]

Local Relevance, however, is a broad suggestion, and the suggestion by itself does not dictate any particular version of the Relevance Approach. And indeed, as an indication of how rapidly the literature is progressing, Aart van Gils and Patrick Tomlin have raised fresh doubts about developing the Relevance Approach in the direction of Local

[42] Tadros, "Localised Restricted Aggregation," pp. 178-81.

Relevance.[43] I will go on to address their objections in Section 5. For now, however, let me go straight to my own preferred version the Relevance Approach.

3.2 The Partial Primacy Account

In developing an adequate version of the Relevance Approach, my account first directs our attention to the individuals who have the greatest need to be helped. In virtue of having the most at stake, these individuals should occupy our moral concern first and foremost. To be sure, the concern we owe to these individuals should *not* lead us to conclude that we ought to always prioritize their interests over the weaker interests of others. As evidenced by the case of *Life vs. Paraplegias*, some admittedly weaker interests should still be relevant to our moral consideration. However, there remains another way for the individuals with the weightiest interests to have pride of place in our process of adjudication, short of according their claims lexical priority — we may let their claims set the *context* and *aim* of the process of offsetting.

Here is the basic idea that I have mind. In adjudicating between the competing interests of individuals, we are to take as our controlling thought that *we should satisfy the strongest individual interests unless these interests can themselves be offset by competing interests that are relevant to them*. This does *not* mean, however, that those claims that are irrelevant relative to the strongest claims do not matter, as is the case under ARC. But it does mean that such claims matter *only insofar as* they bear on whether the strongest claims are in the end offset. In this way, the strongest interests in a given trade-off situation set

[43] Aart van Gils and Patrick Tomlin, "Relevance Rides Again? Aggregate and Local Relevance."

the *context* for determining which competing interests have a role to play in our process of adjudication.

Here is how we can develop the basic idea further. We begin by identifying the strongest individual claims in a trade-off situation — call them C1 — and we ask whether any competing claims are relevant to C1. If so, the claims that are relevant to C1 — call them C2 — have the standing to contribute to the offsetting of C1, *but only if they are not themselves offset*. To check whether C2 are able to be offset, we check to see if there are any competing claims that are relevant to C2. If so, these claims may contribute to the offsetting of C2, but again, only if they are not themselves offset. We continue this "relevance-finding" process until either (a) all the claims present in the trade-off situation are accounted for, or (b) we come upon claims that are not relevant to any of the claims previous ascertained to have the standing to offset other claims. If (b) is the case, then the claims with no offsetting power are deemed to be *inadmissible*. Since inadmissible claims in no way bear on whether the strongest claims are able to be offset, they are completely set aside.

Of the remaining claims previously ascertained to be admissible, we carry out the offsetting process *starting with the weakest of such claims*. This is because the ultimate *aim* in letting claims offset one another, under the current proposal, is to see whether the strongest claims are in the end offset. Thus, although we start the offsetting process from the bottom, so to speak, it is the claims at the top that occupy our ultimate attention. Throughout the process of offsetting, we must make sure that a claim which is irrelevant to a second claim may not contribute to the offsetting of the second claim. Finally, after the process of offsetting is carried to its completion, we should choose the option that

satisfies the strongest non-offset claims. Call this view of interpersonal trade-off the *Partial Primacy Account*, or *PPA* for short.

Although PPA is somewhat complicated to state, its underlying structure is quite simple: Can the strongest individual claims in a trade-off situation be offset by competing claims that are relevant to them, claims which are not themselves offset in turn? If not, we should choose the option that satisfies these strongest claims. If so, we should choose the option that satisfies the strongest, non-offset claims. There is, however, one caveat: Claims that are deemed to be inadmissible are completely set aside.[44]

The underlying rationale for PPA, to repeat, is this. In virtue of having the most at stake, individuals with the strongest claims in a trade-off situation have primacy in our process of adjudication. This primacy, however, is not absolute but only partial; it would be unreasonable for such individuals to insist that their claims should be satisfied no matter the number of relevant claims with which they compete. But such individuals may reasonably insist that their claims should be satisfied *unless* their claims are offset by competing claims that are relevant to them. In other words, the strongest claims in a trade-off situation set a moral benchmark, deviation from which is permissible, but only under the aforementioned "unless" condition. Now, given that the ultimate aim of the process of offsetting is to determine whether the strongest claims are in the end offset, it is only sensible to begin the offsetting process starting with the weakest claims in a trade-off situation. Moreover, given that inadmissible claims exert no impact on the question

[44] This caveat will become important in Section 5.

whether the strongest claims are in the end offset, such claims are set aside in our process of adjudication.

PPA issues the same verdicts in *Life vs. Headaches* and *Life vs. Paraplegias* as does Voorhoeve's ARC. In the former case, the many claims against a minor headache are irrelevant relative to the competing claim against death, so they may not perform any offsetting role. In the latter case, PPA allows the single claim against death to be offset by the many claims against paraplegia, with the remaining non-offset claims against paraplegia (if there are any) dictating which group we ought to save. But in the more complex cases raised by Tomlin and Horton, as we are about to see, the two accounts come apart.

4. Replies to Tomlin and Horton

4.1 A Reply to Tomlin

We are now in a position to address Tomlin's objection. Recall *Case One*:

> Stage 1: We can save either Group A, which consists of one person facing major harm, or Group B, which consists of ten people facing moderate harm.
>
> Stage 2: One person facing minor harm is added to Group A. One million people facing minor harm is added to Group B.

Under PPA, the one claim against minor harm in Group A is offset by one of the competing claims against minor harm in Group B. The many other claims against minor harm in Group

B are then set aside, since they are not relevant to the remaining claim against major harm in Group A. This leaves the claims operative at Stage 1 still in competition. And given that, by earlier supposition, ten claims against moderate harm outcompete one claim against major harm, PPA delivers the verdict that we should save Group B in *Case One*, a verdict that is fully in line with Tomlin's Principle.

The feature of ARC that gives rise to Tomlin's objection is this: Claims of identical strength may nonetheless differ with respect to their status of being relevant under ARC. Tomlin's objection cleverly exploits this feature of ARC by bringing into focus cases in which a greater and a lesser number of claims of identical strength are added to two competing groups, but in which only the lesser number of claims are deemed to be relevant by ARC. The objection then leverages an independent intuition captured by the Tomlin's Principle in order to establish that ARC yields highly counterintuitive implications.

PPA, however, does not share the problematic feature of ARC. Even though the one million claims added to Group B are irrelevant when compared to the competing claim against major harm, they are nonetheless relevant when compared to the single claim against minor harm added to Group A. PPA therefore grants these one million claims an offsetting power which ARC fails to recognize. And by ensuring that the single claim against minor harm added to Group A is completely offset, PPA thereby ensures that Tomlin's Principle is fully conformed to in a situation like *Case One*.

4.2 A Reply to Horton

I now turn to Horton's dilemma. Recall *Case Two*:

Stage 1: We can save either Group A, which consists of one person facing major harm, or Group B, which consists of one million people facing minor harm.

Stage 2: Fifteen people facing moderate harm are added to each group.

PPA yields the following conclusions in *Case Two*. The one million claims against minor harm in Group B completely offset the fifteen claims against moderate harm added to Group A. This leaves the one claim against major harm in Group A and the fifteen claims against moderate harm in Group B still in competition. By earlier supposition, the latter claims outcompete the former. PPA therefore implies that we should save Group B in *Case Two*, an implication that is fully in line with Horton's Principle.

Am I then suggesting that PPA is to embrace the other horn of Horton's dilemma and is thereby in violation of Tomlin's Principle in *Case Two*? Not at all. My contention is rather that Horton's dilemma is a false dilemma because *Case Two* describes a situation to which Tomlin's Principle does not properly apply. The inapplicability of Tomlin's Principle to *Case Two* is in fact explained by PPA itself.

Viewed through the lens of PPA, there is a crucial difference between *Case One* and *Case Two*, a difference which in turn explains why Tomlin's Principle is applicable to the former but not to the latter case. In *Case One*, the addition of claims at Stage 2 does not in any way affect the normative status of the claims already present at Stage 1. This is *not* true in *Case Two*.

In *Case Two*, the one million claims against minor harm lack any offsetting power at Stage 1; they are deemed by PPA to be inadmissible, and they are therefore completely

set aside. However, with the advent of Stage 2, the normative status of the one million claims change. They are now admissible, that is to say, they now have the kind of normative significance that allows them to offset the competing claims relative to which they are relevant. Unlike in *Case One*, therefore, the introduction of new claims in *Case Two* alters more than just the number of claims in the competition. It also alters the normative status of some of the claims already present at Stage 1.

This alternation in status is important, and here is why. The alternation indicates that each of the one million people facing minor harm has a reasonable complaint if her interest is not taken into account in our moral adjudication at Stage 2. We may imagine each of them say, "Look, I understand that, at Stage 1, my interest should be completely set aside, given that a major harm is at stake. But when moderate harms are introduced into the mix at Stage 2, surely my interest should *now* be allowed play an offsetting role. After all, relative to the added moderate harms, my interest in avoiding minor harm *is* relevant!"

The result, then, is that the fifteen claims against moderate harm added to Group A are completely offset. In effect, *Case Two* describes a situation in which, while we should save Group A rather than Group B at Stage 1, at the second stage, only Group B undergoes strengthening after a process of offsetting. In view of this strengthening of Group B, there is nothing strange or bizarre in concluding that we should switch from saving Group A to saving Group B at Stage 2. To put the point more succinctly, Tomlin's Principle fails to be applicable to *Case Two* because the clause "there are no other morally significant changes at t2" fails to be satisfied.

We may say that the addition of the fifteen claims against moderate harm to Group A creates a *bridging effect* in *Case Two*. In the absence of these fifteen claims, the one

million claims against minor harm are wholly without offsetting power. But thanks to the addition of the fifteen claims, the one million claims gain the status of being admissible. The presence of the claims against moderate harm serves as a metaphorical bridge, by means of which each of the one million claims crosses into the realm of normative significance.

Let me stress, in order to avoid confusion, that my reply to Horton is not merely that there is some mechanism internal to PPA which enables us to distinguish between Tomlin's *Case One* and Horton's *Case Two*. If that were my sole reply, then it would simply fail to address Horton's more basic contention that any version of the Relevance Approach that instructs us to save Group B in *Case Two* violates Tomlin's Principle. My reply, rather, is that the reading of Tomlin's Principle which underlies Horton's dilemma is not plausible. It is not plausible because it prohibits claims against minor harm to play an offsetting role even after competing claims against moderate harm are introduced into the calculus in *Case Two*. Such a prohibition, I have argued, can be reasonably rejected by each of those individuals facing minor harm.

5. Replies to van Gils and Tomlin

5.1 Sequential Claims-Matching

My discussion so far shows that PPA is immune to the objections to the Relevance Approach advanced by Tomlin and Horton. But there are further objections I wish to resolve. First, as I earlier mentioned, PPA incorporates the central insight behind Tadros's Local Relevance. Although Tadros himself does not develop Local Relevance in much detail, others have. In a recent article, Aart van Gils and Patrick Tomlin put forward, but

ultimately reject, an alternative way of developing Local Relevance, which they call *Sequential Claims-Matching* (*SCM*). Contrasting PPA with SCM helps to highlight some of PPA's unique features.

SCM advocates the following procedure of adjudication. Starting from the strongest individual claims, we eliminate a claim as soon as it is offset by competing claims that are relevant to it (the offsetting claims are eliminated as well). We then begin this process of elimination anew, focusing now on the strongest claims that remain. Once the process of elimination is carried to its completion, we should choose the option that satisfies the strongest non-eliminated claims. [45]

There are two crucial differences between SCM and PPA. First, SCM employs a process of offsetting that is top-down. By contrast, the process of offsetting under PPA has a bottom-up structure. Second, unlike SCM, PPA deems certain claims to be inadmissible, with the result that the inadmissible claims are completely set aside. These two unique features of PPA flow from the underlying rationale I have supplied in Section 2.2. To repeat, under PPA, it is the individuals with the weightiest interests that should occupy our moral concern first and foremost. The way we show this special concern, I have suggested, is to afford the strongest claims a default status — we should satisfy the strongest claims unless they are themselves offset. Since the aim of offsetting under PPA is to see whether the strongest claims are in the end offset, it only makes sense to start the offsetting process from the bottom. Moreover, since inadmissible claims have no bearing, either directly or

[45] Aart van Gils and Patrick Tomlin, "Relevance Rides Again? Aggregate and Local Relevance," pp. 228-29

indirectly, on whether the strongest claims are in the end offset, PPA tells us to exclude them in our process of adjudication.

These differences between the two accounts allow PPA to avoid the central objection van Gils and Tomlin leverage against SMC. The objection draws on the following case from F. M. Kamm:

Case Six: We can either save A's life or save B's life plus cure C's sore throat.[46]

Intuitively, we are not morally *required* to leave A to die merely because the option of saving B's life has the additional benefit of curing C's sore throat.[47] SCM, however, requires us to save B and C over A. This is because, under SCM, once A's claim and B's claim offset each other, there remains C's claim, which now dictates which group we ought to save.

My PPA is without this problematic implication. It permits the claims of A and B to offset each other, but it deems C's claim against a sore throat to be inadmissible. C's claim is thereby wholly set aside. PPA thus sees the option of saving A's life and the option of saving B's life as equally morally choiceworthy, notwithstanding C's claim against a minor discomfort.

---

[46] Kamm, *Morality, Mortality Volume 1*, p. 101.

[47] There is a further, orthogonal debate over how we should go about deciding between the two options in *Case Six*. Some suggest that we are permitted to simply choose either option, while others think we should give each side a maximal equal chance of being saved. I do not take a stand on the issue here. See, for example, Broome (1990) and Kamm (2007)

5.2 Two Ambiguities

Van Gils and Tomlin moreover identify two general ambiguities that proponents of Local Relevance need to address. The first ambiguity arises from the observation that, within a single trade-off situation, there might be multiple ways for claims to offset one another, which in turn yield incompatible results.[48] Consider, for example:

> ***Case Seven***: We can save either Group A, which consists of one person facing major harm (A1) and one million people facing minor harm (A2), or Group B, which consists of fifteen people facing moderate harm (B1) and one million people facing minor harm (B2).

There are two ways in which the offsetting process could unfold in *Case Seven*. Under *Option One*, we let the claims in A2 and the claims in B2 completely offset one another. This leaves the claims in A1 and the claims in B1 still in competition. Since by earlier supposition, fifteen claims against moderate harm outcompete one claim against major harm, we get the result that we should save Group B in *Case Seven.*

Alternatively, under *Option Two*, we begin by letting the claims in A2 offset the claims in B1. This leaves the claims in A1 and the claims in B2 still in competition. Since claims against minor harm are not relevant to claims against major harm, we get the different result that we should save Group A instead.

---

[48] Aart van Gils and Patrick Tomlin, "Relevance Rides Again? Aggregate and Local Relevance," pp. 245-49.

The way to resolve this ambiguity is to recognize that the offsetting process under Local Relevance is regulated by considerations of fairness. Under *Option Two*, the many claims against minor harm in A2 are allowed to a play a role in determining which group we should save. At the same time, however, the many claims of identical strength in B2 are wholly excluded from participating in the offsetting process. This kind of differential treatment is unjustifiable; each of those individuals in B2 may reasonably argue that the offsetting process under *Option Two* is unfair, given that there is an available option, namely *Option One*, under which the claims in A2 and the claims in B2 are equally allowed to participate in the offsetting process. Stated more generally, when there are multiple ways for claims to offset one another within a single trade-off situation, fairness requires that the ambiguity should be resolved *against* a process of offsetting which excludes some claims that are relevant to claims which are allowed to participate in that process.[49]

The second ambiguity van Gils and Tomlin identify arise in situations where a claim is incompletely offset (248-50). Consider, for example:

> ***Case Eight***: We can save either Group A, which consists of one person facing major harm (A1), or Group B, which consists of one person facing a harm slightly less severe than major harm (B1) and one person facing minor harm (B2).

[49] And here lies another reason why under PPA, the primacy enjoyed by those individuals who have the strongest claims to be helped is only partial. Although the offsetting process described by *Option Two* favors the person who has the strongest claim in *Case Seven*, the special concern we have for the worst-off individual should not lead us to "rig" the offsetting process by unfairly excluding the claims in B2 from participating in offsetting.

In this case, the B1 claim only incompletely offset the competing A1 claim, since the former is slightly weaker than the latter. Now, the B2 claim is not relevant relative to the full strength of the A1 claim, but it may well be relevant relative to what is left of the A1 claim after the latter is incompletely offset. Should we or should we not then take the B2 claim into account in our adjudication of *Case Eight*?

The first thing to note is that, under PPA, the relevance of a claim is determined by comparing the claim to the full strength of some competing claim. This is because the "relevance-finding" exercise under PPA occurs before any offsetting takes place. As a procedural matter, therefore, PPA avoids the ambiguity at issue. The rationale for focusing on the full strength of claims in ascertaining relevance is as follows. In *Case Eight*, while it is true that the claim against minor harm is relevant relative to some *proportion* of the claim against major harm, there is no circumstance under which the person who faces major harm suffers only a proportion of that harm. She will either suffer the full severity of that harm (if her claim is not satisfied) or completely avoid that harm (if her claim is satisfied). Hence, the operative question in *Case Eight* is whether the A1 claim should be satisfied, *not* whether some proportion of the A1 claim should be satisfied. And in answering that operative question, it is plainly the full strength of the A1 claim that needs to be used in ascertaining which competing claims have the standing to defeat the A1 claim.

6. Multi-Option Cases

There is one final objection to the Relevance Approach I want to address. Derek Parfit objects that the approach carries implausible implications in cases where multiple options

are available.[50] Consider the example Parfit introduces to illustrate this supposed implausibility:

> ***Three Options***: We can save either Group A, which consists of one person facing major harm, Group B, which consists of fifteen people facing moderate harm, or Group C, which consists of one million people facing minor harm.

Parfit asks proponents of the Relevance Approach to say what their view implies about which group we ought to save in *Three Options*. Is it to save Group A? But the claims in Group B outcompete those in Group A. Is it to save Group B? But the claims in Group C outcompete those in Group B. Finally, is it to save Group C? But the claims in Group A outcompete those in Group C. It seems that, under the Relevance Approach, we are completely at a loss as to what we should do in the present situation. That, Parfit charges, is an unacceptable conclusion.

Parfit's examples calls for an extension of PPA to multi-option cases. We can do so by recalling the guiding idea behind PPA, which is that the only exception to satisfying the strongest claims in a trade-off situation is that these claims are themselves properly offset. Here is how we can adapt the idea to multi-option scenarios. We begin by identifying the option with the strongest individual claims — call it O1 — and we compare, *in a pair-wise fashion*, each of the available options to O1. If an option "loses" to O1 in

[50] Parfit, "Justifiability to Each Person," p. 384. Norcross raises the same worry in his "Contractualism and Aggregation," pp. 308-09.

accordance with the application of PPA to two-option cases, then that option is eliminated from the competition.

If after the first stage of elimination, there remains multiple options that outcompete O1, then we eliminate O1, and we identify, of the remaining options, the one with the strongest individual claims — call it O2. We then compare, again in a pair-wise fashion, each of the remaining options to O2, and we continue this process of comparison and elimination in the way just enumerated, until either there is a single option that is left standing or there is a tie.

The rationale for applying PPA to multi-option cases in this way has two aspects. First and most obviously, letting the individuals with the strongest claims set the benchmark for elimination expresses the special concern we owe to these individuals, a concern that animates much of the apparatus of PPA. Second, by comparing the options one-on-one to the option with the strongest claims, and eliminating those options that are outcompeted, we ensure that whichever option ultimately emerges victorious, it is one that can successfully offset the option with the strongest claims. In this way, we fully adhere to PPA's central idea that the only exception to satisfying the strongest claims in a trade-off situation is that these claims are themselves relevantly offset by the claims we ultimately satisfy.

We can now respond to Parfit's objection. The individual claims in Group A are the strongest, so we let these claims set the benchmark for elimination. Since the claims in Group C are irrelevant relative to the claims in Group A, the option of saving Group C is eliminated from the competition. This leaves Group A and Group B still in the running. And given that, by earlier supposition, fifteen claims against moderate harm outcompete

one claim against major harm, PPA tells us to save Group B in *Three Options*. Far from being at a loss as to what to say, PPA delivers a clear verdict in Parfit's example.[51]

7. Conclusion

Many of us hold intuitions that suggest there are limits to our acceptance of interpersonal aggregation. In some cases, we are inclined to accept that individual claims may aggregate in order to collectively defeat competing claims. In other cases, we are inclined to reject such interpersonal aggregation. The Relevance Approach is attractive precisely because it offers an initially appealing explanation for our divergent intuitions about the moral significance of numbers. But as the recent critics have charged, the Relevance Approach seems to face devastating counterexamples of its own. This chapter provides a development and defense of the Relevance Approach, one that is animated by the idea that individuals with the strongest claims in a trade-off situation enjoy qualified primacy in our process of interpersonal adjudication.

[51] PPA's treatment of *Three Options* gives rise to a further worry. In a competition between Group B and Group C, PPA instructs us to save Group C. However, in *Three Options*, where Groups A, B, and C are all present, PPA instructs us to save Group B instead. It thus appears that PPA is in violation of a supposed principle of rationality commonly known as the Independence of Irrelevant Alternatives. This is a well-known worry, and proponents of the Relevance Approach have in turn given their response, which is convincing in my view. See footnote 31 above.

Chapter Three: Interpersonal Trade-Off and Risk

1. Introduction

In the previous two chapters, I have focused on the adjudication of interpersonal trade-offs under conditions of certainty. That is, I have focused on cases where the outcomes are certain and where the agent knows both who will be affected by her decision and how they will be so affected. Life, however, is risky. Oftentimes, the choices we face involve balancing risks of harm to some against uncertain benefits for others. How our trade-off decisions should be made in such risky situations is a question to which we should expect moral theories to give adequate guidance.

Recent attempts to develop a nonconsequentialist account of risk, however, have confronted a serious controversy. The disagreement centers around two competing models for assigning significance to the interests of individuals under conditions of uncertainty — an *ex ante* model and an *ex post* model. My central aim in this chapter is to provide a satisfactory resolution of the problem.

I proceed as follows. In the next section, I introduce the two competing models for assessing the significance of individual interests when outcomes are uncertain. Next, in Section 3, I press three lines of objection against the *ex ante* model. In Section 4, I articulate a qualified version of the *ex post* model, and I defend my account against the much-

discussed problem of social risk.[52] My account is qualified in that it introduces a mechanism through which *ex ante* reasoning, or something very much like it, is found to be properly applicable in a range of cases. Seen this way, it may also be said that I am here advocating for a hybrid view of risk. In closing, I address a further objection against the view I defend.

2. *Ex Ante* vs. *Ex Post* Models

How should we assess the interests of individuals in trade-off situations in which the outcomes are uncertain? Let us begin with some uncontroversial cases. Suppose that in:

> ***Disease 1***: Amy is threatened with a serious disease which, if left untreated, would result in her losing an arm. The only way to cure Amy of her disease is to immediately release an odorless gas into her room. Betty is in the room with Amy, and she cannot be removed in time. Releasing the gas would kill Betty.

Suppose, in this and all subsequent cases, that we have to make a decision in isolation, and we cannot consult the wishes of the individuals whose fates we are to decide. I take it that it is clearly morally impermissible for us to release the gas in this case.

Next, consider:

[52] For recent discussions of the problem of social risk, see, for example, Rahul Kumar, "Risking and Wronging," *Philosophy and Public Affairs*, 43 (2015): 27-51; and Johann Frick, "Contractualism and Social Risk," *Philosophy and Public Affairs* 43 (2015): 175-223.

Disease 2: The details are the same as *Disease 1*, except that releasing the gas only imposes a 1/10000 chance of death onto Betty.

I take it that it is clearly morally permissible, though perhaps not morally required, for us to release the gas in *Disease 2*, thereby imposing a small risk of death onto Betty, for the sake of preventing Amy from losing an arm. If one thinks that a 1/10000 chance of death is still too high to make the trade-off permissible, then she should substitute into the example an even lower number, say, one-in-a-billion. Surely, it is not plausible to maintain that we may not permissibly save Amy's arm *no matter* how low the risk of death to Betty becomes.

Finally, consider:

Disease 3: Carl is threatened with a serious disease which, if left untreated, would result in him losing an arm. The only way to cure Carl of his disease is to immediately release an odorless gas into his room. 10000 individuals are in the room with Carl, and they cannot be removed in time. Releasing the gas would cause a lethal reaction in one of the 10000 individuals, though we cannot know in advance who this person will be.

Disagreements arise over the proper analysis of *Disease 3*. According to one line of reasoning, the fact that someone will die if we release the gas in *Disease 3* renders the

situation morally analogous to *Disease 1*. Given that it is morally impermissible for us to release the gas in *Disease 1*, it is similarly impermissible, it is said, for us to release the gas *Disease 3*.[53]

According to another line of reasoning, however, the fact that, by releasing the gas in *Disease 3*, we impose on each particular individual (other than Carl) a 1/10000 risk of death renders each of their situations morally analogous to Betty's situation in *Disease 2*. Since, as we have seen in *Disease 2*, it is permissible to impose on a person a 1/10000 chance of death for the sake of preventing another person from losing an arm, it is concluded that the option of releasing the gas in *Disease 3* is similarly morally permissible.[54]

One point of clarification before we move on. A key difference between *Disease 2* and *Disease 3* is that, in the former case, Betty is the only person who has to bear a small risk of death for the sake of benefitting Amy, whereas in the latter case, 10000 individuals will each bear some risk of death if we decide to release the gas. It might then be suggested that, once we aggregate the complaints of the 10000 individuals against being exposed to some lethal risk, it may well turn out that releasing the gas in *Disease 3* is morally impermissible despite its being permissible to do so in *Disease 2*.

---

[53] For a discussion and defense of this line of reasoning, see, for example, Michael Otsuka, "Risking Life and Limb," in *Identified versus Statistical Lives*, ed. I. Glenn Cohen, Norman Daniels, and Nir Eyal (Oxford: Oxford University Press, 2015), 77-93.

[54] For a discussion and defense of this line of reasoning, see, for example, Frick, "Contractualism and Social Risk." Kumar can be read as endorsing the general spirit behind this line of reasoning. See Kumar, "Risking and Wronging."

At this juncture, it is important to recall the central insight of the Relevance Approach which I have defended in earlier chapters. A person's interest in avoiding a very small risk of death, I submit, is irrelevant in comparison to another person's much stronger competing interest in avoiding the certain loss of an arm. This means that no number of complaints based on a 1/10000 chance of death may collectively defeat a countervailing complaint against the certain loss of an arm. Hence, if the correct analysis of *Disease 3* proceeds from the fact that each of the 10000 individuals is exposed to only a 1/10000 chance of death, and not from the fact that some one of the 10000 individuals is certainly going to die if we release the gas, then the correct conclusion is that releasing the gas in *Disease 3* is morally permissible.

Now, it is sometimes said that the difference between the two lines of reasoning concerns whether individual interests in risky situations should be assessed *ex ante*, based on the *prospects* of harm and benefit the action imposes, or *ex post*, based on the *actual* results that the action has for people.[55] But this is a mischaracterization of the disagreement. Consider:

> ***Villain 1***: An evil villain is about to wreak immeasurable havoc unless you choose between the following:
>
> (a) Plan A, in which the villain cuts off a limb from Debby.

[55] See, for instance, and Joe Horton, "Aggregation, Complaints, and Risk," *Philosophy and Public Affairs* 45 (2017): 56.

(b) Plan B, in which the villain institutes a 70% chance that some one person from a group of seven people will die, though we cannot know in advance who this person will be.

You choose Plan B. Luckily, as things turn out, no one dies.

Suppose, to fix ideas, that under Plane B, a die with ten faces, numbered 1 to 10, will be rolled by the villain. Each of the group of seven individuals is assigned a unique number between 1 and 7. If the die lands on 8, 9, or 10, then everyone survives the situation unharmed. If, on the other hand, the die lands on any number between 1 and 7, then the person who is assigned that number will be killed by the villain. In this way, each particular individual has a 10% chance of being killed by the villain, all the while there is 70% chance that someone from this group of seven individuals will perish at the hands of the villain.

Supporters of the *ex post* model would still register a complaint, based on the 70% prospect that someone will be lethally harmed, against your opting for Plan B in *Villain 1*, even though your decision, in actuality, benefits some and harms no one. Thus, both the *ex ante* model and the *ex post* model are concerned with the prospects of benefit and harm associated with risk-imposing actions, rather than with their actual outcomes. What crucially divides the two models is the thesis that individual interests against a risky action, and the complaints to which they give rise, may be registered against that action based on the prospects that *someone or other* — rather than particular individuals identifiable at the time of decision — would thereby suffer harm. As I shall use the labels, this thesis is affirmed by the *ex post* model (henceforth: *EPM*) and denied by the *ex ante* model (henceforth: *EAM*).

3. Rejecting the Ex Ante Model

I want to now press three objections against *EAM*. The first objection is this: In certain cases involving risk, *EAM* wrongly permits us to secure relatively minor benefits for a few at the cost of imposing severe burdens on many. To see this more clearly, let us return to the plausible observation that there is some degree of lethal risk which may be permissibly traded off against the certain harm of losing a limb. Consider:

> ***Disease 4***: Both Eddie and Fay are threatened by a fatal disease. Two courses of treatment are available for institution, but there is only enough resources to institute one of them.
>
> (a) Under Treatment A, Fay will survive the disease unscathed. Eddie too will survive, but he will certainly lose a limb.
>
> (b) Under Treatment B, Eddie will survive the disease unscathed. Fay has a 9999/10000 chance of surviving the disease unscathed, but there is a one-in-ten-thousand chance that the treatment is ineffective for Fay.

I take it that it is morally permissible for you opt for (b) in this case. In other words, I take it that a 1/10000 risk of death may be permissibly offset by the benefit of preventing the certain loss of a limb. Once again, those who disagree with my calculations may substitute into the example their preferred trade-off calculations, and rework the subsequent examples accordingly.

Now, suppose that:

Disease 5: Gob and 1 million other people are threatened by a fatal disease. Once again, two courses of treatment are available, but only one can be feasibly instituted.

(a) Under Treatment A, everyone will survive the disease, but Gob will certainly lose a limb.

(b) Under Treatment B, Gob, along with all but 100 of the 1 million individuals will survive the disease unscathed. The 100 individuals will die. We cannot know in advance who these one hundred individuals will be.

Clearly, it is morally impermissible for you to opt for (b). *EAM*, however, says otherwise. For the risk of death that any one identifiable individual incurs, under option (b), is at most one-in-ten-thousand. Since, as we have been supposing, a 1/10000 risk of death may be permissibly offset by the benefit of preventing the loss of a limb, *EAM* renders the verdict that you are permitted to save Gob's single limb at the cost of causing the deaths of 100 people.

Matters are still worse for *EAM* once we realize that, at least in principle, the view implies that Gob's single limb may be saved at the cost of *any number* of unidentifiable lives, so long as the risk of death to each identifiable person is kept at or below 1/10000. Thus, if it is onto a group of 10 billion people that lethal risks will be imposed, then *EAM* dictates that we may permissibly spare Gob from losing a single limb at the cost of causing any number between 0 and 1 million deaths. This is deeply implausible.

EPM, let us note, handles cases like *Disease 5* with ease. Since we know in advance that, by opting for (b) in *Disease 5*, some 100 individuals will die, and since complaints based on someone's suffering of lethal harm is clearly stronger than Gob's compliant based on the loss of a single limb, *EPM* delivers the correct verdict that instituting Treatment B in *Disease 5* is morally impermissible.

Let me now move onto my second objection against *EAM*, which is that, intuitively, in certain cases involving risk, uncertainties regarding the identities of potential victims do not matter morally at all. Consider the following case:

> ***Villain* 2**: An evil villain is about to wreak immeasurable havoc unless you choose between the following:
>
> (a) Plan A, in which the villain cuts off a limb from Haley.
>
> (b) Plan B, in which the villain randomly selects one person among Iris and Jake, and cuts off a limb from that person.[56]

EAM emphatically implies that you are morally required to choose (a) over (b). This is so because, according to the *ex ante* line of reasoning, each of Iris and Jake's complaints against losing a limb is discounted by 50%, whereas Haley's complaint is undiscounted. I find this implication counterintuitive. It seems to me that, when it comes to choosing between Plan A and Plan B in this case, facts about whether it is Iris or Jake that will lose a limb under Plan B does not make a moral difference at all. Intuitively, both options are

[56] Horton discusses a similar objection to *EAM*. See Horton, "Aggregation, Complaints, and Risk," 59-60.

morally on a par, and this intuition is best explained by the *ex post* consideration that, under either option, *exactly some one person* will incur the cost of losing a limb.

In saying this, however, I do not wish to suggest that, in the following case, the two available options are morally on a par:

> ***Villain 3***: An evil villain is about to wreak immeasurable havoc unless you choose between the following:
>
> (a) Plan A, in which the villain cuts off a limb from Hayley.
>
> (b) Plan B, in which the villain randomly selects one person among Hayley and Iris, and cuts off a limb from that person.

I accept that in this case, Plan B is morally preferable to Plan A, even though under either option, exactly some one person will incur the cost of losing a limb. In my view, by opting for Plan A over Plan B in *Villain 3*, you would be behaving in a way that is unfair to Hayley. By contrast, no such unfairness is involved in the original *Villain 2* case.

Call the harm of losing a single limb *h*. In *Villain 3*, your choice is between giving Hayley no opportunity at all to avoid *h* and giving Hayley some opportunity to avoid this harm. Since Hayley and Iris are both morally innocent, and since both are equally worthy of moral protection, it would be unfair to Hayley to wholly impose *h* on her, while completely shielding Iris from any risk of harm.

Matters, however, are different in *Villain 2*. There, your choice is between giving Hayley no opportunity at all to avoid *h* and completely shielding her from any risk of harm, at the cost of wholly imposing the harm on either Iris or Jake. Unlike in *Villain 3*, there is

no option available which gives each of Hayley, Iris, and Jake some chance to avoid *h*; it is either onto Hayley or one of Iris and Jake that the harm of losing a limb will befall.

In launching a complaint against your opting for (a) in *Villain 3*, Hayley is not making the self-regarding claim that she should be wholly spared from the harm of losing a limb, while all the harm should instead go to Iris. Rather, she is making the claim that, given that there is a way for she and Iris to share the risk together, and given that she is no less worthy of moral protection than Iris is, it would be unfair to impose the harm on her and her alone. *This* complaint is not available to Hayley in *Villain 2*, for, in objecting to your opting for (a) in Villain 2, Hayley is in effect making the claim that she should be completely shielded from any harm, while the entire harm should be imposed instead on one of Iris and Jake. In the end, what I wish to maintain is simply this: by claiming that the two available options are morally on a par in *Villain 2*, I need not thereby commit myself to the claim that the two available options are morally on a par in *Villain 3*.

Here is my third and final line of objection against *EAM*. It begins with the observation that, in certain cases involving risk, the view seems to imply an implausible stance of indifference toward action alternatives, one of which is clearly morally superior to the other. Consider:

> ***Villain 4***: An evil villain is about to wreak immeasurable havoc unless you choose between which of the following two games of roulette he is to play. Both games employ a roulette wheel with 1000 slots, numbered 1 to 1000:
>
> (a) Under Game 1, if the roulette ball lands on a slot marked by any number between 1 and 500, Nathan will die. Otherwise no one will be harmed.

(b) Under Game 2, if the roulette ball lands on a slot marked by any number between 1 and 500, Ona will die, and if the roulette ball lands on a slot marked by any number between 501 and 1000, Pam will die.

Clearly, opting for Game 1 is a morally superior choice than opting for Game 2. *EPM* can easily account for this judgment. For, under Game 1, the prospect that *someone* will be lethally harmed is 50%, whereas, under Game 2, that prospect increases to 100%. A complaint based on a 100% risk of death is evidently weightier than one based on a 50% risk of death.

Matters, however, are not so clear under *EAM*. On the one hand, Nathan has a complaint against your choosing Game 1, a complaint that is based on a 50% risk of death. On the other hand, each of Ona and Pam has a complaint, based similarly on a 50% risk of death, against your choosing Game 2. And since on the individual complaint model, as it is so far characterized, prospects of harm cannot be aggregated across persons, Ona and Pam's complaints cannot be combined into one. As such, *EAM*, unlike its *ex post* counterpart, cannot produce a complaint against Game 2 that rests on a 100% risk of death.

At this juncture, defenders of *EAM* will likely appeal to a tie-breaking mechanism that roughly functions in the following way: in cases where one must decide whether some harm is to befall each individual in a smaller group or each individual in a larger group, and where the harm each person faces is identical, one is morally required to decide in favour of sparing the larger group of individuals from harm.[57] Defenders of *EAM* may then

[57] For a defense of this tie-breaking mechanism, see Scanlon, *What We Owe to Each Other*, 229-41. There is a worry that this mechanism cannot be justified without resorting to

reason that, since Ona and Pam belong to a group of two, whereas Nathan only comprises a group of one, and since each of Nathan, Ona, and Pam faces the exact same risk of harm, the tie-breaking mechanism dictates that you ought to choose Game 1 over Game 2.

But appealing to such a tie-breaking mechanism is not enough to rid *EAM* of its troubling implications. Consider:

> ***Villain 5***: An evil villain is about to wreak immeasurable havoc unless you choose between which of the following two games of roulette he is to play. Both games employ a roulette wheel with 1000 slots, numbered 1 to 1000:
>
> (a) Under Game 1, if the roulette ball lands on a slot marked by any number between 1 and 200, Quentin will die. Otherwise no one will be harmed.
>
> (b) Under Game 2, Ray is assigned the numbers between 1 and 100; Steve, the numbers between 101 and 200; Tammy, 201 and 300; Udo, 301 and 400; …; and Victor, 901 and 1000. Whichever slot the roulette ball ultimately lands on, the person who is assigned the corresponding number will die.

Clearly, opting for Game 1 is a morally superior choice than opting for Game 2. *EPM* can easily account for this judgment. For, under Game 1, the prospect that *someone* will be lethally harmed is 20%, whereas, under Game 2, that prospect increases to 100%. *EAM*, by contrast, says that you are morally required to opt for Game 2. This is so because Quentin's

---

interpersonal aggregation. See Michael Otsuka, "Scanlon and the Claims of the Many versus the One," *Analysis* 60.3 (2000): 288-93.

complaint against your choosing Game 1, based on a 20% risk of death, is weightier than any one of the complaints from each of the other ten individuals.

Perhaps defenders of *EAM* will now want to resort to a rule that allows for limited interpersonal aggregation in circumstances where the harm involved to each individual is *comparable*. Such a rule is not *ad hoc*, since, in my view, we should allow for trade-offs between, say, some number of deaths and some larger number of permanent full-body paralyses. Defenders of *EAM* may then reason that, since a 20% risk of death is comparable to a 10% risk of death, the harm that befalls each of Ray, Steve, Tammy, Udo, …, and Victor can accordingly be aggregated. And when these harms are summed, they give rise to a complaint against Game 2 that outweighs Quentin's complaint against Game 1.

Unfortunately, however, this defense of *EAM* still falls short. To see this, let us observe, first, that a 0.1% risk of lethal harm is *not* comparable to an 50% risk of lethal harm, just as suffering a day of pain is *not* comparable to suffering two years of pain. Now, consider:

> ***Villain 6***: An evil villain is about to wreak immeasurable havoc unless you choose between which of the following two games of roulette he is to play. Both games employ a roulette wheel with 1000 slots, numbered 1 to 1000:
>
> (a) Under Game 1, if the roulette ball lands on a slot marked by any number between 1 and 500, Wendy will die. Otherwise no one will be harmed.
>
> (b) Under Game 2, each of 1000 individuals are assigned a unique number between 1 and 1000. Whichever slot the roulette ball ultimately lands on, the person with the corresponding number will die.

Plainly, Game 1 is a morally superior choice than Game 2. *EPM*, unsurprisingly, can easily account for this judgment. For, under Game 2, it is absolutely guaranteed that someone will die. Under Game 1, by contrast, the prospect that someone will die is only 50%. *EAM*, however, implies that you are morally required to opt for Game 2, despite knowing that some one person will die as a result. Neither a tie-breaking mechanism nor a rule that allows for limited aggregation can save *EAM* from this counterintuitive implication. I conclude that *EAM* should be rejected.

4. Defending EPM

EPM denies there is any moral significance to the distinction between a complaint based on the fact that someone will suffer some measure of harm and a complaint based on the fact that some particular, identifiable person will suffer the same measure of harm. Under *EPM*, complaints against a risk-imposing action ϕ is assigned based on the full magnitude of the harm that ϕ will likely impose on some one person or another, discounted by the improbability of that harm eventuating.[58]

---

[58] Sophia Reibetanz agrees that, in cases where we are *certain* that someone will be harmed by ϕ-ing, we should assign a complaint against ϕ-ing *ex post*. But she suggests that when we are less than certain that some one person will be harmed by ϕ-ing, we should assess the permissibility of ϕ-ing *ex ante*. See Sophia Reibetanz, "Contractualism and Aggregation," *Ethics* 108 (1998): 303-304. Reibetanz's suggestion, however, is quite implausible. Consider:

> ***Villain 7*:** An evil villain is about to wreak immeasurable havoc unless you choose between the following:
> (a) Plan A, in which the villain cuts off a limb from Xavier.

EPM, however, faces a serious objection, one which has been dubbed by Johann Frick the *problem of social risk*.[59] The objection begins by observing that some large-scale, risk-imposing social policies are characterized by the following features:

1. The policies carry risks of harm that, given how many individuals these policies affect, are virtually guaranteed to eventuate in some one person or another.
2. The loss to each individual who turns out to be burdened by these policies are considerably greater than the gain to each individual who turns out to be benefited by them.
3. It cannot be assumed that those who are burdened by the policies will, in the long run, benefit from the policies' institutions.

Call policies with the aforementioned features *Socially Risky Policies*, or *SRP* for short. *EPM* seem to straightforwardly imply that instituting and maintaining SRP is morally forbidden. The problem, however, is that some SRP are intuitively permissible.

(b) Plan B, in which a 10-faced die, numbered 1 through 10, will be rolled. If the die lands on any number between 1 and 8, some one person from a group of 8 people will die. Otherwise, no one will be harmed.

I take it that supporters of *EPM* would wish to allow the registration of a complaint, based on an 80% chance that someone will be killed, against your choosing Plan B over Plan A.

[59] Frick, "Contractualism and Social Risk," 178-9.

Some have suggested that driving is a type of SRP which *EPM* wrongly deems as morally problematic.[60] But I think we should be wary of this suggestion for at least two reasons. First, it is unclear that driving, as it is being currently practiced in our society, is indeed a permissible type of SRP, since it seems quite likely that there is some alternative set of safety standards, the establishment of which would lower driving-related risks to both drivers and non-drivers. Hence, for all we know, the institution of driving, in its current form, might very well fail to meet the demands of morality, on account of its not having enough precautionary measures in place.

Second, so far as I can tell, driving fails to satisfy the second feature of SRP mentioned above. We should reject the rather unimaginative claim that the only kind of benefit driving brings about is leisurely convenience. It is all but certain that, either directly or indirectly, permitting people to drive has resulted in and will continue to result in the preventions of serious injuries and deaths. Think, for instance, of the benefits of granting emergency vehicles access to roads and of allowing people to transport themselves quickly, to where their presence is urgently needed. If the problem of social risk is to be a decisive objection against *EPM*, then we would do better to focus on less contentious examples of permissible SRP.

To this end, consider public health policies concerning large-scale vaccinations of children. According to Frick,

[60] Kumar, "Risking and Wrong," 27-8.

[I]t is commonly deemed morally unproblematic to systematically inoculate young children against certain serious but non-fatal childhood diseases where there is a remote chance of fatal side effects from the inoculation itself. But this is not because it is unlikely that inoculation will ever lead to disaster for some unlucky children. Given the large number of children inoculated each year, it is a statistical certainty that some number of them will develop fatal complications.[61]

Now, consider a suitably modified version of Frick's central case against *EPM*:

Mass Vaccination: One million children are threatened by a fatal virus. You must choose between the following two vaccines:

(a) Vaccine 1 is certain to save every child's life, but the protection it affords is not complete. If a child receives vaccine 1, the virus will leave her vision seriously but not severely impaired.

(b) Vaccine 2 is risky. It gives every child a 999/1000 chance of surviving the virus completely unharmed, but for every child, there is a 1/1000 chance that vaccine 2 will be completely ineffective against the virus.[62]

---

[61] Frick, "Contractualism and Social Risk," 185-6.

[62] Frick, "Contractualism and Social Risk," 181-2. The case is modified in the following respect: In Frick's original case, vaccine 1 leaves each child partially paralyzed. In my view, complaints against partial paralysis may collectively defeat a countervailing complaint against death. So, even under EPM, option (b) may still be morally permissible. It is worth noting that Frick conducts his discussion within the framework of moral contractualism, and he assumes that contractualist reasoning does not allow for the aggregation of claims across different

Given the Law of Large Numbers, it is virtually guaranteed that at least one child will be killed as the result of your instituting vaccine 2. Intuitively, it is nevertheless permissible for you to opt for vaccine 2 in this case. *EPM*, by contrast, appears to say otherwise, seeing as the *ex post* complaint against instituting vaccine 2, based on the harm of death, is stronger than the *ex post* complaints against instituting vaccine 1, based on the harm of partial vision impairment.

Of course, for the case to be a genuine counterexample to *EPM*, it needs to be the case that the complaints against serious but not severe vision impairment are not sufficient to collectively defeat a countervailing complaint against death in *Mass Vaccination*. I will assume, for the sake of argument, that this is indeed so. Even so, I contend that, despite appearance to the contrary, *EPM can* deliver the intuitively correct verdict that opting for vaccine 2 is morally permissible.

To see this, let me first introduce the notion of a *complaint diminisher*. In my view, some facts are such that, in their presence, the force of someone's complaint against a risk-imposing action can be made weaker than it otherwise would have been. These facts I call complaint diminishers. Let me offer two examples.

Suppose that in:

---

individuals. Hence, under Frick's anti-aggregative understanding of contractualism, EPM rules out option (b), since a complaint against death is stronger than any one countervailing complaint against partial paralysis. I have argued against this anti-aggregative understanding of contractualism in Chapter One.

Virus 1: A group of 1 million people is threatened by a virus X which causes serious but not severe vision impairment if untreated. You must choose between the following:

(a) Let each of this group of 1 million individuals be visually impaired by X.

(b) Subject a different group of 1000 individuals to medical testing, who are not themselves threatened by virus X. We are assured that the testing will be painless for all those involved. The testing will produce a complete cure for those who are threatened by X. We know in advance that one person from this group of 1000 people will die painlessly, though we cannot know in advance who this person will be.

In this case, option (b) carries with it the harm to someone of suffering a painless death. Under *EPM*, this harm — to some one person in the group of 1000 individuals — grounds a complaint against option (b) whose force is greater than that of the complaints grounded in the harm of being visually impaired. *EPM* accordingly deems option (b) to be morally forbidden.

Let us now introduce some additional layers of complexity to our case. Suppose that in:

***Virus 2*:** The details are the same as *Virus 1*, except that each of the group of 1000 individuals has given their informed consent to being the subject of medical testing.

Even in the revised case, option (b) still carries with it *the harm* to someone of suffering a certain death. Moreover, this harm to some one person can still ground *a complaint* against your opting for (b). But *the force* of this *ex post* complaint is clearly diminished in the presence of the fact that people have given their consent, so much so that, in *Virus 2*, option (b) is no longer morally forbidden.[63] We may put this point by saying that:

> *Consent as a complaint diminisher*: the force of a complaint against an action ϕ based on the risk of harm that ϕ imposes on someone (identifiable or not) *A* can be diminished when *A* has given her informed consent regarding incurring the risk of harm ϕ imposes.

Next, suppose that in:

> ***Virus 3***: The details are the same as *Virus 1*, except that the group of 1000 individuals are evil scientists who jointly synthesized virus X with the deliberate aim to cause impairment in others.

Again, option (b) here still carries the same lethal harm as it does in *Virus 1* and in *Virus 2*, and this harm can still ground a complaint against your opting for (b). But the force of

[63] My discussion aims to keep distinct, as a conceptual matter, the following three notions: the harm that ground a complaint against ϕ-ing, the existence of a complaint against ϕ-ing, and the force of a complaint against ϕ-ing. I am happy to accept that, if the force of a complaint against ϕ-ing reduces to zero, the complaint thereby disappears.

this *ex post* complaint is clearly reduced in the light of facts about moral liability, so much so that, in *Virus 3*, option (b) is no longer morally forbidden. We may put this point by saying that:

> *Liability as a complaint diminisher*: the force of a complaint against an action ϕ based on the risk of harm that ϕ imposes on someone (identifiable or not) *A* can be diminished when *A* bears liability for the risk of harm ϕ imposes.

Finally, suppose that in:

> ***Virus 4***: The details are the same as *Virus 1*, except that the group of 1000 individuals are themselves part of the group of 1 million individuals, each of whom is threatened by virus X.

I believe option (b) under *Virus 4* is no longer morally forbidden. This is so, I want to suggest, because some actions are what I shall call, with a harmless air of paradox, *victim-benefiting*, and facts about being victim-benefiting are complaint diminishers. As I understand the notion, a risk-imposing action ϕ is victim-benefiting just in case, *at the time of ϕ's performance* and compared to its alternatives, ϕ gives the best prospects to the individuals who turn out to be the most burdened under ϕ.

What I am here claiming, then, is that, first, some facts are complaint diminishers; second, an option being victim-benefiting can diminish the force of *ex post* complaints against that option; third, in *Mass Vaccination*, the force of *ex post* complaints against your

opting for vaccine 2 is diminished because it is a victim-benefiting option. The result is that opting for vaccine 2 in *Mass Vaccination* is no longer morally impermissible.

The trade-off situation in *Virus 4* differs from that in *Virus 1* in only one respect — the individuals onto whom lethal risk is proposed to be imposed are now the very individuals for whose enhanced benefits the proposal aims to secure. This fact diminishes the force of the *ex post* complaint against opting for (b) in *Virus 4*. Just as, in *Virus* 3, the force of a complaint based on someone's suffering lethal harm is diminished by the fact that that very someone, whoever he or she turns out to be, is liable to be so harmed, the force of such a complaint is similarly diminished by the fact that that very someone, whoever he or she turns out to be, would herself recognize that the risky action under which she is a loser is one which she would reasonably prefer to adopt at the time of deciding between the available alternative options. Losers under a victim-benefiting action are losers of a gamble that the gamblers themselves would have reasonably agreed is worth taking.

Let us not forget, however, that there are trade-off situations do not involve victim-benefiting actions. The complaint diminisher such actions engender, therefore, is not applicable in many cases. Like *Virus 4*, *Mass Vaccination* is a case in which the eventual victims of option (b) are the very individuals whose enhanced benefits the option aims to secure. And even those who turn out to be losers under option (b) can agree that, at the time of decision, option (b) gives them the best welfare prospects compared to the available alternatives. Option (b) under *Mass Vaccination* thus generates a complaint diminisher, one whose presence explains how *EPM* can deliver the verdict that opting for the risky vaccine in *Mass Vaccination* is morally permissible.

My proposed approach to cases like *Mass Vaccination* in certain respects overlaps with the approach of *EAM*. At the time of deciding between vaccine 1 and vaccine 2, opting for the latter imposes on everyone a 999/1000 chance of unharmed survival and a 1/1000 risk of death. *EAM* utilizes these prospects of harm and benefit in its assignment of individual complaints for and against vaccines 1 and 2. I, on the other hand, utilize these prospects to ground a complaint diminisher against the force of *ex post* complaints against opting for vaccine 2. The reason to prefer my *EPM* approach should be clear: cases like *Mass Vaccination* are only a subset of cases involving trade-offs under uncertainty. This is so because not all trade-off situations involve what I call victim-benefiting actions. As I have argued in Section 3, *EAM* carries highly counterintuitive implications in a range of cases involving non-victim-benefiting actions, and it should be rejected on this basis. In my view, the real insight of cases like *Mass Vaccination* is that, in addition to considerations of consent and liability, the force of a complaint against a risk-imposing action based on the fact that someone would thereby suffer harm can be qualified by the consideration that that very someone would have deemed the risky action a gamble worth taking.

I am happy to accept the suggestion that the account of risk I am here advocating is best characterized not as a qualified version of *EPM* but as a hybrid view. Seen this way, my project is to find a principled basis on which to distinguish between cases like *Mass Vaccination*, where we should employ an *ex ante* line of reasoning, and my series of cases involving villains and diseases, where we should employ an *ex post* line of reasoning. The principled basis I offer, then, is this: *Mass Vaccination*, unlike my series of cases involving

villains and diseases, contains as one of its options a victim-benefiting action. It is this feature which explains why it is appropriate to employ *EAM* in *Mass Vaccination*.

5 A Further Objection

Let me, in closing, address a line of objection against what I have argued so far. Consider the following cases, modelled after a pair of cases from Joe Horton, which seek to undermine *EPM*:[64]

> ***Villain (Single Decision) - VSD***: An evil villain is about to wreak immeasurable havoc unless you choose between the following:
>
> (a) Plan A, in which the villain causes serious but not severe vision impairment in each of a group of 1 million individuals.
>
> (b) Plan B, in which the villain randomly selects one person among this group of 1 million individuals and kills that person.

EPM says that opting for (b) in *VSD* in morally impermissible. Now, suppose that in:

> ***Villain (Multiple Decisions) – VMD***: An evil villain is about to wreak immeasurable havoc unless, for each of a group of 1 million individuals in turn, you choose between the following:
>
> (a) The villain causes serious but not severe vision impairment in that person.

[64] Horton, "Aggregation, Complaints, and Risk," 69.

(b) The villain assigns a unique number between 1 and 1 million to that person. At some later time, a die with 1 million faces, numbered 1 to 1 million, will be rolled. Whichever face the die lands on, any person assigned that number will be killed by the villain.

Here, you face a series of 1 million decisions. For each decision, there is an *ex post* complaint against your opting for (a), based on the harm of suffering serious vision impairment, and an *ex post* complaint against your opting for (b), based on the harm of incurring a one-in-a-million risk of death. Since, we can assume, the latter complaint may be permissibly offset by the former, *EPM* seems to say that, concerning the fate of any one individual in this group of 1 million people, it is morally permissible for you to opt for (b) rather than for (a).

Horton thinks what *EPM* has to say about the above two cases violates the following principle of consistency:

Principle of Consistency (PC): if a moral view condemns act X, it should condemn any series of acts that is in all morally relevant respects equivalent to X.[65]

However, we should distinguish between two readings of PC:

[65] Horton, "Aggregation, Complaints, and Risk," 70.

Principle of Consistency 1 *(PC1)*: if a moral view condemns act X, it should condemn, taken as a whole, any series of acts that is in all morally relevant respects equivalent to X.

Principle of Consistency 2 (PC2): if a moral view condemns act X, it should condemn, taken individually, any act that partially constitutes any series of acts that is in all morally relevant respects equivalent to X.

Let us now run Horton's supposed objection against *EPM*.

Your opting for (b) in *VSD* is in all morally relevant respects equivalent to the following series of acts: regarding the fate of each and every innocent person involved in *VMD*, you choose option (b) over option (a). Since *EPM* condemns your opting for (b) in *VSD*, PC1 demands that *EPM* condemns your series of decisions in *VMD* to opt for (b) over (a) taken as a whole. And this is indeed what *EPM* does. For, by choosing option (b) over option (a) for each and every innocent individual involved in *VMD*, you make it certain that someone will die at the hands of the villain. This harm is not permissibly offset by the benefit of avoiding serious but not severe vision impairments. PC1, then, is no objection to *EPM*.

What of PC2? Let us further distinguish between two readings of PC2:

Principle of Consistency 2.1 *(PC2.1)*: if a moral view condemns act X, it should condemn, taken individually and assessed in isolation of the other acts in the series,

any act that partially constitutes any series of acts that is in all morally relevant respects equivalent to X.

Principle of Consistency 2.2 (PC2.2): if a moral view condemns act X, it should condemn, taken individually and assessed against the background of the other acts in the series, any act that partially constitutes any series of acts that is in all morally relevant respects equivalent to X.

EPM does indeed violate PC2.1. Suppose Yolanda is one of the 1 million individuals whose fates you are to decide in *VMD*. The only relevant considerations that bear on your decision about Yolanda, taken in isolation, are: (1) by opting for (a), Yolanda will be visually impaired, and (2) by opting for (b), Yolanda will incur a one-in-a-million risk of death. *EPM* therefore says that, assessed in isolation, it is permissible for you to opt for (b) over (a) in Yolanda's case. But, since your choosing (b) in Yolanda's case partially constitutes a series of acts that is in all morally relevant respects equivalent to your choosing (b) in *VSD*, and since your choosing (b) in *VSD* is condemned by *EPM*, PC2.1, if true, contradicts *EPM*.

But PC2.1 is false. Consider our imperfect duty of beneficence. Someone who has the opportunities to do what this imperfect duty requires, but who, throughout his life, consistently refuses to comply with this duty has done something that is condemned by our best moral view. However, it is not true that, taken in isolation, any particular occasion where the duty of beneficence is applicable but not met is an occasion where the person in

question has done something that is morally condemnable. PC2.1, then, being a false principle, is no objection to *EPM*.

Let us now assess your decision regarding the fate of Yolanda in light of PC2.2. We are here assuming that, for each of the other 999,999 individuals threatened by the villain, you have or will have opted for (b) over (a). Call this the background condition. Now, we ask: given the background condition, is it morally permissible for you to choose option (b) over option (a) in Yolanda's case? PC2.2 says it is not morally permissible. *EPM*, however, says otherwise.

Once we hold fixed the background condition, your opting for (a) in Yolanda's case generates an *ex post* complaint based on the harm of vision impairment, whereas your opting for (b) generates an *ex post* complaint based on the harm of increasing the risk of someone being killed at the hands of the villain from 999,999/1,000,000 to 1. I take it that the latter harm — a marginal increase in the risk of someone being lethally harmed — may be permissibly offset by the benefit of avoiding serious vision impairment. If so, then *EPM* says that, assessed against the background condition, it is permissible for you to choose option (b) over option (a) in Yolanda's case. This result contradicts the dictates of PC2.2.

PC2.2, however, is false. Consider proportionality constraints in the imposition of defensive harm. Suppose a car has been picked up by a strong gust of wind and is now falling toward your daughter Zola. The car will fatally crush Zola unless you hurl a grenade at the car, the explosion of which will obliterate the falling vehicle. However, there are five innocent individuals inside the car, and the explosion from the grenade will certainly kill all five. If the car lands on Zola, we will suppose, the five individuals will survive the incident unscathed.

I take it that our best moral view condemns your throwing of the grenade on the grounds that doing so, even in the name of defending your own daughter, is disproportional. But now suppose that you take out a gun with four bullets, and you shoot and kill four of the five individuals inside the car. Given the new developments, and given that you must now make a decision either way regarding the fate of your own daughter, I think it is now morally permissible for you to hurl a grenade at the car. Granted, your act of killing the four innocent passengers is morally horrendous, and it is strongly condemned by our best moral view. But even so, once we hold these prior acts as fixed and ask the necessary question — are you or are not you not permitted to save your own daughter, who will otherwise be crushed by a falling vehicle — it is clear to my mind that are you morally permitted to do so. PC2.2, by contrast, implausibly says that you are not permitted to save Zola in this scenario.

PC2.2 fails to distinguish between two separate questions we may ask about some decision you are currently facing: (1) Whether some prior series of decisions which led you to this current decision point are morally permissible? And (2) given you have arrived at the current decision point, and given the past is now fixed, what is morally permissible for you now to do? Clearly distinguishing between these two questions will lead us to see that it is morally permissible for you to save Zola when there is only one passenger left in the falling car. It will also lead us to see the falsity of PC2.2. Pc2.2, then, being a false principle, is no objection to *EPM*.

Conclusion

Recent attempts to develop a satisfactory non-consequentialist approach to risk have been divided over the issue of whether to embrace an *ex ante* framework for assigning weight to individual interests or an *ex post* framework. The correct verdict, I believe, is to adopt a broadly *ex post* model, but one which gives voice to the central insight of the *ex ante* model in some cases. I view my paper as proposing a happy and principled reconciliation between the two frameworks. The key question to ask, in deciding between the two approaches, is whether the individuals involved can themselves reasonably agree that the option at issue represents a gamble that, for their own sake, is worth taking.

Bibliography

Anscombe, Elizabeth. “Who is Wronged?” *Oxford Review* 5 (1967): 16-17.

Brennan, Samantha. “Thresholds for Rights*,” Southern Journal of Philosophy*, 33 (1995): 143-68.

Broome, John. “Fairness,” *Proceedings from the Aristotelian Society*, 91 (1990): 87-101.

Frick, Johann. “Contractualism and Social Risk,” *Philosophy & Public Affairs* 43 (2015): 175-233.

Halstead, John. “The Numbers Always Count,” *Ethics* 126 (2016): 789-802.

Horton, Joe. “Aggregation, Complaints, and Risk,” *Philosophy & Public Affairs* 45 (2017): 54-81.

Horton, Joe. “Always Aggregate,” *Philosophy & Public Affairs* 46 (2018): 160-74.

Kamm, F.M. *Morality, Mortality Volume 1* (New York: Oxford University Press, 1993)

Kamm, F.M. "Owing, Justifying, and Rejecting," *Mind* 111 (2002): 323-54.

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

Kumar, Rahul. “Contractualism on the Shoal of Aggregation,” in *Reason and Recognition: Essays on the Philosophy of T.M. Scanlon*, eds. R. Jay Wallace, Rahul Kumar, and Samuel Freeman (Oxford: Oxford University Press, 2011): 129-54

Kumar, Rahul. “Risking and Wronging,” *Philosophy & Public Affairs* 43 (2015): 27-51.

Lazar, Seth. “Limited Aggregation and Risk,” *Philosophy & Public Affairs* 46 (2018): 117-59.

Nagel, Thomas. “Equality,” in *Mortal Questions* (Cambridge: Cambridge University Press, 1979), 106-27

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

Norcross, Alastair. "Contractualism and Aggregation," *Social Theory & Practice* 28 (2002): 303-14.

Otsuka, Michael. "Scanlon and the Claims of the Many versus the One," *Analysis* 60 (2000): 288-93.

Otsuka, Michael. "Saving Lives, Moral Theory, and the Claims of Individuals," *Philosophy & Public Affairs* 34 (2006): 109-35.

Otsuka, Michael. "Risking Life and Limb," in *Identified versus Statistical Lives*, eds. I. Glenn Cohen, Norman Daniels, and Nir Eyal (Oxford: Oxford University Press, 2015): 77-93.

Parfit, Derek. "Equality and Priority," *Ratio* 10 (1997): 202-21.

Parfit, Derek. *On What Matters Volume 2* (Oxford: Oxford University Press, 2011), at 191-212.

Pettit, Philip. "Consequentialism," in *Blackwell's Companion to Ethics*, eds. Peter Singer (Oxford: Blackwell, 1992): 230-40.

Pollock, John. "Defeasible Reasoning," *Cognitive Science* 11 (1987): 481-518.

Raz, Joseph. *Practical Reasons and Norms* (Oxford: Oxford University Press, 1999)

Raz, Joseph. "Numbers, With and Without Contractualism," *Ratio* 16 (2003): 346–67.

Reibetanz, Sophia. "Contractualism and Aggregation," *Ethics* 108 (1998): 303-304.

Scanlon, T.M. "Contractualism and Utilitarianism," in *Utilitarianism and Beyond*, ed. Amartya Sen and Bernard Williams (Cambridge: Cambridge University Press 1982)

Scanlon, T.M. *What We Owe to Each Other* (Cambridge: Harvard University Press, 1998)

Scanlon, T.M. "Replies," *Social Theory & Practice* 28 (2002): 337-58.

Scanlon, T.M. *Being Realistic About Reasons* (Oxford: Oxford University Press, 2014)

Tadros, Victor. "Localized Restricted Aggregation," in *Oxford Studies in Political Philosophy Volume 5*, eds. David Sobel, Peter Vallentyne, and Steven Hall (Oxford: Oxford University Press, 2019): 171-204

Taurek, John. "Should the Numbers Count?" *Philosophy & Public Affairs* 6 (1977): 293-316.

Tomlin, Patrick. "On Limited Aggregation," *Philosophy & Public Affairs* 45 (2017): 223-60.

Van Gils, Aart and Tomlin, Patrick. "Relevance Rides Again? Aggregation and Local Relevance" in *Oxford Studies in Political Philosophy Volume 6*, eds. David So bel, Peter Vallentyne, and Steven Hall (Oxford: Oxford University Press, 2020): 221-55.

Voorhoeve, Alex. "How Should We Aggregate Relevant Claims?" *Ethics* 125 (2014): 64-87.

Voorhoeve, Alex. "Why One Should Count Only Claims with Which One Can Sympathize," *Public Health Ethics* 10 (2017): 148-56.

Wasserman, David, and Alan Strudler, "Can a Nonconsequentialist Count Lives?", *Philosophy & Public Affairs* 31 (2003): 71-94.