

Individualist Theories and Interpersonal Aggregation*

Erik Zhang

This article offers a solution to the numbers problem within an individualist moral framework. Its central aims are as follows: to rescue individualist moral theories, such as moral contractualism, from their long-standing problem with interpersonal aggregation; to demonstrate how, proceeding from an individualist mode of justification, we can nevertheless make the numbers count without directly counting the numbers; to provide an individualist rationale for accepting a partially aggregative criterion of adjudication for resolving interpersonal trade-offs; and finally, to develop an extensionally adequate version of partial aggregation that finds application in more structurally complex cases.

I. INTRODUCTION

Suppose that we can, at 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. How this commonsense verdict is to be given a sound theoretical justification, however, has proven to be a surprisingly difficult issue, raising 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

* This article combines two standalone ancestors, both of which have benefited from the generosity of many people. For their very helpful feedback on earlier drafts, I would like to thank Selim Berker, Tom Dougherty, Nir Eyal, Johann Frick, Liz Harman, Joe Horton, Tom Hurka, Tom Kelly, Rahul Kumar, Sarah McGrath, Liam Murphy, Doug Portmore, Theron Pummer, Gideon Rosen, Tim Scanlon, Peter Singer, Michael Smith, Patrick Tomlin, Daniel Viehoff, Alec Walen, and Jay Wallace. I would also like to thank two anonymous reviewers, as well as Victor Tadros and four other associate editors of *Ethics*, for their insightful and extensive comments.

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 contexts of interpersonal trade-off. For example, in a choice between saving a person's life and preventing a large number of minor headaches (*Life vs. Headaches*), 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 regarded by nonconsequentialists with suspicion. Most memorably, John Rawls remarked that moral theories according to which the losses of some are straightforwardly offset by the greater aggregate gains of others fail to take seriously the separateness of persons.¹ In an effort to provide a clear and principled alternative to consequentialist moral reasoning, leading nonconsequentialists have put forward the suggestion that the fundamental mode of moral justification is *individualist* rather than aggregative in character.² On this suggestion, the moral permissibility of an action depends only on its implications for *single persons*, not on its implications for collections of people.

The trouble, however, is that an individualist moral approach appears to go too far in the opposite direction. No doubt the approach handles well cases like *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. Previous attempts by theorists to resolve this conundrum, including the notable *tiebreaking argument* proposed by T. M. Scanlon,³ have not persuaded critics, and the numbers problem has remained one of the most vexing challenges to the development of nonconsequentialist moral theories to this day.⁴

In this article, I aim to develop a new solution to the numbers problem within an individualist moral framework. The solution I offer unfolds in several stages. I begin, in Section II, by laying the groundwork for an

1. John Rawls, *A Theory of Justice* (Cambridge, MA: Harvard University Press, 1999), 24.

2. See, e.g., Thomas Nagel, *Equality and Partiality* (Oxford: Oxford University Press, 1991); and T. M. Scanlon, *What We Owe to Each Other* (Cambridge, MA: Harvard University Press, 1998).

3. Scanlon, *What We Owe*, chap. 5. Even Scanlon himself later remarks, "When 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 and Practice* 28 (2002): 337–58, 354.

4. See, e.g., 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, 348; David Wasserman and Alan Strudler, "Can a Nonconsequentialist Count Lives?," *Philosophy and Public Affairs* 31 (2003): 71–94; Joseph Raz, "Numbers, with and without Contractualism," *Ratio* 16 (2003): 346–67; Michael Otsuka, "Saving Lives, Moral Theory, and the Claims of Individuals," *Philosophy and Public Affairs* 32 (2006): 109–35; and Derek Parfit, *On What Matters* (Oxford: Oxford University Press, 2011), 2:191–212.

individualist justification for interpersonal aggregation. Among other things, I clarify the central commitments of an individualist moral approach, and I examine Scanlon's attempted solution to the numbers problem. The failures of Scanlon's proposal prove instructive for the presentation of my own account in Sections III and IV. In these sections, I argue that the proper application of individualist moral reasoning leads to the endorsement of a criterion of interpersonal adjudication that is partially aggregative in character, one that instructs us to take numbers into account in some cases of trade-off but not in others, thereby capturing our divergent intuitions in Life vs. Headaches and Life vs. Lives. Finally, in Section V, I briefly discuss some of the recent controversies surrounding the extensional adequacy of partially aggregative views, on which I believe the individualist framework developed here can shed helpful light.⁵ Along the way, I mark and make use of several distinctions whose significance for nonconsequentialist theorizing has not, in my view, been fully appreciated.

If successful, my discussion taken as a whole secures several important results. First, it rescues individualist moral theories, such as moral contractualism, from their long-standing problem with interpersonal aggregation. Second, it demonstrates how, proceeding from an individualist mode of justification, we can nevertheless make the numbers count without directly counting the numbers. Third and finally, it articulates an individualist rationale for accepting a partially aggregative criterion of adjudication for resolving interpersonal trade-offs.

II. CLEARING THE GROUND

Philosophical support for an individualist approach to morality grows out of a long-standing discontent with the aggregative character of utilitarianism and other forms of consequentialism. This discontent has received theoretical articulations from many of the leading nonconsequentialists of our time. It can be found, for instance, in Rawls's aforementioned critique that utilitarian maximization is wholly insensitive to how burdens and benefits are distributed across distinct and separate persons. It can also be found in Thomas Nagel's interpretation that equal and impartial concern for everyone should move us to imaginatively identify with each

5. See, e.g., Alex Voorhoeve, "How Should We Aggregate Relevant Claims?," *Ethics* 125 (2014): 64–87; Patrick Tomlin, "On Limited Aggregation," *Philosophy and Public Affairs* 45 (2017): 232–60; Joe Horton, "Always Aggregate," *Philosophy and Public Affairs* 46 (2018): 160–74; Victor Tadros, "Localized Restricted Aggregation," in *Oxford Studies in Political Philosophy*, ed. David Sobel, Peter Vallentyne, and Steven Wall (Oxford: Oxford University Press, 2019), 5:171–204; and Aart van Gils and Patrick Tomlin, "Relevance Rides Again? Aggregation and Local Relevance," in *Oxford Studies in Political Philosophy*, ed. David Sobel, Peter Vallentyne, and Steven Wall (Oxford: Oxford University Press, 2020), 6:221–55.

individual's point of view separately, rather than combining them into an undifferentiated whole.⁶

More recently, Scanlon has defended a contractualist alternative to consequentialism, which carries with it a signature commitment to the so-called *individualist restriction*.⁷ The restriction limits the basis of deontic assessment to only the *personal reasons* of single individuals, that is, reasons having to do with a person's *own* well-being, interest, or status.⁸ As Scanlon sees it, the individualist restriction flows naturally from contractualism's most central and most attractive ideal, that our interpersonal conduct must be justifiable to *each* of those affected.

Contractualism is likely the best-known development of moral individualism in contemporary ethics. But a defense of individualism need not rest on the acceptance of the general contractualist apparatus, many of whose features prove controversial even among nonconsequentialists. Indeed, Scanlon's own treatment of interpersonal aggregation borrows heavily from F. M. Kamm's noncontractualist, commonsense deontology.⁹ Nevertheless, contractualism's formulation of the individualist restriction helpfully makes explicit that an individualist moral approach bars interpersonal aggregation on two distinct fronts. First, by limiting the basis of deontic assessment to people's personal reasons, the approach bars aggregative considerations from directly figuring in the *content* of people's moral claims. Second, by insisting that only the implications for single individuals matter for determining permissibility, the approach bars the aggregation of claims *across* different individuals.

Thus understood, an individualist approach carries clear advantages as a general framework within which to develop an account of interpersonal trade-offs. In the case of Life vs. Headaches, for example, the approach provides a principled basis for blocking the problematic transition from the *axiological* claim that preventing the many headaches maximizes overall welfare to the *deontic* claim that we should therefore save the greater number. The challenge, of course, is for the approach to explain why the numbers should still count in cases like Life vs. Lives.

Here, two preliminary difficulties need to be addressed. To begin, it is not obvious that the present challenge is one that can be met even in principle. After all, an individualist approach endorses a mode of justification that is wholly antiaggregative across persons, holding that facts

6. Thomas Nagel, "Equality," in *Mortal Questions* (Cambridge: Cambridge University Press, 1979), 106–27; and Nagel, *Equality and Partiality*, chap. 7.

7. As Scanlon himself says, the individualist restriction is "central to the guiding idea of contractualism, and is also what enables it to provide a clear alternative to utilitarianism and other forms of consequentialism" (Scanlon, *What We Owe*, 229).

8. *Ibid.*, 219–20.

9. F. M. Kamm, *Morality, Mortality* (New York: Oxford University Press, 1993), vol. 1, chap. 6; and F. M. Kamm, *Intricate Ethics* (Oxford: Oxford University Press, 2007), chap. 1.

about numbers do not themselves carry any justificatory significance. It then seems to follow that any criterion of adjudication that instructs us to take numbers into account in deciding whom to save straightforwardly runs afoul of the approach's central commitment to antiaggregation.

While this line of reasoning appears persuasive at first glance, it is guilty of what we may call a *level confusion*. An individualist approach is indeed antiaggregative at the level of justification, in the sense that only the personal reasons of single individuals matter for determining permissibility. What the above reasoning overlooks, however, is that individuals might have strong personal reasons in favor of accepting a criterion of adjudication which, as part of its content, instructs us to sometimes take into account the relevance of numbers. We should therefore clearly distinguish between two levels at which aggregative considerations might enter. They might enter at the *level of justification*, as reasons for accepting or rejecting some criterion for adjudicating between competing interests, or they might enter at the *level of adjudication*, as considerations a criterion of adjudication itself tells us to take into account in deciding whom to save. An individualist approach bars aggregation at the level of justification but not at the level of adjudication. Of course, whether or not individuals have strong personal reasons in favor of 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 reasons is not ruled out by moral individualism from the outset.

Next, consider a specific version of Life vs. Lives, wherein we can rescue either A from death or each of B and C from the same fate. It is tempting for us to reason about the situation in the following way. Each of B and C will die if left unaided, but A's well-being is equally at stake. Under an individualist framework, the claims of B and C may not be combined. Nor may we simply appeal to the impersonal worseness of leaving B and C to die as grounds for rescuing the larger group. It then seems that the individual claims on both sides are equally balanced, leaving us without a conclusive basis for saving the greater number.

Precisely this line of reasoning underlies a major part of John Taurek's famous stance that the numbers should not count.¹⁰ Taurek correctly observes that, analyzing the situation from each person's separate concern for her own well-being, no one person has a greater personal claim to be saved than any other person. Taurek concludes that he would resolve the conflict by flipping a coin, thereby giving each person an equal chance of being saved. Taurek's conclusion, however, is in one crucial respect premature. For we should not forget that individual well-being does

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

not exhaust the grounds for determining moral permissibility. 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 against which we may raise legitimate complaints.

Indeed, Scanlon's own proposed solution to the numbers problem appeals to considerations that go beyond a narrow concern for individual well-being. Regarding the case of Life vs. Lives, Scanlon has the following to say:

Either member of the larger group might complain that [a principle which permits us to be insensitive to the 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 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.¹¹

Scanlon's proposal, in effect, 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.

Scanlon's tiebreaking argument, however, confronts several serious worries. For one thing, if our reason for saving the larger group lies in a concern that each individual's presence should make a difference to whom we are to save, then it seems compatible with satisfying this concern that, rather than saving B and C over A outright, we should instead employ the device of a weighted lottery, whereby the presence of each individual makes a difference by affecting people's chance of being saved.¹² For another, it is unclear how Scanlon's proposal can account for the familiar intuition that, as the disparity in numbers increases between the two competing groups, each person in the larger group has an increasingly stronger claim to being saved. For whether we are confronted with a choice between saving one life and two others or between saving one life and one billion others, Scanlon's tiebreaking argument generates claims of individually equal strength against rescuing the smaller group.

11. Scanlon, *What We Owe*, 232.

12. See, e.g., Wasserman and Strudler, "Can a Nonconsequentialist Count Lives?," 82–88; and Otsuka, "Saving Lives," 111–18.

Worse still, Scanlon's proposal runs into trouble when we try to generalize it to certain cases where the harms faced by individuals are unequal in magnitude. Consider a situation in which we can save either a person's life or each of a large number of people from a lesser but still substantial harm, say, permanent paraplegia (*Life vs. Paraplegias*). Intuitively, we should forgo the lifesaving option if sufficiently many paraplegias can be thereby avoided. But in order for the tiebreaking argument to gain a foothold in this case, we must first grant that several people's interests in avoiding paraplegia may *together* create a tie with a competing interest in avoiding death. At this point, we are right to wonder why this way of applying the tiebreaking argument does not similarly extend to the case of *Life vs. Headaches*. In other words, what is to prevent a tie from being created between several people's 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 their own favor?

In response to this worry, Scanlon puts forward the following suggestive proposal:

It 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.¹³

And if we incorporate the distinction between relevant and irrelevant harms into the tiebreaking argument, Scanlon's proposal becomes that only competing harms that are relevant may aggregate to create ties and only harms that are relevant may perform the function of tiebreaking.

I have several misgivings about Scanlon's proposal even in its improved form. First, Scanlon in the above passage simply invokes a relation of relevance between competing harms. But in order for such an invocation to be convincing, we should expect an adequate explication of the notion of relevance at issue. What deeper moral truths, for example, are our intuitions about relevance tracking, and what normative mechanism may we appeal to in explaining how relevance functions? These questions Scanlon does not answer.

13. Scanlon, *What We Owe*, 239–40.

Moreover, it is plainly inadequate for an individualist to simply assert that we should take the number of relevant harms into account in deciding whom to save; 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 that Scanlon fails to discharge. But until such an explanation is given, the distinction between relevant and irrelevant harms remains unavailable to defenders of individualism as a fix to their problem with numbers.

Third and finally, if we do succeed in developing an individualist justification for allowing only relevant claims to aggregate, then the tiebreaking argument simply becomes *redundant*. We may then directly appeal to this justification in vindicating interpersonal aggregation in cases like Life vs. Lives and Life vs. Paraplegias, without resorting, in a roundabout way, to the reasoning of the tiebreaking argument. My goal in the next two sections is to unearth precisely such a justification.

To preview, the argument in what follows develops in two major steps. First, in Section III, I articulate an individualist rationale for rejecting interpersonal aggregation in cases like Life vs. Headaches. The main task is to develop an adequate account of the relation of relevance between competing interests. Next, in Section IV, I put forward what I call the *argument from equal consideration* in favor of saving the greater number in Life vs. Lives and Life vs. Paraplegias. Concerning these cases, my argument seeks to establish that individuals have strong personal reasons, based on a specific interpretation of the ideal of equal consideration, for wanting the larger group to be rescued.

III. NORMATIVE DISABLING

How should a full analysis of Life vs. Headaches proceed under an individualist approach? On the simplest suggestion, we just weigh the competing interests at issue. Given that a person's interest in avoiding death is obviously weightier than any one person's competing interest in avoiding a minor headache, an individualist is straightforwardly led to the conclusion that we ought to choose the lifesaving option, regardless of the number of headaches at stake. But this simple way of analyzing the case lands individualism in a serious problem; after all, a person's interest in avoiding paraplegia, too, is less weighty than a competing interest in avoiding death.

A more promising strategy will begin where Scanlon left off, by appealing to the intuitive distinction between relevant and irrelevant interests. A rough-and-ready suggestion, I take it, goes as follows. In Life vs. Headaches, a person's interest in avoiding a headache is simply too trivial to count against another person's continued survival. Such an interest, therefore, is irrelevant; it ceases to count in such a life-or-death situation.

By contrast, while an interest in avoiding paraplegia is admittedly weaker than a competing interest in avoiding death, it is nevertheless weighty enough to be a relevant interest, one that should still be taken into account in our moral adjudication.

The suggestion seems to me to be on the right track. As it stands, however, it is simply too rough-and-ready. In order to be convincing, the suggestion needs to be developed in at least three major respects. First, we need an account of what it is for an interest to cease to count in some context of interpersonal trade-off. Second, we need to explain how it is that an interest may count in some contexts but not in others. And finally, we need to say more about why it is that an interest can lose its standing to count in interpersonal adjudications.

To these ends, consider the ways in which our practical reasons can conflict. Most familiarly, one of two competing reason-giving considerations may carry less normative force than the other, and the former may in this sense be outweighed by the latter. But the relation of outweighing, as we have already seen, is a poor candidate for underwriting the relation of relevance.

But our practical reasons can conflict in deeper ways. In certain contexts, a consideration which is typically reason-giving can lose its normative force altogether. By way of introducing the phenomenon, consider the following case:

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 funds. After exhausting 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 will allow me to secure my inheritance early.

Do I have a reason to bring about the death of my uncle under the prescribed circumstances? Most of us would answer no, and I think we should take this answer 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 funds 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.¹⁴

14. This claim can be challenged. For example, Mark Schroeder argues that we have good reasons for not taking our negative existential intuitions about reasons at face value. This is because, according to Schroeder, our intuitions about the existence of reasons are colored by a standard presumption that R is a reason to do A only if R is a relatively weighty

In saying this, I do not mean to suggest that one's interest in furthering one's 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 this happens, we may say that the consideration in question is normatively disabled,¹⁵ which differs from its being simply outweighed.¹⁶

It may help our understanding of the phenomenon of normative disabling in the practical domain by considering a seeming analogue in the epistemic domain, in the form of undercutting defeaters.¹⁷ Suppose that as I wander through a gallery I chance upon a room that appears to me to have blue walls. Typically, my perceptual experience 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 member that the lighting in the room makes the walls appear blue regardless of their actual color. In light of this

reason to do A. See Mark Schroeder, *Slaves of the Passions* (Oxford: Oxford University Press), 92–97. Thus, in Rare Stamp, Schroeder might suggest that our intuitive response to the case is misleading; there indeed is a reason for me to kill my uncle, but the reason is simply too weak to be intuitively recognized. In response, we may note two things. First, as Schroeder himself notes, the kind of presumption at issue is removable. We may remove it by more explicitly asking the question, in Rare Stamp, does the improvement in my stamp collection provide any reason at all for me to end my uncle's life? Speaking for myself, I have the clear intuition that the answer is still no. Second, the suggestion that I have a reason, albeit a weak one, to kill my uncle in Rare Stamp confronts difficulties. Imagine that, in a variant of Rare Stamp, my uncle and another family member suddenly fall ill, and I only have enough lifesaving medicine for one of them. There would be something highly objectionable if I invoke the improvement of my stamp collection as a reason in my deliberation for leaving my uncle to die. But why should this be so if the improvement does indeed provide me a genuine reason to bring about my uncle's death? The answer cannot simply be that the reason is too weak, as other weak reasons, such as a slight difference in their likelihood of survival, can appropriately be taken into account in my deliberation. On this point, see Garrett Cullity, "The Context-Undermining of Practical Reasons," *Ethics* 124 (2013): 8–34, 28–30.

15. The idea of normative disabling has been discussed by other philosophers. See, e.g., Jonathan Dancy, *Ethics without Principles* (Oxford: Oxford University Press, 2004), chap. 3; Cullity, "Context-Undermining of Practical Reasons"; and Ralf Bader, "Conditions, Modifiers, and Holism," in *Weighing Reasons*, ed. Errol Lord and Barry Maguire (Oxford: Oxford University Press, 2016), 27–55.

16. Normative disabling also differs from reasons being excluded. An exclusionary reason, according to Joseph Raz, is a second-order reason to refrain from acting on some first-order reason. Joseph Raz, *Practical Reasons and Norms* (Oxford: Oxford University Press, 1999), chap. 1. By contrast, when a consideration is normatively disabled, it is no longer a reason at all under the circumstances. To see how the two phenomena differ, consider my promise to help you move tomorrow. According to Raz, my promise to you gives me an exclusionary reason not to act on certain contrary considerations, such as the consideration that fulfilling the promise is no longer in my own self-interest. But this need not imply that my self-interest is not genuinely reason-providing under the circumstances.

17. See, e.g., John Pollock and Joseph Cruz, *Contemporary Theories of Knowledge*, 2nd ed. (Savage, MD: Rowman & Littlefield, 1999), 36–37.

change in my evidential circumstances, the typical reason-giving force of my perceptual experience is undercut; my experience is no longer a reason for me to believe that the walls are painted blue.

The general point to which I am here calling attention is that the reason-giving force of a consideration is not invariant across contexts. This is true of both reasons for belief and reasons for action. Indeed, this contextual character of a consideration's reason-giving force is made intelligible by reflecting on the reason relation itself. The reason relation, I take it, holds between a fact p , an agent x , an attitude or action a , and a set of circumstances c . The relation obtains just in case p is a reason for x to hold or perform a in c . The specification that whether p constitutes a reason is relative to a set of circumstances helpfully illuminates how it is that a typically reason-giving consideration can lose its normative force. In *Life vs. Headaches*, for example, while I do have a reason to alleviate another's discomfort under normal circumstances, that reason relation fails to obtain in circumstances in which the alleviation comes at the expense of another's life.

The contention thus far is that what it is for an interest to cease to count in some context is for the reason-giving force of that interest to be disabled. The phenomenon of disabling, moreover, takes place at the level of reasons, and the mechanism through which it operates is sourced in the contextual character of the reason relation itself. Now, what deeper explanation may we offer for why it is that disabling occurs in the practical domain? As a first step, here is a proposal that I find attractive: when a consideration which is typically reason-providing is disabled, there is the presence of some value which makes it inappropriate for the consideration to weigh against the value in question.¹⁸

To illustrate, consider, once again, the case of *Rare Stamp*. There, the improvement of my stamp collection runs up against the preservation of my uncle's life. On the present proposal, disabling occurs because the value of human life sets certain limits on which considerations may appropriately weigh against a life's preservation; it is inconsistent with the value of human life for the improvement in the quality of a stamp collection to count in favor of bringing about a person's death.¹⁹ No such inappropriateness exists

18. Here I draw on Victor Tadros's rejection of unrestricted aggregation, to whose discussion I am much indebted. See Tadros, "Localized Restricted Aggregation," 5:172–74.

19. For my purposes, I need not take a committed stance on the priority relations between reasons, values, and appropriateness. Having sympathy for the "reasons-first" approach, however, let me briefly note how my claims in this section could be made compatible with taking reasons to be normatively fundamental. Roughly, we may begin with the suggestion that facts about values are grounded in or explained by facts about reasons. Most familiarly, many suggest that what it is for something to be valuable is for there to be reasons (of the right kind) for us to respond to it in certain ways. We may then add that, in addition to these kinds of response-reasons, a thing's value can also be grounded in another type of

when the option under consideration is instead taking out a short-term loan or calling up a friend for assistance, since these options do not similarly put anyone's life in serious jeopardy.

Of course, we should next ask, why accept this proposal? In particular, why accept that values set the kind of appropriateness conditions that I am here proposing? The answer, as I see it, lies in the intimate connections between the phenomenon of normative disabling, the nature of value, and the conditions of respect.

Consider the following pair of cases. In the first, a parent who is struggling to provide for her family is driven to sell a valuable painting, perhaps an authentic Rembrandt, to a buyer whose desire for the painting's destruction is well-known. Let us suppose that, all things considered, the value of the painting and the reasons against its destruction ultimately outweigh the financial and material gains at issue. Under these stipulations, the decision to sell the painting indicates a familiar kind of mistake about value: it indicates an inaccurate understanding of just how valuable the painting is.

Now, imagine a second case. A person decides to fold up an authentic Rembrandt and use it to stabilize her wobbly dinner table. This person, too, is under a misapprehension of value. Yet it seems natural to say that she is guilty of a more fundamental failing: not only is she mistaken about the exact magnitude of the painting's value, a mistake she shares with the protagonist of our first case, but her decision to use the painting as a table stabilizer also indicates that she does not understand the kind of value that the painting has at all.

Implicit in these commonsensical observations is a distinction we draw between two kinds of mistake about value. In our everyday lives, trade-offs involving values are commonplace, and in deciding these trade-offs, we sometimes go wrong. One kind of mistake about value, we may say, is at bottom a matter of *inaccuracy*. In our first case, a much-needed improvement in the material conditions of one's struggling family is indeed a consideration that may appropriately weigh against the painting's preservation. It is just that, all things considered, an accurate appreciation of the painting's value renders the sale a mistake. But there is another kind of mistake about value, which we may say is fundamentally a matter of *impropriety*. In our second case, a proper understanding of the value of an authentic Rembrandt would simply lead

facts about reasons, namely, facts about which considerations constitute genuine reasons against its preservation. Thus, on this proposal, the value of human life is partly grounded in the fact that, in cases of conflict, a person's interest in avoiding a minor headache is not a genuine reason that counts in favor of the destruction of a life. And it would in this sense be inappropriate for the alleviation of a headache to weigh against the preservation of a person's life.

us to see that the painting's potential to stabilize a wobbly table provides no reason, much less sufficient reason, for its destruction.

This distinction between inaccuracy and impropriety finds further expression in our understanding of the conditions of respect. To be sure, we may elect to understand disrespect quite broadly, as what is shown *whenever* we fail to give accurate weight to the valuable features of the world in our deliberation and action.²⁰ This broad understanding of disrespect, however, carries the uneasy implication that every incorrect decision involving values, even in the most difficult and hard cases, must always be a matter of disrespect. It seems to me, however, that we reserve space in our moral thinking for a clear and distinct form of disrespect that is much narrower in scope. This form of disrespect manifests most evidently in cases in which our failing is one of impropriety and not mere inaccuracy. Hence, in our first case, where the decision to sell is made under admittedly difficult circumstances, it seems strained if not entirely overblown to say of our struggling parent that her decision not only is mistaken but also shows a lack of respect. By contrast, in our second case, we find no hesitation in saying that the decision to use the painting to steady a wobbly table clearly disrespects the value of the painting.

What these reflections point to, I submit, is precisely the idea that certain values set limiting conditions on the kind of considerations that may appropriately weigh against them. What I am calling mistakes of impropriety occur when all that can be said in favor of the option chosen are considerations whose normative force the operative value itself disables. It should, then, come as no surprise that mistakes of impropriety indicate a more basic failing with respect to our appreciation of the valuable aspects of the world: they involve a failure to appreciate that the supposed countervailing considerations, being themselves disabled, do not have the appropriate standing to weigh against the value in question to begin with. And the idea of normative disabling also sheds important light on our understanding of the conditions of respect. When a mistake of impropriety is made, it is not just that the option chosen is supported by considerations with insufficient reason-giving force; it is moreover that the operative value robs these considerations of their reason-giving force altogether. Plausibly, to respect some value, at the very minimum, calls for us to not sacrifice the value for no good reason at all.

If my preceding discussion is broadly on the right track, then a proper individualist analysis of Life vs. Headaches would proceed as follows. It is true that preventing the many headaches generates more well-being in the aggregate. But under an individualist framework, aggregative considerations do not themselves carry any justificatory significance; only personal reasons matter. We should therefore save the life of the single person,

20. See Stephen Darwall, "Two Kinds of Respect," *Ethics* 88 (1977): 36–49.

but not for the simple reason that an interest in avoiding death is the weightiest. Rather, the reason is that, in a life-or-death situation, a person's competing interest in avoiding a minor headache is one whose normative force the value of human life itself disables. The interest therefore provides us no reason for its satisfaction when it directly conflicts with another's interest in avoiding death. The headache-related interests are in this sense irrelevant.²¹ Disabling occurs, moreover, because the value of human life renders it inappropriate for the alleviation of a minor headache to count against another person's survival. A decision to prevent the many headaches, then, would be a mistake of impropriety; it indicates that the decision-maker does not understand the kind of value that human life has at all. And if we nevertheless elect to forgo the lifesaving option, then the person facing death may rightly object that such a decision fails to properly respect the value that her life has.

We may now distinguish Life vs. Headaches and Life vs. Paraplegias in the following way. In the former case, each person's interest in avoiding a minor headache, being normatively disabled, provides us with no reason for its satisfaction. So, we are to ignore headache-related interests in deciding whom to save. By contrast, though a person's interest in avoiding paraplegia is indeed weaker than another person's competing interest in avoiding death, the paraplegia-related interests still retain their reason-giving force. As a result, each person facing paraplegia has a valid claim that, in one way or another, her interest must be taken into account in our moral decision-making; the possibility that it is unjustifiable to her for us to choose the lifesaving option therefore remains open.

I want to conclude the present section by noting one particularly attractive feature of an individualist analysis of Life vs. Headaches. Too often, in thinking about the case, many critics of indiscriminate aggregation

21. Alex Voorhoeve has recently advanced an alternative account of relevance, according to which A's interest is relevant to B's stronger, competing interest just in case, if A herself were called upon to decide between the competing interests, A is morally permitted to favor her own. Voorhoeve, "How Should We Aggregate," 72. The main problem with Voorhoeve's account, in my view, is that facts about relevance and facts about permissible partiality toward oneself are sensitive to distinct considerations, so the latter cannot plausibly ground the former. To see this, imagine that A has made a solemn promise to B that, if B's life is ever threatened, she will save B's life even at the cost of losing both of her arms. Suppose now that we must decide between saving either A's arms or B's life. In virtue of A's earlier promise, it may well be impermissible for A to favor her own interests. But to us, the loss of A's arms is surely not an irrelevant consideration. Indeed, we may dramatize our case even further. Imagine that a large number of people have each made a similar promise to B. Surely, it is implausible to maintain that we ought to save B's life, regardless of the number of people whose limbs are at stake. My account of relevance is not subject to the same difficulty. A's promise to B may well affect the moral permissibility of A's conduct toward B, but it affects the value of neither A's limbs nor B's life to a morally motivated third party.

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 similarly irrelevant. But when this transition is made in the absence of an explicit commitment to an individualist framework, it contains a highly questionable assumption, namely, that if some harm considered on its own is trivial, then an aggregation of such harms must also be trivial.²²

On my present analysis, by contrast, the aggregate well-being generated by the alleviation of the many headaches is irrelevant not because it is trivial but because an individualist approach limits the basis of moral assessment to only the personal reasons of single individuals. For this reason, an individualist can happily accept the axiological claim that a non-trivial amount of aggregate well-being is produced if we prevent the many headaches, while still holding on to the deontic claim that we ought to nevertheless choose the lifesaving option.

IV. EQUAL CONSIDERATION

The central task of this section is to articulate an individualist justification for saving the greater number in cases in which we clearly should. I begin, however, by addressing two questions of larger theoretical significance. First, what exactly is the theoretical ambition of an individualist moral approach in relation to the numbers problem?

Some theorists who are otherwise sympathetic to an individualist approach have counseled that the approach should scale back its theoretical ambition and settle for a kind of moral pluralism. On one suggestion, it is observed, quite correctly, that an individualist approach cannot hope to capture morality in the broadest sense, since some of our reasons for action evidently fall outside of its intended scope. For example, we have reasons not to turn the Grand Canyon into a site for trash disposal, and some of these reasons obtain simply because the Grand Canyon is impersonally valuable, and not just because its destruction would negatively affect the interests of persons. The suggestion then is that defenders of moral individualism should make better use of the kind of pluralism to which they are antecedently committed anyhow; while saving the lesser number in Life vs. Lives would not be wrong under an individualist framework, it would still be wrong all things considered, once impersonal, aggregative considerations are brought into the picture.²³

22. See, e.g., John Halstead, "The Numbers Always Count," *Ethics* 126 (2016): 789–802, 794–95.

23. See, e.g., Parfit, *On What Matters*, 2:216–17; and Johann Frick, "Contractualism and Social Risk," *Philosophy and Public Affairs* 43 (2015): 175–233, 219–33.

This suggestion does have some appeal, but I believe it is not one that individualists should accept. Although moral individualism does not extend to the whole of morality, it should nevertheless be its proper ambition to provide a complete and satisfactory account of interpersonal morality, part of whose content is precisely our duty to aid other people. How our duty to aid operates in cases of interpersonal trade-off should therefore be of central concern to proponents of moral individualism. In short, unless individualists are prepared to make the hugely concessive move that their moral framework is incomplete even with respect to the narrower domain of interpersonal morality, they should not simply rely on consequentialist reasoning to bypass their problem with numbers.²⁴

In saying all this, I might give the impression that an individualist treatment of interpersonal trade-offs should in no way appeal to aggregative considerations. This brings me to the second question I want to raise, namely, what role, if any, may considerations of aggregate well-being play within an individualist moral framework?

One thing is clear: an individualist approach bars considerations of aggregate well-being from directly figuring in the content of people's moral claims. In *Life vs. Lives*, no one may validly claim that she should be saved because doing so saves more lives overall. After all, this is not a claim that she may raise *on her own behalf*.

To be sure, some philosophers deny the very idea of aggregate value.²⁵ They deny that there is any more disvalue, say, in the deaths of a million people than the death of a single person. This view we may call *axiological antiaggregation*. It is crucial, however, to appreciate that axiological antiaggregation and the kind of moral individualism I have been discussing are distinct and orthogonal. In *Life vs. Lives*, for example, a standard consequentialist analysis would proceed in two steps. First, there is the

24. I should note that Voorhoeve's account of relevance, which I examined in n. 21 above, is developed within a more sophisticated kind of pluralism, according to which interpersonal morality itself has two fundamentally competing modes of justification, one individualist and one aggregative, and the correct account of interpersonal trade-offs emerges as a compromise between the two. Voorhoeve, "How Should We Aggregate," 68–70. Voorhoeve's pluralist framework, however, raises three notable worries. First, by introducing a fundamental schism into the very heart of interpersonal morality, the framework compares unfavorably to a thoroughgoing individualism in terms of theoretical unity. Second, since neither the individualist nor the aggregative mode of justification occupies a privileged place within Voorhoeve's framework, it is far from clear why, in analyzing *Life vs. Headaches*, we should focus on the fact that each headache compares so trivially to a life rather than the fact that preventing the many headaches generates much more well-being in the aggregate. Finally, as I will go on to discuss, because Voorhoeve's framework gives a straightforwardly aggregative analysis of *Life vs. Lives*, it cannot explain the appealing thought that each person in the larger group is in a position to complain, on her own behalf, that failing to save her treats her in an unjustifiable manner.

25. See, e.g., Taurek, "Should the Numbers Count?"

claim that saving the greater number generates more well-being overall. Second, there is the claim that we ought to save the greater number precisely because more well-being will be generated as a result. It should be clear that axiological antiaggregation rejects the first claim while moral individualism rejects the second. What an individualist denies, then, is that facts about aggregate well-being by themselves constitute grounds for determining moral permissibility. The distinctively nonconsequentialist element here lies in the denial that an action can be wrong just because it fails to maximize aggregate value. To put the point differently, what an individualist denies is not aggregative axiology but the direct justificatory relevance of aggregative axiology to the morality of interpersonal right and wrong.

But this leaves open the suggestion that aggregate value can play an *indirect* role in an individualist 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 then derive personal reasons in favor of saving the greater number.

At this juncture, it may be instructive for us to briefly consider another context in which impersonal values can play an indirect role in the generation of personal reasons. Imagine a situation in which it is proposed that the Grand Canyon will be turned into a site for trash disposal. While the impersonal value of the Grand Canyon provides a strong basis for rejecting this proposal, there are also personal reasons at play. For part of what it is for the Grand Canyon to be valuable is that visiting and admiring it is worthwhile. Hence, someone who wishes to engage with the Grand Canyon in the appropriate ways has good personal reasons to favor its preservation, reasons that are derived from the impersonal value of the Grand Canyon itself.²⁶ Of course, it remains to be seen how a similar line of reasoning can be applied to the context of interpersonal trade-off.

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 complaints individuals may raise on their own behalf. Scanlon develops this idea in the direction of his tiebreaking argument, claiming that giving everyone's interest equal moral significance in *Life vs. Lives* requires the presence of some individuals to play a tiebreaking role. Scanlon's move, as we have seen, in the end proves unsuccessful. I plan to develop the idea of equal consideration in a new direction.

What does it take to treat the same interests of two different people with equal consideration? A natural answer is that when the interests of two individuals, A and B, are in fact equal in magnitude, one should regard

26. See also Scanlon, *What We Owe*, 200.

the satisfaction of their interests as supported by equally strong reasons. Of course, when the interests of A and B cannot be jointly satisfied, the conflict must be resolved in one way or another. But what equal consideration forbids, absent special justification, is the judgment that there is stronger reason to satisfy the interest of one party 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 the amount of resources we are willing to devote to A. Holding fixed that the lives of A and B do not differ in any morally relevant respect, our decisions reveal an inequality in the consideration we show to the interests of A and B. B may rightly complain, on her own behalf, that her interest in avoiding a lethal harm 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 the parties involved. 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. Let us begin by reframing the situation 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 people. 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. Suppose I decide to save A under the circumstances. Does my decision violate giving equal consideration to the interests of B and C?

Now, in deciding to save A under the circumstances, my decision treats A's life as having the kind of importance that provides me with sufficiently strong reason to forgo the aggregate value of saving two other lives. But if I treat A's life as having this kind of importance, then equal consideration requires that I treat B's life as being equally important. This means that I should treat B's life, too, as giving me sufficiently strong reason to forgo the value of saving two other lives. Next, and continuing with our suppositions, if the importance of B's life gives me sufficiently strong reason to save her even when doing so is opposed by the disvalue of two deaths, then there is even stronger reason for me to save B when saving her is opposed by the disvalue of just one death. In other words, if there is reason of a certain strength to save B even at the cost of two deaths, then there is reason of a comparatively greater strength to save B when the associated 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 whether there is stronger moral reason for us to choose the option of saving A's life at the cost of two deaths, or the option of saving B's life (together with C's) at the cost of just one death. Equal consideration for the lives of A and B implies that the moral reason to choose the latter option is stronger. We may formulate the underlying principle of equal consideration roughly 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 entails that if there is reason of a certain strength to forgo V1 for the sake of X's interest, then 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 personal complaint in the following terms: "If you treat A's life as being so important that you are willing to let two other people die for A's sake, then since my life is just as important, giving equal consideration to my life means that you should treat my life as giving you a comparatively stronger reason to forgo just a single life for my sake. Failing to do so unjustifiably *undervalues* my life in comparison to A's life."²⁷ Or, even more pithily, "Why are you willing to save A at the cost of two lives but no more willing to save me at the cost of just one life?"²⁸

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 in magnitude. 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 complain, again on her own behalf, that while her interest is indeed less weighty than A's interest, we should nevertheless be prepared to allot to her an amount of valuable resources that is proportional to

27. Is the decision to save A also unjustifiable to A because it overvalues her life? This question brings up a larger issue of who can be wronged under an individualist framework. One thing to note is that if we endorse a link between wronging someone and certain reactive attitudes, then on the suggestion that saving A also wrongs A herself, it seems to follow that it would be appropriate for A to blame the person who saves her and that it would be appropriate for the person who saves A to seek forgiveness from A for saving her. It does seem to me that these observations, though far from dispositive, count against the suggestion that A would be wronged by a decision to save her life.

28. I owe this pithy formulation to Theron Pummer.

the relative weightiness of their respective interests. 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 entails that the values we have reason to forgo for the sake of their respective interests are not disproportional to the relative weightiness of their interests.

Now, suppose that I decide to take the lifesaving option in *Life vs. Paraplegias*. My decision treats a person's life as having 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 disproportionately 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 that it may well dwarf the value of saving a single life, then it would be disproportional for me to be willing to forgo this very large amount of value for the sake of saving a life but be no more willing 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 lifesaving option stands to violate the proportionality constraint introduced by EC2. When this happens, each of those individuals facing paraplegia may validly claim that her interest is not given equal consideration.

Let me take care to explain why it is that the foregoing rationale for saving the greater number in *Life vs. Lives* and *Life vs. Paraplegias* does not extend to the case of *Life vs. Headaches*. This is so for two important reasons. First, in order for an individual to validly claim that her interest is not given equal consideration, it must first be true that her interest in fact merits equal consideration. And in order for a person's interest to merit equal consideration, it must first be true that the interest in question at least provides us with some genuine reason for its satisfaction. In other words, a claim of equal consideration must have a normative basis. But, as I argued in Section III, the reason-giving force of an individual's interest in avoiding a minor headache is wholly disabled with respect to the decision whether to save a life or prevent some number of headaches. Because of this, in *Life vs. Headaches*, no one facing a minor headache has the proper normative standing to raise a valid claim based on the ideal of equal consideration, for the simple reason that her interest is not one that we have reason to take into consideration in the first place. To put the point rather bluntly, an interest with no reason-giving force should not be given equal consideration; rather, it should be given no consideration at all.²⁹

29. These remarks apply most straightforwardly to the case of *Life vs. Headaches*, where the issue that needs to be settled is whether to save a life or prevent some number of headaches. But we may give consideration to headache-related interests in settling other

Second and relatedly, while my argument from equal consideration does make indirect appeal to considerations of aggregate well-being, the proper application of the argument calls for the settlement of a prior issue in line with an individualist framework, namely, the aggregation of *which* values may play this kind of indirect role. Here, there is a strong individualist basis for altogether excluding headache-related interests from having even an indirect bearing on our moral adjudication in Life vs. Headaches. This is because, given my earlier argument from normative disabling, no one facing a minor headache is in a position to complain if we simply bar her interest from being part of the aggregate value in the application of the argument from equal consideration. In this way, there is yet another crucial difference between Life vs. Headaches and Life vs. Paraplegias. In the latter case, each person facing paraplegia has a valid claim that we must not simply ignore her interest in our moral adjudication. The way to properly take their interests into account, consistent with the individualist restriction, is to allow each of their interests to play the kind of indirect role I have been advocating. In short, my suggestion that aggregative considerations may play an indirect role in the generation of personal reasons is not a blank check; whether a set of interests may play this kind of role is itself an issue that needs to be settled in line with our commitment to moral individualism.

I think the rationale I provide for saving the greater number is plausible in its own right, but it may 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 in numbers increases between the competing groups, each individual in the larger group has an increasingly stronger claim to being saved. 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 are in effect saying 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 just one life. But when B is in a group of one 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 are saying that, compared to A's life, which we judge to be worthy of forgoing the value of saving one million lives, B's life gives us no stronger reason to

normative questions. For example, suppose that I announce to the world that I am willing to let A die for the sake of preventing just a single headache, but I am only willing to let B die if more than fifty million headaches can be thereby prevented. With respect to the issue *for the sake of how much value am I willing to forgo the lives of A and B*, we may look to the headache-related interests to determine that I am indeed willing to forgo A's life for much less value. And this determination, in turn, may ground a complaint on A's part that her life is even more undervalued by me than B's life.

forgo the value of saving just a single life. B therefore has an even stronger personal complaint, based on the ideal of equal consideration, when she is in a group of one million than when she 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 tiebreaking argument susceptible to this reply is that a weighted lottery allows each person's presence to make a difference. But my proposal has nothing to do with the idea of difference-making. Instead, it relies on the significance we attribute to people's competing interests as revealed by the amount of value we are willing to forgo for their respective sakes, as well as how such attributions relate to the ideal of equal consideration as specified by EC1 and EC2. In general, when one's decision runs contrary to a 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.³⁰

Third, the position I defend here in no way calls into question certain widely held views in axiology. It does not challenge the continuity or transitivity of the "better than" relation, nor does it deny that forgoing the lifesaving option for the sake of preventing sufficient many headaches may generate more well-being overall. Hence, the well-known "spectrum argument," which seeks to establish that there must be some number of headaches the alleviation of which generates more value than the saving of a life, does not in any way threaten my position. What my position instead insists on are two things. First, as mandated by the individualist restriction, conclusions about aggregate well-being do not by themselves carry any justificatory weight in our deontic determinations; only personal reasons matter. Second, aggregative considerations may nevertheless play an indirect role in the generation of personal reasons, but only when that role is consistent with our more basic commitment to moral individualism.

30. Two further observations about the lottery strategy are in order. First, the use of an unweighted lottery, one that gives B and C as a group a 50 percent chance of being saved, is itself subject to an objection from equal consideration. For B and C may rightly complain that, since equal consideration entails that there is comparatively stronger reason to save them than A, this difference in the strength of reason should be reflected in the chance we allot to their survival. Second, a weighted lottery may escape the above objection, but it faces a deeper problem that plagues all lottery proposals. Imagine that we can save either X's life or Y's arm. It should be agreed by everyone that once we arrive at the correct judgment that there is stronger reason to save X, we should simply do so. Running a lottery between X and Y, even a weighted one, is plainly out of place. In the same way, once we arrive at the conclusion that there is comparatively stronger reason to save B and C over A, as the ideal of equal consideration requires, we should simply go ahead and rescue the larger group. Quite generally, the comparative strength of our practical reasons is meant to settle the issue of which option is to be taken.

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 a roundabout way? Why not just admit that we should save the greater number in Life vs. Lives 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 indiscriminate aggregation; accepting the verdict that we should prevent many minor headaches at the cost of leaving someone to die is, I submit, simply too big a bullet to bite.

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 considerations leaves 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?³¹ I think the correct answer to Anscombe's question is that everyone in the larger group is wronged; saving the lesser number under the circumstances would be unjustifiable to each of them. But if the reason for saving the greater number lies directly in considerations of an aggregative kind, then no one in the larger group is in a position to say, on her own behalf, that she in particular has been treated in an unjustifiable manner. For this reason, my individualist rationale for saving the greater number, which embeds aggregative considerations within an ideal of equal consideration, violations of which provide each person the standing to raise a complaint on her own behalf, is well positioned to give the correct answer to Anscombe's question.

Taken together, the argument from normative disabling and the argument from equal consideration reveal a rationale, at the level of justification, for allowing the numbers to count in Life vs. Lives and Life vs. Paraplegias but not in Life vs. Headaches. Crucially, the reason why the numbers count, when they do, is sourced not in an aggregative ideal of maximizing overall well-being but in an individualist ideal of giving equal consideration to those whose interests merit such consideration. In this way, an individualist approach can make the numbers count without directly counting the numbers.

Moreover, at the level of adjudication, the preceding discussion leads to the endorsement of a general view of interpersonal trade-offs which we may call *partial aggregation*, according to which we should take numbers into account in some cases of trade-off but not in others, thereby capturing our divergent intuitions regarding the relevance of numbers in Life vs. Lives and Life vs. Headaches. Partial aggregation, however, remains a quite general, and indeed underdeveloped, suggestion. Recently,

31. Elizabeth Anscombe, "Who Is Wronged?," *Oxford Review* 5 (1967): 16–17.

several authors have expressed skepticism about the prospect of developing an extensionally adequate version of partial aggregation that can handle more structurally complex cases. While I do not have the space to thoroughly engage with all of the extensional objections that have been leveled against partially aggregative views in the recent literature, in the next section I identify some key areas of controversy on which I think the individualist framework developed here can shed some helpful light.

V. PARTIAL AGGREGATION

The three leading cases so far discussed all involve competition between homogeneous groups, to use Patrick Tomlin's helpful term.³² That is, they all involve competition between groups whose sizes stay constant and whose members stand to be burdened in identical ways. Further complexities arise when we consider competitions between nonhomogeneous groups. To streamline our discussion, and in keeping with the recent literature, let us make the following stipulations:

- A claim against MODERATE HARM is relevant to a claim against MAJOR HARM.
- A claim against MINOR HARM is relevant to a claim against MODERATE HARM.
- A claim against MINOR HARM is not relevant to a claim against MAJOR HARM.

Now, consider a case that is central to Tomlin's objection to partial aggregation:

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.³³

How is the relevance status of the competing claims to be determined, especially when different tiers of claims are all present in the same situation? According to the influential Aggregate Relevant Claims (ARC) account developed by Alex Voorhoeve, which has served as the focal point of much of the recent criticisms of partial aggregation, a claim is relevant if and only if it is sufficiently close in strength relative to the strongest competing claim.³⁴

32. Tomlin, "On Limited Aggregation," 238.

33. *Ibid.*, 242.

34. Voorhoeve, "How Should We Aggregate," 66.

ARC, however, gives rise to a serious problem. Tomlin invites us to 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. Given this supposition, ARC implies that we should save Group B at Stage 1. With the advent of Stage 2, however, ARC reverses its verdict and instructs us to save Group A instead. This is because, under ARC, the one million claims against minor harm added to Group B are all irrelevant, while the single claim against minor harm added to Group A is relevant. This implication of ARC is highly counterintuitive. As Tomlin observes, quite plausibly, 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 satisfy 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.

Tomlin's objection to ARC is forceful. But the objection at least leaves open the possibility of developing other versions of partial aggregation that do not suffer from the same shortcoming. Regarding this possibility, however, Joe Horton is not at all hopeful. Building on Tomlin's objection, Horton argues that all versions of partial aggregation inevitably fall prey to a fatal dilemma. Consider a case central to Horton's objection:

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.³⁵

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.

35. Horton, "Always Aggregate," 168.

At Stage 1, any version of partial aggregation should imply that we are to save Group A. In view of Tomlin's Principle, it seems that we should 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 no fewer claims are added to Group A than to Group B.

However, as Horton observes, there is also compelling reason for defenders of partial aggregation to favor the contrary conclusion. This is because, by supposition, the fifteen claims against moderate harm in Group B outcompete the one claim against major harm in Group A, and the one million claims against minor harm in Group B outcompete the fifteen claims against moderate harm in Group A. It is therefore eminently plausible that we should save Group B in Case Two. To state Horton's reasoning more precisely:

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.

And herein lies Horton's dilemma. It appears that, in Case Two, any version of partial aggregation 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 partial aggregation is indeed in deep trouble.

To resolve these challenges, we need to first consider the feature of ARC that gives rise to the troubling implication in Tomlin's Case One. ARC has this implication because, to borrow Victor Tadros's helpful term, it understands relevance as a global matter: the relevance of a claim under ARC is always determined by comparing the claim to the overall strongest claim with which it competes. However, my discussion of the contextual character of normative disabling in Section III motivates a different understanding. True enough, in Case One, the reason-giving force of the claims against minor harm in Group B is normatively disabled relative to the competing claim against major harm. But they still retain their reason-giving force relative to the one claim against minor harm added to Group A. I am therefore in agreement with a suggestion 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 offsetting the second claim. But it may still play a role in offsetting

other competing claims in relation to which it is relevant. Tadros calls this idea *Local Relevance*.³⁶

Local Relevance, however, does not by itself resolve our present challenges. This is because in structurally complex cases, where claims of different tiers are all present, Local Relevance does not yet tell us how exactly the process of offsetting is to be carried out. Here, reasoning within an individualist framework provides distinctive answers. In situations of interpersonal trade-off, moral individualism naturally directs our attention to the individuals who have the greatest need to be helped. In virtue of having the most at stake, these individuals occupy our moral concern first and foremost. Of course, my earlier argument from equal consideration tells us that the individually strongest claims do not enjoy lexical priority. But there is another way for these claims to have the pride of place in our moral adjudication—we may let them set the context and aim of the process of offsetting.

Here is the basic idea I have in mind. In adjudicating between the competing interests of individuals, we are to take as our controlling thought that we should satisfy the individually strongest claims unless they are offset by competing, nonoffset claims that are relevant to them. Here is a quick illustration. 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 offset C1, but only if they are not themselves offset by other competing claims. To see whether C2 can be offset, we check to see whether there are any competing claims that are relevant to C2. If so, these claims may offset 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 previously 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 can in the end be offset, they are completely set aside. In this way, the individually strongest claims set the *context* for determining which claims have a role to play in the offsetting process.

Of the remaining, admissible claims, 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 moral attention. Throughout the

36. Tadros, “Localized Restricted Aggregation,” 5:179.

process of offsetting, we must ensure that a claim which is irrelevant to a second claim may not offset the second claim. Finally, after the process of offsetting is carried to its completion, we should choose the option that satisfies the strongest nonoffset claims. Call this view 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? If not, we should choose the option that satisfies these strongest claims. If so, we should choose the option that satisfies the strongest, admissible, nonoffset claims.

The underlying rationale for PPA, to repeat, is an individualist one. 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 violate my earlier argument from equal consideration 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, nonoffset 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.

We are now in a position to give a full response to the objections from Tomlin and Horton. Consider Tomlin’s Case One. 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 Voorhoeve’s ARC that gives rise to Tomlin’s objection is this: under ARC, claims of identical strength may nonetheless differ with respect to their status of being relevant. Tomlin’s objection cleverly exploits this feature of ARC by bringing into focus a case 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 is deemed to be relevant by ARC. PPA, however, avoids this difficulty. Even though the one million claims added to Group B are irrelevant when compared to the competing claim against major harm, PPA nevertheless allows them to offset the single claim against minor harm added to Group A. In this way, PPA ensures that Tomlin’s Principle is fully conformed to in a situation like Case One.

In Horton's Case Two, PPA yields the following conclusions. 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. Indeed, the inapplicability of Tomlin's Principle to Case Two is illustrated by PPA itself.

Viewed through the lens of PPA, there is a crucial difference between Case One and Case Two. 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. However, with the advent of Stage 2, these one million claims become admissible; they may now 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 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 saying, "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 to 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 t_2 " fails to be satisfied.

We may say that the addition of the fifteen claims against moderate harm to Group A creates a kind of 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 admissibility.

To be clear, my reply to Horton's dilemma is not merely that there is some mechanism internal to PPA which enables us to distinguish Tomlin's Case One from Horton's Case Two. My more basic reply, rather, is that the reading of Tomlin's Principle which underlies Horton's dilemma is not a plausible reading to begin with. 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. Such a prohibition can be reasonably rejected by each of those individuals facing minor harm because it denies their interests the standing to offset competing claims relative to which they are clearly relevant.

My discussion so far indicates that PPA is immune to the objections presented by Tomlin and Horton. Let me conclude my discussion by considering a number of further objections to partial aggregation more recently advanced by Aart van Gils and Patrick Tomlin.³⁷ As I mentioned earlier, PPA incorporates the central insight behind Tadros's Local Relevance. Although Tadros himself does not develop Local Relevance in much greater detail, Van Gils and Tomlin have done so on Tadros's behalf. They put forward, but ultimately reject, an alternative way of developing Local Relevance which they call *Sequential Claims-Matching* (SCM). Contrasting PPA with SCM will help to further highlight some of PPA's distinctive 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 noneliminated claims.

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 inadmissible claims are completely set aside. These two distinctive features of PPA flow from the individualist rationale for PPA that I presented earlier,

37. Van Gils and Tomlin, "Relevance Rides Again?"

which allows the individually strongest claims to set the aim and context of the process of offsetting. And the features also 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 Three: We can either save A's life or save B's life plus cure C's sore throat.³⁸

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. 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. PPA, however, 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 therefore 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 interest in avoiding a minor discomfort.³⁹

Van Gils and Tomlin moreover identify two general ambiguities that proponents of partial aggregation need to address.⁴⁰ 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. Although PPA tells us to start the offsetting process with the weakest admissible claims, this condition by itself does not render the order of offsetting fully determinate in all cases. To see this, consider, for example:

Case Four: 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).

In this case, the A2 claims are among the weakest, so they may serve as the starting point of our offsetting process. But notice that the A2 claims are relevant to both the B1 claims and the B2 claims, so there are now two ways in which we can proceed. Under Option One, we let the A2 claims

38. Kamm, *Morality, Mortality*, 1:101.

39. PPA, as so far developed, does carry the implication that it is permissible to save A's life no matter how many more sore throats we can prevent in Case Three. Speaking for myself, this is an implication that I am fully happy to accept. For I do not think that we are morally required to leave someone to die just because the option of saving another's life produces the additional benefit of curing very many sore throats.

40. Van Gils and Tomlin, "Relevance Rides Again?," 6:245–50.

and the B2 claims completely offset one another, resulting in the verdict that we should save Group B. Alternatively, under Option Two, we let the A2 claims offset the B1 claims, resulting in the verdict that we should save Group A instead.

In my view, the way to resolve the present ambiguity is to once again reason within an individualist moral framework. Under Option Two, while the individuals in A2 are allowed to have their interests play a role in determining which group we ought to save, the many individuals in B2 are wholly excluded from taking part in the offsetting process even though their interests are no less significant. This kind of differential treatment is unjustifiable, and each of those individuals in B2 may reasonably complain that the offsetting process under Option Two treats her unfairly, 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 process of offsetting. 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 but includes others even when these competing claims are relevant to one another.

The second ambiguity Van Gils and Tomlin identify arises in situations where a claim is incompletely offset. Consider, for example:

Case Five: 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).

In this case, we are to imagine that 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 Five?

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. But I think we can say more. In Case Five, while it is true that the B2 claim may be relevant relative to some proportion of the A1 claim, it is not true that the person who faces major harm stands to suffer 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 Five 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, I think, 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.

VI. 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 a proper solution to the numbers problem must meet 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 should attend to Anscombe's question, address the plausibility of the lottery strategy, capture the idea that sometimes individuals have increasingly stronger claims to being saved as the disparity in numbers increases between competing groups, and help us navigate more structurally complex situations of trade-off. I have in this article put forward an individualist 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, it is that we should not expect the numbers problem to admit of an uncomplicated solution.