Relevance and Non-consequentialist Aggregation

J. PAUL KELLEHER
University of Wisconsin – Madison

Interpersonal aggregation involves the combining and weighing of benefits and losses to multiple individuals in the course of determining what ought to be done. Most consequentialists embrace thoroughgoing interpersonal aggregation, the view that *any* large benefit to each of a few people can be morally outweighed by allocating *any* smaller benefit to each of many others, so long as this second group is sufficiently large. This would permit letting one person die in order to cure some number of mild headaches instead. Most non-consequentialists reject thoroughgoing interpersonal aggregation despite also believing it is permissible to let one person die in order to prevent many cases of paraplegia instead. Non-consequentialists defend this asymmetry largely on the basis of intuition, and some rely on the notion of *relevance* to formalize the grounding intuitions. This article seeks to clarify and strengthen the non-consequentialist notion of relevance by engaging with three objections to it.

I. INTRODUCTION

Consequentialists embrace thoroughgoing interpersonal aggregation; most non-consequentialists do not. Interpersonal aggregation *simpliciter* is the combining and weighing of benefits and losses to multiple individuals in the course of determining what ought to be done. (Until section V, I will use 'aggregation' to mean 'interpersonal aggregation'.) For example, many people agree that it is permissible, and can even be obligatory, to help the larger of two groups when all individuals stand to gain the very same benefit. Thus, in cases where we can either save one person from dying or five (and all else is equal), it is not uncommon for consequentialists and non-consequentialists alike to recommend saving the larger number. The same goes for cases where the potential benefit falls short of life-saving, for example where we must choose between restoring the mobility of one person or that of five others. In each of these cases, it seems to matter morally that one group of beneficiaries is larger than the other: when each group's individual benefits are combined and weighed against the other's, the moral scales are tipped in favour of the larger group.

What tends to separate consequentialists and non-consequentialists are cases in which potential beneficiaries stand to gain benefits of different sizes. Again, consequentialists embrace *thoroughgoing* aggregation, which holds that the allocation of *any* large benefit to fewer people can be morally outweighed by the allocation of *any* smaller

© Cambridge University Press 2014
Utilitas Vol. 26, No. 4, December 2014
doi:10.1017/S0953820814000144

https://doi.org/10.1017/S0953820814000144 Published online by Cambridge University Press

benefit, so long as the group of people receiving the smaller benefit is sufficiently large. Hence a proponent of thoroughgoing aggregation will accept that there is some finite number of headaches such that it is permissible (and perhaps even obligatory) to relieve *that* many headaches instead of saving an innocent person's life. My aim in this article is to clarify and strengthen a standard non-consequentialist alternative to thoroughgoing aggregation by engaging with three objections to it. The non-consequentialist alternative I will focus on is built around the technical notion of *relevance*.

Before introducing this notion of relevance, I want to distinguish between two varieties of aggregation – or, perhaps better, between two domains of morality within which aggregation is an issue. The first is *axiological aggregation*, or aggregation within the domain of axiology. Axiology is the study of goodness or value, and it seeks to understand what makes some states of affairs better or worse than others. As Alastair Norcross explains, aggregation within the domain of axiology 'involves the claim that harms and benefits [of different sizes] can be traded off against each other in determining the overall goodness (or badness) of a state of affairs'.[1] The other kind of aggregation is *deontic aggregation*, or aggregation within the domain of *permissibility and obligation*. This is the domain of morality that issues in conclusions about what may or must be done. Since consequentialism is the view that the right thing to do is that which maximizes goodness, consequentialists hold that considerations within the domain of axiology completely determine the facts about moral permissibility and obligation. But non-consequentialists reject the view that the right thing to do is always that which maximizes the good, and so they hold that some moral questions cannot be answered by axiological considerations alone. Now, it is controversial whether thoroughgoing aggregation is the correct approach within the specific domain of axiology. Consider Derek Parfit's so-called 'Repugnant Conclusion':

> For any possible population of at least ten billion people, all with a very high quality of life, there must be some much larger imaginable population whose existence, if other things are equal, would be better, even though its members have lives that are barely worth living.[2]

'Better' here means 'better with respect to the amount of goodness produced by or contained in the state of affairs'. As Parfit's name for this conclusion indicates, many find it very difficult to believe that it would be better to convert a population of 10 billion happy people into a much

[1] A. Norcross, 'Two Dogmas of Deontology: Aggregation, Rights, and the Separateness of Persons', *Social Philosophy & Policy* 26 (2008), pp. 76–95, at 84.

[2] D. Parfit, *Reasons and Persons* (Oxford, 1984), p. 388.

https://doi.org/10.1017/S0953820814000144 Published online by Cambridge University Press

larger population of individuals with lives 'barely worth living'. But note that even if a non-consequentialist concedes that thoroughgoing aggregation operates within the domain of axiology – that is, even if the larger of Parfit's populations really is *better* than the smaller one – the non-consequentialist can still say that it is an open question whether thoroughgoing aggregation is permissible within the domain of deontic assessment. After all, she does not concede that goodness always determines rightness and permissibility.

In any case, my reason for distinguishing axiological from deontic aggregation is simply to stress that this article is concerned with aggregation within the domain of deontic assessment. I will be assuming (but purely for the sake of argument) that thoroughgoing aggregation is perfectly permissible within the domain of axiology.[3]

II. NON-CONSEQUENTIALIST AGGREGATION AND THE IDEA OF RELEVANCE

The non-consequentialist alternative to thoroughgoing aggregation that I seek to explore actually embraces quite a bit of aggregation. In particular, it generally follows Norcross in accepting that for any large benefit b that can be given to each member of group G, there is at least one kind of benefit smaller than b which, when given to sufficiently many others, can morally outweigh giving b to G. For example, this non-consequentialist view can accept that rather than save S's life, one should instead cure paraplegia in, say, 1,000 others. Typically, this stance is not grounded in theory. Many non-consequentialists simply find it implausible that one must always help the group that is worse off or whose members can be helped the most, regardless how many others one can help with less serious but still significant needs.

To avoid inconsistency when saying that aggregation is permissible in some cases but not others, many non-consequentialists have come to rely on the technical notion of *relevance*. To illustrate, suppose that we are presented with the choice between (a) curing one person's one-hour long headache or (b) curing half-hour long headaches afflicting some number of individuals. The non-consequentialist who accepts aggregation in some death-versus-paraplegia cases will also probably accept aggregation in some hour-long-headache-versus-half-hour-long-headache cases. But she may well deny that half-hour headaches can be aggregated to morally outweigh saving a life. She will say, then, that while paraplegia is *relevant* to death (in that paraplegia *can* be aggregated to morally outweigh death), and while half-hour

[3] For an argument against thoroughgoing aggregation within the domain of axiology, see D. Dorsey, 'Headaches, Lives, and Value', *Utilitas* 21 (2009), pp. 36–58.

https://doi.org/10.1017/S0953820814000144 Published online by Cambridge University Press

long headaches are relevant to hour-long headaches, half-hour long headaches are *not* relevant to death. Relevance can thus be thought of as a relation that determines 'when the numbers count'. As T. M. Scanlon puts it, '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.'[4] (In what follows I will refer both to relevance between the harms that we might prevent and relevance between the benefits that we might confer. I take these to be two ways of framing and discussing what is ultimately the same moral relation.)

What could possibly justify belief in the idea that relevance is a genuine relation between types of harm and benefit? I am not certain that anything other than unshakeable intuitions can be marshalled to support it. There are, however, formidable arguments *against* the idea and its commonly associated claim that headaches are not relevant to death, and in the remainder of this article I will engage with three such arguments. My main goal is to assess the force of each argument and to formulate what I take to be the most plausible response on behalf of proponents of the idea of relevance. Owing to space limitations, and to my own reservations about their success, I shall not attempt a full defence of each response. Instead, I seek to understand where and why the conventional notion of relevance needs refinement, and to get that work started. It is my hope that others will find my responses to these three objections useful, and will perhaps develop them further.

The remainder of the article is structured as follows. Section III considers an objection pressed by Alastair Norcross, and argues that the objection is not at all damaging to the idea of relevance. Section IV articulates a risk-based objection to relevance, and argues that the non-consequentialist's most plausible response to it requires the somewhat radical step of foreswearing any appeal to the relative strengths of individuals' claims to assistance. I go on to suggest a sketchy but (I think) still promising account of why it might be that headaches are not relevant to death. Section V considers a powerful objection to the idea of relevance that is due to John Broome. Broome offers an example of an intuitively permissible policy that arguably involves aggregating many prevented headaches to morally (that is, deontically) outweigh saving a life. Here I suggest that the non-consequentialist must lean on a distinction between *inter*personal aggregation and *intra*personal aggregation, and must maintain that Broome's policy is permissible

[4] T. M. Scanlon, *What We Owe to Each Other* (Cambridge, MA, 1998), pp. 239–40. See also F. M. Kamm, *Morality, Mortality: Death and Whom to Save from It*, vol. 1 (Oxford, 1993).

https://doi.org/10.1017/S0953820814000144 Published online by Cambridge University Press

only to the extent that it is defensible in intrapersonal terms. Section VI concludes.

III. NORCROSS'S ARGUMENT AGAINST IRRELEVANT GOODS

Alastair Norcross observes that 'the relation of moral relevance does not obey the principle of transitivity.'[5] Much of the aggregation literature's discussion of transitivity is focused on claims within the domain of axiology. There, the issue is whether the relation 'all things considered better than' is transitive (recall Parfit's Repugnant Conclusion).[6] Norcross's focus, by contrast, is on deontic claims about which options must or may permissibly be chosen. According to Norcross, if one embraces the deontic notion of relevance, one must also accept that there can be cases in which one ought to choose option A over option B, option B over option C, *and* option C over option A. He views this as a serious problem for the claim that relevance is a genuine moral relation between benefits of different sizes. To use his example, suppose for the sake of argument that '(a) the loss of both arms is less serious than but morally relevant to death; [and] (b) that a broken leg is less serious than but morally relevant to the loss of both arms, but not morally relevant to death' (p. 83). Assume also that preventing one death can be morally outweighed by preventing 1,000 people from losing their arms, and that preventing 1,000 people from losing their arms can be outweighed by preventing 1 million people from breaking a leg. Now suppose you must choose between three mutually exclusive options: saving one life, preventing 1,000 people from losing their arms, and preventing 1,000,000 people from breaking a leg. What should be done? One worry is that whichever you choose, you would act wrongly, since for every option there is an alternative that is morally superior.[7] But let us set that issue aside (as Norcross agrees to do) and assume that at least one of your three options must be morally permissible. Even so, Norcross claims that the non-consequentialist's introduction of the notion of relevance has a 'very strange', 'unpalatable consequence'.[8]

To see this, suppose you are morally permitted to choose any of your three options, and you decide to save the one life. Just as you

[5] Norcross, 'Two Dogmas', p. 81.

[6] See, for example, L. S. Temkin 'A Continuum Argument for Intransitivity', *Philosophy & Public Affairs* 25 (1996), pp. 175–210.

[7] Derek Parfit thinks this is a serious problem for a view which invokes the notion of relevance (what Parfit calls 'The Close Enough View'). Whether this is in fact a serious problem turns in part on whether there can be moral dilemmas and whether this might be one. I lack the space to discuss this adequately here. See D. Parfit, *On What Matters* (Oxford, 2011), p. 203.

[8] Norcross, 'Two Dogmas', p. 83.

https://doi.org/10.1017/S0953820814000144 Published online by Cambridge University Press

set off to save the life, however, one of the two other options becomes unavailable when a tree falls and blocks the path to it; in particular, you can no longer consider preventing 1,000,000 broken legs as one of your options. Norcross then states the putative problem: 'You are about to perform the perfectly permissible act of saving a life, when one of your other permissible alternatives becomes unavailable by chance. Now it is no longer permissible to save the life', since we have stipulated that if you are choosing between saving a life and preventing 1,000 people from losing their arms, you must do the latter. According to Norcross, 'We should, if at all possible, avoid having to swallow such an unpalatable consequence.'[9] Since this consequence follows from the 'intransitivity' engendered by the notion of relevance, we should abandon that notion. This would appear to force non-consequentialists to choose between thoroughgoing aggregation (which in effect makes headaches relevant to death) and thoroughgoing anti-aggregation (which denies that paraplegia is relevant to death), both of which most non-consequentialists wish to reject.

In response to Norcross, David Lefkowitz has argued that a nuanced account of relevance can be used to convert the problematic three-option case into a non-problematic two-option case. According to Lefkowitz's 'orbital' notion of relevance, 'the option of preventing a given harm provides an agent with both a first-order reason to prevent that harm, and a second-order reason not to consider as part of his deliberation the option of preventing irrelevant harms.'[10] The idea of a harm's 'orbit' is just the idea of there being a certain range of alternative harms that are relevant to it. So far, this adds nothing new to the notion of relevance as I have described it. What is new is Lefkowitz's claim that when one has an opportunity to prevent a given harm, one thereby acquires two practical reasons: a first-order reason to prevent the harm, and a second-order reason to ignore opportunities to prevent harms that are not within the orbit of the first harm. Lefkowitz claims that this 'two-level theory of practical rationality' can be used to avoid the dilemma Norcross constructs for those who embrace the idea of relevance.[11] Since (by hypothesis) a broken leg is not within the orbit of an imperilled life, the option of preventing many broken legs should not even be on the table in the first place, according to Lefkowitz. If that is correct, then one is left with just the two options of saving the life and preventing 1,000 lost pairs of arms, and one must therefore prevent the lost pairs of arms.

[9] Norcross, 'Two Dogmas', p. 83.

[10] D. Lefkowitz, 'On the Concept of a Morally Relevant Harm', *Utilitas* 20 (2008), pp. 409–23, at 416.

[11] Lefkowitz, 'On the Concept', p. 418.

https://doi.org/10.1017/S0953820814000144 Published online by Cambridge University Press

I do not believe that Lefkowitz's response to Norcross is successful. It is of course clear why the orbital notion of relevance provides a reason to ignore the prospect of preventing the broken legs. Since a broken leg is not within the orbit of a death, we have a second-order reason to ignore each and every broken leg whose prevention conflicts with preventing the death. What is not so clear, however, is why this second-order reason automatically trumps the second-order reason there is to *consider* the broken legs given that a broken leg *is* within the orbit of a lost pair of arms. Lefkowitz seems to assume that once a second-order reason exists to ignore a harm as irrelevant, that harm can never be heard from again. But Lefkowitz gives no reason to think that second-order 'exclusionary' reasons are *ipso facto excluding* reasons. That is, he does not address the possibility that a second-order exclusionary reason can be counterbalanced or neutralized by another second-order reason that pulls in the other direction. Instead, he asserts that since death is the worst harm one can prevent in Norcross's case, 'Therefore, according to the orbital conception of relevant harms, the agent should only consider in his deliberation the possibility of preventing harms that are relevant to' death.[12] But this is not an automatic consequence of the orbital notion of relevance, and thus it must be supported by further argument. It is true that given the prospect of preventing a death, one has a second-order reason to ignore the prospects one has for preventing broken legs. But for all Lefkowitz has argued, it is also true that given the prospect of preventing a lost pair of arms, one has a second-order reason to give concomitant consideration to all those harms that are within the orbit of a lost pair of arms. And if that is correct, then we are arguably where Norcross left us, namely in a three-option case that invites a strange consequence when one's permissible option is rendered impermissible by the removal of an option one permissibly avoided in the first place.

In my view, a better way to rebut Norcross is simply to show that the non-consequentialist should not be embarrassed by the consequence that Norcross finds strange and unpalatable. Thus, consider the following case. Suppose that (a) Anne and Bill have played each other in Scrabble 1,000 times, and Anne has always won; (b) Bill and Clara have played 1,000 times, and Bill has always won; and (c) Clara and Anne have played 1,000 times, and Clara has always won. You have promised your rich uncle that you will place a bet on his behalf at the National Scrabble Tournament's final three-person championship game between Anne, Bill and Clara. Your only guidance is to place the wisest bet possible. Since you are morally bound by your promise to bet, and since the intransitivity between the finalists leaves none

[12] Lefkowitz, 'On the Concept', p. 418.

https://doi.org/10.1017/S0953820814000144 Published online by Cambridge University Press

as 'the wisest' choice, let us assume you permissibly blindly draw a name out of a hat, selecting Anne. However, just as you are about to place your bet on Anne, you learn that Bill has been disqualified (it turns out he is Canadian and ineligible for the US Tournament). If this happens, then plainly you must now bet on Clara, since she has never lost to Anne. But this is precisely the phenomenon that Norcross called 'unpalatable' and claimed must be avoided: a permissible option is rendered impermissible by the removal of an option one permissibly avoided in the first place.

As far as I can tell, the only thing strange about the Scrabble case is the empirical implausibility of the fictional players' head-to-head Scrabble history. But that was not Norcross's concern. Rather, he claimed the serious problem arises once Bill is disqualified. But *pace* Norcross, there is nothing all that bizarre about what you are required to do once Bill is disqualified. I therefore agree with John Broome that:

> Cycles may occur in what one ought to choose; it may be that, given a choice between A and B one ought to choose A, and given a choice between B and C one ought to choose B, and given a choice between C and A one ought to choose C.[13]

Deontic intransitivity is simply not *itself* the theoretical deal-breaker Norcross thinks it is.[14] Thus, in order to criticize the non-consequentialist's reliance on the notion of relevance, I believe one is better served by focusing more directly on the claim that headaches are not relevant to death. Putting concerns about deontic intransitivity aside, then, is it really true that headaches are irrelevant to death? The next two sections consider powerful arguments for scepticism about that claim.

IV. CLAIMS TO ASSISTANCE AND A RISK-BASED ARGUMENT AGAINST IRRELEVANT GOODS

Non-consequentialists are famous for thinking that morality sometimes imposes constraints against maximally promoting the good. Often these constraints are said to forbid harming some to provide benefits to others. But non-consequentialists who reject thoroughgoing aggregation evidently must embrace a constraint of a different kind. For none of the cases I've discussed involves harming; they are all

[13] J. Broome, *Weighing Lives* (Oxford, 2004), p. 60.

[14] It is worth noting that what Norcross calls deontic 'intransitivity' is arguably not really a form of intransitivity at all. As Daniel M. Hausman explains, '[T]ransitivity concerns preferences over a single set of alternatives. Transitivity says nothing about how preferences over one set of alternatives {apple pie, blueberry pie} should be related to preferences over a different set of alternatives {apple pie, blueberry pie, cherry pie}.' See D. M. Hausman, *Preference, Value, Choice and Welfare* (Cambridge, 2012), pp. 15–16.

https://doi.org/10.1017/S0953820814000144 Published online by Cambridge University Press

cases of deciding whom to aid and whom to leave unaided. This may seem to make the idea of relevance all the more mysterious. After all, the idea that human agents are to be regarded as inviolable against bodily aggression has deep roots in ordinary moral thought. Yet no comparable violation seems to occur when *aid* is allocated to many people with minor ailments rather than to fewer people in imminent peril. Why, then, should we believe in constraints against thoroughgoing aggregation?

A common non-consequentialist answer to this question invokes the idea of a *claim* to assistance. There is more than one account of claims in the literature, but at least since Thomas Nagel's seminal discussion of aggregation, many non-consequentialists have expressed sympathy with what Nagel calls the 'individual acceptability' model (IA) of claims and of how conflicting claims should be reconciled. (I will briefly discuss a different model in section V.) IA holds that the strength of an individual's claim to assistance is commensurate with the urgency of her needs (and perhaps also with the degree to which she can be helped), and that conflicts between claims are to be adjudicated by finding the outcome 'that is least unacceptable to the person to whom it is most unacceptable'.[15] Nagel's discussion suggests something like the following relation between an outcome's unacceptability to an individual and the strength of her claim: the strength of her claim to assistance (which she possesses before any allocation choice is made) is equal to the strength of the complaint she could lodge if she is not helped, and the strength of the complaint she could lodge if not helped is equal to that outcome's degree of unacceptability to her. For example, in a two-way conflict case pitting one person in mortal peril against many others who are about to suffer mild headaches, the person with the dire need possesses a much stronger claim to assistance, and she will therefore have a much stronger complaint if the others are helped instead of her. This in turn means that the outcome involving cured headaches will be more unacceptable to the person in mortal peril than the outcome involving life-saving would be to those who are forced to endure headaches. Since IA instructs helping agents to pursue the outcome that is least unacceptable to the person to whom it is most unacceptable, IA in this case requires the life to be saved. We can articulate the same choice procedure in terms of claims: the person in mortal peril has a much stronger claim to assistance than any of the others, and IA directs agents in a position to help to satisfy claims in order of their strength. As Nagel puts it, 'Each individual

[15] T. Nagel, 'Equality', in his *Mortal Questions* (Cambridge, 1991 [1979]), pp. 106–27, at 123.

https://doi.org/10.1017/S0953820814000144 Published online by Cambridge University Press

with a more urgent claim has priority, in the simplest version of such a view, over each individual with a less urgent claim.'[16] Following Nagel, this method of adjudicating conflicts between claims has been called 'pairwise comparison', since it compares each individual's claim with each other individual's claim, rather than (for example) comparing the strength of one person's claim with the *aggregated* strength of many other people's *combined* claims.

Plainly, a constraint against thoroughgoing aggregation falls out of IA's procedure of pairwise comparison. The problem, however, is that IA seems to rule out *any* form of aggregation. Take, for example, the two-way conflict case involving one person in mortal peril and a great many (e.g. a billion) others who will be paralysed if not helped. Since there is no one among the group who has as strong a claim as the person facing death, there is no person among the group who will find the outcome involving life-saving more unacceptable than the worst off person would find the outcome in which she is left to die. And since IA adjudicates claims (and prospective complaints of unacceptability) using the method of pairwise comparison, the verdict will always come down in favour of the individual with the strongest claim. Nagel understood this and believed it was unacceptable:

> It seems to me that no plausible theory can avoid the relevance of numbers completely. There may be some disparities of urgency [of claims] so great that the priorities persist whatever numbers are involved. But if the choice is between preventing severe hardship for some who are very poor and deprived, and preventing less severe but still substantial hardship for those who are better off but still struggling for subsistence, then it is very difficult for me to believe that the numbers do not count, and that priority of urgency goes to the worse off however many more there are of the better off.[17]

I have already said that most non-consequentialists agree with Nagel that numbers can be decisive in at least some cases. The hard part, as Nagel saw, is accounting for this in a theoretically satisfying way. One cannot, for example, simply admit that the dictates of IA can in principle be outweighed by a large enough amount of aggregate value, since we are assuming that any amount of value can be generated by curing sufficiently many headaches. But nor can selective forms of aggregating smaller benefits fit within a framework that requires the pairwise comparison of claims. The only option open to non-consequentialists who are sympathetic to IA is to say that IA is right to stress the moral importance of claims, but then to concede that sometimes relatively weaker claims *can* be aggregated to outweigh

[16] Nagel, 'Equality', p. 118.
[17] Nagel, 'Equality', p. 125.

https://doi.org/10.1017/S0953820814000144 Published online by Cambridge University Press

stronger claims. This of course amounts to saying that some relatively weaker claims are nonetheless relevant to comparatively stronger claims. Such a view would still impose some constraints on promoting the good through aiding, as it would remain impermissible to satisfy a great number of very weak claims when one could instead satisfy one very strong claim. This is more or less the position reached by T. M. Scanlon in what is arguably the most careful and sustained attempt to fit the selective aggregation of weaker claims into an IA model. Scanlon holds that IA (which, following Parfit, he calls the 'complaint model') 'avoids implausible cases of aggregation in what seems, intuitively, to be the right way'.[18] But he realizes that IA also rules out *plausible* forms of aggregation. He therefore endorses a mixed view on which weaker claims can be aggregated if and only if they are deemed relevant to stronger ones.[19] Scanlon does not offer an account of why some weaker claims are relevant to stronger ones while others are not. He suggests only that claims are to be regarded as relevant when ignoring them would not give 'proper consideration' to the individuals whose claims they are.[20]

I worry, however, that things are much worse for Scanlon's account than he realizes. This is because a focus on claims seems to open the door to an argument in support of aggregating many very weak claims over fewer very strong ones. That of course would undermine the rationale for introducing the notion of relevance in the first place. To see the problem, consider what I shall call the 'case of tiny reductions in risk'.[21] Suppose you come upon a very large island nation inhabited by 300,000,001 people. Suppose also that you possess a very versatile medication. In particular, a full dose of the medicine (which is all you have) will cure one person who is about to die of a fatal disease, X, while 1/300,000,000th of the medicine will reduce a healthy person's risk of acquiring and dying from X by one-in-10,000,000 (from a five-in-10,000,000 risk to a four-in-10,000,000 risk). As it happens, there is one person in the nation who is fatally ill with X right now, while the other 300,000,000 are currently healthy but still face a five-in-10,000,000 chance of acquiring and dying from X. You therefore have the choice between (1) curing the one sick person now or (2) preventing thirty (expected) deaths from X in the future by reducing each healthy person's small risk by a further tiny amount. (To give you some

[18] Scanlon, *What We Owe*, p. 241.

[19] Scanlon, *What We Owe*, pp. 239–40.

[20] Scanlon, *What We Owe*, p. 240.

[21] This case was inspired by a case presented by Norman Daniels in 'Reasonable Disagreement about Identified vs. Statistical Victims', *Hastings Center Report* 42 (2012), pp. 35–45, at 41. For a related risk-based case, see S. Reibetanz, 'Contractualism and Aggregation', *Ethics* 108 (1998), pp. 296–311, at 302.

https://doi.org/10.1017/S0953820814000144 Published online by Cambridge University Press

perspective, the risk-reduction that you can give to each healthy person is of the same magnitude as the risk of death incurred by the average American on an average car trip. This suggests that few of the currently healthy people would be willing to make a great sacrifice to achieve that level of risk reduction. Indeed, I will assume that each of them is indifferent between achieving that level of risk reduction and curing the short but unpleasant headache they currently have.) Now let us ask: would it be wrong to use your medicine to reduce many people's already low risk by a further tiny amount instead of saving the life of the one person currently dying of X? I do not think this would be wrong; indeed, I think you might very well be morally *required* to provide the tiny reductions in risk in order to prevent thirty (or so) deaths from X. But if that is true, *then it is permissible after all to promote the good by neglecting one very strong claim in order to satisfy a great many very weak claims*. And if that is right, then it seems we must admit that very weak claims are indeed relevant to very strong claims, and thus that it can be permissible to cure or prevent a great many headaches instead of saving a life. This result would render appeals to relevance pointless.

I can see two ways for a non-consequentialist like Scanlon to evade this line of argument. The first is to deny that this case involves 300 million tiny claims being aggregated to outweigh one very strong claim; instead (the response continues) this case involves 30 (or so) very strong claims conflicting with one very strong claim. This response requires one to hold that the strength of an individual's claim to assistance is a function of the harm he will *in fact* suffer if not helped, rather than a function of his antecedent risk of suffering the harm. Following Johann Frick, I will call this the *ex post* approach to assessing the strength of claims.[22] Since we can say with confidence that in a population of 300 million there will *in fact* be 30 (or so) people who develop X, it will follow on the *ex post* approach that there are 30 (or so) people who have a claim that is *just as strong* as the claim possessed by the one person already suffering from X. On this view, if we assume (as virtually all non-consequentialists do)[23] that claims of the same strength are relevant to one another, then one should provide the 300 million tiny reductions in

[22] See Johann Frick, 'Treatment vs Prevention in the Fight against HIV/AIDS and the Problem of Identified vs Statistical Lives', *Identified v. Statistical Lives: Ethical, Legal, and Medical Perspectives*, ed. I. Glenn Cohen, Norman Daniels and Nir Eyal (Oxford, in press). Unaware of Frick's work at the time, I discussed the *ex post* view in 'Prevention, Rescue, and Tiny Risks', *Public Health Ethics* 6 (2013), pp. 252–61, under the name of 'Brock and Wikler's view' of how to assess the strength of claims.

[23] A notable exception is J. Taurek, 'Should the Numbers Count', *Philosophy and Public Affairs* 6 (1977), pp. 293–316.

https://doi.org/10.1017/S0953820814000144 Published online by Cambridge University Press

risk, as that is your only way to satisfy the largest group of very strong claims.

I do not think this response succeeds, however, because I think that the *ex post* approach to claims should be rejected. To see why, consider the following example.[24] Suppose Alice currently has a 99 per cent chance of dying and that one other person has a headache. And suppose we must choose between curing the other person's headache and eliminating Alice's risk. On the *ex post* view of how to assess the strength of claims, the strength of Alice's claim to assistance is not a function of her antecedent risk, but rather depends entirely on what *will in fact* happen to her if she is not helped. So if we now choose to cure the headache and Alice miraculously does not die, the proponent of the *ex post* view must say that Alice *in fact had no claim at all* to our assistance. But that is not plausible. Surely if anyone had a claim to assistance, it was Alice.[25] I have elsewhere offered other arguments against the *ex post* view (as has Frick),[26] but I think this example suffices to show that in cases involving risks of harm, the strength of one's claim to assistance is a function of one's *antecedent* risk of suffering a given harm. And if that is correct, then in preventing 30 (or so) deaths by providing tiny reductions in risk to 300,000,000 people, one is satisfying 300,000,000 very weak claims, rather than 30 (or so) very strong claims. The case of tiny reductions in risk is, therefore, a case in which it seems permissible to neglect one very strong claim in order to satisfy a great many very weak claims instead. But then that is precisely what happens when one cures a great many headaches instead of saving a life. So in putting the focus on claims and their relative strength, non-consequentialists like Scanlon invite the very sort of aggregation they originally invoked the ideas of claims and relevance to rule out.

The only other way I can see of blocking the conclusion that headaches are relevant to death is to abandon the reliance on claims altogether. Since claims were introduced to help place constraints on promoting the good, and since the case of tiny reductions in risk seems to show that it is permissible to satisfy many very weak claims instead of satisfying one very strong claim, it appears that claims are unable to serve the function for which they were intended. My suggestion, then, is that non-consequentialists should not invoke them at all – or

[24] I deploy this case to make similar point in 'Prevention'.

[25] Of course, both defenders and opponents of the *ex post* view can agree that Alice's antecedent risks gave us good *epistemic* reason to believe that Alice had a strong claim. But that is the *only* role that proponents of the *ex post* view give to antecedent risks, and this is what makes that view so implausible.

[26] See again Kelleher, 'Prevention', and Frick, 'Treatment vs Prevention'.

https://doi.org/10.1017/S0953820814000144 Published online by Cambridge University Press

at least they should not invoke them in cases of aiding (perhaps claims work differently in contexts that involve harming). But with claims out of the picture, the idea of relevance will have to be interpreted as a relation between the *discrete units of goodness* that different actions bring about.[27] On this view, the permissibility of aggregating benefits would turn on two equally crucial factors: (1) the *amount* of good one can produce; (2) the *internal structure* of that quantity of goodness. Thus, if one is able to prevent a death or prevent many headaches, it will not matter how much good one can do by preventing headaches, for that amount of good will simply not have the requisite internal structure: since *the-good-of-a-prevented-headache* is not relevant to *the-good-of-a-prevented-death*, it does not matter how much good one can do by preventing headaches. This view would capture the permissibility of catering to individuals with very weak claims in the case of tiny reductions in risk, and the impermissibility of catering to individuals with very weak claims in a life-versus-headaches case. In providing tiny reductions in risk, one is promoting the good *by preventing deaths*. The resulting amount of good therefore has an internal structure suitable to be a candidate for receiving priority over saving another person from imminent peril. The same cannot be said when one chooses to prevent many headaches instead of saving the life.

Can non-consequentialists live with this proposal? I think they may have to. Of course, many non-consequentialists will not like the suggestion that constraints on aggregation are not grounded in facts about what beneficiaries are entitled to claim, but rather in facts about how benefactors must go about promoting the good. But that seems a small concession, especially since non-consequentialists who believe in relevance are already willing to run roughshod over the strongest claims in order to satisfy many weaker yet still relevant ones. Moreover, consider a two-way conflict case in which we must choose between saving a group of individuals or saving a larger number of others. Many non-consequentialists will already agree with Frances Kamm that 'as the number of people in the larger group who outnumber those in the smaller group increases, there is a *greater* wrong done if we do not save the greater number.' But as Kamm points out, 'there seems to be no individual complaint that can capture this.' So even in the most uncontroversial cases of aggregation – where lives are aggregated to outweigh other lives – the reasonable non-consequentialist must, it seems, be willing to take up 'some point of view outside that of any

[27] The argument from tiny risks therefore provides reasons to join Kamm in talking about (ir)relevant *goods*, rather than (ir)relevant claims. See Kamm, *Intricate Ethics: Rights, Responsibilities, and Permissible Harm* (Oxford, 2007), p. 34.

https://doi.org/10.1017/S0953820814000144 Published online by Cambridge University Press

individual'.[28] My proposal to treat claims as morally inert, and to impose more direct constraints on how the good is to be promoted, simply takes this line of reasoning a bit further.

It is of course standard for consequentialists to hold that the only plausible moral perspective outside that of any individual's is a maximizing perspective from which thoroughgoing aggregation is viewed as unproblematic. But there are alternatives worth exploring. Nagel, for example, recommends a perspective from which agents (or perhaps moral communities understood as collective agents) are to display equal and impartial concern for all others. Here concern for others is construed as concern for them *as individuals*, as 'a separate concern for each person' rather than a concern for 'a conglomeration of them'. He then adds that this form of other-regarding concern 'is realized by looking at the world from each person's point of view separately and individually, rather than by looking at the world from a single comprehensive point of view'.[29] But this last step can be questioned. Consider a parent who loves her five children equally. Suppose she is told that one child, Sarah, has a fatal condition, and that two of the other four will also die soon, with each of the four having an equal probability (<1.0) of ending up in the group of two that will die. If, tragically, the parent must choose between preventing Sarah's death or preventing two deaths among the other four, surely a choice to minimize deaths is consistent with her showing equal concern for each of her children. Yet in light of her greater antecedent risk, Sarah has a stronger claim on her mother's assistance than does any of her four siblings (at least on the view of claims that I have endorsed). So I think Nagel is wrong to hold that a perspective of equal concern requires one to 'look at the world from each person's point of view separately and individually, rather than by looking at the world from a single comprehensive point of view'.

It would, however, be nice if the non-consequentialist could say a bit more about why his or her preferred 'comprehensive point of view' treats paraplegia, but not headaches, as relevant to death. Here I think non-consequentialists would do well to expand upon an observation made by Roger Crisp in the course of defending an 'impartial spectator' approach to distributive ethics. After criticizing traditional utilitarianism for embracing thoroughgoing aggregation (which he calls aggregation 'all the way up'), Crisp writes:

[28] F. M. Kamm, 'Owing, Justifying, and Rejecting', *Mind* 111 (2002), pp. 323–54, at 348; emphasis in original.

[29] Nagel, 'Equality', p. 127.

https://doi.org/10.1017/S0953820814000144 Published online by Cambridge University Press

It is important to note that the objection often made against utilitarians – that they are concerned merely with happiness and not with individuals – cannot be made against the spectator. He or she recognizes the distinction between persons, and, to that extent, [his or her] compassion is here to be understood as a personal, rather than an impersonal virtue.[30]

Here Crisp draws on work by Kai Draper, who writes:

It has often been observed that classical utilitarians value well-being too impersonally. Because they value for its own sake happiness per se, they value persons merely as a means to the end of maximizing aggregate happiness . . . It is safe to say that a genuine compassionate individual will value well-being in a more personal way . . . A genuine compassionate concern for well-being must be distinguished from [a] sort of benefit fetishism.[31]

I think this is a view that can be adopted by the non-consequentialist who now wishes to appraise allocation decisions from a point of view outside that of any individual beneficiary's. She can now say that headaches are not relevant to death because, in entertaining aggregation in that context, one reveals that one is more concerned with (or for, or about) benefits *per se* than she is with (or for, or about) individuals. Something like this idea does seem to capture my own hesitation about letting headaches outweigh death. At any rate, with the appeal to claims no longer on the table, non-consequentialists would do well to explore and expand upon Crisp's and Draper's suggestion that thoroughgoing aggregation is inconsistent with a proper concern *for persons*, in contrast to a concern for bits of welfare.[32] The resulting view will still need to appeal to intuitions, and I admit that it is hard to imagine a theoretical explanation for why one cannot be properly concerned for persons while aggregating headaches when the alternative is preventing one death, while one *can* be properly concerned for persons while aggregating headaches when the alternative is preventing one earache. But the non-consequentialist is obviously not afraid to rely on intuitions or to say that the relevance of a harm or benefit is context-dependent.

V. JOHN BROOME'S ARGUMENT AGAINST IRRELEVANT GOODS

I have argued that the proponent of relevance should declare that claims to assistance (assuming there are such things) are in effect

[30] R. Crisp, 'Equality, Priority, and Compassion', *Ethics* 113 (2003), pp. 745–63, at 757.
[31] K. Draper, 'The Personal and Impersonal Dimensions of Benevolence', *Noûs* 36 (2002), pp. 201–27, at 208.
[32] Although he only makes passing reference to aggregation, Richard Yetter Chappell provides a valuable related discussion in 'Value Receptacles', *Noûs* (forthcoming).

https://doi.org/10.1017/S0953820814000144 Published online by Cambridge University Press

morally inert when it comes to deciding whether and when aggregation is permissible. Claims therefore cannot do the aggregation-blocking work that many non-consequentialists ask of them. Whatever work they *appear* to be doing must be explained by some other theoretical element or principle. Perhaps, however, this apparent uselessness of claims stems from a faulty account of how claims work. Such an objection can be found in John Broome's work. According to Broome, claims *never* work as constraints blocking all appeals to aggregative considerations. Broome therefore rejects the notion that some goods or claims are ever irrelevant. According to him, there is always a *pro tanto* requirement of fairness that individuals' claims be 'satisfied in proportion to their strength'. More specifically, Broome holds 'that equal claims require equal satisfaction, that stronger claims require more satisfaction than weaker ones, and also – very importantly – that weaker claims require some satisfaction. Weaker claims must not simply be overridden by stronger ones.'[33] How, then, can claims be satisfied in proportion to their strength in the sort of two-way conflict situation on which I have been focusing? Broome's answer is that they cannot, since by hypothesis only one of the groups can be helped. In such cases, Broome says that the best we can do is to provide everyone with 'a sort of surrogate satisfaction' in the form of a lottery.[34] By holding a suitable lottery, we can ensure that those with equally strong claims receive equal chances of being helped, and that those with claims of different strengths receive differently sized chances. This is not, however, to say that holding suitably weighted lotteries is always the *right* thing to do, only that holding such lotteries would be the *fair* thing to do. Whether one must do the fair thing depends, Broome claims, 'on how important fairness is in the circumstances'.[35] And Broome is explicit that the aggregate sum of value associated with helping one group can be so large that it justifies directly helping that group – even if that sum is generated by curing a great many headaches:

> I believe that a lot of small benefits can add up to be as important as one large benefit. When a patient in a United Kingdom hospital gets a headache, he or she is given an analgesic. Over a few years, the UK Health Service gives out a few million analgesics to cure headaches. The cost of all these pills adds up, and eventually it will amount to more than enough to save someone's

[33] J. Broome, 'Fairness', in his *Ethics out of Economics* (Cambridge, 1999), pp. 111–22, at 117.
[34] Broome, 'Fairness', p. 119.
[35] Broome, 'Fairness', p. 120.

https://doi.org/10.1017/S0953820814000144 Published online by Cambridge University Press

life. Evidently, the Health Service thinks that curing all those headaches is as valuable as saving a life. I agree.[36]

Broome's discussion and his Health Service example raise two important questions for the approach to aggregation that I sketched in the previous section. First, is Broome correct to retain a focus on claims by arguing that fairness (but not necessarily rightness) requires their proportional satisfaction, rather than satisfaction in order of strength (as the competing IA approach has it)? Second, is Broome correct that the Health Service example illustrates that 'all goods are relevant' to life-saving, even the good of headache-alleviation? In my view, these questions must be sharply distinguished, for I believe that Broome's account of claims is implausible, whereas I think that his Health Service example is perhaps the strongest argument in the literature for the view that all goods, including headache-alleviation, are relevant to life-saving.

The main problem with Broome's account of fairness and of how claims 'work' is nicely illustrated by Brad Hooker:

> Suppose your claim on the medicine comes from the fact that you need it to save your life, and my claim on it comes from the fact that I need it to save my little finger. Suppose an average life is something like a thousand times more important than my little finger. So should the matter of who gets the medicine be decided by a lottery in which you have a 999/1000 chance of winning and I have a 1/1000 chance? Given that your claim is so much stronger than mine, how could it be right to take *any* risk that I rather than you might end up with the good?[37]

Here Broome might agree with Hooker that the life should be saved straight away (without holding a lottery); Broome could simply say that since a life is so much more valuable than a mere finger, saving the life straight away is the *right* thing to do all things considered (even if it is not the *fair* thing to do). But as Hooker goes on note, 'Letting the stronger claim win [without first holding a weighted lottery] seems *completely* fair.'[38] And that is what is so difficult to accept about Broome's account of claims and fairness: he must say that when choosing between saving a life and preventing even *one* mild headache, it is *unfair* to save the life straightaway without first holding a weighted lottery. But that is hard to believe. It is not plausible to say that in directly saving the life, one has done *the unfair* thing. If Hooker

[36] Broome, 'All Goods Are Relevant', *Summary Measures of Population Health*, ed. C. J. L. Murray, J. A. Salomon, C. D. Mathers and A. D. Lopez (Geneva, 2002), pp. 727–9, at 728.

[37] B. Hooker, 'Fairness', *Ethical Theory and Moral Practice* 8 (2005), pp. 329–52, at 349.

[38] Hooker, 'Fairness', p. 349; emphasis added.

https://doi.org/10.1017/S0953820814000144 Published online by Cambridge University Press

and I are right about that, then it means that fairness does *not* require proportional satisfaction of claims (at least not in every case), and thus that Broome's attempt to vindicate the indispensability of claims fails.

What, then, should we say about Broome's Health Service example? The force of that example does not hinge on the success of Broome's view of claims and fairness. Even if one were sympathetic to a Crisp- and Draper-inspired 'proper concern for persons' view of why headaches are not relevant to death, Broome's Health Service example should still give one pause. For I suspect many find the Health Service's policy permissible, and it is hard to see why it should be permissible to aggregate headaches in that context but not in the two-way conflict cases I have been discussing. Consider, for example, a rather unsatisfying explanation from Frances Kamm. In response to the related observation by Dan Brock that many people find it permissible to fund public Botanical Gardens for the sake of minor aesthetic pleasures instead of an extra ICU bed for life-saving, Kamm writes:

> [Brock] suggests that I might make my claim normative rather than descriptive; what we do is just wrong. But there may be other explanations for our actual investments. We may not exclude even the most trivial concerns from receiving some of our resources, because the resources are *divisible* . . . So, problems may receive divisible resources in proportion to their importance, but also in proportion to the probability of our resources successfully dealing with the problem . . . Even with divisible resources, I still think it should require a much greater number of people experiencing headaches and a much higher probability of curing them to appropriately invest as much in headache cures as in cures for fatal diseases.[39]

It is not clear to me why the (in)divisibility of one's resources for assistance should make a difference to whether the headaches one can cure are relevant to the deaths one can prevent. But let us nevertheless confine our attention to cases in which resources are divisible. Here, Kamm suggests that headaches are in fact *not* irrelevant: 'So, if we do not face the choice of either giving all of our money to curing a fatal disease that hits a few people or saving many from headaches or withered arms, we could give some to each cause.'[40] And yet in direct response to Broome's Health Service case, Kamm writes:

> [S]uppose that we did argue even for the permissibility of investing in cures for truly minor problems affecting many, such as headaches, rather than in a cure of a rare fatal disease . . . This does not imply that here and now we should not save someone from dying from the rare fatal disease, if we could, rather than cure millions of headaches. For example, suppose that, surprisingly, giving someone who develops the fatal disease all of the aspirin that has been

[39] F. M. Kamm, 'Replies', *Philosophy and Phenomenological Research* 58 (1998), pp. 969–74, at 973–4.

[40] Kamm, *Intricate Ethics*, p. 47 n. 74. In context it is clear that by 'could' Kamm means 'could permissibly'.

https://doi.org/10.1017/S0953820814000144 Published online by Cambridge University Press

produced to cure headaches could still now save him. It could be wrong to leave him to die . . . It is here and now that the irrelevant utilities of headache cures do not aggregate to override saving a life.[41]

It is hard for me to see why Kamm's view is not straightforwardly incoherent. After all, if it really is permissible to *fund* aspirin for headaches in the way the Health Service does, surely that means it is permissible to *actually use* that aspirin to cure headaches even when it is possible (say) to sell the aspirin to pay for an extra dialysis machine that could save a life now.

In my view, since the Health Service policy of funding and providing aspirin seems entirely permissible, the only way to rebut Broome's conclusion that 'all goods are relevant' is to argue that the Health Service case does not actually involve the sort of aggregation that the non-consequentialist rejects. To see how such an argument might go, consider one real-world context in which it is obviously permissible to balance death against the prevention or cure of many headaches: viz. when one drives to the store to buy a bottle of aspirin. By driving, one risks death. But, typically, one accepts this risk as the price one must pay to be in a position to relieve relatively minor aches and pains. Following Scanlon, let us call this form of aggregation – i.e. the aggregation of many relieved aches and pains within one's own life – *intrapersonal* aggregation.[42] Intrapersonal aggregation is aggregation within a single life, rather than aggregation between or across different lives (which I have been calling *inter*personal aggregation). The two-way conflict case we have been discussing quite clearly involves interpersonal aggregation. There, we are asking whether it is permissible to let the headaches of many different people outweigh saving one other person's life. I believe that if one is to deny that the Health Service case justifies the interpersonal aggregation of headaches over lives, one must argue that the Health Service case instead involves intrapersonal aggregation. More specifically, one must argue that funding aspirin is permissible *only to the extent that* it involves permissible intrapersonal aggregation. Otherwise the example really would establish the permissibility of thoroughgoing interpersonal aggregation.

How might such an intrapersonal aggregation argument go? We might imagine that the Health Service reasonably assumes that each citizen would be willing *ex ante* to trade guaranteed headache-relief (when in hospital) against the small risk that he or she will be denied whatever life-saving treatment the aspirin funds displace. Because not

[41] Kamm, *Intricate Ethics*, pp. 36–7.
[42] Scanlon, *What We Owe*, p. 237.

https://doi.org/10.1017/S0953820814000144 Published online by Cambridge University Press

much life-saving could be achieved with the aspirin budget, and because several different life-saving programmes would be good candidates for the funds if the budget were redirected, no single individual who faces death in the hospital can say with any confidence that *his* chances of living would have been appreciably greater if aspirin were not funded. So the Health Service might reasonably suppose that it is rational, from each British citizen's standpoint, to trade guaranteed headache cures against a very small increase in the risk of dying in the hospital.

This rationale is obviously not the rationale Broome ascribes to the Health Service in his formulation of the example. As he presents it, the Service's goal is straightforwardly to maximize the value it can achieve with its limited funds. But if the non-consequentialist is right to reject thoroughgoing interpersonal aggregation, the Health Service cannot permissibly invoke the rationale Broome has in mind. Its rationale will instead have to make essential reference to intrapersonal aggregation. In doing so, it will have to draw upon the value of autonomy – of giving people what they themselves prefer (or perhaps what they reasonably prefer) – rather than exclusively drawing upon a benevolent concern to promote well-being (where that can be interpreted in either the maximizing consequentialist manner or in the 'proper concern for persons' manner suggested by Crisp and Draper). Of course, many non-consequentialists might be perfectly happy to let policy be shaped both by the demands of benevolence and by a requirement to respect autonomy. But the hard part comes in explaining how intrapersonal aggregation and respect for autonomy can justify a policy that would be impermissible if it were defensible solely by reference to thoroughgoing interpersonal aggregation.

To appreciate the challenge, consider a case in which the demands of benevolence arguably trump the requirement to respect autonomy. (I borrow the following science-fiction case from Tom Dougherty, who reports that it originally comes from Michael Otsuka.)

> There is a giant roulette wheel in the sky. It has a billion slots that are connected to chutes. The chutes are directed at individuals who have swallowed poison. The poison will make each individual have a headache and die, if untreated. To make the wheel work, you have to press a button. This will give the wheel an indeterministic spin, and then release the contents of each slot into a chute. Before you do so, you have to make an antidote to put in the slots. You only have enough resources to choose between these options:
>
> Vials for Everyone. You make a billion vials of an antidote that will certainly save each recipient's life, but do nothing for each recipient's headache.
>
> Vials That Cure Headaches. You make a billion minus one vials of an antidote that will certainly save each recipient's life and cure this person's headache.[43]

[43] T. Dougherty, 'Aggregation, Beneficence and Chance', *Journal of Ethics & Social Philosophy* 7 (2013), pp. 1–19, at 17.

https://doi.org/10.1017/S0953820814000144 Published online by Cambridge University Press

It would not be surprising to learn that each of the billion individuals prefers Vials That Cure Headaches to Vials for Everyone. After all, virtually everything we do in life carries *some* risk of death, and a 1/1,000,000,000 chance of dying does not seem too high a price to pay for headache-relief. (If one does find it too high, one can simply adjust the example to involve more people, more vials and more chutes.) This suggests that if individual autonomy were the only value at stake here, one should choose Vials That Cure Headaches. But I suspect many people will share my (and Dougherty's) intuition that it would be wrong to choose Vials That Cure Headaches. Here, a concern for the fate of each individual – and especially for the individual who will die under Vials That Cure Headaches – seems to pull decisively in favour of Vials for Everyone. So here we have a case in which the demands of benevolence conflict with demands of autonomy, and benevolence arguably wins. How, then, could a non-consequentialist defend the Health Service's aspirin policy by invoking the idea of intrapersonal aggregation? Why should the Health Service's policy not also be criticized as flouting an overriding demand of benevolence?

I believe the answer will have to do with the fact that the Health Service's aspirin protocol can be embedded in a larger framework of policy that strikes a reasonable balance between benevolence and respect for autonomy. I take it as uncontroversial that the Health Service makes many allocation decisions that are quite clearly motivated by benevolent concern for the lives and well-being of the individuals it serves. Among these is the initial choice to fund any ICU beds at all, as well as the choice to pay for various forms of cancer screening and treatment. (If we moved to an even higher-order level of policymaking, we could also cite here various public health measures such as sanitation, clean water and pollution controls.) With this background of benevolent policy in place, it seems to me permissible for the Health Service to turn *some* of its attention to the risk-taking preferences of the individuals it serves. That is, when the Health Service can make a good case that it is displaying *enough* benevolent concern for the lives and well-being of those it serves, it can then permissibly add a layer of policy aimed at satisfying individuals' autonomous preferences for risk-taking. Compare this appraisal of the Health Service case to the Roulette Wheel case. Unlike the Health Service's policymakers, the decision-maker in the Roulette Wheel case cannot point to any background choice or policy that embodies her benevolent concern for the billion individuals who are now the targets of her choice. Rather, *this* is her opportunity to display proper benevolence. Here it seems reasonable for her to choose benevolence over respect for preferences, both because a life is at stake and because

https://doi.org/10.1017/S0953820814000144 Published online by Cambridge University Press

her choice of Vials for Everyone does not seriously injure or curtail the autonomy of any individual. After all, each will continue to have his own life to lead after her decision is made, and they can take all the risks they want after she is out of the picture. Admittedly, if the decision-maker were repeatedly faced with this choice day in and day out, her situation might well become more and more like the Health Service's situation; at some point it will be reasonable for her to say, 'I have saved these people's lives time and again; it is time to give them what they want.' But until that point arises, she is obligated by demands of benevolence to neglect the individuals' preferences and prevent death. That is simply what benevolence, properly balanced against respect for autonomy, requires, and that is why the intrapersonal rationale can license headache-alleviation in the Health Service example but not in the Roulette Wheel example.

Let me stress that I am not claiming that the Health Service *in fact* justifies its aspirin policy by citing reasons of intrapersonal aggregation; nor am I claiming that others judge the policy to be permissible because they implicitly see that it can be presented as a case of intrapersonal aggregation (rather than interpersonal aggregation). I am instead seeking to explain why the Health Service's policy is in fact permissible. It is permissible, I claim, only to the extent that it is a case of permissible intrapersonal aggregation. I understand that some people will judge the Health Service policy to be (1) permissible and (2) a case of interpersonal aggregation, and will therefore accept that (3) there is some number of headaches that outweigh saving a life in an interpersonal two-way conflict case. But in light of how counter-intuitive (3) is, I suggest we have strong reason to hold that independently reasonable intrapersonal considerations offer the best account of the permissibility of the Health Service's policy. I therefore think Broome is wrong to offer his example as establishing that 'all goods are relevant' or (what comes to the same thing) that the notion of relevance is pointless.

VI. CONCLUSION

My goal in this article has been to clarify and refine the non-consequentialist notion of relevance by engaging with three lines of argument against it. Obviously I am very sympathetic to the view that benefits and harms can be relevant and irrelevant to one another, and that there are indeed some hard constraints against certain kinds of interpersonal aggregation. But throughout my main aim has been to articulate the non-consequentialist's best response to each of the three challenges I have considered. I have been surprised to see where this has led, especially with regard to the idea that claims to assistance are

https://doi.org/10.1017/S0953820814000144 Published online by Cambridge University Press

morally inert in the context of deciding whom to help. That conclusion in particular seems very difficult for a non-consequentialist to accept, but I am willing to entertain it because I do not see another way to block the route to thoroughgoing interpersonal aggregation that is opened up by what I called the case of tiny reductions in risk. In that case, it does seem perfectly permissible to promote the good by catering to those with exceedingly weak claims, even as the very strong claim of someone in dire peril is ignored. I think it is important here to stress that most non-consequentialists already accept a great deal of aggregation that does not easily mesh with any plausible framework that cares deeply about giving competing claims a fair hearing. That, after all, is why non-consequentialists needed the idea of relevance in the first place. One tentative conclusion, then, is that if non-consequentialists wish to hold onto a meaningful notion of relevance, they must focus more of their attention on developing a decidedly 'supply-side' morality concerned with the proper promotion of goodness by benefactors, and less attention on developing 'demand-side' principles concerned with the fair or even-handed treatment of prospective beneficiaries. I think this would certainly drive non-consequentialist moral views further in the direction of consequentialism, and that will of course make many non-consequentialists nervous. But I think that if there is one thing I have shown decisively, it is that we non-consequentialists should already be rather nervous in embracing an idea as slippery and mysterious as relevance. Perhaps a little more consequentialism is a price worth paying for a little less mystery.[44]

jkelleher@wisc.edu

[44] For helpful comments and discussion, I thank Paul Audi, Thomas Dougherty, Molly Gardner, Dan Hausman, Gordon Hull, Justin Klocksiem, Win-chiat Lee, John C. Moskop, Peter Nichols, Robert Streiffer, and audiences at Amherst College, the University of Nebraska-Omaha, the 2012 Rocky Mountain Ethics Congress in Boulder, Colorado, and the 2013 Junior Scholars in Bioethics Workshop at Wake Forest University.

https://doi.org/10.1017/S0953820814000144 Published online by Cambridge University Press