Aggregation and Two Moral Methods

F. M. KAMM
Harvard University

I begin by reconsidering the arguments of John Taurek and Elizabeth Anscombe on whether the number of people we can help counts morally. I then consider arguments that numbers should count given by F. M. Kamm and Thomas Scanlon, and criticism of them by Michael Otsuka. I examine how different conceptions of the moral method known as pairwise comparison are at work in these different arguments and what the ideas of balancing and tie-breaking signify for decision-making in various types of cases. I conclude by considering how another moral method that I call virtual divisibility functions and what it helps reveal about an argument by Otsuka against those who do not think numbers count.

Does the number of people we can help count morally in deciding what to do in conflict situations when we cannot help everyone? In this article, I shall revisit this question. I shall first present some arguments I have presented elsewhere, but expand on them in various ways. Then I shall consider new arguments for and against numbers counting in the course of which two moral methods – pairwise comparison and virtual divisibility – will be examined.

I

Among those who argue that numbers do not count is John Taurek. That is, he claims that if we can save either A or a group of B, C, D, and E, there is no reason to save the greater number *per se*. (I shall assume, unless noted otherwise, that there are no morally relevant differences between individuals in the cases to be discussed and that groups between which we must choose have no overlapping members.) Two major premisses in his argument are that (1) no one of the people in the larger group will suffer a greater loss than the one person would; and (2) it is not true that we produce a better outcome if the greater number survive, as there is no impartial perspective from which to judge this; it is better for the one for him to survive and better for the five for them to survive. (It is only better for each of the five that the five survive as this is a way to his own survival, assuming they are strangers.) His conclusion is, if we wish to show equal concern and respect, we should give everyone an equal chance to survive by, for example, tossing a coin.

Among the implications of Taurek's view that numbers do not count that have not, to my knowledge, been emphasized are the following.

© 2005 Cambridge University Press
DOI: 10.1017/S0953820804001372
Utilitas Vol. 17, No. 1, March 2005
Printed in the United Kingdom

https://doi.org/10.1017/S0953820804001372 Published online by Cambridge University Press

(a) Suppose a trolley is headed towards killing five people and it is possible for a bystander to turn it away from them and towards a track where one person will be killed instead. If numbers do not count, this should be treated as if it were a case of one person on each track. If we do not favour letting die over killing, we might then toss a coin to determine if we let five die or turn the trolley to one. Similarly, if we do not favour letting die over killing, it should be just as permissible to toss a coin and, if called for, turn a trolley that is headed towards one person away from him and onto a track where five will be killed. Those who think numbers do not count may, of course, combine that view with a moral distinction between killing and letting die and then avoid these results. But suppose a trolley is at a crossroads and if it stays there it will set off a nuclear weapon that will destroy civilization. We could avoid this by turning the trolley onto a track where one person will be killed or onto a track where five people will be killed. Those who do not think mere numbers of people count could still think that a qualitative difference such as destruction of civilization counts. Hence, it seems they should toss a coin to decide whether to kill one or to kill five if the trolley were at a crossroads.[1]

While those who think that numbers count would say that someone would make the morally wrong decision if he saved one person instead of a different group of five, I do not think that they would necessarily say that we must force such a person to save the five just because he is going to be saving someone anyway. However, I believe that those who think numbers count would say that we should interfere, even by force, with someone who is about to turn a trolley away from killing one person to killing five instead or about to send a trolley to kill five when he could send it to kill one instead, even if in doing this he is carrying out a fair toss of a coin. Dealing with numbers in the context of killing rather than saving may make clearer the contrast between the numbers and anti-numbers camps.[2]

(b) Suppose that in one part of the country someone is about to save a group of five people rather than save one different person when it is impossible to save everyone, and he does this without giving both groups an equal chance. He merely counts the numbers. In another part of the country someone is about to save a group of five people rather than save two different persons when it is impossible to save

[1] I am grateful for discussion of my cases here to Michael Otsuka.

[2] As we shall see, Elizabeth Anscombe thinks that numbers can be a reason to save one group over another, but not a reason that one *must* act on, so long as one will save someone in a non-overlapping group. (See her 'Who Is Wronged?', *The Oxford Review* 5 (1967), p. 16.) Would she also think that there is no definitive reason to harm fewer people so long as one will harm someone in a non-overlapping group?

https://doi.org/10.1017/S0953820804001372 Published online by Cambridge University Press

everyone, and he does this without giving both groups an equal chance. He merely counts the numbers.

Suppose numbers should not count in these two circumstances. Then these two decision-makers each behave incorrectly. Suppose we can go only to one but not both of them in time to reason with him so that he gives everyone he is dealing with an equal chance. If the first decision-maker does not give an equal chance, one person will have been treated unfairly in a way that is possibly bad for him. If the second decision-maker does not give an equal chance, two people will have been treated unfairly in a way that is possibly bad for them. But if numbers do not count, this cannot be a reason to go to the second decision-maker. Rather we should toss a coin to decide where fairness will occur.

I do not raise these points as arguments against those who do not think numbers count, only to point out some implications of their position that I think are striking.

Now, let us evaluate Taurek's argument. Taurek's Premiss (2) should be eliminated as a premiss in an anti-numbers argument, I believe. One reason is that the view that there is no objective perspective from which to judge outcomes yields results that even those who reject counting numbers would find unacceptable. For, as Taurek notes, Premiss (2) also implies that from A's perspective the saving of his leg could be a better state of affairs than the saving of B's life. Yet, I think, even those who think that numbers do not count are not thereby committed to tossing a coin between saving A's leg and B's life. This implies that they are imposing, from some impartial perspective, at least some objective comparison of what A and B stand to lose and so of the outcome in which A does not lose a leg versus the one in which B's life is saved. It is only if A and B stand to lose something equal or comparable (and perhaps attain something equal or comparable) that they will toss the coin, other things being equal. We can all still recognize that from A's perspective it is not as good for A if B survives as if A does and this is part of what leads us to want to give A his equal chance. I have called such a mixing of objective and subjective perspectives 'Sobjectivity'.[3]

Indeed, Taurek's own Premiss (1) suggests that we are engaging in a pairwise comparison, individual by individual, to see if anyone will suffer a loss greater than anyone else, for if someone will, we would not toss a coin. This presupposes the more objective view of losses I have just described and is compatible with the rejection of Premiss (2).

However, this degree of commitment to objective comparison does not, by itself, I think, commit one to the rejection of Premiss (2). As

[3] I first described Sobjectivity in 1987 ('The Choice between People, Common Sense Morality, and Doctors', *Bioethics* 1 (1987), p. 255), and again in *Morality, Mortality*, vol. 1 (New York, 1993).

https://doi.org/10.1017/S0953820804001372 Published online by Cambridge University Press

noted, taking into account whether one person will suffer a loss greater than another is consistent with taking to heart that when the losses are equal, A still prefers that he not suffer his loss rather than that B not suffer his equal loss and vice versa. For this reason one might toss a coin between them. And if no one of a greater number will suffer a greater loss than A would, we cannot refuse to toss the coin simply on the ground that A's loss would be much less than anyone else's, because it will not be less.

However, the following argument, which I call the Argument for Best Outcomes, suggests how we might argue that it is worse if more people die and so reject Premiss (2): (1) Using Pareto Optimality, we see that it is worse if both B and C die than if only B dies, even though it is not worse for B. That is $B + C < B$. (2) A world in which A dies and B survives is just as bad as a world in which B dies and A survives. This is true from an impartial point of view, even though the worlds are not equally good for A and B. That is, there is moral equivalence in the death of A or B. (3) Given (2), we can substitute A for B on the right side of the moral equation in (1) and get the result that it is worse if B and C die than if A dies. That is, if $B + C < B$ and $A = B$, then $B + C < A$. Alternatively, we can substitute A for B on the left side of the moral equation in (1) and get that $A + C < B$.

Even if Taurek's Premiss (2) is wrong and it would be worse that B and C die than that A dies, this does not necessarily mean that it is right for us to save B and C rather than A. We cannot automatically assume that it is morally permissible to produce the best state of affairs for this may violate justice or fairness. Some might claim that if we save B and C on the basis of (3), we abandon A to save the greater number without giving him a chance and not giving equal chances is unfair. They might object that we could generate an intransitivity: where '>' means 'clearly ought to be saved', and '=' means 'equally permissible to save', $B + C > B$, $A = B$, but $-(B + C > A)$ (or $-(A + C > B)$) because it could be unfair to deprive someone of his equal chance to be saved, a factor not present if we act on $B + C > B$ or toss a coin because $A = B$. I shall return to this issue below.

II

Unlike Taurek, Elizabeth Anscombe argued that numbers can be a reason for saving some people rather than others but not a reason one must act on. We may save A rather than B and C or we may save B and C rather than A. She believes that A's need gives us a reason to save him, and saving the greater number can also give one a reason to save B and C. But when one has a reason to do X and a reason to do Y, one need not have a further reason to do X rather than Y, she claims. Let us consider

https://doi.org/10.1017/S0953820804001372 Published online by Cambridge University Press

Anscombe's argument more carefully. The claim that one need not have a reason to do X rather than Y, when one has a reason to do X and a reason to do Y, suggests that when someone could save A and B rather than just save A alone, he does nothing wrong if he just saves A, as he has a good reason (satisfying need) to do that. This seems clearly wrong. It is not merely that satisfying need implies that it is wrong to save no one. It also implies that it is wrong to save A alone when one could also save B. But nothing Anscombe says explains this and what she does say, literally, suggests that one does not need a justification for ignoring all the reasons that there are so long as one has some reason to act. Further, the reason of satisfying need is present for both saving A and for saving the greater number. But if this reason is present for saving either side, someone might ask, why not satisfy it along with a reason (even if it is not a conclusive reason) for saving the greater number?

This brings me to the next point: Anscombe's view that taking numbers into account is a possible reason that we could allow to determine our choice implies that it is not unfair (or otherwise wronging) to A not to give him an equal chance when there is more than one person on the other side. Anscombe explicitly says that she thinks no individual considered as an individual in the larger group is wronged if we save A instead of him. We now also see that she thinks that A is not wronged, if we do not give him an equal chance. (This contrasts with the view we have examined that the reason numbers might not be any reason at all when conflicts are present is that it conflicts with giving everyone an equal chance.)

If it is wrong to save only A when one can save A and B, this implies that satisfying more need is better, at least when no one loses out on an (maximal) equal chance. But, as argued above, if numbers can be a possible reason in conflict cases, this implies that failure to give equal chances is not a wrong-making feature. When these two claims are combined, it might seem that they imply that satisfying more need is better (i.e. we should save the greater number) in a way that rules out saving the one. Only if Anscombe thought that giving someone an equal chance were at least another possible reason for action, would one still have permission to save the one rather than the greater number.

III

Hence, it is important to see if it really is wrong to produce the best outcome in conflict cases because someone is wronged in not being given an equal chance. It is also important (especially in answering Anscombe) to see if someone is wronged, if one does take giving equal chances as a possible reason on which one can act. Here are two arguments against its being wrong not to give equal chances, the second

https://doi.org/10.1017/S0953820804001372 Published online by Cambridge University Press

of which also implies that it is wrong to give equal chances (unless the greater number requests this) because someone in the group with more people will be wronged.

(i) The Consistency Argument indirectly shows that in saving the greater number of people, we need not be overriding fairness or justice. In many other cases, non-consequentialists who think that numbers count will not violate justice in order to save the greater number. For example, they will not ordinarily kill one person even in order to save a thousand. Moreover, they (arguably) would not deprive a teacher of a chance for an organ transplant simply because a doctor, who alone can save the lives of four others, also needs the organ. Why would such non-consequentialists, who refuse to sacrifice justice or fairness in order to save more lives in these other cases, override justice or fairness even to save two people rather than save one? It is most reasonable to believe that they would choose to save the two rather than the one because fairness or justice is not being overridden in such a case. That is, those who are sensitive to issues of fairness and justice do not think fairness or justice requires that we give A a chance in a conflict with B and C.

(ii) Second, what I call the Balancing Argument claims that in a conflict, justice demands that each person on one side should have her interests balanced against those of one person on the opposing side; those whose interests are not balanced out in the larger group help determine that the larger group should be saved. Those who are balanced in the larger group are saved instead of those in the smaller group and thus substituted for them. If we instead toss a coin between one person and any number on the other side, giving each person an equal chance, we would behave no differently than if it were a contest between one and one. If the presence of each additional person would make no difference, this seems to deny the equal significance of each person. Thus, justice does not conflict with producing the best outcome when different numbers of lives are at stake.

To answer Anscombe, the Balancing Argument claims that any individual who remains unbalanced can complain that he as an individual is wronged, if his presence does not make a difference to the outcome (when this is in his interest) when all individuals stand to lose their lives. This argument also implies that Anscombe would be wrong to say that one may take giving equal chances even as a possible reason on which to act, for this would conflict with the balancing that we owe each person. Hence, aggregating the interests of many people, thereby producing the greatest good, might be required, but for the distinctly non-consequentialist reason that what we owe each individual is to weigh him against an opposing equal and allow the remaining person's weight to count as well, rather than because we have a duty to produce the greatest good.

https://doi.org/10.1017/S0953820804001372 Published online by Cambridge University Press

IV

Thomas Scanlon is a contractualist.[4] He insists on what he calls the Individualist Restriction (IR) in his version of contractualism, namely that only an individual, on behalf of himself or another individual, can register complaints against a principle proposed for governing our relations with each other. However, he wants to be able to argue that when we have a choice between saving a smaller group of people and saving a larger group of people whose members do not overlap, we should save the larger group, other things being equal. (While Scanlon rejects Taurek's second premiss (that we would not produce the better outcome in saving the greater number in a conflict situation), he does not wish to rely on the fact that the outcome is better, within a contractualist framework, as an argument for saving the greater number. Similarly, the Balancing Argument should work independently of considering that we produce the best outcome in saving the greater number.) He claims that he can adhere to the IR and still count numbers of people because counting numbers need not involve considering the complaints of a group against a principle that ignores numbers. Rather, this justification for counting number is what he calls the Tiebreaker Argument.

Consider a conflict between saving A and saving B and C, each of whom would otherwise lose his life. Scanlon claims that if A were alone, A's weight would be recognized by saving A. If he were in conflict only with B, we would have to recognize B's weight as well, and hence modify our allocation procedure by tossing a coin. When C is present, if we still only tossed a coin giving everyone an equal chance, C's presence would make no difference to what we do. But, Scanlon thinks, C as an individual could complain about this on his own behalf. This is because the weight of A and B could be recognized and fully taken account of, if we allow them to balance each other, and C's weight as another person needing to be saved could be taken account of by making him the tiebreaker between A and B. The Tiebreaker Argument is like the Balancing Argument[5] (as Scanlon has acknowledged).

Notice that because A's presence is said to be dealt with if it is balanced against B's (even though A gets no chance of having his life saved), it turns out that in the context of A tying with B, it is C (and not A or B) who would ultimately have a complaint based in fairness if he did not affect the outcome. (Though B would be another beneficiary of the satisfaction of C's fairness complaint, he could not

[4] Scanlon's complete theory is laid out in his *What We Owe to Each Other* (Cambridge, Mass., 2000).

[5] First in 'Equal Treatment and Equal Chances', *Philosophy & Public Affairs* 14 (1985), p. 177, and then in *Morality, Mortality*, vol. 1.

https://doi.org/10.1017/S0953820804001372 Published online by Cambridge University Press

directly complain on his own behalf, on grounds of fairness, if C does not affect the outcome when it is in C's interest to do so.) Notice also that if we think that the Balancing and Tiebreaking Arguments show that numbers count, we still need not think that in all situations when life and death are at stake for everyone and we have to choose between a non-overlapping smaller and larger group, we should choose the larger group. For example, suppose that we have a scarce organ to distribute. It is necessary either for transplantation to save the life of a teacher or for creating a serum that alone can cure a terrible headache of a doctor who is the only one who can do surgery to save four patients. Even though we face a choice of doing what will save four lives or doing what will save only one, the fact that the teacher would be excluded on instrumental grounds (i.e. he cannot do surgery) unrelated to the best direct use of the particular scarce resource we have might be a reason not to do what saves the greater number.

V

Scanlon's IR is a descendant of the moral method of pairwise comparison used by many non-consequentialists and also suggested by Taurek's Premiss (1). But there seem to be different understandings of how pairwise comparison should operate, though this is not commonly noticed. One view suggests that in conflict scenarios we should search to see whether someone is worse off than anyone else and also whether someone would get a bigger benefit than someone else. Some might say that in doing pairwise comparisons we should search only for baselines (i.e. how badly off someone would be if not helped) and not consider how big a benefit is at stake for people at the same baseline. This is one difference in views about pairwise comparison. Let us suppose both baseline and benefit should be compared one person at a time, pairwise in conflict situations.

According to the next view of pairwise comparison, which I shall call the Context-Aware View (for reasons to be explained below), once we find a match of baseline and benefit, the person on one side for whose match we were searching is no longer compared with anyone else. He has met his match and is in a sense silenced relative to others. An alternative view, which I shall call the Blinker View, involves pairwise comparing a person from one side until we find someone on the other side who will be worse off and/or get a bigger benefit.[6] It is only when we find *more* than someone's match that he is no longer pairwise compared with individuals on the other side. This is the second difference in views about pairwise comparison. On the latter view, if one is equal to anyone

[6] US English for 'blinker' is 'blinder'.

https://doi.org/10.1017/S0953820804001372 Published online by Cambridge University Press

on the other side, and all on the other side are equal, one is equal to everyone on the other side. Hence a coin should be tossed. When there are four equal persons (same baseline, same benefit), two on each side, the Context-Aware View of pairwise comparison tells us to toss a coin because each is matched with an equal and opposite number. The Blinker View also tells us to toss a coin, but only because at least one person (and, in fact, each person) on one side should have an equal chance against *both* of those on the other side. I think that the Context-Aware View of pairwise comparison is involved in Scanlon's Tiebreaker Argument and my Balancing Argument and that the Blinker View is involved in Michael Otsuka's criticism of these two arguments.[7]

Let us consider these two views in more detail. This should make clear why they bear these names.

Imagine a group of a thousand individuals who are in competition with another group of 900 individuals for an education, where all are equally needy and capable of as good an outcome. We pair off each of the 900 on one side with each of the 900 on the other. There are people on one side with no matches on the other, each as needy and capable of benefit as the original 900. Because of this, according to the Balancing and Tiebreaker arguments, we should balance the two sets of 900 individuals and break the tie between them not by tossing a coin, but by doing what, in addition to helping the 900 on one side, will help 100 more people. Notice that in this way of doing pairwise comparison, when we are comparing any two equal sets of individuals, we allow ourselves to be aware of the context in which they are situated, hence the title Context-Aware View. That is, we are aware that there are further people as well, and *this is what leads us to balance out the members of a set rather than make a random choice (e.g. toss a coin) between them*. Balancing 'silences' the people balanced; they are not further compared with any other individuals.

Contrast this understanding of pairwise comparison with the interpretation of how to do pairwise comparison offered by the Blinker View. We take individual A from one side, and pairwise compare him with someone on the other side, but we do so 'with blinkers on' as to the context of these two individuals. If they are equal in all morally relevant respects, we would see no reason not to toss a coin to decide between them. We then take the blinkers off before tossing the coin, becoming aware of the context, and if we find another individual on the side opposite to A's with whom to compare A, we again compare

[7] See his 'Scanlon and the Claims of the Many versus the One', *Analysis* 60 (2000), p. 288.

https://doi.org/10.1017/S0953820804001372 Published online by Cambridge University Press

the two with blinkers on as to their context.[8] If they are equal in all morally relevant respects, we still see no reason not to toss a coin. We follow this procedure until A is compared with all individuals on the opposite side. A is not balanced and silenced by having met his match. He would only be silenced by meeting more than his match, that is, someone needier or perhaps likely to produce a better outcome. Then we compare anyone else on A's side with all the same individuals on the opposite side in the same way, putting blinkers on and off. This way of doing pairwise comparison would eventually result in our having to toss a coin between a smaller group and a larger group of persons of equal need and outcome, because each person on one side would be owed an equal chance against any number of other individuals on the other side.

Notice that balancing which silences occurs as the result of a pairwise comparison match only when we are aware of a context in which there are individuals who could be helped in addition to those who would be balanced. Given that the Tiebreaker Argument makes use of balancing which silences, this will affect, I believe, how we should understand Scanlon's IR. It should be understood as implying: (1) Several individuals' complaints together cannot create a tie with another individual's complaint; only an individual's complaint can do that. (2) But individuals' complaints may be evaluated in the context of another individual's complaint and this may change the relation between the individuals' complaints. (3) The flipside of (2) is that an individual complaint may be evaluated in the context of how other individuals' complaints relate to each other. It may also be evaluated in order to see how it relates to and changes the relation of those other individual complaints.

Scanlon clearly enunciates (1). However, I believe that his presentation conceals (2) and (3). This is because he emphasizes comparing the complaints of individuals, and this may suggest comparing each individual's complaints with blinkers on. To make clearer the role of (2) and (3) in a tiebreaker argument, consider the following Three Islands Case. Suppose that we could save the life of either (1) A on one island or (2) B on another island, or (3) C on a third island. It would be a mistake to think that the first two people balance and silence each other and then conclude that we should help C. Rather, we should give each person a maximum equal chance to be saved (e.g. pick one by using

[8] Michael Otsuka believes that the fact that we do take the blinkers off and make sure to compare A with any other individual on the other side means that the decision procedure is affected by the presence of another person. Hence, it is not true, he says, that we decide in the same way as we would if it were a contest between A and just one other person. See his 'Scanlon and the Claims of the Many', p. 288.

https://doi.org/10.1017/S0953820804001372 Published online by Cambridge University Press

the three-straws method). It is only when an additional person will be saved along with A or B that balancing which 'silences' those balanced is allowed to occur. In that sort of case, proponents of the Tiebreaker or Balancing Argument not only pairwise compare A and B, but they see what I would call a dynamic relation between the complaints of A, B and C, that is, they see that B balances and silences the weight of A only when C joins B. (This is an instance of (2).) Further, C's complaint is an individual (versus group) complaint in sense (1), but not in the sense of the significance it would have in isolation (as in the Three Islands Case) or when pairwise compared with A's life with blinkers on. Rather, C's complaint is considered with blinkers off in a context where not only does A match B (the relation between other individuals' complaints) but B would be saved if C is. (These are instances of (3).)

But if all this is true, it may be said, do not the Balancing and Tiebreaker Arguments which are meant to lead to the conclusion that numbers count assume that numbers count? For they assume that a person (like C) being helped *as well* (that is, in addition to B) makes a difference (that is, leads to balancing which silences rather than tossing a coin or drawing straws). Are not the arguments, therefore, circular?

I believe that one of Michael Otsuka's criticisms of the Tiebreaking and Balancing arguments[9] is that to work as arguments for numbers counting they must already assume that numbers do count. I think his point is based on noticing that according to these two arguments someone like C has a complaint not considered in isolation, as in the Three Islands Case, but only in the context where B *in addition* has a life that would be saved when C's is.

My answer to this challenge is as follows: Pairwise comparison, combined with an impartial perspective on the individual's losses and gains, implies that there is equal moral value in helping either person who is equal in morally relevant respects. We recognize that the satisfaction of B's concern morally speaking gives us all that would have been achieved by the satisfaction of A's concern and this grounds the substitution of his concern with that of B's.[10] This implies that at the point where we consider that C will be helped *in addition to* B, the counting of numbers is no more controversial than the claims that we get the better outcome

[9] Ibid.

[10] Though, of course, it does not give A all the satisfaction of A's concern that saving him would have given him. It also does not give to the fact that A's concern would not be as well satisfied when B survives as when he does the recognition that giving equal chance does. Thus, a premiss in these arguments caters even less to the subjective perspective of each person than a procedure that refuses to toss a coin between A's leg and B's life. As noted above, rejecting a fair toss between A's leg and B's life is not quite enough to generate a commitment to the further impartial perspective of the Tiebreaker and Balancing arguments.

https://doi.org/10.1017/S0953820804001372 Published online by Cambridge University Press

if we save B and C rather than B alone (Pareto Optimality) and also that we should save B and C rather than B alone.[11] Those who object to counting numbers do not reject this view that it is better if B and C are saved than if B alone is and that we ought to save B and C rather than B alone. So, the Balancing Argument for counting numbers does not assume that numbers count in the controversial sense that in conflict situations we should help the greater number. It tries to prove this without assuming it: the appropriateness of the impartial perspective, the moral equivalence of A and B from that perspective, and substituting saving one for the other, bear most of the weight of the argument.

VI

Now, suppose we can save either A's life or B's life but not both, and also save C's legs only if we save B. (Call this the Legs Case.) If we just pairwise compare the potential loss and gain of every individual with blinkers on, losing or gaining legs pales before losing or gaining life and would be irrelevant to our saving a life. Yet, proponents of Balancing and Tiebreaker arguments could (correctly, I think) treat the legs as a tiebreaker, given the way they do pairwise comparison with context awareness. This is a case in which a tiebreaker between two lives is less than another life. Let us consider how such cases shed light on the idea of tie-breaking.

Typically when we think of breaking a tie, our focus is on (a) the individuals who are tied and (b) how to settle things between *them*. Factor (b) emphasizes that there is a tie and that we must break it, but does not yet tell us to use balancing that silences, as opposed to tossing a coin to do it. Factor (a) emphasizes that we do not focus on how it would be wrong to forego consideration of *that* which would break the tie for its own sake. Taking this perspective, if A ties B when lives are at stake and we could also cure a sore throat in C (who is otherwise fine) only if we save B, it might be said that it would be permissible to use C as a tiebreaker (call this the Sore Throat Case). That part of Scanlon's argument which focuses on how A's weight has been taken account of by requiring that someone who has as strong a complaint, such as B, balance him out, would lead one to conclude that A could have no complaint if C's sore throat broke the tie. After all, A's complaint has been given all the recognition it deserves, it might be said, when it is balanced against B's equal complaint. But unlike the additional person who needs his life saved, C with the sore throat cannot do what Scanlon says C could do if his life were also at stake. He cannot compare himself to B and say to A that just as the decision

[11] I shall have reason to modify this claim below.

https://doi.org/10.1017/S0953820804001372 Published online by Cambridge University Press

procedure was changed in recognition of B's weighty concern, so it should be changed in recognition of C's comparably weighty concern. If the sore throat were used as a tiebreaker, because we understand tie-breaking arguments as focusing on those in the tie, it seems that someone who has a much smaller concern could have an effect on the procedures used as great as would be had by someone who has an equally weighty complaint. That is, the outcome will be no different in a tie-breaking argument whether C has a sore throat or whether C's life is at stake.

Consideration of the Sore Throat Case shows that there are (at least) two separate issues in tie-breaking arguments: (i) is A in a tie with B; (ii) would the tie breaker have a complaint for his own sake, based on the seriousness of his own need, if he does not break the tie. Only if the latter is true should we proceed to engage in balancing and silencing as opposed to tossing a coin as a way of breaking the tie, I think. My own view[12] is that in the Sore Throat Case the sore throat should not be a tiebreaker. I believe it is inappropriate for A to lose his 50 per cent chance of being saved if we toss a coin between him and B simply for the sake of getting the additional gain of curing C's sore throat. If we did balance and silence, A could complain, even though his weight has been balanced by B's. So, sometimes treating each person properly will stand in the way of producing the best outcome. Hence, Scanlon's Tiebreaker argument is too simple, if it implies that we should not give A and B equal chances in the Sore Throat Case. On the other hand, if Scanlon's argument truly emphasizes the complaint on his own behalf of the individual who is to be a tiebreaker, it may here give the right result, but only because *it is more than a simple tie-breaking argument*.[13]

The Sore Throat Case suggests that we owe each individual something in addition to pairwise matching or balancing with an equal, namely, to balance and silence each with his equal only when the nature of the need of any additional person on one side is serious enough relative to the needs of those in a tie. This really means that those who are balanced and silenced will have no complaint if this is done only when the additional people on one side would themselves have a complaint as individuals if it were not done. This is one reason why I think it is misleading to describe the balancing argument as

[12] A view that I discussed at length in *Morality, Mortality*, vol. 1.

[13] Notice that it is not clear that whenever lives are at stake, we should not break ties with a minor concern such as a sore throat. For suppose there are a billion (or more if necessary for argumentative purposes) individuals on a billion separate islands waiting to be saved and we can save only one, and one of them is someone whose sore throat would also be cured. Here each person has only a very small chance of being saved. It might be that the loss to him of a one-in-a-billion chance can be outweighed by saving someone and also curing someone else's sore throat.

https://doi.org/10.1017/S0953820804001372 Published online by Cambridge University Press

a tie-breaking argument (as Thomas Scanlon does). For when we must break a tie, our focus is on the individuals tied, not on how the refusal to use someone else's need as a tiebreaker will wrong the potential tie-breaking individual.

Now return to the Legs Case, where saving C's legs is what would be a tiebreaker between A's or B's life. If A should not complain if C's legs are a tiebreaker, it is presumably because C, in the context where B will be saved, is appropriately asking us to focus on his complaint if he loses his legs as well as the fact that B can balance A. He does not just say that we need somehow to break the tie. But it is still true that he cannot say to A (as Scanlon emphasizes C could if his life were at stake) that just as the decision procedure was changed when B arrived on the scene, it should be changed now that he, C, is on the scene. For he cannot claim that his loss and gain are anywhere as important as B's or that he could stand in the same pairwise relation to A as B did. The argument for non-fairness to C in not considering him, which I think is correct, can neither be based on the mere fact that he can break a tie between two people when the focus is on them, nor on the fact that he has the same standing based on his substantive complaint as they have (and for this reason the focus should be on him when they create a tie). In the Legs Case as in the Sore Throat Case, someone with less of a complaint than A would be allowed to determine an outcome just as much as if his life were at stake. (This is one reason why even those who think that numbers of lives should count, may want to move back from tie-breaking here and instead give only a higher proportional chance to the side with B and C. I myself would not.) Hence, the Legs Case shows that the additional people on one side can themselves have a complaint as individuals if we do not balance and silence, even when they would not be as badly off as those who will be balanced and silenced.

In sum, tie-breaking is a somewhat misleading title if it leads us only to think of the equal weight of those between whom the tie exists. But the remedy of focusing on the complaint of the tiebreaker on his own behalf (which leads to balancing and silencing as a way to break the tie) is also misleading, if it implies that the tiebreaker must say that he is another person with a complaint equal to that of those tied. Morally appropriate consideration of the good of another in tie-breaking is both narrower and broader than these approaches imply.

So far, I have been considering the bearing on the idea of tie-breaking on cases in which what is at stake for an additional individual is less than what is at stake for the people who create the tie. But the Sore Throat Case also bears on the claim, made above, that taking account of another person in a balancing argument employs a form of aggregation that is no more controversial than the claim that we ought to save B

https://doi.org/10.1017/S0953820804001372 Published online by Cambridge University Press

and C rather than B alone.[14] We said that if A and B are tied, from an impartial view saving each is of equal moral value. So saving B is like saving A, and as it is clear that we ought to save A and C instead of just A, so we ought to save B and C instead of just A. But the Sore Throat Case presents a challenge to this simple view. For while we ought to save A and cure C's sore throat rather than just save A, I claimed that it is not true that we ought to save B and cure C's sore throat rather than (toss a coin between A and B and if he wins the toss) save A. (This is consistent with it still being true that we would produce a better outcome if we save B and cure C's sore throat than if we save A.) For purposes of deciding what we ought to do (rather than whether we produce the best outcome), it is not true that once we get to the stage of the Balancing or Tiebreaker arguments of considering B and C together versus A, the aggregation of B and C is always no more controversial than the aggregation of A and C. But this should not lead us to doubt the Balancing and Tiebreaker arguments for saving the greater number of lives, I think. For the objection to directly saving B and curing C's sore throat rather than tossing a coin between A and B does not stem from an objection to the moral equivalence of saving A and B or to balancing. Rather it stems from a concern about when balancing (rather than tossing a coin) is triggered. It is a concern over when the weight of the additional person is great enough so that the case should be treated as one involving anyone besides A and B at all; if the case should be treated as if only A and B are involved, their moral equivalence does not involve just saving B but involves tossing a coin to give each an equal chance.

VII

Above I noted that in Scanlon's view only individuals can be tie-makers. (This was (1) in my analysis of how we should understand Scanlon's IR.) It has been pointed out by Derek Parfit[15] that if Scanlon also accepts that in pairwise comparison we should compare not only baselines (how badly off someone would be) but also how much he can be benefited, then his views will have highly implausible implications. For example, suppose we can either help one person, A, who would otherwise die, to live twenty years or help two people, B and C, who would otherwise die, to live twelve years and eight years, respectively. If only individuals can create ties, and pairwise comparison is for baseline and benefit, then there is no tie in this scenario, and we should not save the greater

[14] As Michael Otsuka emphasized to me.

[15] In 'Rediscovering Reasons' (unpublished), and in 'Justifiability to Each Person', *Ratio* 16 (2003), p. 368.

https://doi.org/10.1017/S0953820804001372 Published online by Cambridge University Press

number according to the Tiebreaker Argument. (There is no tie in the sense that we would not toss a coin between A and any of the others alone in deciding whom to aid.) Not saving the greater number, however, seems wrong.

I have elsewhere suggested how Scanlon might deal with the problem.[16] I shall briefly repeat then expand on that earlier discussion here. It may help to think about Parfit-type cases, what we would do if the good we could give to A were, counterfactually, divisible. So imagine that he could first get a portion of it that another person, B, might get instead, if we tossed a coin between them. For example, A and B could be tied with respect to how badly off they will be if not helped, and also with respect to the possibility of getting a benefit of living up to twelve years. There would be two ways to break the tie. One way is vertically (as I call it), by imagining that if A gets the first benefit, he could also get another benefit of eight more years which B cannot have. A second way to break the tie is horizontally (as I call it), by imagining that if B gets the first benefit, another person C, who would be as badly off as A and B if not helped, will benefit to the level of eight years of life. If we give greater weight to tie-breaking on the horizontal dimension, it could be because it is morally more important to help another person who will be as badly off as the others (so long as he is helped significantly) than it is to provide someone (such as A) with an additional eight years on top of twelve.[17] (I think this could be true even if the benefit given to C did not, when combined with that given to B, amount to twenty years.)

In the actual case, of course, we cannot divide the potential benefit to A in this way and decide where to put the parts of it. But the reasoning we have just used, which I call Balancing Argument II, could, I suggest, still be used to justify saving the greater number in our actual case. In essence, we can reason in the actual case with another moral method that I call the Method of Virtual Divisibility (see Figure 1). That is, when we must choose between saving A who would, in fact, receive twenty years, or B who would get twelve and C who would get eight, we imagine creating a tie between A and B by marking off the point (twelve years) at which they would be tied and see whether the additional eight years we can produce are better allocated by saving C or by saving A for additional years. If helping the worst off takes priority and there is 'virtual' tying between A and B at twelve years, we should save the greater number. So the idea is that if Scanlon would distribute an actually divisible benefit a certain way, he should be committed to saving the greater number even when it is a choice between giving a

[16] See my 'Owing, Justifying, and Rejecting', *Mind* 111 (2002), p. 323.
[17] I argued for the greater moral value of distributed rather than concentrated benefits in *Morality, Mortality*, vol. 1.

https://doi.org/10.1017/S0953820804001372 Published online by Cambridge University Press

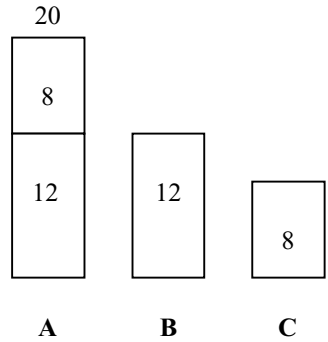


Figure 1

greater benefit to one person or smaller benefits to others who would otherwise be as badly off as that one person.

But the use of the Balancing Argument II and the Method of Virtual Divisibility can be extended, and its extension casts doubts on Scanlon's insistence that only an individual confronting another individual can create a tie. Suppose A could be saved for thirty years of life, or we could save B for only three years and C for only three years. It might indeed not make sense to break the 'virtual tie' between A and B at three years (imaginatively) by giving C three years instead of achieving in A an additional twenty-seven years. That is, the complaint of the worst off person to be taken care of to a significant degree when B would be helped as much as A up to the virtual tie point might be trumped by the size of the additional benefit possible for A. But suppose we add ten other people (D, E, etc.) to the side of B and C, and each of those people needs as badly to be saved and can live for three years. Now it seems more reasonable to save the greater number who would each be as badly off as A would be if not helped. Here it becomes clear that if there is any tie-breaking that justifies doing this, it is neither a tie created by A and one person actually equal to him in baseline and benefit, nor a tie created by any one additional person via 'virtual divisibility' and balancing. At best, it is a tie created by a group of people via virtual divisibility and balancing. That is, A and B are virtually tied for three years, A and C are virtually tied for the second three years that A might have, A and D are virtually tied for the third three years that A might have, etc. Only once we are (imaginatively) left with a much smaller additional gain possible to A can one single person complain that helping another worst off person is more important than giving this additional gain to A. (Perhaps it is even best not to speak of Scanlon-type tie-breaking at all as occurring in such scenarios where numbers

https://doi.org/10.1017/S0953820804001372 Published online by Cambridge University Press

count, if what is really going on is that components of a group of people rather than any individual are creating ties.)

Let us now consider possible problems with the Method of Virtual Divisibility. Recall the case where either A gets twenty years or B gets twelve and C gets eight. The Method is somewhat misleading because it suggests that one of our options is giving (a) additional resources, and (b) in multiple acts of aid to A moving him first to twelve years and then beyond to twenty. It also suggests that these are resources and attention that might otherwise go to a greater number of people who still lack even basic significant life extensions. But, in fact, if we help A, then it is just in helping him once, with no more resources going to him than are needed to help *each* one of B and C, that we will produce a bigger benefit. Indeed, if we wanted aid to help A live only to the twelve-year level, we might have to give him fewer resources than would be given to each of B and C on the other side, as he is such an efficient consumer of resources.

We must be careful, therefore, not to commit what might be called 'the divisive fallacy' when concluding by the Method of Virtual Divisibility that we should help the greater number each to a lesser degree rather than help a smaller number each to a greater degree (when they are all equally needy). The 'divisive fallacy' involves thinking that because we would not produce benefits to one set of people when distribution of divided resources is involved, we are thereby committed to not producing benefits in the same set when a single (smaller) resource is to be distributed. That is, we must not think that we would be undertaking multiple acts of allocation with greater resources for A and, therefore, unfairly helping him a lot, for we would really only be undertaking one allocative act to A with no more than his equal share of resources. In general, aiding a person with resources that amount to, for example, twenty times his share must be distinguished from helping one person with his equal share of resources even when this results in our not being able to give each of twenty others his equal share of resources. The former behaviour has a consideration against it from the point of view of fairness that the latter lacks.

Hence, if we decide to aid a great number of people each to achieve less than one person who would get more, when all share the same baseline, this should be just because achieving a greater good in one person (which is like adding goods to ones already obtained) is not as important as helping more people achieve significant goods relative to the original baseline.[18]

[18] The Method of Virtual Divisibility can also mislead if one mistakenly thinks that A in our example will actually get at least twelve years of life regardless of our decision about whether to give him or C eight additional years. This is not true.

https://doi.org/10.1017/S0953820804001372 Published online by Cambridge University Press

VIII

Those who deny that numbers count reject, unlike Scanlon, the idea of tying in the sense that involves thinking that we morally get everything we would have got if we had saved A if we instead save B, and so we may sometimes balance A and B and save B and C, when everyone faces the same bad outcome if not saved. So Balancing Argument II that I gave in Section VII to deal with a problem Parfit raises for Scanlon would not be acceptable to them (even if it were to Scanlon). This is for the same reason that the original Balancing Argument is not acceptable to them. It depends on the permissibility of giving A zero if we can get just as much moral value in helping someone else as we would get in helping A, and we can thus also significantly help other people as badly off. If there were a tie between individuals, those who reject counting numbers would only toss a coin to choose between them. And this would be true if it were a choice between saving A and saving B, whether the two are tied in the most obvious sense or in the sense that involves virtual divisibility, even when B is accompanied by C.

Michael Otsuka, however, would appear to have developed an argument against numbers not counting that does not rely on a balancing-of-equals step that numbers opponents reject.[19] He seems to give an argument for counting numbers that bypasses the issue of how to deal with ties. However, I do not think that his argument does, in fact bypass the need for some sort of tying and balancing argument, as I hope to show. I see a relation between my Balancing Argument II and the Method of Virtual Divisibility just described and Otsuka's argument, and so I shall present his argument in a way that highlights that relation. (I do not mean to deny his originality.) His argument also works by imagining what someone would do when a resource is divisible and then claiming that this has implications about what one must do when a resource is not, in fact, divisible. But he makes novel use of the assumption that those who think numbers do not count accept pairwise comparison that checks for both baselines and (crucially) degrees of benefit, so that they would not toss a coin between people who are at the same baseline but get very different benefits. He asks us to imagine that four people will be completely paralysed if we do nothing. We have three pills. In three of the people, one pill will relieve paralysis in two limbs to a maximum of two unparalysed limbs per person. In one person, one pill will relieve paralysis in one limb, but recovery can be achieved to a maximum of three limbs (with three pills).

[19] See his 'Skepticism about Saving the Greater Number', *Philosophy & Public Affairs* 32 (2004), p. 413.

https://doi.org/10.1017/S0953820804001372 Published online by Cambridge University Press

Otsuka claims that in distributing the pills, the anti-numbers counting person need not think it right to give all three pills to the person who would thereby get three working limbs. Instead of doing this, he would prefer (in increasing order) to: (1) give two pills that provide two working limbs in the fourth person and one pill that provides two limbs in one of the three; (2) give one pill that gives one working limb to the fourth person and two pills that each give two working limbs to two of the three; or (3) (best of all) give three pills that each provide two limbs in three people. The reasoning here relies on the fact that even an anti-numbers theorist can believe that we should put the pill where we can produce the bigger benefit, given that each person would be as badly off if not helped. Because the pill will produce more good in each of the three people than it will in the fourth person, pairwise comparison (of the sort that involves baseline and benefit) never yields a tie between any of the three who can get bigger benefits and the lone fourth person. (A tie, on an anti-number counting view, would imply that we should toss a coin between the greater and the lesser number, if we want to show equal concern.) Pairwise comparison in Otsuka's case with his assumptions, therefore, results in the willingness of an anti-number theorist to leave one person with zero and help the greater number seemingly without relying on the idea that in a tie we get as much moral value if we save one person or another and so we may balance, silence, and count numbers.

But then, Otsuka asks, will it not be inconsistent with it being the anti-numbers theorist's least favoured option to distribute the three pills so as to help one person get three limbs, for him (given his anti-numbers position) to favour choosing to help one person in the following case: If he is helped, one person will get three working limbs, while if three people are helped, each of them will get two working limbs. The anti-numbers position seems to commit the theorist to helping the single person on the following grounds: Numbers of people do not count, we should *not* toss a coin between only two people who would be at the same baseline when one will get a significantly bigger benefit than the other, and three limbs working is a significantly bigger benefit than two limbs working. But if the anti-numbers theorist would make choices stepwise in the first pill distribution case that lead to someone who could have had three working limbs having none so that three people may each have two, why should he not choose in making one choice in the second case to treat three people so that they each have two working limbs instead of treating one person to have three?[20] Doing this seems

[20] Using the Method of Virtual Divisibility that I discussed above, we could imagine that each of the working limbs that would be present in the single person, when compared with two limbs in one of the three, would not tie with or dominate the two limbs. (Suppose

https://doi.org/10.1017/S0953820804001372 Published online by Cambridge University Press

to Otsuka to be required by consistency, but it involves deciding to give smaller benefits to each of several simply on the basis of counting numbers of people. Hence, Otsuka concludes, the view that numbers of people do not count leads to inconsistent choices in cases, and this is a ground for rejecting the view.[21]

Does Otsuka's argument succeed against the anti-numbers theorist and does it bypass the need for moral equivalence, tying and balancing? Here are some reasons to think it does not: (I) Otsuka's argument may only show that those who are against counting numbers must be committed to a form of pairwise comparison that compares only baselines (i.e. how badly off someone will be) but not the amount of benefit to be gained. That is, they must not refuse to toss a coin between a person who avoids total paralysis by getting three working limbs and one who avoids total paralysis by getting two or one working limb. This is consistent with their refusing to toss a coin between a person who will be totally paralysed if not aided and someone who would only be paralysed in one limb if not aided. Why should they be committed to this? Because, I shall argue, pairwise comparison of degree of benefit depends on a hidden form of tying that involves balancing and substitution of morally equivalent benefits that anti-numbers people reject. And if pairwise comparison of benefits depends on such balancing, then Otsuka's argument, appearances to the contrary, does not provide an argument against the anti-numbers theorist that does not rely on the sort of balancing and substitution that is at the heart of the Balancing and Tiebreaker arguments.

To see why all this might be true, consider in a conflict case that part of a greater good that a person A could potentially get which B, who is at the same baseline, could also potentially get. This is the most crucial part of the good, as it is the part that lifts a person from his lowest baseline. From the point of view of someone who rejects the view that balancing is permissible because it is just as good if A or B gets that part of the good, it is morally inappropriate to take the fact that A could also get an additional unit of good that B would not get as a reason to favour A. We must first settle whether A or B will get the first, more crucial part of the good, and that would require tossing a coin. So, where A

one thinks the difference between getting two limbs and getting one is not significant enough to interfere with a tie between them, and so we should toss a coin. Then one could create another example, perhaps using years of life, where each one year in a candidate who will live thirty years is to be compared via virtual divisibility with twenty years in each of thirty other candidates that could be got instead. It is harder to say there is a tie between one year and twenty.)

[21] Otsuka notes that Scanlon faces the same problem even though he thinks numbers do count, for he would respond to Otsuka's two cases in the same way as the anti-numbers theorist given that he considers both baseline and benefit to create a tie.

https://doi.org/10.1017/S0953820804001372 Published online by Cambridge University Press

stands to get two working limbs if not totally paralysed, and B stands to get only one working limb, we must first settle who gets at least one limb. A's winding up with two limbs if he wins the toss is only a means to his not being deprived of a chance to get one limb. In other words, if we let the greater benefit that one person can achieve decide who gets the pill in Otsuka's case, we will have assumed that it is equally good whether A or B gets the first unit (e.g. one limb) of good, and only that allows us to decide things by considering the additional unit of good that A would get. If Otsuka's argument relies on virtual balancing at the level of one limb and virtual breaking of the tie with an additional working limb that will be present in one person and not in another, we may as well use the original Balancing and Tiebreaker arguments given by Kamm and Scanlon against the anti-numbers position.

Hence, Otsuka seems to have shown at most that those rejecting counting numbers because they reject a balancing or tie-breaking argument are committed to what many would also think is another implausible view, namely do not select among those who share the baseline merely because some can achieve significantly greater benefits. (I shall argue below that he does not actually show this.)

(II) How does Otsuka's argument relate to the divisive fallacy? In Otsuka's first case, to produce three working limbs in someone, we would have to give him three pills, pills that could individually help others to a lesser degree. It may be this *method* of bringing about the person with three working limbs that the anti-number theorist rejects. Hence, when Otsuka imagines in the second scenario the anti-number theorist preferring to help someone who will have three limbs rather than help each of three people have two limbs, making there be a three-limbed person in existence could involve giving him only one pill, the same number any one of the three would get. This is because, for example, his system employs the pill more efficiently. As noted above, in discussing Scanlon and problems with the Method of Virtual Divisibility, it is not inconsistent with refusing to give someone multiple resources, when others could use a portion of the total, to give someone a single portion of a resource because we thereby produce a bigger benefit in him, even if doing this means that we are not able to give the rest of the resources to other people.

Consider other cases like this. I might choose to distribute intelligence genes in a conflict case so that many have adequate intelligence, leaving one person with nothing rather than make the one a genius by giving him all the genes. But this alone (all by itself) does not prove that I have a reason to save several people of normal intelligence rather than someone who is already a genius. I might also choose to give several people each reasonable life spans rather than create one very long-lived person, but this alone does not imply that if I can save either someone

https://doi.org/10.1017/S0953820804001372 Published online by Cambridge University Press

who already has long-life-span genes or else several with shorter-life-span genes, my first choice commits me to saving the latter rather than the former.

Suppose that the numbers sceptic might in this way distinguish different methods by which we help to make there be a person with three working limbs. Then Otsuka's argument does not – contra (I) above – even show that the numbers sceptic is committed to the form of pairwise comparison that does not compare benefits but only baselines. The numbers sceptic can avoid inconsistency by distinguishing two different ways in which the three-working-limb person can come about, favouring his coming about when one method is used and disfavouring it when the other method is used.

What if the benefit to be gained by each person is equal (e.g. if no one could get more than one limb from a pill), but only one person could benefit to the extent of getting three working limbs? In such a variant of Otsuka's case, the anti-numbers theorist could commit straightforwardly to giving one person rather than another a pill on the grounds of helping a worst off person before helping a better off person. Therefore, he would prefer that three people each get one pill than that one person gets three. He should pick a straw from a set of four, giving each person an equal chance to get a pill, leaving one without anything. But suppose that for some reason a pill given to three of the people (A, B and C) only works if all three take it. However, D's pill will work if he alone takes it. Then the anti-numbers theorist faces a choice of giving D a pill or giving A, B and C one each. Giving it outright to the three would deprive D of his equal chance. So, the anti-numbers theorist should toss a coin between D and the three. If D wins, he gets his pill. But because two other pills will now go to waste if we give D only one pill, it becomes permissible to give D all three pills rather than give one to each of the other three people.[22]

frances_kamm@harvard.edu

[22] This article was presented at the conference on Thomas Scanlon's *What We Owe to Each Other*, University College, London, June 2004. I am grateful to Professor Véronique Munoz-Dardé for inviting me to speak. For comments, I am grateful to Michael Otsuka and members of the audience at UCL, including Thomas Scanlon, and to Alexander Friedman.

https://doi.org/10.1017/S0953820804001372 Published online by Cambridge University Press