

Contractualism and Aggregation

In *What We Owe to Each Other*¹ Tim Scanlon attempts to have his cake and eat it too. He claims both that his contractualist theory provides an intuitive way to block the aggregative feature of consequentialist theories that leads to counterintuitive results, and that it allows for at least the appearance of aggregation in those situations in which aggregative reasoning is intuitively acceptable. His argument is interesting and challenging, but, I will argue, ultimately unconvincing.

1. Scanlon's Argument for Limited Aggregation

So how does Scanlon's contractualism block aggregative moral reasoning? "According to contractualism, when we address our minds to a question of right and wrong, what we are trying to decide is, first and foremost, whether certain principles are ones that no one, if suitably motivated, could reasonably reject" (189). His argument that contractualism blocks aggregation depends on a feature that, he claims, is central to the guiding idea of contractualism: "its insistence that the justifiability of a moral principle depends only on various *individuals'* reasons for objecting to that principle and alternatives to it" (229). The "mode of justification" of consequentialist theories, on the other hand, "is, at base, an aggregative one: the *sum* of a certain sort of value is to be maximized" (230). This leads to the allegedly implausible implication that "in principle ... imposing high costs on a few could always be justified by the fact that this brought benefits to others, no matter how small these benefits may be as long as the recipients are sufficiently numerous" (230). Thus, to use a familiar example, a consequentialist would have to admit that it could be permissible to kill an innocent person in order to prevent some, perhaps astronomically high, number of people from suffering minor headaches. Contractualism, claims Scanlon, rejects this kind of justification:

A contractualist theory, in which all objections to a principle must be raised by individuals, blocks such justifications in an intuitively appealing way. It allows the intuitively compel-

¹T.M. Scanlon, *What We Owe to Each Other* (Cambridge, Mass.: Harvard University Press, 1998). All page references in the text are to this book.

ling complaints of those who are severely burdened to be heard, while, on the other side, the sum of the smaller benefits to others has no justificatory weight, since there is no individual who enjoys these benefits and would have to forgo them if the policy were disallowed. (230)

However, as Scanlon points out, there are situations in which aggregative reasoning seems to be called for. For example, if faced with a choice between saving one person from death and saving two different people from death, it seems that we ought, other things being equal, to save the two. Moreover, it seems that we ought to save the two, *because* there are two of them, as opposed to one. But what does the contractualist say about this choice? Consider the different principles that might govern our choice in this case: (a) we ought to save the larger number; (b) we ought to save the smaller number; (c) we are permitted to save either the larger or the smaller number. If we simply focus on the individual reasons for objecting to any of these principles, it seems that they are of equal strength. Scanlon sums up the problem as follows:

It therefore seems that as long as it confines itself to reasons for rejection arising from individual standpoints contractualism will be unable to explain how the number of people affected by an action can ever make a moral difference. (230)

In response to this problem, Scanlon suggests a ground on which a member of the larger group could reasonably reject (c), which is unavailable to a member of the smaller group:

In such a case, either member of the larger group might complain that this principle did not take account of the value of saving his life, since it permits the agent to decide what to do in the very same way that it would have permitted had he not been present at all, and there was only one person in each group. The fate of the single person is obviously being given positive weight, he might argue, since if that person were not threatened then the agent would have been required to save the two. And the fact that there is one other person who can be saved if and only if the first person is not saved is being given positive weight to balance the value of saving the one. The presence of the additional person, however, makes no difference to what the agent is required to do or to how she is required to go about deciding what to do. This is unacceptable, the person might argue, since his life should be given the same moral significance as anyone else's in this situation ... The conclusion ... is that any principle dealing with cases of this kind would be reasonably rejectable if it did not require agents to treat the claims of each person who could be saved as having the same moral force. (232)

It is important to be clear on just what the reasoning is here. The complaint of a member of the larger group is that his life is not being given the same moral significance as that of the member of the smaller group. Despite what Scanlon says, the contrast cannot be with the life of the other member of the larger group, since the very same complaint is available to both members. Furthermore, the complaint is justified not simply by the fact that the offending principle permits the agent to decide what to do "in the very same way that it would have permitted had he not been present at all." Such a counterfactual test would also render unacceptable the principle

that the agent should save the larger group. If such a principle were employed, the single member of the smaller group could also complain that his presence makes no difference to what the agent is required to do. What is more, if the groups were more unequal, say one versus three, the counterfactual test could be used by a member of the *larger* group to reject the principle that the agent should save the larger group. The correct counterfactual test is something more like the following: if, apart from my presence, the reasons on each side were balanced, would the principle permit the agent to decide in the same way as it would have done if I had not been present? In Scanlon's example, of course, the first part of the test is already satisfied. But the test, if it is to be plausible, must also apply to more unequal groups. The crucial concept, then, somewhat obscured by the particular choice of example, is that of reasons for and against a particular action being *balanced*.

A consequentialist would, of course, agree that if the reasons for and against a particular action are balanced, the introduction of a further reason on one side makes a difference to how the agent is required to go about deciding what to do. Furthermore, a consequentialist would agree that any principle that fails to take any individual's welfare into account is unacceptable, and thus that it is reasonable for that individual to reject such a principle. It might be objected that consequentialism grounds the reasonable rejection of a principle in the fact that it doesn't take everyone's welfare equally into account, whereas Scanlon's approach grounds the unacceptability of a principle in its reasonable rejection by an individual. But this would be misleading. Even for a contractualist, the fact that an individual can reasonably reject a principle is not a basic moral datum in no need of explanation or grounding. What makes it reasonable to reject the principle that it is permissible to save either the larger or the smaller group is that it doesn't require agents "to treat the claims of each person who could be saved as having the same moral force." Talk of balancing reasons is a mere heuristic to help identify those principles that are reasonably rejectable for the relevant reason. Contractualism departs from consequentialism in its account of precisely which reasons can balance which other reasons. For the consequentialist, any kind of morally relevant reasons can, in principle, balance any other kind. Scanlon, on the other hand, claims that some moral reasons are simply not "relevant" to other moral reasons.

To see in more detail how Scanlon's contractualism is supposed to depart from consequentialism, we need to examine an example in which the balancing of reasons approach is available to a consequentialist, but unavailable to a contractualist. Scanlon presents such an example as follows:

Jones has suffered an accident in the transmitter room of a television station. Electrical equipment has fallen on his arm, and we cannot rescue him without turning off the transmitter for fifteen minutes. A World Cup match is in progress, watched by many people, and

it will not be over for an hour. Jones's injury will not get any worse if we wait, but his hand has been mashed and he is receiving extremely painful electrical shocks. Should we rescue him now or wait until the match is over? Does the right thing to do depend on how many people are watching ... ? (235)

Scanlon's answer is that, of course we should rescue Jones now, no matter how many people are watching the match. The principle to which he appeals to ground this judgment is that "if one can save a person from serious pain and injury at the cost of inconveniencing others or interfering with their amusement, then one must do so no matter how numerous these others may be" (235). So why couldn't one of the sports fans reasonably reject such a principle on the grounds that it fails to take her interests into account? The answer, apparently, is that such a claim would rely on the possibility of there being a balance of moral reasons that could be tipped by the addition of the reason to prevent inconvenience to one person. In the case involving loss of life on both sides, it is easy to see how the reasons for and against an action could be exactly balanced. If you remove one member from the group with two, you have the *same reason*, not just the same strength reason, to save one group as the other. However, when the reasons supporting different options (or principles) are themselves different, the tie-breaking assumption can't be applied:

But when the harms in question are unequal, we cannot create ... a tie simply by imagining some of the people in the larger group to be absent. To claim that there is a tie in such a case would be already to claim that the fact that there are more people in one group makes it reasonable to reject a principle requiring one to help the smaller number, each of whom would suffer the greater harm. So we cannot use this "tie-breaking" argument to justify the selection of a principle requiring one to save the greater number in such cases. (235)

There are two obvious lines of objection to this attempt to incorporate into contractualist reasoning a limited role for the numbers of people affected by different options. (i) Even if we reject the possibility of balancing a large harm, such as death, against very small harms, such as headaches, there are intermediate harms, such as permanent paralysis or mutilation, that can be added together to balance the harm of death. (ii) There are many, intuitively acceptable, real-life examples of summing large numbers of small benefits to offset small numbers of large harms. I will consider Scanlon's discussions of these objections in turn.

2. Moral Relevance and the Denial of Transitivity

In response to (i), Scanlon suggests a further modification of his approach to incorporate the notion of moral relevance between harms:

[I]t seems that our intuitive moral thinking is best understood in terms of a relation of "relevance" between harms. If one harm, though not as serious as another, is nonetheless serious

enough to be morally "relevant" to it, then it is appropriate, in deciding whether to prevent more serious harms at the cost of not being able to prevent a greater number of less serious ones, to take into account the number of harms involved on each side. But if one harm is not only less serious than, but not even "relevant to," some greater one, then we do not need to take the number of people who would suffer these two harms into account in deciding which to prevent, but should always prevent the more serious harm. (239-40)

This suggestion, admittedly not explored in detail by Scanlon, is vulnerable to a potentially fatal objection, as I will explain. It is fairly clear that the relation of moral relevance does not obey transitivity. To see why, suppose, first, that it does. Consider now a descending scale of finitely many different harms, from the most serious, such as death, all the way down to the most trivial, such as a minor temporary headache. The difference in seriousness between any two adjacent harms is no larger than is necessary for the lesser harm to be clearly less serious than the greater harm. Suppose, also, that for every harm on the scale above the most trivial, there is some lesser harm that is relevant to it. Call this second assumption the "continuity assumption." Transitivity and continuity together entail that the most trivial harm is relevant to the most serious harm, precisely the result that the notion of moral relevance is intended to avoid. Can we preserve transitivity by rejecting continuity? This would involve finding a break (or breaks) in the scale between two harms, such that the harm directly below the break is not morally relevant to the harm directly above the break. Given that the difference between any two adjacent harms is as small as is compatible with the harms being morally distinct, the postulation of a break in the scale would run directly counter to the intuition that suggested the notion of moral relevance in the first place. Where could such a break plausibly occur? The most likely candidate would be just below death. There is, we might think, something special about death. As Clint Eastwood says in "Unforgiven," "It takes away all a man has, and all he's gonna have." Unpleasant as even severe mutilation is, perhaps it's still worse that one person die than that any number are mutilated. This might be the view of death espoused by those students in introductory classes who claim that life is "invaluable" or "infinitely valuable," but is it really plausible? Can anyone who really considers the matter seriously honestly claim to believe that it is worse that one person die than that the entire sentient population of the universe be severely mutilated? Clearly not. Perhaps the break in the sequence of harms could occur at some later point. Perhaps there is some harm short of death, that is worse than any number of any lesser harms. This seems even more implausible, though, than the claim that death is worse than any number of any lesser harms.

We must, therefore, conclude that the relation of moral relevance, if it is to do the work intended for it by Scanlon, does not obey strict transitivity. So what? If the notion of moral relevance were supposed to constrain

our judgments of all-things-considered betterness, this would be a serious problem. Although some brave souls have seriously entertained the possibility that "all-things-considered better than" is not a transitive relation,² the sheer implausibility of the suggestion makes the standard objections to utilitarianism, Kantianism, or contractualism appear trivial by contrast. However, Scanlon suggests the notion of moral relevance as part of an account of which principles are reasonably rejectable, and thus of which options are permissible, obligatory, or forbidden. To demonstrate that, even in this context, the failure of transitivity leads to highly implausible results, I need to consider an example.

Suppose, for the sake of argument, (a) that the loss of both arms is less serious than but morally relevant to death; (b) that a broken leg is less serious than but morally relevant to the loss of both arms, but not morally relevant to death; (c) that in a choice between saving one life and preventing one thousand people from losing both arms, it is obligatory to aid the larger group; (d) that in a choice between preventing one thousand people from losing both arms and preventing one million people from breaking a leg, it is obligatory to aid the larger group. (The choice of examples is unimportant.) Consider now three different choices: (i) Save one life or prevent one thousand people from losing both arms. (ii) Prevent one thousand people from losing both arms or prevent one million people from breaking a leg. (iii) Save one life or prevent one million people from breaking a leg. From (b), (c) and (d) it follows that it is obligatory to aid the larger group in (i) and (ii), and the smaller group in (iii). So far, so good. But what happens when we are faced with all three options in one choice? No answer here seems satisfactory.

Consider the possibility that one of the options, say saving the life, is obligatory. But now suppose that, just as you are about to save the life, it becomes impossible for you to prevent the million people from breaking a leg. Perhaps the largest group is further away than the other two, and your fuel tank is punctured by a jagged rock on the road to the one person. You are still able to save either the one or the thousand, but you can't reach the

²See L. Temkin, "Intransitivity and the Mere Addition Paradox," *Philosophy & Public Affairs* 16 (1987): 138-87, and "A Continuum Argument for Intransitivity," *Philosophy & Public Affairs* 25 (1996): 175-210; W. Quinn, "The Puzzle of the Self-Torturer," *Philosophical Studies* 59 (1990): 79-90; S. Rachels, "Counterexamples to the Transitivity of 'Better Than'," *Australasian Journal of Philosophy* 76 (1998): 71-83. In his 1987 paper Temkin uses the same central example as Rachels, but Temkin's explanation for the supposed intransitivity is the same as the one he provides in his earlier paper. Quinn doesn't explicitly claim that "better than" is intransitive, but his arguments, if successful, entail that a utilitarian should deny the transitivity of "better than." I discuss Temkin's 1987 paper in my "Intransitivity and the Person-Affecting Principle," *Philosophy and Phenomenological Research* 59 (1999): 769-76. I discuss Temkin's 1996 paper and Quinn's in my "Comparing Harms: Headaches and Human Lives," *Philosophy & Public Affairs* 26 (1997): 135-67.

million in time. Now you find yourself faced with choice (i), in which it is obligatory to save the thousand and forbidden to save the one. But this is very strange. You were about to do your duty, virtuously eschewing both forbidden alternatives, when one of the forbidden alternatives by chance becomes unavailable, as a result of which the other forbidden alternative becomes obligatory, and the previously obligatory alternative becomes forbidden. We should, if at all possible, avoid having to swallow such an unpalatable consequence. The same reasoning applies, *mutatis mutandis*, to the hypothesis that either of the other alternatives is obligatory in the three-option choice. Perhaps, then, each option is permissible in the three-option choice. But the implausibility of this can be demonstrated by the very same thought experiment. 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. A further possibility is that each option is forbidden in the three-option case. But this is even more unpalatable than the previous suggestions. Not only would we have to accept that a previously forbidden alternative can become obligatory by the chance deletion of another forbidden alternative, but we would also have to accept the existence of situations in which an agent, through no fault of her own, cannot help but do wrong. What is more, such situations may be very common. Both through the agency of charities and through our own efforts, many of us are able to bring many different types and levels of aid to others.

I very much doubt that the notion of moral relevance can do what Scanlon needs it to do, without causing worse problems than those it is intended to solve. I have not, of course, proved that the notion of moral relevance cannot be incorporated into a contractualist theory in such a way as to provide for limited aggregation. I have shown that the notion cannot obey strict transitivity, if it is to serve its purpose. I have also demonstrated the kinds of problems that would arise from an intransitive relation playing the role envisioned for it by Scanlon.

3. Acceptable and Unacceptable Tradeoffs

But why should a contractualist, or any other moral theorist, want to rule out the kinds of aggregative reasoning that allow for tradeoffs between large numbers of small benefits and small numbers of large harms? Don't we already accept such tradeoffs in our everyday thinking? As examples of this, Scanlon considers public building projects, designed to bring added convenience to many, but undertaken in the knowledge that they will inevitably result in serious injury or even death for a few. How, then, do such projects differ morally from the situation involving Jones in the TV transmitter room? This is what Scanlon has to say about the difference:

It is important in understanding our reaction to these cases to note that they involve failing to prevent accidental injuries rather than either intentionally inflicting serious harm on a few people, or withholding aid from people who need it, in order to bring small benefits to others. They differ in this respect from my original television studio example, and if they did not differ in this way our reaction to them would be very different. (236)

But it is by no means clear that the distinction between failing to prevent accidental injuries, on the one hand, and withholding aid from people who need it, on the other, can bear the kind of moral weight that Scanlon is suggesting here. If we refuse to rescue Jones before the match is over, we are withholding aid from someone who needs it. But what if the situation were slightly different? Suppose that we see that the electrical equipment is about to fall on Jones's arm, and the only way we can prevent the accident involves turning off the transmitter for fifteen minutes. If we are allowed to aggregate small benefits in cases involving preventing accidents, but not in cases involving withholding aid from people who need it, we might be permitted, or even required, to allow the equipment to fall on Jones, even at the cost of permanent injury, and then required to turn off the transmitter for fifteen minutes to rescue him. I hope the absurdity of this result is clear. Whatever grounds our intuition that we must help Jones in Scanlon's version of the example would also ground an intuition that we must prevent the accident in my version.

Another difference between public building projects, on the one hand, and either version of the television studio example, on the other, is that the identity of the victim is known in advance of making any decision in the latter cases, but not in the former. When we consider whether to go ahead with a large building project, we may well know that *some* people will be injured or killed in the course of completing it, but we don't know *who* they are. In either version of the television studio example, we know that it is Jones who needs to be saved. This may well make a psychological difference in our reactions to the different cases, but it is clear that it cannot bear any moral weight. To see why, consider another variation on the television studio example. There are two workers in the transmitter room, Smith and Jones. An explosion in the room kills one, and traps the other. We don't know who is still alive, but we do know that another explosion will kill the survivor in an hour. However, in order to complete the rescue, we must shut down the transmission of the World Cup match for fifteen minutes. If we wait until the match is over, it will be too late. Both Smith and Jones will be dead, and we will never know which one could have been saved (the second explosion will destroy the evidence). It is clear that our lack of knowledge of the victim's identity makes no difference to the morality of the situation. Whatever we should do in this case would not be altered, if we were to discover that the survivor of the initial explosion was Jones. Perhaps this is because we have already narrowed down the identity

of the victim to one of two people. In the example of a large building project, the victim(s) could be any of a much larger number of people. But a little reflection suffices to show that this can't be a morally significant factor, either. If there were a hundred, a thousand, or even a million workers in the transmitter room, all but one of whom were killed in the initial explosion, our duty to the one unknown survivor would be the same as our duty to Jones in any of the other versions of the example.

Scanlon also suggests that our willingness to permit certain risky activities "depends crucially on the assumption that precautions have been taken to make the work safe and that, in addition, workers have the choice of whether or not to undertake the risks involved" (236). The latter condition may be satisfied in the kinds of projects to which Scanlon refers, though even in these cases, some of the risks may affect passers-by. But there are other, also commonly accepted, examples of imposing risks of large harms to a few in order to bring small benefits to many. For example, thousands of people die in automobile accidents every year in the United States. It is highly probable that the number of deaths is positively correlated with the speed limits in force on highways, at least within a certain range. One of the effects of raising speed limits is that there are more accidents, resulting in more deaths and injuries. One of the effects of lowering speed limits is that there are fewer accidents. Higher vehicle safety standards also affect both the numbers of accidents and the severity of the injuries sustained when accidents do occur. Another effect of raising speed limits is that more gasoline is consumed, which raises the level of particulate pollution, which also leads to more deaths.³ Stricter standards for fuel-efficiency also affect the amount of gasoline consumed. There are, then, many different measures that we, as a society, could take to lower the number of automobile-related deaths, only some of which we do take. There are also many measures we could take that would *raise* the number of such deaths, some of which we do take. Furthermore, it is not obvious that we are wrong to fail to do all we can to reduce the number of deaths. Consider our failure to impose a national speed limit of 50 mph. If there were a national speed limit of 50 mph, it is overwhelmingly likely that many lives would be saved each year, as compared with the current situation. One of the costs of our failure to impose such a speed limit is a significant number of deaths. The benefits of higher speed limits are increased convenience for many. Despite this, it is far from obvious that the failure to impose a 50 mph speed limit is wrong.⁴ Can we confidently say

³According to a study by the Natural Resources Defense Council, "Some 64,000 Americans may die prematurely each year because of air pollution" (Reuter, 9 May 1996).

⁴There are those who react to what I say about the 50 mph speed limit by declaring that I have convinced them that it *is* wrong not to impose it. But what I say about the 50 mph speed limit can also be said about a 40 mph speed limit, or a 30 mph speed limit, or even

of the victims of higher speed limits that they have freely chosen to undertake the risks involved? This is clearly not applicable to the victims of higher levels of particulate pollution. Nor does it seem to be true of *all* the victims of road accidents. Many children are killed on the roads. Many of these may have had no say over whether to travel that way. It seems clear, then, that the appeal to freely assuming risks will not serve to distinguish intuitively acceptable tradeoffs from intuitively unacceptable ones.⁵

4. The Appeal to Intrapersonal Aggregation

Scanlon suggests that a contractualist may be able to distinguish between the television studio example and the building projects example by employing what he calls "intrapersonal aggregation." Consider first "a principle that allows projects to proceed, even though they involve risk of serious harm to some, provided that a certain level of care has been taken to reduce these risks" (236). Clearly, the strongest objection to such a principle would come from the standpoint of someone who is seriously injured or killed as a result of the project. How, we might ask, could there be an individual generic reason supporting the principle that could be considered reasonable when contrasted with the generic reason of avoiding death? Scanlon's answer is that an individual could aggregate within her own life all the relatively minor inconveniences imposed by adopting a more stringent requirement:

In meeting the level of care demanded by the principle, they might argue, they have done enough to protect others from harm. Refusing to allow activities that meet this level of care would, they could claim, impose unacceptable constraint on their lives. (237)

Although Scanlon does not argue for any particular risk-permitting principle, we are clearly supposed to agree that there is an acceptable one that would permit at least some of the public projects we find intuitively permissible. It seems clear, though, that any such principle would be reasonably rejectable, so long as we are restricted to intrapersonal aggregation of inconveniences. How could even severe constraints on a life be anything like as bad as violent death? Furthermore, if we rule out interpersonal aggregation, we must do it both for the inconveniences suffered by the many under the stricter principle, and for the deaths suffered by the few under the more permissive principle. (It might be tempting to smuggle in inter-

about abolishing private automobiles altogether. Very few are hardy enough to follow their respect for life to such extremes.

⁵For a more detailed discussion of the speed limits example, and of other attempts to distinguish it from intuitively unacceptable tradeoffs, see Norcross, "Comparing Harms"; M. Ridge, "How to Avoid Being Driven to Consequentialism: A Comment on Norcross," *Philosophy & Public Affairs* 27 (1998): 50-58; and A. Norcross, "Speed Limits, Human Lives, and Convenience: A Reply to Ridge," *Philosophy & Public Affairs* 27 (1998): 59-64.

personal aggregation of deaths under the guise of deciding what counts as an adequate level of care to make the project safe. However, it is clear that consistency requires that we not simply reduce this question to the question of how many deaths we are prepared to live with. If we did, then what would prevent us from factoring in the numbers of people constrained in our assessment of what constitutes “unacceptable constraint”? If we can’t directly take the numbers of the dead into account, we couldn’t distinguish morally between a principle that imposes severe constraints on many to prevent a few deaths, and one that imposes severe constraints on a few to prevent many deaths. I am fairly confident that it is not simply the unreconstructed consequentialist in me that insists on a significant moral distinction between these principles.

Suppose that intrapersonal aggregation will suffice to accept a principle allowing a significant number of intuitively permissible risky projects. Why, then, wouldn’t such a principle also allow us to leave Jones to suffer for an hour while millions enjoy the World Cup? Scanlon’s reply is that the rejection of a principle allowing us to let Jones suffer would not impose unacceptable constraints on anyone’s life. Such a rejection involves accepting something like a principle requiring “an agent to save one person from an hour of extreme pain even at the cost of inconvenience to others, regardless of the number of people so inconvenienced” (238). Scanlon argues that we already follow something like this principle, and that “the occasions to which it applies seem sufficiently rare that the costs on each of us are not very significant” (238). Do we really follow anything like this principle? Consider a governor of a state deciding whether to veto a law raising speed limits, or deciding whether to press for legislation lowering speed limits. Either action could save many people from far more than an hour of extreme pain, at the cost of inconvenience to many others. Do we accept that the governor is always required to opt for the pain- (or death-) saving option? If we do accept anything like Scanlon’s principle, we would have to understand saving someone from pain to be different from preventing someone from suffering pain. Perhaps we would also have to understand the someone to be identifiable in advance. However, as I argued above, neither the distinction between failing to prevent harm and withholding aid, nor between a victim who is identifiable and one who is not can bear any moral weight. If a contractualist approach that appeals to intrapersonal, but not interpersonal, aggregation treats as morally different two principles that differ only on one (or both) of these dimensions, that is a serious strike against such an approach.

I have a further worry about the appeal to intrapersonal aggregation to judge certain principles unacceptably constraining. Our judgments of what kinds or degrees of restrictions would be unreasonable are, no doubt, largely shaped by culture and propaganda. In the United States, for exam-

ple, with its strongly libertarian streak, some restrictions may be judged unacceptable that would be accepted as common sense in certain Western European nations. Attitudes to tradeoffs between freedom and safety also change with time. When seat-belt laws were first introduced, they were widely viewed as unacceptable restrictions on freedom. In the course of one generation, that view has all but disappeared. These worries don't, of course, prove that there isn't a fact about what restrictions would be unacceptably constraining. Perhaps seat-belt laws really are unacceptably constraining, and we have been brainwashed into believing otherwise. Perhaps, on the other hand, the denial of all forms of personal motorized transport would not constitute an unacceptable constraint on our freedom. The difficulty of supporting such judgments should make us wary of accepting intrapersonal aggregation as sufficient to justify the intuitively acceptable tradeoffs between great harms and small benefits. We are, naturally, eager to find a principle that both explains and justifies our intuitive moral judgments. This eagerness may lead us to embrace the appeal to intrapersonal aggregation, without carefully considering whether it can really do the considerable moral work required of it. I suspect that it cannot. We must at least consider the possibility that the reason we commonly accept higher speed limits and risky building projects, even if we are wrong to do so, is that the benefits accrue to large numbers of people.

Conclusion

I have argued that Scanlon's attempt to accommodate limited aggregation within a contractualist theory fails. Where does this leave Scanlon's contractualist approach? If we maintain reasonable rejectability, *from a generic individual standpoint*, as the test of a moral principle, either we should reject even the limited aggregation involved in choosing whether to save either one or two people from death, or we should accept the intuitively troublesome aggregation involved in killing one person to prevent many headaches. I have argued elsewhere⁶ that embracing the latter course is far more palatable than rejecting it. If we do embrace this option, though, we have lost a significant motivation for opting for a contractualist moral theory as opposed to a consequentialist one. This strikes me as a virtue of the suggestion, but I suspect that Scanlon would disagree.

Alastair Norcross

Department of Philosophy
Southern Methodist University
anorcross@mindspring.com

⁶Norcross, "Comparing Harms," and A. Norcross, "Great Harms from Small Benefits Grow: How Death Can Be Outweighed by Headaches," *Analysis* 58 (1998): 152-58.

Copyright of Social Theory & Practice is the property of Florida State University / Dept. of Philosophy and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.