

Sequential Dominance and the Anti-Aggregation Principle

Johan E. Gustafsson*

According to the widely held anti-aggregation principle, it is wrong to save a larger number of people from minor harms rather than a smaller number from much more serious harms. This principle is a central part of many influential and anti-utilitarian ethical theories. According to the sequential-dominance principle, one does something wrong if one knowingly performs a sequence of acts whose outcome would be worse for everyone than the outcome of an alternative sequence of acts. The intuitive appeal of the sequential-dominance principle should be obvious; everyone is knowingly made worse off if it is violated. In this paper, I present a number of cases where one is forced to violate either the anti-aggregation principle or the sequential-dominance principle. I show that these principles conflict regardless of whether one accepts a counterfactual or a temporal, worsening view of harm. Moreover, I show that this result holds regardless of how much worse a harm has to be in order to count as a much more serious harm.

According to a widely held view, it is wrong to severely harm someone to save a large number of people from a minor harm. T. M. Scanlon illustrates the underlying intuition with the following television-studio example:

Suppose that 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 [p. 1594] 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—whether it is one million or five million or a hundred million? It seems to me that we should not wait, no matter how many viewers there are, and I believe that contractualism can account for this judgment [...]¹

* I would be grateful for any thoughts or comments on this paper, which can be sent to me at johan.eric.gustafsson@gmail.com.

¹ Scanlon (1998, p. 235).

More generally, Scanlon endorses *the anti-aggregation principle*, which

does not require, or even permit, one to save a larger number of people from minor harms rather than a smaller number who face much more serious injuries.²

John M. Taurek makes similar claims in his case against letting the numbers count.³ And a similar principle follows, given consequentialism, from an evaluative view which Larry S. Temkin claims most people accept.⁴ In this paper, I shall argue that every theory that forbids one from saving a greater number of people from minor harms rather than a lesser number from more serious harms will in some situations force one to knowingly choose a sequence of acts that is worse for everyone than another available sequence of acts.

I shall present an example where one must either violate the anti-aggregation principle or violate a sequential dominance principle. This dominance principle concerns not just acts but also sequences of acts. A sequence of acts is a set of acts one can perform during a time interval without performing any other acts. We say that one performs a certain sequence of acts during a time interval $t-t'$ if and only if one performs all and only these acts during $t-t'$. Because the dominance principle I shall propose might, at first blush, seem unduly complex, it will help to first consider a simpler, tentative principle. According to

the sequential-dominance principle (tentative version),
if one can perform a sequence of acts during $t-t'$ with outcome x , it is wrong to perform a sequence of acts during $t-t'$ whose outcome would be worse for everyone than x .

This principle has at least three potential problems.

First, performing a dominated sequence of acts—that is, a sequence of acts with an outcome that is worse for everyone than the outcome of an alternative sequence of acts—might seem permissible if one did not know from the start what the [p. 1595] outcomes of these sequences would be. So, to sidestep this issue, we weaken the principle to cases where one knows the outcomes of the sequences of acts one can perform.

Second, one might object that it need not be wrong to perform a dominated sequence of acts during a time interval if one cannot, at the start

² Scanlon (1998, p. 238). The name is due to Raz (2011, p. 207).

³ Taurek (1977, pp. 307–310).

⁴ Temkin's (2012, p. 32) 'second standard view' yields that any outcome where a larger number of people are saved from minor harms is worse than any (otherwise similar) outcome where instead a smaller number of people are saved from much more serious injuries. Given that it is wrong to perform any act whose outcome is worse than the outcome of some alternative act, we get the anti-aggregation principle.

of the interval, form an intention to perform the whole of a dominating sequence of acts—that is, a sequence of acts whose outcome would have been better for everyone. There might be sequences of acts such that, even though one can intentionally perform each act in the sequence individually, one cannot at the start of the sequence form an intention to perform the whole sequence. We say that one intentionally performs a sequence of acts during $t-t'$ if and only if one performs it during $t-t'$ and for all times t'' such that $t \leq t'' \leq t'$, one forms or has formed at t'' an intention to perform all the at t'' still unperformed acts in the sequence.⁵ Hence we further weaken the principle to cases where one can intentionally perform a dominating sequence of acts.

Third, one might think that it is a category mistake to predicate wrongness to sequences of acts rather than to individual acts. Nonetheless, this worry may also be sidestepped. It need not be the whole dominated sequence of acts that is wrong. We need not claim more than that one performs at least one wrong act if one performs a dominated sequence of acts.

Hence, to avoid these potential problems, we shall instead rely on a weaker, revised version of the principle. According to the revised and final version of

the sequential-dominance principle,

a person P performs at least one wrong act during $t-t'$ if

- (i) P knows at t which sequences of acts P can perform during $t-t'$,
- (ii) P knows at t the outcome for each of these sequences,
- (iii) P can intentionally perform a sequence of acts during $t-t'$ with outcome x , and
- (iv) P performs a sequence of acts during $t-t'$ whose outcome is worse for everyone than x .

The idea is that, if you knew all along that what you did was worse for everyone than some other sequence of acts you could have intentionally done instead, then you must have done something wrong. The intuitive appeal of the sequential-dominance principle should be obvious.⁶ Hence

⁵ See Carlson (2003, pp. 182–183) for a similar account of performability.

⁶ Nevertheless, some people—for example, Peterson (2013, pp. 24–25)—object to dominance principles in cases where the dominating outcome is more unequal than the dominated one. But we can set aside this problem for the purposes of this paper, because the dominating outcomes improve well-being by the same amount for everyone over the dominated outcomes in all cases where we apply the sequential-dominance principle in this paper. Moreover, a standard view in the literature on inequality is that inequality matters less at higher levels, see e.g. Temkin (2001, pp. 346–348). So the inequality in the dominating outcome should plausibly matter less than the inequality in the dominated one, given that everyone's well-being is improved by the same amount.

it should at least count as a [p. 1596] drawback for theories if they require one to violate it. But, apart from that, the relevant moral principles on Scanlon's version of contractualism are those that no one could reasonably reject.⁷ And, since *everyone* is knowingly made worse off if the sequential-dominance principle is violated, it seems likely that no one could reasonably reject it. Thus violations of the sequential-dominance principle should be especially problematic for contractualist theories of Scanlon's type.

For the example, let o_1 , o_2 , and o_3 be three outcomes that do not differ in other respects than the well-being of persons P_1 , P_2 , and P_3 . Each person's well-being in these outcomes is given in the following table:⁸

	P_1	P_2	P_3
o_1	1	2	3
o_2	3	1	2
o_3	2	3	1

Furthermore, let o_1^- , o_2^- , and o_3^- be three outcomes just like o_1 , o_2 , and o_3 except that everyone has slightly less well-being. The well-being of P_1 , P_2 , and P_3 in these outcomes can be given as follows, where ϵ is an arbitrarily small positive number:

	P_1	P_2	P_3
o_1^-	$1 - \epsilon$	$2 - \epsilon$	$3 - \epsilon$
o_2^-	$3 - \epsilon$	$1 - \epsilon$	$2 - \epsilon$
o_3^-	$2 - \epsilon$	$3 - \epsilon$	$1 - \epsilon$


Consider first a case where one has a single choice between all six of o_1 , o_2 , o_3 , o_1^- , o_2^- , and o_3^- . This non-sequential case does not seem to pose any serious problems for the anti-aggregation principle. Although it is unclear how one should calculate harms in non-binary choices, it seems clear that it is wrong to choose one of o_1^- , o_2^- , and o_3^- , since they are each dominated in this choice by an alternative outcome that would be better for everyone. And it seems clear that it is permissible to choose any one of the remaining outcomes o_1 , o_2 , and o_3 , since their well-being distributions are simply permutations of each other.

But consider instead the following decision tree:⁹ [p. 1597]

⁷ Scanlon (1998, p. 197).

⁸ For our discussion, we assume cardinal measurability of individual well-being and interpersonal comparability of gains and losses of well-being. We need to be able to make interpersonal comparisons of differences in the seriousness of harms in order to apply the anti-aggregation principle.

⁹ Cantwell (2003, p. 389) uses a decision tree of this form in a money-pump argument against cyclic preferences.



In this case, one knows at t_1 which sequences of acts one can perform during t_1 – t_3 , and one also knows their outcomes. Moreover, one can at t_1 intentionally perform each of these sequences of acts. Their respective outcomes are o_1 , o_2 , o_3 , o_1^- , o_2^- , and o_3^- . Since o_1^- , o_2^- , and o_3^- are worse for everyone than o_1 , o_2 , and o_3 respectively, the sequential-dominance principle yields that one does something wrong if one performs one of the sequences with outcomes o_1^- , o_2^- , and o_3^- . Regardless of what one does in S_1 , one will face one of the situations S_2 , S_3 , and S_4 at t_2 . Should one, for instance, face S_2 at t_2 , one has a choice between o_1 and o_2^- . The loss in well-being for P_1 if one chooses o_1 over o_2^- is almost twice as great as the loss in well-being for P_2 and P_3 if one chooses o_2^- over o_1 . It seems then that by choosing o_1 in S_2 one saves a greater number of people from minor harms rather than a lesser number from more serious harms. We assume here the standard counterfactual view of harm, where an act harms a person P if and only if P would be worse off if the act were performed than P would be if the act were not performed—we shall, however, consider lifting this assumption later.¹⁰ It follows that the anti-aggregation principle does not permit choosing o_1 in S_2 . We get, *mutatis mutandis*, that the anti-aggregation principle permits neither choosing o_2 in S_3 nor choosing o_3 in S_4 . So following the anti-aggregation principle forces one to choose one of o_1^- , o_2^- , and o_3^- in S_2 , S_3 , and S_4 . But, if one chooses one of o_1^- , o_2^- , and o_3^- in S_2 , S_3 , or S_4 , one performs a sequence of acts that one knew would be worse for everyone than an alternative sequence one could have intentionally performed during t_1 – t_3 . Thus, in this case, following the anti-aggregation principle forces one to do something wrong according to the sequential-dominance principle.

One might object that the one who is harmed the most in each pairwise choice at t_2 is only harmed about twice as much as the others. Hence

¹⁰ See e.g. Kavka (1982, p. 96) and Parfit (1984, p. 69).

one might think that the anti-aggregation principle does not apply, since the more harmed person does not face *much* more serious harm. Nevertheless, we can easily raise the stakes, at the cost of a more complicated example.

For example, we can add one more person and one more alternative outcome as follows: [p. 1598]

	P_1	P_2	P_3	P_4
o_1	1	2	3	4
o_2	4	1	2	3
o_3	3	4	1	2
o_4	2	3	4	1

We then get that P_1 is harmed almost three times more if one chooses o_1 over o_2^- than the other people are if one chooses o_2^- over o_1 , and so on, where again o_2^- is like o_2 but everyone has slightly less well-being.

More generally, we can get an example where the harm for the one is arbitrarily as many times greater than the avoided harms for the many with the following scheme:

	P_1	P_2	P_3	P_4	...	P_n
o_1	1	2	3	4	...	n
o_2	n	1	2	3	...	$n-1$
o_3	$n-1$	n	1	2	...	$n-2$
o_4	$n-2$	$n-1$	n	1	...	$n-3$
\vdots	\vdots	\vdots	\vdots	\vdots	\ddots	\vdots
o_n	2	3	4	5	...	1

As before, we let o_i^- be just like o_i except that everyone has slightly less well-being. In this version, one knows at t_1 that one can perform $2n$ alternative sequences of acts during t_1-t_3 with outcomes $o_1, o_2, \dots, o_n, o_1^-, o_2^-, \dots$, and o_n^- . And each of these sequences is intentionally performable at t_1 . Then one faces a choice at t_2 between o_1 and o_2^- , between o_2 and o_3^- , ... , or between o_n and o_1^- . At t_2 , one is not permitted to choose one of o_1, o_2, \dots , and o_n , because that would be to save a large number of people from minor harms rather than one person from an almost $n-1$ times as serious harm. So, in order to follow the anti-aggregation principle, one has to choose one of o_1^-, o_2^-, \dots , and o_n^- at t_2 . But, in that case, one performs a sequence of acts that one knew would be worse for everyone than an alternative sequence of acts one could have intentionally performed during t_1-t_3 . Since, in this example, the harm for the one might definitely be much more serious than the avoided harm for the many, the anti-aggregation principle applies and yields thus the conflict with the sequential-dominance principle.

Another possible objection is that someone who knows that she will follow the anti-aggregation principle in all future situations can predict that she would not stick to a plan to perform the sequence of acts with outcome o_1 . And, if so, she might not be able to form an intention to perform that sequence of acts, because she predicts that it will not be performed. Hence someone like this could not face the type of example I have presented.¹¹ But note that the anti-aggregation principle is a general [p. 1599] principle that should also apply to people who are able to at least form intentions to do acts that violate the principle. Such people might face situations like those in my examples, even if they follow the anti-aggregation principle on every occasion. And, in such cases, they will be forced to violate either the sequential-dominance principle or the anti-aggregation principle.

Furthermore, one might object that the anti-aggregation principle is perhaps not meant to be applied to sequential choices. If so, my argument might seem to miss its mark. In my argument, however, we only need to apply the anti-aggregation principle in the final choice nodes of the decision tree. And, in these final choice nodes, the choices are not sequential, because their outcomes do not depend on any further choices. Accordingly, since my argument only applies the anti-aggregation principle to non-sequential choices, the principle cannot be saved by limiting its application to such choices.

But one might instead have the reverse reaction and try to save the anti-aggregation principle by limiting its application so that it only applies to initial choices—and not to later, succeeding choices—in sequential-choice situations.¹² While this revision might block my above argument, it has at least two problems.

The first problem is that the initial choice of one sequential-choice situation might be one of the succeeding choices in an earlier sequential-choice situation. Hence one would have to complement the proposal with some criterion for picking out the relevant sequential-choice situation of a given choice. To rule out counter-examples like the ones above in this manner, it seems that one cannot allow that the anti-aggregation principle is applied to sequential-choice situations whose initial choice is part of an earlier sequential-choice situation. And, since the initial choice in most sequential-choice situations is part of an earlier sequential-choice situation, this would in most cases require that the anti-aggregation principle is applied to a sequential-choice situation whose initial choice was made a long time ago and includes all possible succeeding choices at

¹¹ Cf. R. H. Strotz's (1955–1956, p. 173) sophisticated-choice proposal and Levi's (2002, p. S241) similar objection to sequential money pumps.

¹² Cf. McClennen's (1990, p. 13) resolute-choice proposal.

least up until one's present choice. This, however, would make it very hard to get any action guidance from the anti-aggregation principle, because one would have to consider a potentially enormous decision tree going all the way back to some initial choice in the distant past. And then it might not be clear whether the principle would, for example, permit one to put off the rescue in Scanlon's television-studio example, since this might depend not only on one's available acts in that situation but also on the sequences of acts one could have performed starting from some earlier choice situation.

The second problem is that, in order to avoid the conflict with the sequential-dominance principle in my examples, this revision must not at t_2 forbid one from saving a larger number of people from minor harms rather than a smaller number who face much more serious injuries. Since the situations at t_2 in my examples could be of the same type as the television-studio example, it follows that this revised principle does not forbid some options that are analogous to putting off the rescue until the match is over. Hence Scanlon's intuition that one should not put off the rescue is at odds with, rather than support for, this revision of his principle. [p. 1600]

Finally, one might object that one can block my argument by adopting an alternative view of harm where someone is harmed only if they suffer a decrease in well-being. Instead of the counterfactual view we assumed above, we could adopt the worsening view of harm where, roughly, one is harmed by an act if and only if the act causes one to have a lower well-being after some time t than one had before t .¹³ On this view, there need not be anyone who is harmed in my examples. Hence my argument would be blocked. The worsening view, however, does not remove the conflict between the anti-aggregation principle and the sequential-dominance principle. To see this, consider a case where you get four successive opportunities to either do an act that harms three of four people or an act that does more serious harm to the remaining person. In each of these four choice situations, the person who faces a more serious harm does not face a more serious harm in the other choices. So the first choice is between either seriously harming P_1 or much less seriously harming P_2 , P_3 , and P_4 , the second choice is between either seriously harming P_2 or much less seriously harming P_1 , P_3 , and P_4 , and so on. Furthermore, suppose that from the time of the first choice you know that you are going to face these four choices and you also know the outcomes of all alternative acts. More specifically, you know each act's effect on the well-being of each person. In the table below, the number for a combination of an act and a person

¹³ Feinberg (1986, p. 149) discusses and rejects a similar version of the worsening view. See, however, Perry (2003) for a defence.

represents how the act affects the person's well-being.

		P_1	P_2	P_3	P_4
First choice	a_1	0	-1	-1	-1
	a_2	-2	0	0	0
Second choice	a_3	-1	0	-1	-1
	a_4	0	-2	0	0
Third choice	a_5	-1	-1	0	-1
	a_6	0	0	-2	0
Fourth choice	a_7	-1	-1	-1	0
	a_8	0	0	0	-2

Given the worsening view of harm, each one of a_2 , a_4 , a_6 , and a_8 is wrong according to the anti-aggregation principle.¹⁴ The trouble is that, if one were to refrain from doing these acts and thus instead do a_1 , a_3 , a_5 , and a_7 , everyone would be worse off in the end than if one were to do a_2 , a_4 , a_6 , and a_8 .¹⁵ Hence, even on the worsening view of harm, one cannot both follow the anti-aggregation principle and avoid doing something wrong according to the sequential-dominance principle. [p. 1601]

I wish to thank Arif Ahmed, Per Algander, Gustaf Arrhenius, Marc Fleurbaey, Iwao Hirose, Christopher Jay, Martin Peterson, Christian Piller, Mozaffar Qizilbash, Wlodek Rabinowicz, Daniel Ramöller, Jussi Suikkanen, and an anonymous referee for valuable comments.

References

- Cantwell, John (2003) 'On the Foundations of Pragmatic Arguments', *The Journal of Philosophy* 100 (8): 383–402.
- Carlson, Erik (2003) 'Dynamic Inconsistency and Performable Plans', *Philosophical Studies* 113 (2): 181–200.
- Feinberg, Joel (1986) 'Wrongful Life and the Counterfactual Element in Harming', *Social Philosophy & Policy* 4 (1): 145–178.
- Kavka, Gregory S. (1982) 'The Paradox of Future Individuals', *Philosophy & Public Affairs* 11 (2): 93–112.
- Levi, Isaac (2002) 'Money Pumps and Diachronic Books', *Philosophy of Science* 69 (3): S235–S247.

¹⁴ Like before, if one does not think that a twice as serious harm is a *much* more serious harm, we can, in a similar manner, easily raise the stakes by increasing the number of choices and people.

¹⁵ Parfit (2003, p. 385) presents a similarly structured case.

- McClennen, Edward F. (1990) *Rationality and Dynamic Choice: Foundational Explorations*, Cambridge: Cambridge University Press.
- Parfit, Derek (1984) *Reasons and Persons*, Oxford: Clarendon Press.
- (2003) ‘Justifiability to Each Person’, *Ratio* 16 (4): 368–390.
- Perry, Stephen (2003) ‘Harm, History, and Counterfactuals’, *San Diego Law Review* 40 (4): 1283–1313.
- Peterson, Martin (2013) *The Dimensions of Consequentialism: Ethics, Equality and Risk*, Cambridge: Cambridge University Press.
- Raz, Joseph (2011) *From Normativity to Responsibility*, Oxford: Oxford University Press.
- Scanlon, T. M. (1998) *What We Owe to Each Other*, Cambridge, MA: Harvard University Press.
- Strotz, R. H. (1955–1956) ‘Myopia and Inconsistency in Dynamic Utility Maximization’, *The Review of Economic Studies* 23 (3): 165–180.
- Taurek, John M. (1977) ‘Should the Numbers Count?’, *Philosophy and Public Affairs* 6 (4): 293–316.
- Temkin, Larry S. (2001) ‘Inequality: A Complex, Individualistic, and Comparative Notion’, *Noûs* 35 (s1): 327–353.
- (2012) *Rethinking the Good*, Oxford: Oxford University Press.