

Ex-Post Average Utilitarianism Can Be Worse for All Affected, even if No Contingently Existing Person Is Affected

Johan E. Gustafsson and Dean Spears*

Draft: October 29, 2024 at 5:59 p.m.

ABSTRACT. According to *Ex-Post Average Utilitarianism*, prospect X is at least as good as prospect Y if and only if the expected average well-being is at least as great in X as in Y . Relative to the *ex-ante* approach of taking the average of peoples' expectations, this *ex-post* approach has the advantage of not needing well-defined expectations of well-being for contingent people — people who exist in some but not all states of nature. Nevertheless, we show that *Ex-Post Average Utilitarianism* can oppose the interests of all affected persons. Moreover, we show this without relying on any comparisons of expectations of well-being for contingent people: Our objection can be made with cases in which no contingently-existing person is affected. Finally, we show that our objection can be made even if lifetime well-being has only an ordinal structure (in which case prior objections to Average Utilitarianism would not apply).

Average utilitarianism is the view that the value of a final outcome is equal to the average well-being in that outcome.¹ This view about how to evaluate final outcomes can be extended in several ways to also evaluate prospects — that is, probability distributions over possible final outcomes. The most straightforward way to do so is to let the value of a prospect be equal to the prospect's expected value. This view is called *Ex-Post Average Utilitarianism*:²

Ex-Post Average Utilitarianism Prospect X is at least as good as

* We would be grateful for any thoughts or comments on this paper. They can be sent to johan.eric.gustafsson@gmail.com.

¹ Sidgwick (1907, p. 415) distinguished average and total utilitarianism, although he (1907, pp. 415–16) favoured the latter. The name 'average utilitarianism' comes from Rawls 1971, p. 599.

² Harsanyi 1985, p. 44; 1986, p. 57. Nevertheless, in his correspondence with Ng (1983, p. 168), Harsanyi seems to defend *Ex-Ante Average Utilitarianism*. This conflicts with Harsanyi's (1955, p. 313) commitment to the expectation taking for the social ranking.

prospect Y if and only if the expected average well-being is at least as great in X as in Y .

A rival to *Ex-Post* Average Utilitarianism is *Ex-Ante* Average Utilitarianism, which assess prospects by the average of each person's expected well-being conditional on that person existing:³

Ex-Ante Average Utilitarianism A prospect x is at least as good as prospect y if and only if the average conditional-on-existence expected well-being for possible people in x is at least as great as the average conditional-on-existence expected well-being for possible people in y .

An advantage of *Ex-Post* Average Utilitarianism over *Ex-Ante* Average Utilitarianism is that the former maximizes an expectation, which means it satisfies expected utility theory for general betterness.⁴ Nevertheless, we show that *Ex-Post* Average Utilitarianism can oppose the interests of all affected persons. Moreover, we can show this without assuming any well-defined expectations of well-being for contingent people — indeed while leaving contingent people entirely unaffected.

³ Ng 1983, p. 168.

⁴ *Ex-Ante* Average Utilitarianism violates expected utility theory. For instance, *Ex-Ante* Average Utilitarianism violates statewise dominance (even in non-sequential choices). Consider the following case where there are two possible states of nature S_1 and S_2 with an equal probability.

	Prospect A		Prospect B	
	S_1	S_2	S_1	S_2
	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$
Ann	5	Ω	6	2
Bob	5	1	Ω	2

In S_1 , the average well-being is 5 in A and 6 in B , and, in S_2 , the average well-being is 1 in A and 2 in B . So B is better than A in every state of nature according to average utilitarianism. But, according to *Ex-Ante* Average Utilitarianism the value of A is 4 and the value of B is 3; so A is better than B . (The same counterexample also works against other *ex ante* approaches that would aggregate peoples' conditional-on-existence expected well-being while ignoring probabilities of their existence — such as an approach that values each prospect according to the minimum, among possible people, of individual expected well-being conditional on existence.)

1. Prior objections to average utilitarianism

The literature already contains many strong objections to average utilitarianism. Without denying these objections, an advantage of our contribution is that it requires less assumptions about value than the standard objections to *Ex-Post* Average Utilitarianism require: Though we present our core example using real-valued lifetime well-being levels for simplicity, we will show that our objection would work if lifetime well-being were merely ordinally ranked (in which case prior prominent objections to Average Utilitarianism would not apply). So, while we claim that it's a dialectical advantage of our argument that it cannot be met by rejecting these assumptions, we don't deny that these objections are compelling.

One standard objection concerns the addition of lives of negative well-being. Like every form of average utilitarianism, the *ex-post* variant is vulnerable to the Sadistic Conclusion:⁵

The Sadistic Conclusion When adding people without affecting the original peoples' well-being, it sometimes can be better to add a number of people with negative well-being, rather than a number of people with positive well-being.

A potential response to this objection is that the average utilitarian can reject the meaningfulness of the concept of negative well-being. Negative well-being levels are plausibly those levels that are less good than the neutral well-being level.⁶ But how are we to understand the neutral well-being level? One compelling idea is that the neutral well-being level is the level that is equal to the well-being level of non-existence. Yet some prior research has questioned whether non-existing people would have any well-being at all.⁷ And, crucially, *Ex-Post* Average Utilitarians are not committed to any comparisons of well-being between existence and non-existence.⁸ They may consistently deny the underlying assumption of the existence of negative well-being in these objections.⁹ There are, of course,

⁵ Arrhenius and Bykvist 1995, p. 85 and Arrhenius 2000b, p. 251. Parfit (1984, p. 406) also puts forward an objection based on the addition of lives with negative well-being in his two-hells case.

⁶ Chisholm and Sosa 1966, pp. 247–8.

⁷ Williams 1973, p. 87, Parfit 1984, p. 487, and Broome 1993, p. 77.

⁸ See Harsanyi in Ng 1983, pp. 168–9.

⁹ Note further that denying the existence of negative well-being also blocks many other conditions in population ethics, such as the Mere-Addition Principle (Parfit 1984,

other accounts of the neutral well-being level.¹⁰ But those accounts, too, are contestable.¹¹

It may be objected that, even without any notions of negative well-being, *Ex-Post* Average Utilitarianism still faces a strong objection. The theory entails that adding a small number of people with lives barely worth living can sometimes be better than adding a large number of people with magnificent lives—this would happen if the background population to which lives were being added had an average that was even more magnificent. But this variation of the objection still makes assumptions about absolute well-being. Lives that are barely worth living seem to be lives at a well-being level that is barely above the neutral well-being level, and magnificent lives seem to be lives at a well-being level that is much higher than the neutral level of well-being.

Another standard objection to average utilitarianism is the Egyptology objection.¹² It is the objection that, given average utilitarianism, the evaluation of acts today may depend on how well off people were in ancient Egypt (or in other distant and unaffected subpopulations). If the ancient Egyptians were very well off, then the addition of a person with a certain level of well-being could be worse than not adding them (because they would lower the average) even though, if the ancient Egyptians were, instead, less well off, the addition of the person would be better (because they would raise the average). A potential response to this objection is that the average utilitarian could maintain that whether the well-being of the ancient Egyptians were high or low is morally relevant, so sensitivity to such facts is not a drawback.

A final standard objection to average utilitarianism is the utility monster.¹³ The objection is that, even supposing that everyone in history would live very good lives, it would be better according to average utilitarianism if there had only been a single person that is just slightly happier. A weakness of this objection, in a dialectic between total and average utilitarian-

p. 420 and Ng 1989, pp. 537–8), which rules out Average Utilitarianism, and the Repugnant Conclusion (Parfit 1984, p. 388), which, in a variant that compares *additions* to populations, rules out Average Utilitarianism (Anglin 1977, p. 746 and Spears and Boudolfson 2021, p. 574).

¹⁰ See, for example, Arrhenius 2000a, pp. 17–27.

¹¹ Gustafsson 2020, pp. 96–100. Note, again, that we are only claiming that the assumption of a neutral level of well-being is a dialectical disadvantage; we are not denying that there is such a level.

¹² McMahan 1981, p. 115 and Parfit 1984, p. 420.

¹³ Nozick 1974, p. 41.

ism, is that much the same objection also applies to total utilitarianism (if the single person is sufficiently happy to have more total well-being). To assess this objection, notice that, for the single-person population to be better on total utilitarianism, their welfare has to be much higher — not just slightly higher than the average. But, rather than categories of absolute value, this objection relies on categories of differences of value. That is, some differences in well-being are slight and some very large. Of course, the size of the differences are morally relevant, but this objection needs these categories of differences to be different in kind in their moral relevance, which is questionable.

Although many of these objections are strong, they are not conclusive. Our new objection blocks defences from an average utilitarian who is willing to insist that the structure of utility is merely ordinal. And it illustrates how *Ex-Post* Average Utilitarianism quickly encounters problems in a setting of social risk. These problems are conceptually distinct from these prior objections.

2. Our objection

In this paper, we present a new objection to *Ex-Post* Average Utilitarianism, which avoids the drawbacks of the earlier objections.¹⁴ Accordingly, the above responses to those objections do not apply to this new objection.

Consider the following prospects, where columns are risky states of nature, rows are people (including contingent people), and cells are outcomes which could include ordered well-being levels, represented by numbers, or non-existence, represented by Ω :

	Prospect A		Prospect B	
	S_1	S_2	S_1	S_2
	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$
Ann	1	Ω	1	Ω
Bob	1	7	9	1

There are two possible states of nature S_1 and S_2 with an equal probability. In both prospects, Ann exists with a well-being of 1 in S_1 but she does not exist at all in S_2 . Bob, on the other hand, exists in all states of nature in

¹⁴ Or, for a reader persuaded by these prior objections to Average Utilitarianism, our new example adds to their strength.

both prospects. In A , Bob has a well-being of 1 in S_1 and a well-being of 7 in S_2 . In B , Bob has a well-being of 9 in S_1 and a well-being of 1 in S_2 .

Part of the motivation for average utilitarianism is that it does not rely on comparisons of well-being between existence and non-existence. Indeed it is unclear whether a prospect could be better-for a person who would exist only in some possible outcomes and would not exist in others. We will say that a person is *affected* in a comparison between two prospects X and Y if and only if there is a state of nature in which that person's outcome (which could be either existence with a well-being level or non-existence) is different in X than in Y ; otherwise, the person is *unaffected* in that comparison. Given this distinction, we can rely on a weakened form of *ex-ante*-Pareto dominance that only applies to people who are affected in a comparison and who are necessary in a comparison (that is, people who exist in all possible states):

Necessity-Restricted All-Affected Stochastic Pareto If (i) all contingent people are unaffected in a comparison between prospects X and Y and (ii), for each affected person, X stochastically dominates Y , then X is better than Y .

Notice that Necessity-Restricted All-Affected Stochastic Pareto does not invoke expectation-taking even for necessary people: It merely relies on stochastic dominance.

It may be helpful to compare Necessity-Restricted All-Affected Pareto with Anteriority:¹⁵

Anteriority If every potential person faces the same prospect in X as in Y , then X and Y are equally good.

Anteriority, combined with some further assumptions, rules out average utilitarianism.¹⁶ But an advantage of relying on Necessity-Restricted All-Affected Stochastic Pareto rather than Anteriority is that the former (unlike the latter) does not assume that we are indifferent between prospects which are the same except that some contingent person (who exists with the same well-being and probability) exists in different states of nature. Contingent people do not have well-defined expectations, because they

¹⁵ McCarthy 2017, p. 226. We have added 'potential' to highlight that the involved people need not exist.

¹⁶ Thomas 2022, pp. 280–2.

do not have a well-being level in all potential final outcomes.¹⁷ So Anteriority cannot be supported by the claim that each person has the same expectation in the prospects. Given a variant of our example that replaces the 9 with a 7, Anteriority rules out *Ex-Post* Average Utilitarianism without further assumptions.

Returning to our example, Ann is unaffected by a choice between *A* and *B*. But, for Bob, *B* stochastically dominates *A*. Hence, for the only affected person (in either existence or well-being), *B* is clearly better than *A*. So, according to Necessity-Restricted All-Affected Stochastic Pareto, *B* would be better than *A*. And yet, *Ex-Post* Average Utilitarianism entails that *A* is better than *B*, because the expected average well-being is 4 in *A* but just 3 in *B*.¹⁸ Therefore, in the comparison of *A* and *B*, *Ex-Post* Average Utilitarianism opposes the interests of all affected persons.

This counter-example to *Ex-Post* Average Utilitarianism avoids the drawbacks of the earlier objections. It does not rely on negative well-being. Unlike the unaffected ancient Egyptians in the Egyptology objection who could be relevantly better or worse off, this counter-example does not rely on any morally relevant change to the unaffected Ann — as her prospect is exactly the same in both prospects. Finally, unlike the utility-monster objection, this counter-example does not work against total utilitarianism.¹⁹

¹⁷ Blackorby et al. 1998, p. 10.

¹⁸ Note, moreover, that *Ex-Post* Average Utilitarianism would still favour *A* even if Ann's well-being in S_1 of *B* were 2.

¹⁹ By the same reasoning as in our example, Necessity-Restricted All-Affected Stochastic Pareto is also incompatible with some notable *ex-post* extensions to social risk of other non-separable population axiologies defined elsewhere in the literature, such as variable-value utilitarianism (Hurka 1983, Ng 1989, pp. 244-250, formalized as 'number-dampened utilitarianism' by Blackorby et al. 2005, pp. 144-7); rank-discounted utilitarianism (Asheim and Zuber, 2014, pp. 632); variable-population extensions of equally-distributed-equivalent egalitarianism (Fleurbaey 2010, pp. 657-658 for fixed-population cases; Spears and Zuber forthcoming); and variable-population expected maximin, which would value a prospect according to the expectation across states of the minimum well-being among people alive in that state. We explore these extensions further by using a principle related to Necessity-Restricted All-Affected Stochastic Pareto to characterize an additively-separable family of variable-population social welfare functions in a companion paper for the economic theory literature that cites the priority of this paper.

3. A generalization to ordinal utility

Our objection against average utilitarianism can be made even if lifetime well-being levels or utilities are merely ordinally ranked so no categories (like neutral, zero, barely worth living, negative, or excellent) are meaningful and so no differences in well-being levels are meaningful. To see this, consider a set U that contains all the possible lifetime utilities and is ordered, so lifetime well-being level u is better, worse, or just as good as lifetime well-being level v , but that is all one can say.

This is enough to define a generalized average. A generalized average takes any population — here meaning any set of lifetime well-beings in U — and returns an element of U that represents them, independently of size. A few rules would be necessary for any generalized average function g . First, the generalized average of a population in which every person has the same well-being level is that well-being level (this includes one-person populations).

$$(1) \quad g(u, u) = g(u) = u.$$

Second, the generalized average ignores the order in which people are listed, so the generalized average of u and v is the same as the generalized average of v and u .

$$(2) \quad g(u, v) = g(v, u).$$

The third rule also amounts to a requirement that U be a rich enough set for an average to make sense. This third rule says that if u is a better lifetime well-being than v , then there exists a u^+ that is better than u such that the generalized average of u^+ and v is worse than u .

$$(3) \quad \text{If } u \succ v, \text{ then there exists } u^+ \text{ such that } u^+ \succ u \text{ and } u \succ g(u^+, v).$$

Now, consider two prospects:

	Prospect C	Prospect D
	S_1 S_2	S_1 S_2
	$\frac{1}{2}$ $\frac{1}{2}$	$\frac{1}{2}$ $\frac{1}{2}$
Ann	v Ω	v Ω
Bob	v u	u^+ v

An *ex post* approach using any g would reduce C to $\langle v, u \rangle$ and would reduce D to $\langle g(u^+, v), v \rangle$. By social stochastic dominance, C would be judged better, even though the only affected person is better-off in D .

We wish to thank Kacper Kowalczyk for valuable comments.

References

- Anglin, Bill (1977) ‘The Repugnant Conclusion’, *Canadian Journal of Philosophy* 7 (4): 745–754.
- Arrhenius, Gustaf (2000a) *Future Generations: A Challenge for Moral Theory*, Ph.D. thesis, Uppsala University.
- (2000b) ‘An Impossibility Theorem for Welfarist Axiologies’, *Economics and Philosophy* 16 (2): 247–266.
- Arrhenius, Gustaf and Krister Bykvist (1995) *Future Generations and Interpersonal Compensations: Moral Aspects of Energy Use*, Uppsala: Uppsala University.
- Asheim, Geir and Stéphane Zuber (2014) ‘Escaping the Repugnant Conclusion: Rank-Discounted Utilitarianism with Variable Population’, *Theoretical Economics* 9 (3): 629–650.
- Blackorby, Charles, Walter Bossert, and David Donaldson (1998) ‘Uncertainty and Critical-Level Population Principles’, *Journal of Population Economics* 11 (1): 1–20.
- (2005) *Population Issues in Social Choice Theory, Welfare Economics, and Ethics*, Cambridge: Cambridge University Press.
- Broome, John (1993) ‘Goodness Is Reducible to Betterness: The Evil of Death Is the Value of Life’, in Peter Koslowski and Yuichi Shionoya, eds., *The Good and the Economical: Ethical Choices in Economics and Management*, pp. 70–84, Berlin: Springer.
- Chisholm, Roderick M. and Ernest Sosa (1966) ‘On the Logic of “Intrinsically Better”’, *American Philosophical Quarterly* 3 (3): 244–249.
- Fleurbaey, Marc (2010) ‘Assessing Risky Social Situations’, *Journal of Political Economy* 118 (4): 649–680.
- Gustafsson, Johan E. (2020) ‘Population Axiology and the Possibility of a Fourth Category of Absolute Value’, *Economics and Philosophy* 36 (1): 81–110.
- Harsanyi, John C. (1955) ‘Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility’, *The Journal of Political Economy* 63 (4): 309–321.

- (1985) ‘Does Reason Tell Us What Moral Code to Follow and, Indeed, to Follow Any Moral Code at All?’, *Ethics* 96 (1): 42–55.
- (1986) ‘Utilitarian Morality in a World of Very Half-hearted Altruists’, in Walter P. Heller, Ross M. Starr, and David A. Starrett, eds., *Social Choice and Public Decision Making: Social Choice and Public Decision Making: Essays in Honor of Kenneth J. Arrow, Volume I*, pp. 57–73, Cambridge: Cambridge University Press.
- Hurka, Thomas (1983) ‘Value and Population Size’, *Ethics* 93 (3): 496–507.
- McCarthy, David (2017) ‘The Priority View’, *Economics and Philosophy* 33 (2): 215–257.
- McMahan, Jeff (1981) ‘Problems of Population Theory’, *Ethics* 92 (1): 96–127.
- Ng, Yew-Kwang (1983) ‘Some Broader Issues of Social Choice’, in Prasanta K. Pattanaik and Maurice Salles, eds., *Social Choice and Welfare*, pp. 151–173, Amsterdam: North-Holland.
- (1989) ‘What Should We Do about Future Generations? Impossibility of Parfit’s Theory X’, *Economics and Philosophy* 5 (2): 235–253.
- Nozick, Robert (1974) *Anarchy, State, and Utopia*, New York: Basic Books.
- Parfit, Derek (1984) *Reasons and Persons*, Oxford: Clarendon Press.
- Rawls, John (1971) *A Theory of Justice*, Cambridge, MA: Harvard University Press.
- Sidgwick, Henry (1907) *The Methods of Ethics*, London: Macmillan, seventh edn.
- Spears, Dean and Mark Budolfson (2021) ‘Repugnant Conclusions’, *Social choice and welfare* 57:567–588.
- Spears, Dean and Stéphane Zuber (forthcoming) ‘Foundations of Utilitarianism under Risk and Variable Population’, *Social Choice and Welfare*.
- Thomas, Teruji (2022) ‘Separability and Population Ethics’, in Gustaf Arrhenius, Krister Bykvist, Tim Campbell, and Elizabeth Finneron-Burns, eds., *The Oxford Handbook of Population Ethics*, pp. 271–295, New York: Oxford University Press.
- Williams, Bernard (1973) *Problems of the Self: Philosophical Papers 1956–1972*, Cambridge: Cambridge University Press.