Bounded Utilities and Ex Ante Pareto

Petra Kosonen

petra.kosonen@austin.utexas.edu

11/17/2022

Bounded Expected Totalism

- I will investigate the compatibility of two standard theories: Total Utilitarianism and Expected Utility Theory with a bounded utility function.
- Let's call the combination of these views *Bounded Expected Totalism*.
- I will argue that Bounded Expected Totalism violates *Ex Ante Pareto*, the principle that what is in expectation better for everyone is better overall.

Bounded Expected Totalism

Both Total Utilitarianism and Expected Utility Theory with a bounded utility function are true.

Ex Ante Pareto

For all prospects X and Y, if X is at least as good as Y for everyone, and X is better than Y for some, then X is better than Y.

Total Utilitarianism

- Total Utilitarianism states that a population is better than another just in case the total quantity of well-being it contains is greater.
- Let $X \succsim Y$ mean that X is at least as good as Y.
- Also, let $\mathrm{W}(A)$ denote the total quantity of well-being in the state of affairs A and let $\mathrm{w}(S_i)$ denote the well-being of individual S_i.
- Then, more formally, Total Utilitarianism states the following:

Total Utilitarianism

For all states of affairs A and B (in which n and m individuals exist, respectively), $A \succsim B$ if and only if $\mathrm{W}(A) \geq \mathrm{W}(B)$, where

$$\mathrm{W}(A) = \sum_{i=1}^{n} w(S_i) \text{ and } \mathrm{W}(B) = \sum_{i=1}^{m} w(S_i).$$

Expected Utility Theory

- Next, let $\mathrm{EU}(X)$ denote the expected utility of prospect X.
- Also, let O be the set of possible outcomes, $p_X(o)$ the probability of outcome o in prospect X and $u(o)$ the utility of o.
- Then, Expected Utility Theory states the following:

Expected Utility Theory

For all prospects X and Y, $X \succsim Y$ if and only if $\mathrm{EU}(X) \geq \mathrm{EU}(Y)$, where

$$\mathrm{EU}(X) = \sum_{o \in O} p_X(o) u(o).$$

Boundedness

- What does it mean for utilities to be bounded?
- If utilities are real-valued, then boundedness means the following:

Boundedness

There is some $M \in \mathbb{R}$ such that for all outcomes x, $|u(x)| < M$.

- In other words, Boundedness rules out arbitrarily and infinitely good outcomes.
- Standard axiomatizations of expected utility maximization require utilities to be bounded.[1]

[1] See for example Kreps (1988, pp. 63–64), Fishburn (1970, pp. 194, 206–207), Hammond (1998, pp. 186–191) and Russell and Isaacs (2021).

Continuity

- Consider, for example, the von Neumann-Morgenstern axiomatization of Expected Utility Theory.
- Let $X \succ Y$ mean that X is better than Y.
- Also, let XpY be a risky prospect with a p chance of prospect X obtaining and a $1 - p$ chance of prospect Y obtaining.
- Then, if prospects are compared by their expected utilities, Boundedness follows from the following von Neumann-Morgenstern axiom:

Continuity

If $X \succ Y \succ Z$, then there are probabilities p and $q \in (0, 1)$ such that $XpZ \succ Y \succ XqZ$.

- To see why Continuity implies Boundedness, let's consider the two ways in which Boundedness might be false.

Continuity implies Boundedness

- First, Boundedness might be false because there is an infinite sequence of prospects A_1, A_2, A_3, … such that A_2 is at least twice as good as A_1, A_3 is at least twice as good as A_2, and so on, with respect to some baseline.
- Let A be a mixed prospect that assigns probability $1/2^k$ to prospect A_k.
- Then, we have that

$$\mathrm{EU}\left(A\right) = \sum_{i=1}^{\infty} p\left(A_i\right) u\left(A_i\right) = \infty.$$

- Next, choose some prospects B and C such that $\infty > \mathrm{EU}\left(B\right) > \mathrm{EU}\left(C\right) > -\infty$.
- Then, we have that A is better than B, which is better than C.
- However, for all probabilities $q \in (0, 1)$, $\mathrm{EU}\left(AqC\right) = \infty$.
- Therefore, AqC is better than B for all probabilities $q \in (0, 1)$. This is a violation of Continuity.[2]

[2] This is a modified argument from Kreps (1988, pp. 63–64).

Continuity implies Boundedness

- Secondly, and more generally, Boundedness is false if some prospect A is infinitely better than another (good) prospect B.
- This leads to a violation of Continuity because the mixed prospect ApC (where C certainly gives nothing) is better than B for all probabilities $p \in (0, 1)$.
- So, the supposition that Boundedness is false leads to violations of Continuity.
- Thus, it follows from Continuity that Boundedness is true.

Probability Fanaticism

- Boundedness has been discussed as a possible alternative to Probability Fanaticism:[3]

Probability Fanaticism (Positive)

For any probability $p > 0$, and for any (finitely) good outcome o, there is some great enough outcome O such that probability p of O (and otherwise nothing) is better than certainty of o.

- If one accepts Boundedness, then for any tiny probability of a great outcome, there is still some certain modest positive outcome that is worse.
- However, it is not the case that for any certain modest positive outcome, an *arbitrarily* small probability of a sufficiently great outcome is better.
- So, Boundedness avoids Probability Fanaticism.

[3] Wilkinson (2022, p. 449).

Bounded Expected Totalism

- Next, I will present Bounded Expected Totalism in more detail.
- Let *well-being* refer to how good some outcome is for an individual.
- And, let *social utility* refer to how good some outcome is overall, from an axiological point of view.
- Also, let *expected individual utility* represent how good some prospect is for an individual, and let *expected social utility* represent how good some prospect is overall.
- In the context of Expected Utility Theory, I will denote these by $\mathrm{EU}_{\mathrm{Ind}}$ and $\mathrm{EU}_{\mathrm{Soc}}$, respectively.
- In general, I will use *individual betterness* to refer to betterness from an individual's point of view.
- Similarly, I will use *overall/impersonal betterness* to refer to betterness from a moral point of view.

The social transformation function

- To combine Total Utilitarianism and Expected Utility Theory, we need a *social transformation function* that takes the total quantity of well-being as input and gives social utilities as output.
- This transformation function must be non-linear if an infinite or arbitrarily large number of happy individuals might exist, as then the total sum of individuals' well-being might be infinite or arbitrarily large (and similarly for negative well-being).
- But, as Bounded Expected Totalism requires expected social utilities to be bounded, the expected social utilities assigned to prospects that might result in an infinite or arbitrarily large number of happy individuals must be bounded.

The social transformation function

- One might object that the total quantity of well-being cannot be infinite or arbitrarily large because there is an upper limit to how many individuals might exist.
- This upper limit might be due to, for example, the Universe being finite.
- However, this may not be true, so we need a decision theory that can also handle these possibilities. If there is even a tiny probability that an infinite or arbitrarily large number of individuals exist, then the transformation function must be non-linear for utilities to be bounded.
- And, even if we were certain that there is an upper limit to how many individuals might exist, the total quantity of well-being might still be very large.
- So, Bounded Expected Totalism could still prescribe what might be considered fanatical choices in cases that involve tiny probabilities of huge outcomes.

The social transformation function

- Suppose that the social transformation function is non-linear. It will also have the following qualities:
- First, more well-being is always better, so the social transformation function must be strictly increasing with the total quantity of well-being; it must assign greater utilities to outcomes that contain more well-being.
- Secondly, because utilities are bounded above, similar increases in well-being must (after some point at least) matter less and less. Consequently, the social transformation function must be strictly concave on some subset of its domain. (Similarly: strictly convex if utilities are bounded below).
- Lastly, for utilities to be bounded, the social transformation function must be sufficiently concave with positive total well-being and sufficiently convex with negative total well-being; the contribution of additional (positive or negative) well-being to social utility must tend to zero.

Bounded Expected Totalism

- Let f be this transformation function.
- Also, let $p(A_i)$ denote the probability of state of affairs A_i, $\mathrm{W}(A_i)$ the total quantity of well-being in A_i and $\mathrm{w}(S_{ij})$ the well-being of individual S_j in A_i.
- Then, we can state Bounded Expected Totalism formally as follows:

Bounded Expected Totalism

Bounded Expected Totalism: For all prospects X and Y, $X \succsim Y$ if and only if $\mathrm{EU}_{\mathrm{Soc}}(X) \geq \mathrm{EU}_{\mathrm{Soc}}(Y)$, where

$$\mathrm{EU}_{\mathrm{Soc}}(X) = \sum_{i=1}^{n} p(A_i) f\left(\mathrm{W}(A_i)\right) = \sum_{i=1}^{n} p(A_i) f\left(\sum_{j=1}^{m} w(S_{ij})\right).$$

Summary

1. On Bounded Expected Totalism, when calculating the value of a prospect, one first calculates the total quantity of well-being in every possible state of the world.
2. Then, one transforms each state's total quantity of well-being into social utilities.
3. Finally, to get the expected social utility of a prospect, one multiplies the social utility of each state with that state's probability and sums these up.

- To summarize, social utilities might be bounded if the total quantity of well-being is itself necessarily bounded.
- However, this is not true; therefore, Bounded Expected Totalism requires a social transformation function that takes the total quantity of well-being as input and outputs social utilities.

The cardinal structure of well-being

- As mentioned above, the social transformation function takes the total quantity of well-being as input.
- To make sense of 'total quantity of well-being', we need well-being to have a 'cardinal structure', which allows us to make statements about *how much* more well-being an individual has in some outcome compared to another outcome.
- There are two ways of deriving the cardinal structure of well-being:

1. First, the cardinal structure of well-being might be understood in a 'primitivist' sense, according to which it can be defined independently of the individual betterness relation on gambles.
2. Alternatively, the cardinal structure of well-being might be understood in a technical sense as, for example, von Neumann-Morgenstern utilities. On the technical understanding, if the individual betterness relation satisfies a set of axioms, it can be represented by an expectational utility function.

Risk neutrality and Bernoulli's hypothesis

- If von Neumann-Morgenstern utilities represent the cardinal structure of well-being, then individual betterness is, by definition, risk-neutral with respect to well-being.
- It might still be risk-averse with respect to money or happy years of life.
- But it cannot be risk-averse with respect to well-being because well-being just is the quantity whose expectation the betterness relation can be represented as maximizing.
- This view satisfies the following principle:[4]

Bernoulli's hypothesis

One alternative is at least as good for a person as another if and only if it gives the person at least as great an expectation of their well-being.

[4] Broome (1991, p. 142). I have replaced 'good' with 'well-being'.

Risk-aversion with respect to well-being

- If Bernoulli's hypothesis is false, then individual betterness might be risk-averse with respect to well-being.
- For example, agents might be represented as maximizing risk-weighted expected utility.
- Alternatively, well-being could be understood in a primitivist sense.
- The primitivist view requires that quantities of well-being have meaning independently of how much they count when evaluating uncertain prospects.[5]
- But if such a metric of well-being is available, then individual betterness might be risk-averse with respect to this (non-technical) well-being.
- Note that this view is compatible with Expected Utility Theory (but not with Bernoulli's hypothesis).

[5] Broome (1991, p. 217).

Unbounded well-being and bounded utilities

- Next, let an *agent's transformation function* be a function that takes that person's well-being levels as input and outputs their individual utilities (to be used in decision-making under risk).
- If individual betterness over prospects is sufficiently risk-averse with respect to well-being, such that the agent's transformation function approaches asymptotically some upper bound with more well-being, then well-being itself can be unbounded without leading to unbounded utilities.

Prudential Fanaticism

- Finally, individual betterness might be risk-neutral with respect to well-being.
- And, happy days of life might not contribute less to well-being the more happy days the agent already has (and similarly for unhappy days).
- Given that individuals might live arbitrarily long at a constant positive well-being level, this view implies that both well-being and utilities are unbounded.
- This leads to a prudential analogue of Probability Fanaticism.

Prudential Fanaticism (Positive)

For any probability $p > 0$, and for any (finitely) good outcome o, there is some great enough outcome O such that probability p of O (and otherwise nothing) is prudentially better than the certainty of o for some individual S.

Summary

- To summarize, the social transformation function uses the 'total quantity of well-being' as input. To make sense of this notion, well-being must have a cardinal structure.
- This structure could be primitive, that is, given independently of individual betterness relation on gambles.
- Alternatively, it could be defined in a technical way, as for example von Neumann-Morgenstern utilities.
- If the cardinal structure is defined using the von Neumann-Morgenstern axioms, then individual betterness is risk-neutral.
- But if it is primitive, or defined in some other way, then it is at least initially an open question whether individual betterness is risk-neutral, risk-averse, or what, with respect to well-being.
- Next, I will show that Bounded Expected Totalism violates Ex Ante Pareto if individual betterness is risk-neutral with respect to well-being.

The risk-neutral case

- Let well-being levels be represented by real numbers.
- As argued above, the social transformation function f must be strictly concave on some subset of its domain.
- For the sake of argument, let's suppose it is strictly concave at 1.
- Then, there must be some positive constants δ and ϵ such that $f(1) - f(1-\delta) > f(1+\delta+\epsilon) - f(1)$.
- This is because the smaller benefit (δ) contributes more when added to a population at a lower well-being level than the greater benefit ($\delta + \epsilon$) when added to a population at a higher well-being level.

The risk-neutral case

Consider the following prospects:

The Risk-Neutral Case:

Risky Gives a 0.5 probability of a $1 + \delta + \epsilon$ well-being level; otherwise, it gives a well-being level of $1 - \delta$.

Safe Surely gives a well-being level of 1.

Suppose that the betterness relation of some agent, Alice, is risk-neutral with respect to her well-being.

Then, Risky is better than Safe for Alice (with all positive values of δ and ϵ), as Risky gives her a higher expectation of well-being than Safe does.

The risk-neutral case

- But is Risky also better than Safe impersonally?
- The answer is no.
- Given that the constants δ and ϵ are such that $f(1) - f(1-\delta) > f(1+\delta+\epsilon) - f(1)$, Safe is impersonally better than Risky.
- The situation is illustrated by the following graph:[6]

[6] Gustafsson (2022) presents this case to illustrate that *Ex-Post* Prioritarianism violates Ex Ante Pareto, a fact that goes back at least to Rabinowicz (2002).

The risk-neutral case

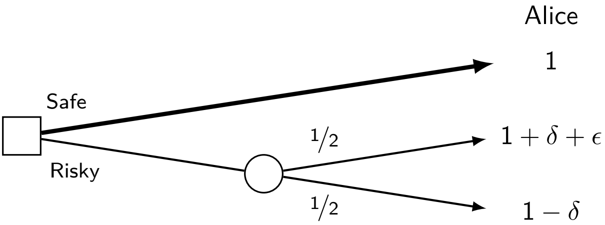

- Here, the square represents a choice node, while the circle represents a chance node.
- Going up at the choice node means accepting Safe, and going down at the choice node means accepting Risky.
- Thus, if we go up, Alice gets a well-being level of 1.
- On the other hand, if we go down, there are two possible states of the world, each with a 0.5 probability.
- In state 1, Alice gets a well-being level of $1 + \delta + \epsilon$. And, in state 2, Alice gets a well-being level of $1 - \delta$.

The risk-neutral case

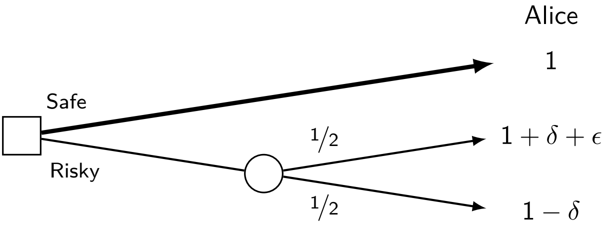

- The expected social utility of going up is $\text{EU}_{\text{Soc}}(\text{Safe}) = f(1)$.
- And, the expected social utility of going down is $\text{EU}_{\text{Soc}}(\text{Risky}) = \frac{1}{2} \cdot f(1+\delta+\epsilon) + \frac{1}{2} \cdot f(1-\delta)$.
- Given that $f(1) - f(1-\delta) > f(1+\delta+\epsilon) - f(1)$, $\text{EU}_{\text{Soc}}(\text{Risky})$ is less than $\text{EU}_{\text{Soc}}(\text{Safe})$.
- Thus, going up is impersonally better than going down, according to Bounded Expected Totalism.
- However, going down is better than going up for Alice (and equally good for everybody else). So, we have a violation of Ex Ante Pareto.

Summary

- To summarize, Bounded Expected Totalism violates Ex Ante Pareto if individual betterness is risk-neutral with respect to well-being.
- So, if well-being is understood as von Neumann-Morgenstern utilities, or in a primitive way and individual betterness is risk-neutral with respect to well-being, then the combination of Expected Utility Theory and Total Utilitarianism prescribes prospects that are better for none and worse for some.
- This happens because the social transformation function is concave on some subset of its domain.
- Consequently, Bounded Expected Totalism is at least sometimes risk-averse with respect to (positive) well-being.

The risk-averse case

- Next, I will show that Bounded Expected Totalism violates Ex Ante Pareto even if individual betterness is risk-averse with respect to well-being.
- It is already known that individual risk attitudes incompatible with Expected Utility Theory can cause tensions with Ex Ante Pareto.[7]
- However, the violation of Ex Ante Pareto discussed in this section happens even if the risk-aversion is of the kind that is compatible with Expected Utility Theory.

[7] See for example Nebel (2020) and Mongin and Pivato (2015).

The risk-averse case

- If individual betterness is risk-averse with respect to well-being, then it may no longer be true that Risky is better than Safe for Alice.
- So, Bounded Expected Totalism might not violate Ex Ante Pareto in the earlier case.
- If Alice's transformation function corresponds to the social transformation function when Alice is the only person who exists, then Risky is at least as good as Safe for Alice if and only if Risky is at least as good as Safe impersonally (and vice versa).
- However, how much Alice's well-being contributes to social utility depends on how many individuals exist and what their well-being levels are.
- The greater the total quantity of well-being, the smaller the contribution of additional well-being is.

Alice and Others

- Suppose that, when Alice is the only person who exists, Alice's loss of δ would reduce social utility by x units, and her gain of $\delta + \epsilon$ would increase it by more than x units.
- Then, in the one-person case, Risky is better than Safe (both impersonally and, by Ex Ante Pareto, for Alice).
- Now change the case; suppose that, besides Alice, there is a large number N of other, unaffected people.

Alice and Others: A large number N of other people have very good lives in state 1 ($p = 0.5$) and neutral lives in state 2 ($p = 0.5$).

Risky Gives Alice a well-being level of $1 + \delta + \epsilon$ in state 1 and a well-being level of $1 - \delta$ in state 2.

Safe Gives Alice a well-being level of 1 in states 1 and 2.

Alice and Others

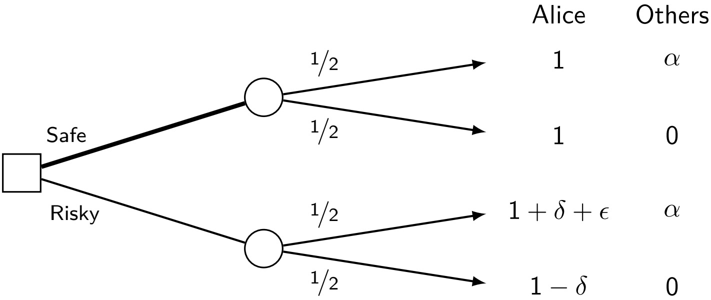

- In the state where Alice would lose δ (state 2), the other people have neutral lives. It follows that, no matter how large N is, her loss of δ would still reduce social utility in that state by x units.
- On the other hand, in the state where Alice would win $\delta + \epsilon$ (state 1), the N people have very good lives. Let α denote the total quantity of well-being of the N people with very good lives.

Alice and Others

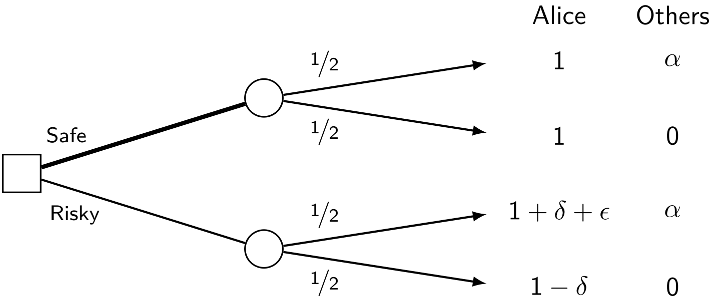

- As we increase N, the social utility in state 1 approaches the upper limit of utilities until it comes within x units of the upper limit. Then, increasing Alice's well-being by $\delta + \epsilon$ contributes less than x to social utility in that state.
- So, the $\delta + \epsilon$ increase in Alice's well-being in state 1 is no longer sufficient to compensate for the possible loss of δ well-being (and x units of utility) in state 2.
- It follows that, with a sufficiently large N, Safe is impersonally better than Risky.

Alice and Others

- This contradicts Ex Ante Pareto since Risky is better than Safe for Alice, and Safe and Risky are equally good for each of the N additional people.

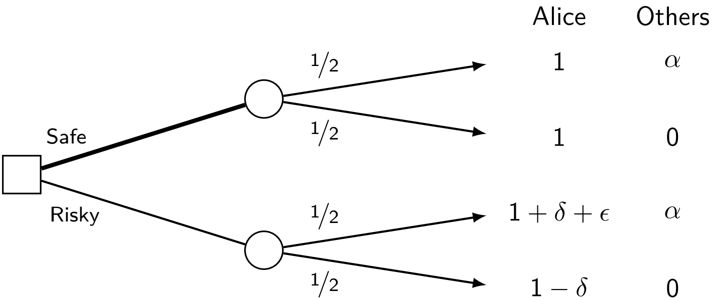

- To summarize, Bounded Expected Totalism violates Ex Ante Pareto even if individual betterness is risk-averse with respect to well-being.
- So, the combination of Expected Utility Theory and Total Utilitarianism prescribes prospects that are better for none and worse for some.

Harsanyi's social aggregation theorem

- Next, I will discuss how the earlier examples relate to a famous result in this area, namely, Harsanyi's social aggregation theorem.
- Harsanyi's social aggregation theorem shows that if both individual and social betterness relations can be given an expected utility representation, and the overall betterness relation satisfies Ex Ante Pareto, then social utilities are weighted sums of individual utilities.[8]
- Premises:

1. Each individual's betterness relation obeys the von Neumann-Morgenstern axioms.
2. The overall betterness relation obeys the von Neumann-Morgenstern axioms.
3. Ex Ante Pareto.[9]

[8] Harsanyi (1955).

[9] Harsanyi (1955) uses Pareto Indifference in the original formulation of the theorem, while Harsanyi (1977, p. 65) uses Weak Ex Ante Pareto.

Harsanyi's social aggregation theorem

- The conclusion of Harsanyi's theorem is that social utilities are weighted sums of individual utilities.
- Thus, overall betterness can be represented as maximizing the expectation of a weighted sum of individual utilities.
- If, in addition, we assume equal weighting for all individuals, then this theorem shows that the social utility function must be a sum of individual utilities.
- Harsanyi's theorem shows, in other words, that if individual and overall betterness relations are represented by expectational utility functions, then in order to satisfy Ex Ante Pareto, the social utility function must be a linear combination of individual utilities.

Harsanyi's social aggregation theorem

- I showed that Total Utilitarianism combined with Bounded Expected Utility Theory violates Ex Ante Pareto.
- Therefore, if one accepts Bounded Expected Totalism, that premise of Harsanyi's theorem fails.
- The reason that led to its failure was that a non-linear social transformation function is needed because the number of individuals might be infinite or arbitrarily large.
- In fact, it is unsurprising that one of Harsanyi's premises must be rejected; if the number of individuals might be infinite or arbitrarily large, then social utilities cannot be weighted sums of individual utilities because this could lead to unbounded social utilities.
- So, given that a bounded expected totalist rejects Harsanyi's conclusion, they cannot accept all his premises.

Harsanyi's social aggregation theorem

- This is worrying because Harsanyi's theorem is often considered one of the best arguments for utilitarianism.
- The conclusion of Harsanyi's theorem is that, for any fixed and finite population, social utility is an affine (or linear) function of total individual utility.
- However, once we consider the possibility of an infinite or arbitrarily large population, we find that social utility must be non-linear if social utilities are bounded and additive with individual utilities.
- And this leads to violations of Ex Ante Pareto.

Harsanyi's social aggregation theorem

- All this can be taken to support *Average Utilitarianism*, namely, the view that one population is better than another if and only if the average well-being it contains is greater.[10]
- Alternatively, these cases might be taken to undermine Boundedness (and Continuity). One might accept, for example, *Unbounded Expected Totalism*, namely, the view that combines Total Utilitarianism and Expected Utility Theory with an unbounded utility function.[11]
- Finally, the arguments might be taken to indirectly support alternatives to Boundedness, such as discounting small probabilities (also violates Ex Ante Pareto).

[10] Average Utilitarianism does not require a non-linear social transformation function; if individual utilities are bounded, then the average of those must also be bounded.

[11] See for example McCarthy et al. (2020).

Conclusion

- I have shown that Bounded Expected Totalism violates Ex Ante Pareto. Separate examples of Ex Ante Pareto violations were given for risk-neutrality and risk-aversion.
- A general argument to the effect that Bounded Expected Totalism must violate Weak Ex Ante Pareto was also given.
- Lastly, the implications of these cases for Harsanyi's social aggregation theorem were discussed.
- To conclude, combining two standard theories, Total Utilitarianism and Expected Utility Theory with a bounded utility function, results in violations of Ex Ante Pareto: The combination of these views implies that a prospect can be impersonally better than another prospect even though it is worse for everyone who is affected by the choice.[12]

[12] This talk focused on the compatibility of Expected Utility Theory and Total Utilitarianism, but the problem with Ex Ante Pareto arises for, for example, Critical-Level Utilitarianism in exactly the same way.

Appendix: General proof

- Next, I will give a general proof for why Bounded Expected Totalism must violate (Weak) Ex Ante Pareto if social utilities are bounded above and below.
- This proof shows that a violation of (Weak) Ex Ante Pareto must happen regardless of whether individual utilities are bounded or unbounded and whether individual betterness is risk-neutral, risk-averse or risk-seeking.
- Consider the following prospects:

  Risky* vs. Safe*:

  *Risky** Gives a 0.5 probability of δ additional well-being; otherwise, it decreases well-being by $-\delta$.

  *Safe** Does not increase or decrease well-being.
- The general idea is that if social utilities are bounded above, then (at least at some point) the social transformation function is concave with a positive total quantity of well-being.
- This means that, at least sometimes, the overall betterness relation is risk-averse with respect to well-being.
- So, with some positive total quantity of well-being W, Safe* is impersonally better than Risky*.

General proof

- So, whether Risky* is overall better than Safe* (or vice versa) depends on the total quantity of well-being.
- However, whether Risky* is better than (or equally as good or worse than) Safe* for Alice does not depend on the total quantity of well-being.
- Thus, there must be a counterexample to Weak Ex Ante Pareto:

> **Weak Ex Ante Pareto:** For all prospects X and Y, if X is at least as good as Y for everyone, then X is overall at least as good as Y.

General proof

- The proof for why Bounded Expected Totalism violates Weak Ex Ante Pareto goes as follows:
- Fix any $\delta > 0$.
- The following two claims are true:
    (1) If Risky* is impersonally at least as good as Safe* no matter how much total well-being there is in the background population, then social utility is unbounded above.
    (2) If Safe* is impersonally at least as good as Risky* no matter how much total well-being there is in the background population, then social utility is unbounded below.

General proof

- If social utility is bounded above and below, there must be a counterexample to Weak Ex Ante Pareto.
- Suppose, for example, that Risky* is at least as good as Safe* for Alice.
- This could be because Alice's betterness relation is risk-neutral with respect to her well-being and Risky* is therefore equally as good as Safe* for Alice.
- Alternatively, Alice's betterness relation might be risk-seeking.
- Either way, if social utilities are bounded above, then (1) shows that Risky* cannot be impersonally at least as good as Safe* no matter how much total well-being there is in the background population.
- So, with some total quantity of well-being, Safe* is impersonally better than Risky*—which contradicts Weak Ex Ante Pareto.

General proof

- Similarly, suppose that Safe* is at least as good as Risky* for Alice.
- Again, this might be because Alice's betterness relation is risk-neutral with respect to her well-being.
- Alternatively, it could be because her betterness relation is risk-averse.
- However, given that social utilities are bounded below, (2) shows that Safe* cannot be impersonally at least as good as Risky* no matter how much total well-being there is in the background population.
- So, with some total quantity of well-being, Risky* is impersonally better than Safe*, contrary to Weak Ex Ante Pareto.

General proof

- Proof of (1) goes as follows: Consider background populations with total well-being levels of 0, δ, 2δ, 3δ, and so on. Let $x = f(\delta) - f(0)$.
- If Risky* is impersonally at least as good as Safe* with respect to all these background populations, then the difference between $f(n\delta)$ and $f((n-1)\delta)$ is at least as great as the difference between $f((n-1)\delta)$ and $f((n-2)\delta)$, for each $n > 2$.
- It follows that $f(n\delta)$ is at least as great as nx.
- Thus, f is unbounded above.
- One can give a similar proof for (2).
- So, if social utilities are bounded above and below, there must be a counterexample to Weak Ex Ante Pareto, regardless of whether individual utilities are bounded or unbounded and whether individual betterness is risk-averse, risk-neutral or risk-seeking.

References I

Broome, J. (1991), *Weighing Goods: Equality, Uncertainty and Time*, Blackwell, Oxford.

Fishburn, P. C. (1970), *Utility Theory for Decision Making*, Wiley, New York.

Gustafsson, J. E. (2022), 'Ex-ante prioritarianism violates sequential ex-ante Pareto', *Utilitas* **34**(2), 167–177.

Hammond, P. J. (1998), Objective expected utility: A consequentialist perspective, *in* S. Barberà, P. J. Hammond and C. Seidl, eds, 'Handbook of Utility Theory Volume 1: Principles', Kluwer, Dordrecht, pp. 143–211.

Harsanyi, J. C. (1955), 'Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility', *Journal of Political Economy* **63**(4), 309–321.

Harsanyi, J. C. (1977), *Rational Behavior and Bargaining Equilibrium in Games and Social Situations*, Cambridge University Press, Cambridge.

References II

Kreps, D. M. (1988), *Notes on the Theory of Choice*, Westview Press, Boulder.

McCarthy, D., Mikkola, K. and Thomas, T. (2020), 'Utilitarianism with and without expected utility', *Journal of Mathematical Economics* **87**, 77–113.

Mongin, P. and Pivato, M. (2015), 'Ranking multidimensional alternatives and uncertain prospects', *Journal of Economic Theory* **157**, 146–171.

Nebel, J. M. (2020), 'Rank-weighted utilitarianism and the veil of ignorance', *Ethics* **131**(1), 87–106.

Rabinowicz, W. (2002), 'Prioritarianism for prospects', *Utilitas* **14**(1), 2–21.

Russell, J. S. and Isaacs, Y. (2021), 'Infinite prospects', *Philosophy and Phenomenological Research* **103**(1), 178–198.

Wilkinson, H. (2022), 'In defence of fanaticism', *Ethics* **132**(2), 445–477.