

Discussion Paper Series

Institute-Wide Joint Research Project

Methodology of Social Sciences: Measuring Phenomena and Values

Repugnance at the Limit

Susumu Cato
University of Tokyo

Ko Harada
University of Tokyo

Ken Oshitani
University of Tokyo

August 17, 2023

E-23-004

Institute of Social Science, University of Tokyo

[<https://web.iss.u-tokyo.ac.jp/methodology/en/dp/dp/>]

Repugnance at the Limit

Susumu Cato

susumu.cato@gmail.com
Institute of Social Science,
University of Tokyo

Ko Harada

ko.harada11@gmail.com
Graduate School of Economics,
University of Tokyo

Ken Oshitani

kenoshitani@gmail.com
Institute of Social Science,
University of Tokyo

Abstract

The implications of the repugnant conclusion for consequentialist theories, such as total utilitarianism, have been extensively discussed since the work of Derek Parfit. Notably, these discussions make the paradoxes of population ethics depend on essentially welfarist conditions, represented by the dominance condition (or the Pareto Principle). This might lead us to think that the repugnant conclusion does not pose as much of a practical problem for deontologists, who deny that we always ought to do what produces the most aggregate welfare. In this study, we offer two impossibility results using what we call *the limit aggregation property*. This states that, when there is a conflict between maximizing overall welfare and non-welfaristic considerations, there is a threshold number of persons such that it is morally preferable to prioritize welfaristic considerations above this threshold. We argue that this property should be accepted even by deontologists, insofar as they do not commit themselves to the implausible absolutist position that fails to assign any moral weight to considerations of aggregate welfare. Our results therefore state that any normative population theory that is not absolutist in this sense entails (a variant of) the repugnant conclusion or some other implausible conclusion. We conclude that the repugnant conclusion must be taken seriously even by non-absolutist deontologists, not just consequentialists.

Keywords: variable population, repugnant conclusion, sadistic conclusion, aggregation, threshold deontology.

1. Introduction

It is often assumed that the paradoxes of population ethics, such as the mere addition paradox, pose a challenge primarily for welfaristic normative theories, which hold that it is always morally preferable to do what best promotes overall welfare, but that they do not pose as much of a practical problem for non-welfaristic normative theories, such as deontological ethics. The repugnant conclusion holds that it follows from given certain widely accepted assumptions that for any state which contains some number of persons with a high level of well-being, there is a state that is morally preferable containing more people at a level of well-being that is barely worth living. The repugnant conclusion is standardly taken as posing a problem for normative population theories that direct us to maximize aggregate welfare, such as total utilitarianism (Rawls 1971, 162–163; Glover 1977, 31; Parfit 1984, 378; Cowen 1996, 755; Zuber et al., 2021). On the other hand, philosophers have argued that normative theories that do not take a welfarist structure can avoid the problematic practical implications of the repugnant conclusion, since they deny the premise that it is always normatively preferable to bring about the best consequences (Boonin-Vail 1996, 291; Meijers 2017, 253–254). One reason for believing that non-consequentialists can simply sidestep the repugnant conclusion is their rejection of *the dominance condition* (or *the Pareto condition*), which holds that a state that contains higher welfare is always morally preferable to a state that contains less welfare.¹ Thus, non-consequentialists can hold that non-welfaristic considerations (such as those pertaining to the rightness of acts) can plausibly override considerations of welfare in determining the outcomes that we ought morally to bring about.

¹ For a more precise statement of the dominance condition, see Arrhenius (2000), 261.

This paper offers a novel take on the paradoxes of population ethics that allows us to see why they pose a problem even for non-consequentialists who deny the dominance condition and who therefore deny that we always have normative reasons to do what is best. Most notably, our argument is based, not on the dominance condition, but a plausible assumption that even many deontologists would accept: *the limit aggregation axiom*. To put it briefly, the limit aggregation axiom states that, even if there is a conflict between maximizing overall welfare and non-welfaristic considerations, there is a threshold number of persons such that, above this threshold, it is morally preferable to prioritize welfaristic considerations (more specifically, the level of aggregated well-being). We argue that the limit aggregation axiom is a relatively plausible condition that ought to be accepted by anyone who is not an ethical absolutist, or in other words, someone who refuses to accord *any* weight to considerations of aggregate well-being. A significant implication of our argument, therefore, is that it shows that the standard paradoxes of population ethics have a far wider applicability than previously thought.²

The rest of this paper is organized as follows. In Section 2, we show our new impossibility results, which presupposes the limit aggregation axiom in place of dominance. Section 3 explains how the limit aggregation axiom is presupposed by many deontological theories, most notably threshold deontology. This shows how our new impossibility results have implications

² An axiologist might reply that the deontologist's response to the repugnant conclusion simply misses the point of the repugnant conclusion. The repugnant conclusion is an axiological thesis concerning the comparative goodness of states of affairs: the conclusion is puzzling because it intuitively seems that a state containing less people who are more well-off is *better* than a state containing more people who are less well-off. If so, then moving the focus away from the axiological question of what state of affairs is best to the normative question of what we ought to do simply changes the subject without addressing the original axiological puzzle (Parfit 1984, 429). While this point is well-taken, it does not negate the significance of our thesis. Our claim is that *even if* the deontologist understands the repugnant conclusion as a normative puzzle regarding what we ought to do, she cannot avoid facing a normative version of the same problem, insofar as she is committed to certain plausible assumptions, including the limit aggregation axiom.

not only for consequentialist theories but also for non-consequentialist theories as well. Section 4 concludes.

2. Two impossibility results with the limit aggregation property

This section offers two impossibility results in population ethics. Our framework is the standard setting for population ethics. We assume that persons' well-being levels are cardinally measurable and interpersonally comparable. For each population, information about individual welfare is represented by a list of individual well-being levels, (w_1, \dots, w_n) where w_i is the individual i 's well-being level.³ Note that non-welfaristic/non-consequential information is not excluded in this framework.⁴ The "zero" level of well-being represents neutrality. Populations are denoted by $A, B, C, \dots, P, Q, \dots$, and so on.

A theory for population ethics is represented by an "ordering" over the set of populations. That is, populations are compared by the ordering. We call this ordering a *population principle*. Note that such an ordering represents a moral/normative *theory*, which tells us what kind of populations we morally ought to choose. Importantly, for the purposes of this paper, we will understand the population principle as a normative "ought-to-be-chosen" relation (Arrhenius forthcoming, 359; Greaves, 2017). Some population A *ought to be chosen rather than* another

³ In this paper, we assume that individual well-being levels are continuous values. Arrhenius (2009, forthcoming) considers the framework where well-being levels are restricted to discrete numbers. Our results can be easily extended to Arrhenius's framework.

⁴ We do not exclude non-consequential information because each population, say P , can be considered as a combination of a well-being distribution (w_1, \dots, w_n) and descriptions of non-consequential information. In this case, each population describes a comprehensive social situation, including welfarist information and non-welfarist information. A typical P becomes $((w_1, \dots, w_n), \phi)$, where ϕ includes descriptions of non-welfarist/non-consequential information. Notably, our formal statements hold for this general framework. This relation is of practical importance for consequentialists and non-consequentialists alike.

population B if and only if it would be wrong to bring about B rather than A . Note that any “ought-to-be-chosen” relation is considered asymmetric (i.e., if A ought to be chosen rather than B , then it is not the case that B ought to be chosen rather than A). We will say that it is *permissible to choose* population A against B if and only if it would not be wrong to bring about A rather than B . Thus, it is permissible to choose population A against B when it is not the case that B ought to be chosen rather than A . We assume that completeness and transitivity on population principles. That is, all populations are comparable, and “ought-to-be-chosen” relations and permissibility relations are transitive.⁵

In this paper, we focus exclusively on the normative relation concerning what we ought to choose, since this relation is of practical importance for both consequentialists and non-consequentialists. Of course, consequentialists hold that the normative relation of “ought to be chosen” is reducible to an axiological relation—namely, the goodness (“betterness”) relation. More specifically, they believe that we ought to choose, or are required to choose, a given population rather than another if that population is better than the other. The corresponding “at-least-as-good-as” relation is induced from the betterness relation as follows: A is at least as good as B if and only if B is not better than A . By contrast, according to non-consequentialists, the normative relation is independent of the goodness relation, since it is not the case that we always ought to choose the option that realizes the best outcome. Despite these differences, however, the non-consequentialist and the consequentialist can both rely in the end on some

⁵ Our framework can be formally expressed by using a binary relation \succ as follows. We write $A \succ B$ when A ought to be chosen rather than B (or it would be wrong to bring about B rather than A). We define \succcurlyeq by letting $A \succcurlyeq B$ if and only if $\neg B \succ A$. Then, it is permissible to choose A against B if and only if $A \succcurlyeq B$. We impose completeness and transitivity on \succcurlyeq . Completeness requires that, for all populations A, B , $A \succcurlyeq B$ or $B \succcurlyeq A$. Transitivity requires that if $A \succcurlyeq B$ and $B \succcurlyeq C$, then $A \succcurlyeq C$. This logically implies that if $A \succ B$ and $B \succ C$, then $A \succ C$.

kind of normative ordering to determine the population that we morally ought to choose to bring about.

With these preliminaries in place, let us move on to our two formal results. The formal statement of the first impossibility result is as follows.

For any population principle that satisfies the limit aggregation axiom, the repugnant addition holds or the mere addition property is violated.

We offer definitions of the properties that we mentioned in the statement. The first property is a variation of the well-known statement of the repugnant conclusion. According to the original version, “[f]or any possible population of at least ten billion people, all with a very high quality of life, there must be some much larger imaginable population whose existence, if other things are equal, would be better even though its members have lives that are barely worth living” (Parfit, 1984). This is repugnant because we have the intuition that we could not be morally required to bring about a larger population with very high well-being levels rather than a population with very low positive well-being levels. We offer a more precise statement as follows; see Blackorby, Bossert, and Donaldson (2005) for a mathematically sophisticated expression of this.

Repugnant conclusion. For any perfectly equal population A with very high positive well-being levels, and for any small positive well-being level α , there exists a very large number m , which is larger than the population size of A , such that we ought to choose to bring into being m persons at the level of α rather than A .

The second property is the mere addition property. Note that a person who obtains a positive well-being level is worth living because the “zero” utility level corresponds to “neutrality.” Having more persons who are worth living is never morally impermissible. This is the idea of this property. More formally, we have the following:

Mere addition property. For any population A , if another population B includes only positive well-being levels, then it is morally permissible for us to bring about the population where B is added to A rather than A .

This has been widely used in the literature since the work of Parfit (1976, 1982, 1984). Its formal analysis is undertaken by Ng (1989) and Carlson (1998), who establish impossibility results within a consequential framework.

Finally, we introduce the limit aggregation axiom. Compared to the two properties mentioned above, this one may be relatively unfamiliar. For example, let us consider one individual, say, Smith. Assume that Smith obtains 1,000 units of well-being in world ω , but he obtains $-1,000$ units in world ω^* . The difference is huge for him. Having ω^* instead of ω is devastating to Smith. Now, Smith is not the only person in these worlds. Consider an additional person who obtains 0 in world ω while he/she obtains 1 in world ω^* . This person prefers ω^* , but this does not matter a lot to this additional person. Given this, it is intuitively plausible to say that if there are only the two persons (Smith and the additional person), ω is better than ω^* . Then, what if there are n persons, other than Smith, who obtain 0 in ω and 1 in world ω^* ? If n is below a certain threshold, many theories would say that ω is better than ω^* . Indeed, it could be held that ω is better than ω^* even if n exceeds 2000. This might be the case if Smith has a moral claim against being subject to excessive losses in order to bring about gains in the

aggregate welfare of others (Scanlon 1998, 235). In other words, non-welfaristic considerations such as Smith's moral claim can plausibly override welfaristic considerations, making it better all-things-considered to choose ω even though ω^* is slightly better in welfaristic terms. But what if n is an extremely large number, such as a billion? Beyond a certain threshold, it might seem that small gains for a sufficiently large number of people can override significant losses for a fewer number of people. This is what the following axiom states.

The limit aggregation axiom. Let A, B be two populations with the same population size. For any positive numbers δ, ε , there exists a threshold integer n^* such that we ought to choose to bring about the addition of n persons at the level of $\delta + \varepsilon$ to A rather than the addition of n persons at the level of δ to B as long as n is larger than n^* .

The three properties are enough to establish the aforementioned impossibility result. We will prove the claim by showing that the repugnant conclusion holds if the mere addition property is satisfied. Take any large population P with very high well-being levels. The size of this population is denoted by m . We now take a very small, positive well-being level, ε . Let us consider $P^+(n) = (\frac{\varepsilon}{2}, \dots, \frac{\varepsilon}{2}, P)$, where $P^+(n)$ is the population in which n persons with $\frac{\varepsilon}{2}$ are added to P . We compare it with the following population:

$$Q(m + n) = (\varepsilon, \dots, \varepsilon).$$

$Q(m + n)$ is the population consisting of $m + n$ persons, who obtain ε . Since $Q(m + n)$ and $P^+(n)$ have the same population size for any number n , one can apply the limit aggregation property. Each person in P obtains much higher utility levels than ε , which is a very small number. Thus, m persons lose a huge amount of well-being in Q . On the other hand, n persons gain some well-being in Q ; to be precise, the gain is equal to $\frac{\varepsilon}{2}$. By the limit aggregation property, there exists a number n^* such that bringing about $P^+(n)$ rather than bringing about

$Q(m + n)$ is morally wrong as long as n is larger than n^* . By the mere addition property, it is guaranteed that it is permissible to bring about $P^+(n)$ rather than P . By transitivity, it would be impermissible to bring about P rather than $Q(m + n)$. This implies that the repugnant conclusion holds. The impossibility result is established.

The formal statement of the second impossibility result is as follows.

For any population principle that satisfies the limit aggregation axiom, the repugnant addition or the sadistic conclusion holds.

Now, we offer definitions of the repugnant addition and the sadistic conclusion. The repugnant addition is a version of the repugnant conclusion. This essentially states that adding very high well-being levels is worse than adding very low positive well-being levels. More formally:

Repugnant addition. There exists a population A such that, for any perfectly equal population B with very high positive well-being levels, and for any small positive well-being level α , there exists a very large number m , which is larger than the population size of B , such that we ought to choose to bring about the addition of m persons at the level of α to A rather than the addition of B to A .

There are many versions of the repugnant conclusion; Spears and Budolfson (2021) listed them and discussed the differences among them. This version of the repugnant conclusion was imposed by Arrhenius (2009, forthcoming) in a slightly different framework. Spears and Budolfson (2021: 575) claim that if the fact that the repugnant conclusion holds for an ordering this is sufficient to disqualify the ordering normatively, then one can disqualify any (other)

ordering for which the repugnant addition holds. That is, the repugnant addition is considered “repugnant” as long as the original repugnant conclusion is “repugnant.” On this ground, Arrhenius (2009, forthcoming) and Spears and Budolfson (2021) offered impossibility results using a variation of the repugnant addition.⁶

The next property is the sadistic conclusion. This idea was proposed in a series of works by Arrhenius (2000, 2009). We use a very standard version. Consider a population with positive well-being levels on the one hand and another population with negative well-being levels on the other hand. If either of them is added to the third population, which is the one that we morally ought to choose? The sadistic conclusion holds if there is a case where we morally ought to add negative welfare. More formally:

Sadistic conclusion. There exist three populations A , B , C , where all persons in B obtain positive well-being levels, all persons in C obtain negative well-being levels, and we ought to choose to bring about (A, C) rather than (A, B) . In other words, we ought to choose to bring about the addition of persons with positive well-being levels to A rather than the addition of persons with negative well-being levels to A .

Recently, Franz and Spears (2020) have shown that the mere addition principle and the avoidance of the sadistic conclusion are logically related to each other under certain conditions.

We now show the result. By way of contradiction, both the sadistic conclusion and repugnant addition are avoided. Take any large population P with very high well-being levels. The size of this population is denoted by m . We now consider a very small, positive well-being level,

⁶ They imposed what we can call the “very repugnant addition.”

ε . Let us compare $P^+(n) = (\frac{\varepsilon}{2}, \dots, \frac{\varepsilon}{2}, P)$ with $P^- = (-\varepsilon, P)$, where $P^+(n)$ is the population in which n persons with $\frac{\varepsilon}{2}$ are added to P , and P^- is the population where one persons with $-\varepsilon$ is added to P . Since the sadistic conclusion is avoided, it is permissible to bring about $P^+(n)$ rather than P^- for any number n . Now, let us consider the following population:

$$Q^*(m + n - 1) = (-\varepsilon, \varepsilon, \dots, \varepsilon).$$

Under $Q^*(m + n - 1)$, one individual obtains a negative utility level $-\varepsilon$, while $m + n - 1$ persons obtain ε . We note that $Q^*(m + n - 1)$ and $P^+(n)$ have the same population size for any number n . By the limit aggregation property, it holds that we ought to choose to bring about $Q^*(m + n - 1)$ rather than $P^+(n)$ if n is sufficiently large. By transitivity, we ought to choose to bring about $Q^*(m + n - 1)$ rather than P^- . This means that we ought to choose to add $(\varepsilon, \dots, \varepsilon)$ instead of adding P . This implies that the repugnant addition holds. The impossibility result is established.

3. What is the meaning of the limit aggregation property?

We must stress that the limit aggregation property is not entirely new. There are existing works that employ similar properties to derive the impossibility theorem. For example, Blackorby, Bossert, Donaldson, and Fleurbaey (1998) and Spears and Budolfson (2021) show versions of the mere addition paradox relying on variations of the limit aggregation property. Moreover, Arrhenius (2000, 2009) and Bossert, Cato, and Kamaga (2022) prove a tension between the repugnant and sadistic conclusions using such properties.

These authors' results are mathematically closer to ours. However, there is a crucial difference: these existing works also use a property of dominance, which states that bringing about one population rather than another is always better than another—or in normative terms, that we always ought to choose the bring about the former rather than the latter—if everyone is better off in the former.⁷ The dominance property plays an indispensable role in existing derivations of the impossibility results.⁸

The dominance property is plausible if one assumes a consequentialist normative theory, according to which we always morally ought to bring about the best consequences.⁹ This is true for other works that offer impossibility results, which do not necessarily impose properties like the limit aggregation property. For example, non-anti-egalitarianism, an essential property for deriving the paradox proposed by Parfit (1984) and Ng (1989), essentially implies the dominance property. On the other hand, the dominance property is less plausible if we assume a non-consequentialist normative theory, such as a theory that also takes into consideration deontological constraints. Roughly speaking, prior derivations of the impossibility results

⁷ This is the Paretian dominance property (the Paretian condition). Some works use the egalitarian dominance property, which is logically weaker than the Paretian condition. The egalitarian dominance property requires that, for any two populations with the same size, in which all persons obtain the same utility level, a population with high utility levels is better.

⁸ Strictly speaking, Spears and Budolfson (2021) do not use dominance condition. They consider two closely related axioms. One of them requires that a population with positive well-being levels is always morally preferable to a population with negative well-being levels regardless of their population sizes, and hence as dominance condition, it requires $(0.001, 0.001, 0.001)$ to be always morally preferable to $(-0.001, -0.001, -0.001)$. On the other hand, the other requires a population with positive well-being levels is always morally preferable to a population with positive but lower well-being levels as long as the former population has a larger population size. Note that, it does not necessarily require $(0.001, 0.001, 0.001)$ is morally preferable to $(-0.001, -0.001, -0.001)$. However, since only a slight increase in population size (at least one person) is needed for a population principle to be sensitive to extremely small differences in well-being levels, the objection to dominance condition can be applied to both axioms.

⁹ Kaplow and Shavell (2001) and Adler (2018) show that the Pareto condition excludes the use of non-consequential information for social evaluations under certain conditions.

impose properties that are plausible on the ground of consequentialism. For the purposes of establishing axiological impossibilities, it is quite natural to rely on consequentialist properties. However, when we recast the impossibilities in normative terms (as we have done above), the assumption of dominance can be challenged on non-consequentialist or deontological grounds.

More specifically, we should emphasize that deontological theories can avoid previous versions of the repugnant conclusion by rejecting welfarism. This means that A may not be morally more plausible than B even if A welfaristically dominates B . Dominance and related properties assumed in the existing works suggest that we ought to choose to bring about A in all cases where persons are better off. For example, dominance requires that we ought to choose to bring about $(5, 5, 5, 5)$ than $(4.999, 4.999, 4.999, 4.999)$, whatever the other non-consequentialist features of our situation are. Deontologists would plausibly deny this. Since the two numbers are very close, it is reasonable for deontologists to say that there can be some non-welfaristic differences between the two situations, which can override 0.001 units of gains of the four individuals. Since dominance does not allow this possibility, it is likely to be rejected by deontological theories. Therefore, a possible response from deontologists or non-welfarists to the existing impossibility results is to argue that these results are irrelevant to them because dominance is implausible.

A novel contribution of this paper is to show that the impossibility results that plague population ethics do not need to assume dominance: one faces repugnance or sadism under only the limit aggregation property.¹⁰ Notably, the limit aggregation property is respected not only by consequentialist theories such as utilitarianism and prioritarianism but by many deontological theories. This is because the limit aggregation axiom is compatible with the claim

¹⁰ We stress that violation of the mere addition property can be regarded as a certain type of sadism.

that we need not always do what produces the best outcome. The limit aggregation axiom merely says that we ought to do what produces the best outcome when the amount of good at stake is sufficiently large (that is, above a certain threshold). This means that our result has significance even for deontological theorists.

Of course, we do not claim that *all* deontological theories accept the limit aggregation property. Indeed, absolutist deontology, which assigns unequivocal priority to deontological concerns, does not meet the property. On the other hand, if deontological theories are non-absolutist, in the sense that they do take people's welfaristic gains into account in *some* way, it is natural to consider that the limit aggregation property is satisfied by them. Note that the property can be reasonable even for deontologists because it does not specify the number of persons whose gains override someone's sacrifice. It requires the existence of such a number.

To illustrate this point, take the following example, adapted from a paper by Alastair Norcross. Let us say that lowering the speed limit by a certain amount could save one person from being killed in a car accident but would impose a mild degree of inconvenience on 1000 drivers. Further, let us say that in welfaristic terms, the badness of the inconvenience endured by the 1000 drivers outweighs the badness of the death of the one person. The deontologist could still plausibly maintain that it would be impermissible to let one person die for the convenience of a 1000 people, even if this would maximize overall welfare. However, it is less plausible to maintain that a mild degree of inconvenience imposed on *any* number of drivers (for instance, a billion) would not suffice to make it permissible to let one person die of a car accident. This accords with common-sense morality: we do not believe that allowing cars to drive at the current speed limit is morally impermissible, even if lowering the speed limit by 10 mph could result in saving, say, tens of thousands of lives a year (Norcross, 1997, 159). If the deontologist

accepts this judgment, however, then she also accepts the limit aggregation axiom, since she believes that there is some threshold number of persons such that a lesser sacrifice to these persons can override the stronger claims of some fewer number of persons.

In more general terms, the limit aggregation property is presupposed by threshold deontologists, who hold that non-welfarist considerations (such as our duties to avoid intentionally causing harm) can be normatively overriding up to a certain point, but that, when a sufficient amount of good is at stake, it is permissible to choose the option that produces the best outcome (Nagel 1979, 62; Moore 1997, 723; Alexander 2000, 894). It is important to note that many, if not most, deontologists are in fact threshold deontologists: for instance, even Robert Nozick, who holds that individuals have inviolable rights against being used as means for promoting the social good, concedes that deontological considerations can give way if the good at stake is sufficiently great (Nozick 1974, 30 n*).

Conversely, there can be normative theories that reject aggregative reasoning of any kind, and which consequently are able to avoid the impossibility results shown above. To see this point, let us take the maximin criterion, which has been substantially discussed by Bossert (1990) and Zuber (2018). Note that, for any population, one can identify the well-being level of the worst-off. Thus, according to the maximin criterion, a population is better than another if and only if the worst-off obtains a higher well-being level under the former. As a distinguishable nature of the maximin criterion, it absolutely rejects any form of aggregation. One can say that this is “absolutely anti-aggregation.”¹¹ Any gain of people other than the worst-off, does not matter

¹¹ Asheim and Zuber (2014) proposed what are called *rank-discounted principles*; these principles are similar to *geometrism* by Sider (1991), but it is anti-non-egalitarian, unlike Sider’s principle. Rank-discounted principles assign moral weights depending on the relative positions of persons’ well-being levels in the entire population. Because of this nature, the limit aggregation property is violated. However, this does not imply that the impossibilities do not hold for Asheim and Zuber’s principle.

to the maximin criterion. Because of this, this criterion avoids both the repugnant and sadistic conclusions (more precisely, it can even avoid any version of the repugnant addition and any version of the sadistic conclusion).¹²

4. Concluding remarks

In this paper, we argued that, contrary to common wisdom, the standard paradoxes of population ethics pose an important practical challenge to non-consequentialist theories, not just to consequentialist theories. While previous demonstrations of the paradoxes assumed the dominance condition, which is a controversial premise that non-consequentialists could plausibly reject, our argument is based on a relatively weaker premise (namely, the limit aggregation axiom) that many deontologists would accept—as long as they are not absolutists who reject the moral relevance of aggregative welfare altogether.

References

- Adler, M. D. (2018). Prioritarianism: Room for desert? *Utilitas*, 30(2), 172–197.
- Alexander, L. (2000). Deontology at the threshold. *San Diego Law Review* 37, 893.
- Arrhenius, G. (2000). An impossibility theorem for welfarist axiologies. *Economics and Philosophy*, 16, 247–266.
- Arrhenius, G. (2009). One more axiological impossibility theorem, in L. Johansson, J. O’sterberg, and R. Sliwinski (eds.), *Logic, Ethics, and All That Jazz: Essays in Honour of Jordan Howard Sobel*. *Uppsala Philosophical Studies* 57, 23–37. (Uppsala: Uppsala University)
- Arrhenius, G. (forthcoming). *Population Ethics: The Challenge of Future Generations*. Oxford: Oxford University Press.

Although does not avoid the sadistic conclusion, although it avoids the repugnant addition. First, it violates the mere addition property. Second, although it avoids the very/strong sadistic conclusion and the repugnant addition, it does not avoid the sadistic conclusion. Note that Asheim and Zuber (2014) uses the very the sadistic conclusion under the name of the sadistic conclusion. But our version of the sadistic conclusion holds for it; see Bossert, Cato, and Kamaga (2022).

¹² Note that this maximin avoids the sadistic conclusion, but it does not satisfy the mere addition property. Hence, the maximin criterion is a counterexample to the second impossibility theorem but not to the first one. For the first one, we can consider Sider’s (1991) geometrism, which exhibits anti-egalitarianism.

- Asheim, G. B., and Zuber, S. (2014). Escaping the repugnant conclusion: Rank- discounted utilitarianism with variable population. *Theoretical Economics*, 9(3), 629– 650.
- Blackorby, C., Bossert, W., and Donaldson, D. (2005). *Population Issues in Social Choice Theory, Welfare Economics, and Ethics*. New York: Cambridge University Press.
- Blackorby, C., Bossert, W., Donaldson, D., and Fleurbaey, M. (1998). Critical levels and the (reverse) repugnant conclusion. *Journal of Economics*, 67, 1–15.
- Boonin-Vail, D. (1996). Don't stop thinking about tomorrow: Two paradoxes about duties to future generations. *Philosophy & Public Affairs*, 25(4), 267-307.
- Bossert, W. (1990). Maximin welfare orderings with variable population size. *Social Choice and Welfare*, 7, 39–45.
- Bossert, W., Cato, S., and Kamaga, K. (2022). Revisiting variable-value population principles. *Economics & Philosophy*, 1–17. doi:10.1017/S0266267122000268
- Carlson, E. (1998). Mere addition and two trilemmas of population ethics. *Economics & Philosophy*, 14(2), 283-306.
- Cowen, T. (1996). What do we learn from the repugnant conclusion? *Ethics*, 106(4), 754–775.
- Franz, N., and Spears, D. (2020). Mere Addition is equivalent to avoiding the Sadistic Conclusion in all plausible variable-population social orderings. *Economics letters*, 196, article no. 109547.
- Glover, J. (1977). *Causing Death and Saving Lives*. London, UK: Penguin Books.
- Greaves, H. (2017). Population axiology. *Philosophy Compass*, 12, e12442.
- Kaplow, L., and Shavell, S. (2001). Any non-welfarist method of policy assessment violates the Pareto principle. *Journal of Political Economy*, 109(2), 281–286.
- Meijers, T. (2017). Citizens in appropriate numbers: evaluating five claims about justice and population size. *Canadian Journal of Philosophy*, 47(2-3), 246–268,
- Moore, M. (1997). *Placing Blame: A General Theory of Criminal Law*. Oxford, UK: Oxford University Press.
- Nagel, T. (1979). *Mortal Questions*. New York: Cambridge University Press.
- Norcross, A. (1997). Comparing harms: headaches and human lives. *Philosophy & Public Affairs*, 26(2), 135-167.
- Nozick, R. (1974). *Anarchy, State, and Utopia*. Oxford, UK: Blackwell.
- Ng, Y.-K. (1989). What should we do about future generations? Impossibility of Parfit's theory X. *Economics & Philosophy*, 5, 235–253.
- Parfit, D. (1976). On doing the best for our children. In: M.D. Bayles (ed.). *Ethics and Population*, pp. 100–102. Cambridge, MA: Schenkman.
- Parfit, D. (1982). Future generations, further problems. *Philosophy and Public Affairs*, 11, 113–172.
- Parfit, D. (1984). *Reasons and Persons*. Oxford, UK: Oxford University Press.
- Rawls, J. (1971). *A Theory of Justice*. Cambridge, MA: Harvard University Press.
- Sider, T. R. (1991). Might theory X be a theory of diminishing marginal value? *Analysis*, 51(4), 265–271.
- Scanlon, T.M. (1998). *What We Owe to Each Other*. Cambridge, MA: Harvard University Press.
- Spears, D. and Budolfson, M. (2021). Repugnant conclusions. *Social Choice and Welfare*, 57, 567–588.
- Zuber, S. (2018). Population-adjusted egalitarianism. CES Working Paper, halshs- 01937766.
- Zuber, S., Venkatesh, N., Tännsjö, T., Tarsney, C., Stefánsson, H. O., Steele, K., ... and Asheim, G. B. (2021). What should we agree on about the repugnant conclusion? *Utilitas*, 33(4), 379–383.