

Calibrating variable-value population ethics: Implications for the effort to avoid the repugnant conclusion

August 17, 2021

Abstract

Population ethics has long focused on the effort to avoid Parfit’s *Repugnant Conclusion*. But a growing literature suggests that this effort has been misunderstood, and may therefore have been misplaced. Here, we add to this literature by exploring Variable-Value axiologies’ proposed solutions to the challenges of population ethics. These views avoid Parfit’s Repugnant Conclusion, while satisfying some weak instances of the *Mere Addition* principle (for example, at small population sizes). We apply calibration methods to Variable-Value views while assuming: first, some very weak instances of Mere Addition, and, second, some plausible empirical assumptions about the size and welfare of the intertemporal world population. We find that Variable-Value views imply conclusions that should seem repugnant to anyone who opposes Total Utilitarianism due to the Repugnant Conclusion. So, any wish to avoid repugnant conclusions is not a good reason to choose a Variable-Value view. More broadly, these calibrations teach us something about the effort to avoid the Repugnant Conclusion. Our results join a recent literature arguing that prior efforts to avoid the Repugnant Conclusion hinge on inessential features of the formalization of repugnance — and therefore may be less normatively significant than is traditionally assumed.

1 Introduction

Much research in population ethics — as studied by both philosophers and economists — is motivated by the quest to avoid what Parfit (1984) called the *Repugnant Conclusion*, one version of which states that:¹

¹Parfit’s own formulation of the Repugnant Conclusion states that: “For 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, p. 388) Our formulation is closer to Arrhenius’s (forthcoming).

The Repugnant Conclusion (Informal version). *For any perfectly equal population of very well-off people, there is a better population consisting entirely of lives that are barely worth living.*

Total Utilitarianism, according to which a population is better the greater sum of welfare it contains, is widely recognized to entail the Repugnant Conclusion. No matter how well-off people are in some population *A*, and independently of *A*'s size, there is some (potentially much bigger) imaginable population *Z* that contains a greater sum of welfare than *A* does — even though people in *Z* have lives that are each barely worth living (understood as having barely positive welfare).

Most paths to avoiding the Repugnant Conclusion begin by abandoning what Parfit called the *Mere Addition principle*, which can be stated thus:

Mere Addition (Informal version). *By adding any life worth living to any population, without making anyone else worse off, we do not make the population worse.*

Total Utilitarianism implies the Mere Addition principle. But this principle is violated by *Average Utilitarianism*, according to which a population is better the greater average welfare it contains. And Average Utilitarianism avoids the Repugnant Conclusion: Population *Z*, whose members all have lives that are barely worth living, contains lower average welfare than *A*. So, *A* is better than *Z*, according to Average Utilitarianism.

Somebody who abandons Mere Addition argues that merely adding a life worth living, without making anyone worse off, can make a population worse. But what about adding a life *well* worth living? Consider merely adding a person who lives a very good life by modern standards: say, a professor living in a developed country in 2020. Surely by adding a person like that to any population, without thereby making anyone else worse off, we have made the population better? Not according Average Utilitarianism. To see this in an

Spears and Budolfson (2021) have argued that formalizations of the Repugnant Conclusion should be broader — including, for example, additions to unaffected, intersecting populations — but for this paper we ignore that proposal and focus on what they call a “restricted” formalization.

absurd example: adding our professor to a single-person “population” whose member is only a tiny bit better-off than the professor makes the resulting population *worse*, according to Average Utilitarianism. In fact, if the future of humanity is as long and wonderful as some hope (Ord, 2020), then adding a person like this to the *actual* intertemporal world population² makes the resulting population worse, according to Average Utilitarianism. This violates what we shall call *Weak Mere Addition* (which we state formally in section 2).

In light of the above counterintuitive implications of on the one hand Total Utilitarianism (in particular, the Repugnant Conclusion) and on the other hand of Average Utilitarianism (violating Weak Mere Addition), some theorists have been attracted to a family of views that are often called *Variable-Value views*.³ Some views within this family avoid the Repugnant Conclusion altogether while capturing the intuition that adding a well-off person to a *small* population makes the resulting population better. More specifically, these views hold that the quantity that added persons (with a fixed level of welfare) contribute towards the overall value of a population decreases as the size of the population increases, *cumulatively* contributing only a bounded amount, which is how such views escape the Repugnant Conclusion. Other views within this family merely limit the reach of the Repugnant Conclusion, in the sense that, compared to Total Utilitarianism, they less often judge large populations were people are not so well off to be better than small populations of well-off people.

Various versions of Variable-Value views have been rigorously formalized.⁴ These formalizations and the ensuing analysis has focused on *qualitative* properties of Variable-Value population ethics: with which *axioms* do these proposals comply? However, there

²By “intertemporal world population” we mean the totality of humanity (see fn. 6) throughout history.

³Hurka (1983) coined the term, and was probably the first to suggest such a view in response to Parfit’s Repugnant Conclusion, but views in this family have since been proposed or investigated by Ng (1989), Sider (1991), Asheim and Zuber (2014), and Pivato (2020), although not all of these authors endorsed the Variable-Value axiology that they identified or explored.

⁴Examples include Ng (1989), Asheim and Zuber (2014), and Pivato (2020).

has not been a similar focus on the *quantitative* implications of these Variable-Value views. In particular, one might wonder *how fast* the quantity that an added person contributes towards the overall value of a population diminishes as e.g. the size and average welfare of the population increases, and what implications that will have for various trade-offs between increasing the size and the average welfare of a population. Similarly, one might wonder precisely which weakenings of the Mere Addition principle these views can accommodate without implying seemingly repugnant instances of the Repugnant Conclusion.⁵

Our aim in this paper is to fill the above gap in the population ethics literature. In particular, we shall assume some very weak instances of Mere Addition and calibrate what Variable-Value axiologies, that satisfy such weak instances of Mere Addition (but violate the stronger Mere Addition principle that Total Utilitarianism entails), imply under what we take to be plausible empirical assumptions about the future. Informally, the weak Mere Addition that we assume ensures that merely adding people who are very well-off by modern standards, such as professors in the developed world, does not make the population worse. The empirical assumption we make is that the future of humanity is long and prosperous, such that, in particular, the average welfare of the total intertemporal world population⁶ is higher than the average welfare of the world population up to 2020.⁷

⁵Our aim is not to examine *all* Variable-Value views. In particular, we shall not be concerned with those variable-value views that satisfy the strong version of Mere Addition (i.e., the version entailed by Total Utilitarianism), such as the theory examined in Sider's (1991). Instead, the aim is to examine those views that (unlike Average Utilitarianism) imply some weak instance of Mere Addition, without implying the strong version of Mere Addition.

We note also that a normative reason for excluding from our examination the view in Sider (1991) is that it implies what Arrhenius's (forthcoming) calls "The Very Anti Egalitarian Conclusion: For any perfectly equal population of at least two persons with positive welfare, there is a population which has the same number of people, lower average (and thus lower total) welfare and inequality, which is better." In fact, Sider himself rejects the view due to implications like this (Sider, 1991, 270).

⁶In what follows, we focus on *human* populations. This is not because we think that the welfare of animals is unimportant. But how precisely to integrate animal welfare into population ethics is far from evident. Therefore, for the sake of simplicity, we focus on human welfare and populations of people.

⁷If the reader finds our empirical assumptions implausible, then she can of course read our argument and conclusion as being *merely conditional* on these assumptions.

Our main observation is that, when combined with the above two assumptions, Variable-Value axiologies imply countless *prima facie* problematic instances of the Repugnant Conclusion. (By an “instance” of the Repugnant Conclusion, we mean the judgement that some particular population consisting only of lives that are barely worth living is better than some particular perfectly equal population of very well-off people.) Of course, they do not imply the *qualitative* Repugnant Conclusion stated above — which holds for *all* populations of very well-off people. But these implications, we argue, should nevertheless seem every bit as repugnant to those who oppose to the Repugnant Conclusion.⁸ So, on the face of it, these results would seem problematic for any arguments that the Repugnant Conclusion requires us to reject Total Utilitarianism in favor of alternatives that avoid the qualitative Repugnant Conclusion, or that escaping the qualitative Repugnant Conclusion is a goal of core methodological importance for population ethics. Ultimately, such problems are important for the methodology of population ethics even if Variable-Value axiologies prove inadequate for other reasons.

It might be worth providing some additional remarks to motivate our approach.⁹ First, we assume that some (but perhaps not all) of those who are happy with giving up the traditional Mere Addition principle will nevertheless find it hard to reject some very weak instances of the principle. After all, we seem to have stronger reasons to think that a mere addition of a very well off person does not make the world worse than we have to think that a mere addition of a life barely worth living does not make the world worse.¹⁰ Therefore, there is, we think, something to be gained from exploring what happens when we replace Mere Addition with a weaker principle that only applies to people who are

⁸In fact, according to the principle of “unrestricted instantiation” (Tännsjö, 2020), these implications *must* be seen as repugnant if the Repugnant Conclusion is to be an argument against Total Utilitarianism.

⁹Many thanks to [blinded] for making us see the need to address the motivation.

¹⁰If one is, say, a convinced total utilitarian, then one might nevertheless have equally strong (subjective) reason to think that a mere addition of a very well off person *could* make the world worse as one has to think that a mere addition of a life barely worth living *could* make the world worse. (Cf. Lenman, 2000, Greaves, 2016.)

very well-off.

Second, we think that valuable lessons can be learnt from exploring what population axiologies imply given reasonable empirical assumptions, as opposed to merely exploring what these axiologies imply in theory. In particular, our finding that Variable-Value views have counterintuitive implications, given empirical assumptions that we accept for our actual world population, provides a valuable lesson that is not learnt from simply learning that these axiologies have counterintuitive implications given assumptions about the world population that we take to be false. For that shows that Variable-Value views do not only have counterintuitive implications in hypothetical scenarios; they also have counterintuitive implications in empirically plausible scenarios.

We proceed as follows. In Section 2 we lay out the formal framework of the paper, which allows us to state more formally the views and conditions we informally describe above, and introduce the reader to calibration methods in decision theory. Then, in Section 3, we begin with two simple, illustrative examples: we use such calibration methods to examine what two prominent Variable-Value views imply when they have been calibrated to the current world population and what we take to be reasonable assumptions about the future population. Section 4 contains our main formal result: we use the same calibration methods to present a more general result, that is, a result for a large family of Variable-Value views that do not satisfy the strong version of Mere Addition. Finally, in Section 5 we ask what the upshot of our arguments is for population ethics and in particular for the focus in the population ethics literature on avoiding the Repugnant Conclusion.

2 Formal framework for population ethics

Our framework, terminology, and notation follow closely that of Asheim and Zuber (2014). Let \mathbb{N} denote the set of natural numbers and \mathbb{R} the set of real numbers, while $\mathbb{R}^+ \subset \mathbb{R}$

denotes the set of strictly positive real numbers. Let $\mathbf{X} = \bigcup_{n \in \mathbb{N}} \mathbb{R}^n$ denote the set of possible finite distributions of lifetime well-being. More formally, $\mathbf{X} = \bigcup_{n \in \mathbb{N}} \mathbb{R}^n$ is a set of vectors of real numbers, where each number represents the lifetime well-being of some person. A generic such vector for a population of m people is denoted $\mathbf{x} = (x_1, \dots, x_m)$, where x_i denotes the lifetime well-being of individual i . The size of the population given by \mathbf{x} is denoted by $\mathcal{N}(\mathbf{x})$ (and will, as mentioned, always be finite). For any vector \mathbf{x} , we write the average lifetime well-being of its members as \bar{x} .

Let \succsim on \mathbf{X} denote a (weak) better-than relation on \mathbf{X} , such that for any $\mathbf{x}, \mathbf{y} \in \mathbf{X}$, $\mathbf{x} \succsim \mathbf{y}$ means that \mathbf{y} is at least as good as \mathbf{x} . Throughout the discussion we shall assume that the better-than relation is transitive, reflexive, and complete,¹¹ which means that the relation generates a better-than *order*. The strict relation, \prec , and indifference, \sim , are respectively the asymmetric and symmetric counterparts of \succsim .

Built into our framework is an *anonymity* axiom, which holds that the “better-than relation” we study is invariant under permutations of the vectors in \mathbf{X} . For instance, let \mathbf{x}' be the vector that results when the lifetime well-being of i and j in \mathbf{x} are switched. Then the better-than relations that we shall consider are all indifferent between \mathbf{x} and \mathbf{x}' , that is, they deem these two distributions to be equally good. Intuitively, this means that it does not matter *who* receives what welfare; all that matters is how many people have each welfare level. This assumption rules out some person-affecting views.

For any $\mathbf{x} \in \mathbf{X}$ with m members, let $\mathbf{x}_{\square} = (x_{[1]}, \dots, x_{[r]}, \dots, x_{[m]})$ be the nondecreasing reordering of \mathbf{x} . In other words, in \mathbf{x}_{\square} the elements of \mathbf{x} have been put in a nondecreasing order, such that for each rank $r \in \{1, \dots, m\}$, $x_{[r]} \leq x_{[r+1]}$, meaning that individual with rank $r + 1$ is at least as well off as individual with rank r . The anonymity assumption

¹¹Although the assumption of completeness is standard in the population economics literature, some population ethicists have made attempts to avoid the Repugnant Conclusion by relaxing it. (See e.g. attempt by Parfit, 2016; see also discussion in Arrhenius, 2016.) However, since the aim of this paper is to explore Variable-Value views which imply a complete better-than relation, the assumption of completeness is harmless here.

ensures that when two or more individuals are equally well-off, how they are ranked relative to each other does not affect the ranking of populations.

Finally, $(z)_n \in \mathbb{R}^n$ denotes the perfectly-equal distribution where all n individuals have lifetime well-being z . And let $(\mathbf{x}, (z)_n)$ denote distribution $\mathbf{x} \in \mathbf{X}$ with n added individuals that all have lifetime well-being z . When only one individual with well-being level y is added to \mathbf{x} , we denote this by (\mathbf{x}, y) .

With this formalization, different axiological views, such as those discussed above, can be seen as different views about the structure of \succsim . This allows for convenient formal statements of the views and conditions we informally discussed in the last section. For instance, Total Utilitarianism can be formulated thus:

Total Utilitarianism (TU). *For any $\mathbf{x}, \mathbf{y} \in \mathbf{X}$:*

$$\mathbf{x} \succsim \mathbf{y} \Leftrightarrow \sum_i x_i \leq \sum_i y_i$$

We can now also state the Repugnant Conclusion more formally:¹²

The Repugnant Conclusion (Formal version). *For all $y, z \in \mathbb{R}$, where $y > z > 0$, and for any $k \in \mathbb{N}$, there is a $n \in \mathbb{N}$ such that $(y)_k \prec (z)_n$.*

It is easy to verify that Total Utilitarianism implies the Repugnant Conclusion.¹³

The Variable-Value views that we later discuss will be contrasted with both Average Utilitarianism (to be formally defined) and Critical-Level Generalized Utilitarianism

¹²This formalization is slightly different from that of Blackorby et al. (2005), who require that $n > k$. See Spears and Budolfson (2021) for a discussion of heterogeneity in formalizations of the Repugnant Conclusion in the prior literature.

¹³Nebel (forthcoming) has recently formulated a version of totalism that avoids the Repugnant Conclusion, by including a lexical threshold in the conception of individual welfare. As our aim here is not to discuss the extent to which Total Utilitarianism implies the Repugnant Conclusion — but rather the extent to which Variable-Value views imply the repugnant conclusions — we will not discuss Nebel’s or other totalist views that avoid repugnance. (We do however explore lexical views in [blinded].)

(CLGU).¹⁴ CLGU is a family of generalized total utilitarian views, which include for instance Critical-Level Utilitarianism, Prioritarianism, and Total Utilitarianism as special cases. The general view can be stated thus:

Critical-Level Generalized Utilitarianism (CLGU). *For any $x, y \in \mathbf{X}$:*

$$x \succsim y \Leftrightarrow \sum_i [g(x_i) - g(c)] \leq \sum_i [g(y_i) - g(c)]$$

where g is non-decreasing and non-convex, and $c \geq 0$ is the critical level at which adding a life becomes a social improvement.

Some views in this CLGU family avoid the standard formalization of the Repugnant Conclusion, at the cost of entailing another counterintuitive result, such as Arrhenius's (2000) family of sadistic conclusions (Franz and Spears, 2020). CLGU is a significant family for our purposes because it is fully additively separable. That is, a CLGU better-than relation satisfies same-number, different-number, and existence independence axioms, always with a constant critical level (Blackorby et al., 2005). So, because our calibration results depend on additive separability being violated, no CLGU view is subject to the calibration arguments of our paper.

Average Utilitarianism can now simply be stated as:

Average Utilitarianism (AU). *For any $x, y \in \mathbf{X}$:*

$$x \succsim y \Leftrightarrow \bar{x} \leq \bar{y}$$

It can be easily verified that Average Utilitarianism does not imply the Repugnant Conclusion. However, Average Utilitarianism is well-known to violate the Mere Addition principle, which we can now formally state as:

¹⁴CLGU was introduced by Blackorby and Donaldson (1984) and subsequently explored in depth by Blackorby et al. (2005) and Bossert (2017).

Mere Addition (Formal version). *For any $x \in X$, and for any $z \in \mathbb{R}^+$, $x \succsim (x, z)$.*

Denying Mere Addition, for a complete ordering, is equivalent to entailing what we call the Anti-Natalist Conclusion:

Anti-Natalist Conclusion. *There exists a $z \in \mathbb{R}^+$ and an $x \in X$ such that $(x, z) \prec x$.*

In the remainder of this paper, we examine a novel weakening of Mere Addition that we argue is highly plausible (and that we believe will seem plausible to even some of those who are willing to reject the strong Mere Addition principle). To state the principle, let us stipulate that there is well-being level beyond which lives at that level are excellent by the standards of 21st-century developed countries; and let $\mathbb{E} \subset \mathbb{R}^+$ be the set of well-being levels that are excellent by this same standard. For concreteness, let's set that level at 97.5th percentile of lifetime well-being in our current global population. To make things even more concrete, we shall occasionally assume that a typical professor in a developed country is at that level.¹⁵ And let X_R be the set of vectors that could realistically represent the lifetime well-being of the intertemporal (human) world population.¹⁶ We can now finally state:

Weak Mere Addition. *For any $x \in X_R$, and for any $y \in \mathbb{E}$, $x \succsim (x, y)$.*

While denying Mere Addition, for a complete ordering, “only” implies accepting the Anti-Natalist Conclusion, denying Weak Mere Addition in addition implies accepting a Strong Anti-Natalist Conclusion:

Strong Anti-Natalist Conclusion. *There exists a well-being level $y \in \mathbb{E}$ and a population $z \in X_R$ such that $(z, y) \prec z$.*

¹⁵ Note that since \mathbb{E} is a subset of \mathbb{R}^+ , the concreteness assumption implies that a typical professor in a developed country, and more generally people in the top 2.5 percent of the current global well-being distribution, have lives worth living. While some may question this assumption (e.g., Benatar, 1997), we hope that it will strike most readers as innocent.

¹⁶The exact content of this set is not important for our argument. We merely introduce the set to emphasize that our argument does not require the use of any arbitrary world population, but only (what we take to be) realistic populations.

We ourselves are sceptical of the Anti-Natalist Conclusion. But we think that there is even stronger reason to reject the Strong Anti-Natalist Conclusion. The latter is of course logically stronger than the former (since \mathbb{E} is a strict subset of \mathbb{R}^+). Moreover, we think that, intuitively, the lives of people in the top 2.5 percent of the current global well-being distribution are all *much* better than barely worth living. Some may question this intuition (e.g., Benatar, 1997), but we shall simply take it for granted. But then the Strong Anti-Natalist Conclusion will be correspondingly “much” stronger than the Anti-Natalist Conclusion.

As we show in the next two sections, however, unless Variable-Value views imply the Strong Anti-Natalist Conclusion, when these views have been calibrated to plausible empirical assumptions about the size and welfare of the future world population, they must imply many instances of the Repugnant Conclusion; and, in fact, many instances that would seem problematic for those who recommend Variable-Value views over Total Utilitarianism due to the Repugnant Conclusion.

This follows from our novel application of a familiar logic in decision theory: calibration of variable-value objective functions to reveal tensions between intuitions for large-quantity decisions and intuitions for small-quantity decisions. The leading result in this literature is Rabin’s (2000) celebrated argument about expected utility theory. Formally, we merely extend Rabin’s argument about choice under risk to analogous functional forms in population ethics.

Rabin established that an expected utility maximizer can only be moderately risk averse when relatively small sums of money are involved—e.g. always turning down 50-50 gambles between losing \$100 and winning \$105—if she is extremely risk averse when larger sums of money are involved—e.g. turning down 50-50 gambles between losing \$2,000 and winning any (including infinite) amount of money. So, the lesson of Rabin’s argument is that an expected utility maximiser is either surprisingly risk averse when

stakes are large or surprisingly risk neutral when stakes are small.¹⁷

Our calibration result is that Variable-Value views are surprisingly anti-natalist when few extra lives are at stake (more specifically, they entail the Strong Anti-Natalist Condition) unless they are surprisingly totalist when more lives are at stake. (By “totalist” we mean making choices that should be found objectionable by those who reject Total Utilitarianism because of the Repugnant Conclusion.) In other words, our result is that Variable-Value views can only accommodate an extremely weakened Mere Addition principle if they also imply countless and *prima facie* problematic instances of the Repugnant Conclusion.

3 Two illustrative examples

This section turns to two algebraically simple examples of Variable-Value population axiologies. Both of these are well-known in the literature to either completely avoid traditional formalizations of the Repugnant Conclusion or at least limit the reach of the conclusion. We ask what these views recommend in repugnance-type tradeoffs between large and small populations, once calibrated to satisfy Weak Mere Addition and facts or predictions about the world population. Because they are algebraically tractable and well-studied, it is instructive to see why these particular versions of Variable-Value views imply seemingly problematic instances of the repugnant conclusions, before considering (in the next section) a more general argument.

¹⁷Nebel and Stefánsson (2020) apply a similar logic to inequality averse views about how to order populations of a fixed size, in particular, to Prioritarianism and Rank-Discounted Utilitarianism, and find that such views can only be moderately inequality averse when small differences in welfare are at stake if they are extremely inequality averse when larger welfare differences are at stake.

3.1 Number-dampened generalized utilitarianism: The power-function case

The first view in the Variable-Value family that we shall consider can be stated as follows:

Number-Dampened Generalized Utilitarianism (NDGU). *There is a concave (and increasing) real-valued function f such that for any $\mathbf{x}, \mathbf{y} \in \mathbf{X}$:*

$$\mathbf{x} \succsim \mathbf{y} \Leftrightarrow \bar{x} f(\mathcal{N}(\mathbf{x})) \leq \bar{y} f(\mathcal{N}(\mathbf{y}))$$

The concavity of f means that NDGU reduces the value of additions to the population as population size grows. Moreover, if (but only if) f is bounded, then NDGU avoids the Repugnant Conclusion; otherwise it merely limits the reach of the conclusion, as previously explained.

For illustrative purposes, we assume in the below example that

$$f(\mathcal{N}(\mathbf{x})) = \mathcal{N}(\mathbf{x})^\alpha$$

for some $\alpha \in (0, 1)$. (But, as previously mentioned, we provide a more general argument in the next section.) Note that this simplifying assumption too means that the value of additions to a population diminishes as the population size grows. One way to see this, and connect it to our formulation of the Repugnant Conclusion, is that $(y)_k \prec (z)_n$ just in case $\frac{y}{z} < \left(\frac{n}{k}\right)^\alpha$. From this we can see that as α approaches 0, the difference between y and z needed for NDGU to imply that $(z)_n \succsim (y)_k$ becomes smaller; which is just another way to say that the closer α comes to 0, the closer NDGU gets to Average Utilitarianism, and the fewer instances of the Repugnant Conclusion it implies. However, when $\alpha > 0$ — and, more generally, when f is unbounded — NDGU implies the Repugnant Conclusion, and, moreover, for *some intuitively large* difference between y and z , and for any k , there will be

an n such that $(y)_k \prec (z)_n$.

Moreover, while NDGU violates the strong version of Mere Addition that Total Utilitarianism implies (in the simple case: since $\alpha < 1$), it does imply some weak instances of Mere Addition. Our aim now is to explore what happens when we assume that it satisfies some particular (very plausible) instances of Weak Mere Addition.

The above algebraic formulation applies a particular, one-parameter family of functional forms to the concave-transformation proposal of Hurka (1983) and Ng (1989). Here we expand our historical horizons: The relevant fact is that adding a well-off (by today's standards) professor lowers average lifetime well-being because there will be many future people who will be even better off. In making this assumption, we follow recent literature on the *possible* long-term prosperity of humanity, such as by Greaves and MacAskill (2019) and Ord (2020). If humanity reaches its full potential, both in terms of quality and quantity, then α must be high if the addition of a well-off professor is not to be a worsening.

If there will be 2×10^{12} people overall, for example, and if the average lifetime well-being will be three times as high as that of the modern professor's, then that implies an α of at least 0.67.¹⁸ If the future is even better than that and the average person will be five times as well-off as our professor, then we have an α of 0.80 — greater because today's extra happy professor pulls the intertemporal average down by more.

These, too, imply seemingly problematic repugnant-like quantitative consequences. For instance, $\alpha = 0.67$ implies that a population of 100,153,521 people with lifetime well-being of 0.1 is better than a population of 100,000 people with lifetime well-being of 10 (i.e., a population where each person is 100 times better off than in the larger low quality population). And $\alpha = 0.80$ implies that a population of 31,683,480 people with lifetime well-being of 0.1 is better than a population of 100,000 people with lifetime well-being of 10.¹⁹ Both of these better-than judgements (i.e., those entailed by $\alpha = 0.67$ and $\alpha = 0.80$)

¹⁸Recall, from fn. 6, that we focus in this paper on human populations.

¹⁹The reason why the number of people with 0.1 lifetime well-being that is needed to outweigh 100,000

would, we suspect, be found repugnant by many of those who oppose Total Utilitarianism due to the Repugnant Conclusion.

3.2 Rank-discounted generalized utilitarianism

The second Variable-Value view that we shall consider can be stated as follows:²⁰

Rank-Discounted Generalized Utilitarianism (RDGU). *There is a $\beta \in (0, 1)$ such that for any $\mathbf{x}, \mathbf{y} \in \mathbf{X}$:*

$$\mathbf{x} \succsim \mathbf{y} \Leftrightarrow \sum_r \beta^r g(x_{[r]}) \leq \sum_r \beta^r g(y_{[r]})$$

where g is increasing and weakly concave.

A version of this view is defended by, for instance, Asheim and Zuber (2014). It avoids the (universally quantified) Repugnant Conclusion because $\beta^1 + \beta^2 + \beta^3 \dots$ is a convergent series, which ensures that the aggregated value of a perfectly-equal population is bounded and remains finite, no matter how large it becomes. Therefore, if k , in our formal statement of the Repugnant Conclusion, is sufficiently large, and if y is sufficiently larger than z , then there is *no* n such that, by RDGU, $(y)_k \prec (z)_n$. In other words, the (universally quantified) Repugnant Conclusion does not follow from RDGU. But, if β is sufficiently close to 1, then even large y could be part of an instance of the Repugnant Conclusion with a z that is small enough to capture the purportedly repugnant features of the Repugnant Conclusion. (Note that, by definition, $\beta < 1$. However, if β were equal to 1, then RDGU would be equivalent to CLGU with $c = 0$.)

RDGU does not satisfy the strong Mere Addition principle that Total Utilitarianism entails. This is because adding a life to a population lowers the weights of any otherwise-existing higher-utility lives, which may worsen the population by more than the additional

people with lifetime well-being of 10 is smaller in the second case, is that as α becomes larger, the resulting axiology becomes closer to a totalist one.

²⁰Recall that $\mathbf{x}_{[]} = (x_{[1]}, \dots, x_{[r]}, \dots, x_{[m]})$ is the nondecreasing reordering of \mathbf{x} .

life improves it. However, RDGU must satisfy *some* Weak Mere Addition principle, since $\beta > 0$, which means that *some* mere additions are valuable. And, in fact, the closer β is to 1, the closer RDGU comes to implying the strong Mere Addition principle, in the sense of implying stronger instances of Mere Addition. We want to examine what RDGU implies if we assume that it satisfies particular (very plausible) instances of Weak Mere Addition.

To that end, we assume that a typical professor in the developed world is at the 97.5th percentile of lifetime well-being in our current world (human) population,²¹ so only 2.5 percent of people are better-off. That is around 182 million people. Under RDGU, adding such a professor reduces the weight on those 182 million very well-off people. The implication is that, if adding such a professor is not a worsening, β must be very close to 1. In particular, β must be greater than 0.99999995 if the 182 million people are all *no more than 1.0001 times as well-off* as our well-off professor, and even closer to 1 if the better-off people are even better-off than that (which is of course more realistic). β reaches 0.999999999 if the better-off people are a little more than six times as well-off as the professor, for example.²²

Now consider what these high β s imply for repugnant-like tradeoffs. How many people each with a life at 0.1 would be needed to be better than a population with 100,000 people at 10 (i.e., a population where each person is 100 times better off than those in the low quality population)? The answer ranges from 13,813,089 for $\beta = 0.99999995$ down to just above ten million as β becomes closer to 1.^{23,24} But those who believe that the Repugnant Conclusion must be avoided at all theoretical cost will presumably not be happy to have 13.8 million people at 0.1 instead of 100,000 people at 10. And yet, such is the consequence

²¹For simplicity we here ignore past and future generations, but the large number of future people who would, we assume, be better off than today's happy professors only increases the force of this argument.

²²Note that our argument can accommodate a positive critical level and a prioritarian transformation: simply subtract the critical level and/or do the prioritarian transformation before the rank-based weighting.

²³The sum of a geometric series is $\frac{1-\beta^n}{1-\beta}$. Because well-being is constant in both populations, this requires solving for n such that $0.1 \times \frac{1-\beta^n}{1-\beta} > 100 \times \frac{1-\beta^{100}}{1-\beta}$.

²⁴Here we follow the convention of normalizing the g measure around the 0 welfare level, that is, we assume that $g(0) = 0$.

of choosing RDGU and maintaining that creating the happy professor is not a worsening given our actual world population.

4 A more general argument

The quantitative results of Section 3 depended upon two specific functional forms. Here we present a more general argument which applies to any variable-value goodness function of the form

$$\text{for any } \mathbf{u} \in \mathbf{X} : W(\mathbf{u}) = \bar{u} \times f(\mathcal{N}(\mathbf{u})),$$

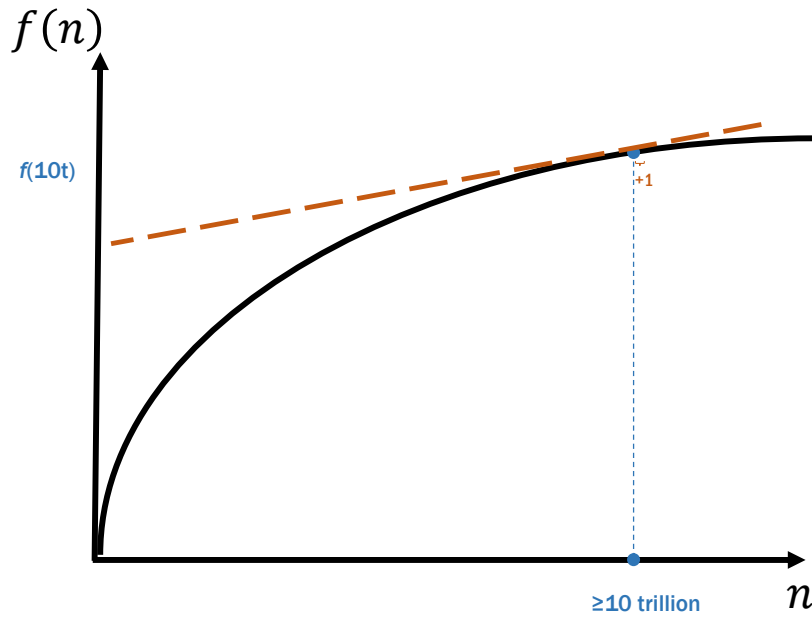
where f is positive, increasing, and concave.²⁵

Figure 1 illustrates the argument. First, fix a lifetime well-being level for a weak mere addition assumption: the lowest lifetime well-being level that, if added to our intertemporal population, would not decrease the goodness of the population. So far, we have been using the typical lifetime well-being of a developed-country professor, or more specifically the 97.5th percentile of the 2020 global socioeconomic distribution. But here, to get the strongest possible argument, we want the *lowest* lifetime well-being level for which mere addition reasoning can be safely applied. We expect that, for most readers, this is a level below that of most lives in 2020 developed countries. We denote this a well-being level of 1.

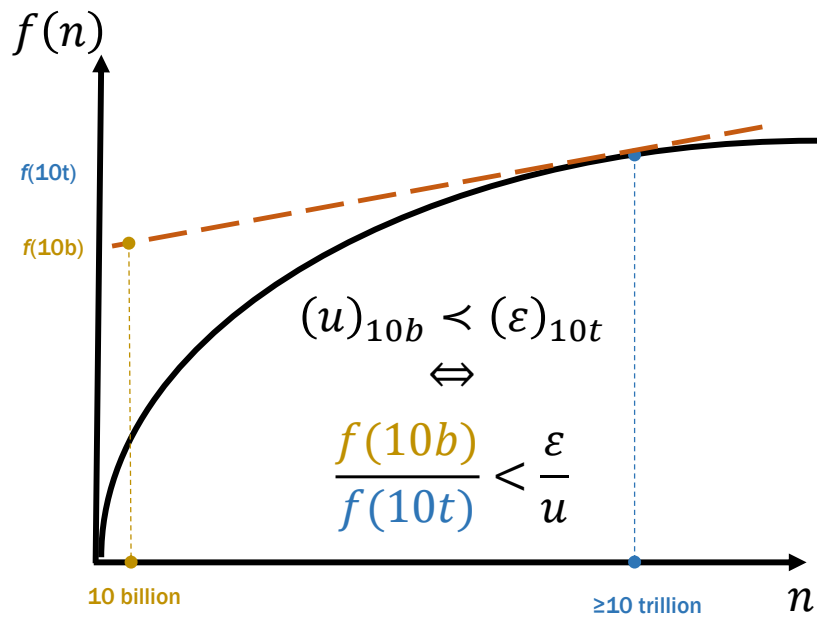
Next fix a lifetime well-being level of what we expect the business-as-usual average lifetime well-being level to be for the overall intertemporal human population. By “business-as-usual” we mean the situation that will occur if the one life at well-being level 1 is *not* added. We label this average lifetime well-being level γ for “good” and assume that the long-run future of humanity is such that γ is greater than 1.

²⁵Note that such a view violates the strong Mere Addition principle entailed by Total Utilitarianism. Hence, the view examined in Sider (1991) does not have this form. But, as mentioned in fn. 5, there are strong normative reasons for excluding from consideration the view in Sider (1991).

Figure 1: A graphical representation of concave variable-value calibration
 Panel a. Addition of a good life bounds the slope of f at 10 trillion (10t)



Panel b. This slope bounds the ratio of f at 10 trillion and 10 billion (10b)



For an explanation of the choice of population sizes, see fn. 26.

We must next make an approximate assumption about the size of the business-as-usual population. We assume 10 trillion²⁶, but the below logic can be applied to a different assumption; it would merely change our example quantities. Because 1 is the lowest well-being level such that adding it to a population does not make the population worse, and because \bar{u} is continuous in utilities, given the functional form of W we can conclude that:

$$\frac{\gamma \times 10 \text{ trillion} + 1}{10 \text{ trillion} + 1} f(10 \text{ trillion} + 1) = \gamma f(10 \text{ trillion}).$$

This inequality fixes the slope b of f at 10 trillion: $b = \frac{f(10 \text{ trillion}) \times (\gamma - 1)}{\gamma \times 10 \text{ trillion} + 1}$.

10 billion is a commonly-used size of the high-welfare population (A) in the Repugnant Conclusion literature. By the concavity of f , the slope at 10 billion can be no less positive than the slope at 10 trillion. Extending this linearly towards zero provides an upper bound on f at 10 billion — but if concavity is steep the actual value may be well below this bound, resulting in even more repugnant-like implications. In particular, we can conclude by concavity that $f(10 \text{ billion}) < f(10 \text{ trillion}) - b \times (10 \text{ trillion} - 10 \text{ billion})$, so the ratio of $f(10 \text{ billion})$ to $f(10 \text{ trillion})$ is bounded above.

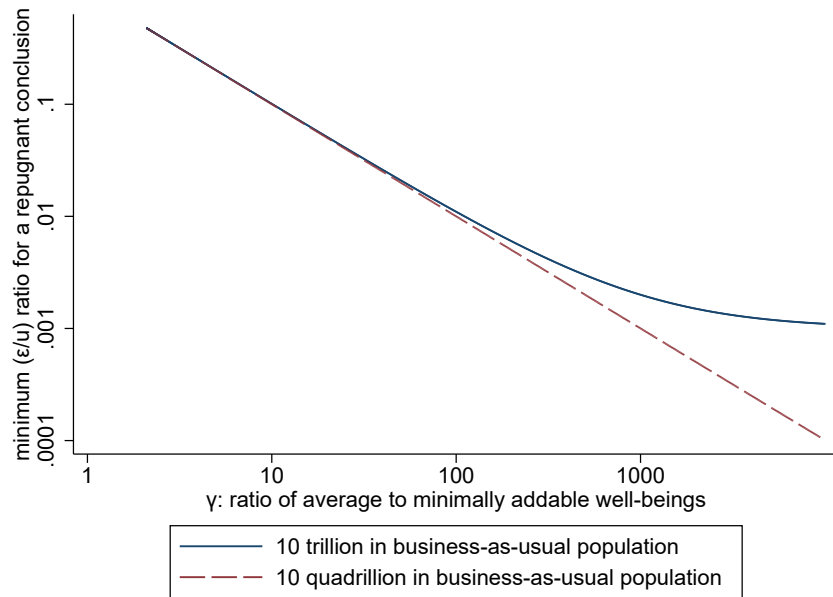
Now we are in a situation to draw quantitative repugnant conclusions. Let ε be a low positive utility level and u be a much higher utility level. A case where 10 trillion lives at ε is better than 10 billion lives at u would be an instance of the repugnant conclusion. As 1 depicts, this depends on the ratio that we have just bounded. The larger population will be better if $\frac{\varepsilon}{u} > \frac{f(10 \text{ billion})}{f(10 \text{ trillion})}$. By the bounding argument above, we can know this is true if

$$\frac{\varepsilon}{u} > 1 - \frac{(10 \text{ trillion} - 10 \text{ billion}) \times (\gamma - 1)}{\gamma \times 10 \text{ trillion} + 1}.$$

Therefore, there is a repugnant conclusion if $\frac{\varepsilon}{u}$ exceeds a threshold that is a function of γ , which is plotted in Figure 2. Notice that both axes of Figure 2 are *ratios* of well-

²⁶10 trillion is an Ord (2020)-type estimate of the plausible size of the intertemporal human population.

Figure 2: Result of more general argument: minimum $\frac{\epsilon}{u}$ threshold for a repugnant conclusion as a function of γ



beings, which underscores that this argument does not depend on any precise assumptions about well-being levels. The figure includes the case we have been considering, where the business-as-usual intertemporal population has 10 trillion members, as well as an alternative case with 10 quadrillion members, as a robustness check to verify that it makes little difference. The larger γ , the smaller ϵ can be as a ratio of u and the more “repugnant” the example can seem.

The wide horizontal axis in Figure 2 can presumably accommodate many readers’ beliefs about the γ and the plausible goodness of the future. The implications for large enough γ seem to constitute the very same repugnance that motivates some population ethicists to avoid Total Utilitarianism. To be sure, our argument uses an empirical premise about the size and well-being of the future population, formalized in the assumption that $\gamma > 1$. We believe that it is a strength of our argument that it speaks to a plausible calibration of the actual world. After all, that means that the implications we derive are not

merely theoretical possibilities, but rather results that we would get if we were to apply the theories under examination in actual policy evaluation (under certainty).²⁷ But any reader uncomfortable with these empirical premises can read our argument as a *conditional* one, where the results are conditional on a plausible and relevant hypothetical future.²⁸

5 Lesson and concluding remarks

Recall that the intuitive appeal of Variable-Values views was supposed to be that they could avoid (or at least limit the reach of) the Repugnant Conclusion while satisfying at least some weak instance of the Mere Addition principle. We have now seen, however, that if these views satisfy what we take to be a very plausible, and certainly weak, instance of Mere Addition, and if in addition we make plausible empirical assumptions about the intertemporal world population, then these Variable-Value views have implications that, we suggest, those who oppose Total Utilitarianism due to the Repugnant Conclusion will find repugnant.²⁹

What should we conclude from our results? Most narrowly, a lesson of our results is that when calibrated to the real world — that is, the actual world population and what we think are plausible empirical assumptions about the future population — Variable-Value views would substantially agree with Totalist views on how to rank most policies that affect a relatively small number of people. Moreover, if we assume that policy choices typically affect only a relatively small number of people — that is, small in relation to the total intertemporal world population — then the implication is that these Variable-Value views and Totalist views typically recommend the exact same courses of action (especially

²⁷Actual policy evaluation is of course an exercise in decision-making under uncertainty, which makes it even harder to avoid repugnant-like conclusions, as Arrhenius and Stefánsson (2018) point out.

²⁸See, e.g., Ord (2020), Greaves and MacAskill (2019).

²⁹In fact, given Tännsjö's principle of unrestricted instantiation (recall fn. 8), these implications *must* be deemed repugnant if the Repugnant Conclusion is to be used as an argument against Total Utilitarianism.

if the menu of possible options is coarse). The only escape would be for these Variable-Value views to be strikingly anti-natalist, such that they do not even satisfy weak instances of Mere Addition that would involve the lives of well-off readers of this paper.

More broadly, these results teach us something about the effort to avoid the Repugnant Conclusion. Population ethicists have long understood that escaping undesirable or unintuitive implications is impossible. But this paper adds to a growing recent literature — including Spears and Budolfson (2021) on additions to an unaffected population and Arrhenius and Stefánsson (2018) on risky choice between uncertain populations — that finds repugnant conclusions even under approaches to population ethics commonly understood to avoid repugnance. Collectively, these results suggest that the effort to avoid the Repugnant Conclusion has, in some ways, hinged on questionable features of the formalization of repugnance (such as the features that exclude the cases documented in this paper); that some of this effort may therefore be misplaced; and that perhaps avoidance of the Repugnant Conclusion should not be a core goal of population ethics research. We therefore suggest that our arguments here support the position in Zuber et al. (2021), a recent statement of agreement by many authors from diverse perspectives who argue that avoiding the Repugnant Conclusion has been overemphasized by population ethics research.

References

- Arrhenius, Gustaf.** 2000. “An impossibility theorem for welfarist axiologies.” *Economics & Philosophy*, 16(2): 247–266.
- 2016. “Population Ethics and Different-Number-Based Imprecision.” *Theoria*, 82(2): 166–181.

- Arrhenius, Gustaf and H Orri Stefánsson.** 2018. "Population ethics under risk." working paper, IFFS.
- Arrhenius, Gustaf.** forthcoming. *Population Ethics: The Challenge of Future Generations*: Oxford University Press.
- Asheim, Geir B and Stéphane Zuber.** 2014. "Escaping the repugnant conclusion: Rank-discounted utilitarianism with variable population." *Theoretical Economics*, 9(3): 629–650.
- Benatar, David.** 1997. "Why It Is Better Never to Come Into Existence." *American Philosophical Quarterly*, 34(3): 345–355.
- Blackorby, Charles and David Donaldson.** 1984. "Social criteria for evaluating population change." *Journal of Public Economics*, 25(1-2): 13–33.
- Blackorby, Charles, Walter Bossert, and David J Donaldson.** 2005. *Population issues in social choice theory, welfare economics, and ethics*: Cambridge University Press.
- Bossert, Walter.** 2017. "Anonymous welfarism, critical-level principles, and the repugnant and sadistic conclusions." working paper, University of Montréal.
- Franz, Nathan and Dean Spears.** 2020. "Mere Addition is equivalent to avoiding the Sadistic Conclusion in all plausible variable-population social orderings." *Economics Letters*, 196: 109547.
- Greaves, Hilary.** 2016. "Cluelessness." *Proceedings of the Aristotelian Society*, 116(3): 311–339.
- Greaves, Hilary and William MacAskill.** 2019. "The case for strong longtermism." working paper, Global Priorities Institute.
- Hurka, Thomas.** 1983. "Value and population size." *Ethics*, 93(3): 496–507.

- Lenman, James.** 2000. "Consequentialism and Cluelessness." *Philosophy and Public Affairs*, 29(4): 342–370.
- Nebel, Jacob M. and H. Orri Stefánsson.** 2020. "Calibration Dilemmas in the Ethics of Distribution." working paper, University of Southern California and Stockholm University.
- Nebel, Jacob M.** forthcoming. "Totalism Without Repugnance." in Tim Campbell, Jeff McMahan, and Ketan Ramakrishnan eds. *Festschrift for Derek Parfit*.
- Ng, Yew-Kwang.** 1989. "What Should We Do About Future Generations?: Impossibility of Parfit's Theory X." *Economics & Philosophy*, 5(2): 235–253.
- Ord, Toby.** 2020. *The Precipice: Existential Risk and the Future of Humanity*: Hachette Books.
- Parfit, Derek.** 1984. *Reasons and Persons*: Oxford.
- 2016. "Can we avoid the repugnant conclusion?" *Theoria*, 82(2): 110–127.
- Pivato, Marcus.** 2020. "Rank-additive population ethics." *Economic Theory*, 69(4): 861–918.
- Rabin, Matthew.** 2000. "Risk Aversion and Expected-utility Theory: A Calibration Theorem." *Econometrica*, 68(5): 1281–1292.
- Sider, Theodore R.** 1991. "Might theory X be a theory of diminishing marginal value?" *Analysis*, 51(4): 265–271.
- Spears, Dean and Mark Budolfson.** 2021. "Repugnant Conclusions." working paper; prior version is IZA Discussion Paper 12668.
- Tännsjö, Torbjörn.** 2020. "Why Derek Parfit had reasons to accept the Repugnant Conclusion." *Utilitas*: 1–11.

Zuber, Stéphane, Dean Spears, Johan Gustafsson, Mark Budolfson, and others. 2021.

“What should we agree on about the repugnant conclusion?” *Utilitas*: 1–5.