Intergenerational Population Ethics

Paolo G. Piacquadio*

October 2025

Abstract

This paper characterizes a novel class of welfare criteria for endogenous population models. The core innovation is to treat parents' fertility preferences as ethically relevant, deferring to them on the quality-quantity trade-off under conditions of equality. The framework disentangles aggregation over existing generations from the ethical choice of when to add people, yielding a simple, recursive, and time-consistent formulation that nests discounted utilitarianism as a special case. The resulting fertilitarian welfare criteria resolve long-standing normative dilemmas that afflict existing alternatives, including total and average utilitarianism. I illustrate the criteria in stylized examples and a Barro–Becker environment.

JEL Code: D60; D70.

Keywords: Intergenerational justice; fertility preferences; exponentially-discounted utilitarianism; social welfare.

*Department of Economics, University of St. Gallen, and CESifo Research Network. Email: paolo.piacquadio@unisg.ch. Address: Rosenbergstrasse 20, 9000 St. Gallen, Switzerland.

The author is grateful to Jonathan Ho, Rolf Aaberge, Geir Asheim, Yaroslav Chechel, Maya Eden, Bard Harstad, Paul Kelleher, Antony Millner, Flore Richard, Itai Sher, Dean Spears, Richard Tol, Christian Traeger, Kristen Vamsaeter, and Peter Wakker for comments, in addition to seminar and conference audiences in Copenhagen, Munich, NE&EP, Oslo, Oxford, Santa Barbara, Stanford, and St. Gallen. This project has received funding from the European Research Council under the European Union's Horizon 2020 research (Grant 804104).

1 Introduction

Population change plays a central role in shaping long-run economic outcomes. It affects resource scarcity, innovation, and intergenerational equity. For example, if climate change makes Earth a harsher place to live, households may choose to have fewer children. A smaller population may ease pressure on natural resources, but may also slow innovation. Yet, despite the centrality of population change, normative economics lacks a coherent and widely accepted criterion for evaluating policies that influence population size.[1] As a result, population dynamics is generally treated as exogenous and excluded from welfare analysis.

To illustrate the challenge, consider Robinson Crusoe deciding how many children to have. What principles determine whether the world in which Robinson has more children is better than the one in which he has fewer or no children?

Standard welfare criteria evaluate outcomes by aggregating the utilities of all individuals who will ever exist. Under total utilitarianism, social welfare increases with the sum of utilities. Under average utilitarianism, it increases with the mean. Both approaches lead to counterintuitive and often unacceptable conclusions (Blackorby et al., 2005; Arrhenius, 2000; Greaves, 2017).

Under total utilitarianism, Crusoe should have a child whenever the child's utility exceeds his own utility loss. This implies that multiple children with moderate utility can justify extreme suffering of the parent.[2] Moreover, a sufficiently large population living barely tolerable lives can be ranked above a smaller population enjoying excellent lives—this is the well-known repugnant conclusion (Parfit, 1984).

Average utilitarianism fares no better. Again, multiple children with sufficient

[1] The problem of evaluating population-sensitive allocations is widely recognized as one of the most difficult and unresolved issues in both economics and moral philosophy. See, among others, Blackorby, Bossert, and Donaldson (2005) for a systematic axiomatic treatment in economics, and Parfit (1984), Broome (2004), and Greaves (2017) for discussions in philosophy.

[2] In a companion paper (Piacquadio, 2025), I name this dilemma the "tyranny of parenthood" and show all standard criteria face it. More in detail, the tyranny of parenthood holds whenever the addition of children with modest utility—possibly a dynasty of additional persons (children, grandchildren, etc.)—can justify extreme suffering of the parents, regardless of their preferences.

utility may justify great parental suffering. Worse, a world with a much smaller population is ranked higher even if the gain in average utility is arbitrarily small.

To avoid these dilemmas, I make two departures from standard practice. First, I disentangle the two tasks of welfare criteria for endogenous population settings: (i) how to trade off the well-being of those who exist; and (ii) how to decide when to add people. Second, I propose the following desiderata to inform these tasks:

- Consistency with standard welfare criteria when population is fixed (exogenous), governing aggregation over existing generations; and
- Respect for parental fertility preferences: absent inequalities, adding a child increases social welfare only if it makes the parents better off.

The *fertilitarian* social welfare functions proposed here extend exponentially discounted utilitarianism to settings with endogenous population and defer the quality-quantity trade-off to parental preferences.[3] This paper provides their formal axiomatic foundation and shows how these criteria avoid the paradoxes that afflict total and average utilitarianism.

I consider an infinite-horizon overlapping-generations model in which generation t are the children of $t-1$ and the parents of $t+1$. The framework is standard and flexible, covering the endogenous-fertility models used in growth and climate applications.

I numerically illustrate the criteria in a stylized Barro-Becker model (Barro & Becker, 1989). Starting from a Pareto efficient steady-state, the fertilitarian criterion characterized here recommends no policy interference, whenever individual and social discounting are the same. Intuitively, at such steady-state there are no inequalities, so disagreement on the social discount rate is the only accepted justification for distorting parents' fertility choices. Barro and Becker also recognize

[3]More precisely, the type of exponentially-discounted utilitarianism considered here is one where: *i)* it is possible to add inequality aversion beyond the concavity of the individuals' utilities (as for prioritarianism); and *ii)* other-regarding preferences are excluded, in line with Dworkin (1977)'s argument against double-counting.

that parents may not place enough weight on future generations and suggest (Barro & Becker, 1989, p.497): "A social planner who is more sympathetic to children and other future generations than the current generation will place relatively more weight on numbers of people and consumption in the future."

In contrast to fertilitarianism, average utilitarianism fails time consistency and runs into the "Egyptology objection" (McMahan, 1981; Parfit, 1984): if previous generations were worse off (dragging the average well-being below the steady-state one), it is beneficial to increase the population beyond the steady-state level. Unsurprisingly, total utilitarianism is much more prone to population growth. It recommends dedicating almost all resources to childbearing and increasing fertility more than ninefold, despite the consequent 90% drop in per-capita consumption.

This paper primarily contributes to the literature on axiomatic population ethics by enriching the framework with a dynastic structure, disentangling aggregation from population choice, and relying on parents' fertility preferences. The present criteria sidestep the repugnant conclusion (Parfit, 1984). The trade-off between the quality of life and the number of individuals is established through fertility preferences. Thus, if a much larger population with very low well-being is considered repugnant, this will be reflected in parents' fertility preferences and, thus, incorporated in the fertilitarian social welfare function. However, the criteria may violate the mere addition principle: when fertility preferences of parents are satiated (and absent inequalities), social welfare respects their views. In this case, social welfare would not increase when children are added, even if these would enjoy a very high level of well-being.

A further contribution is to the literature on intergenerational justice: the separation yields a simple, recursive, and time-consistent extension of exponentially discounted utilitarianism to endogenous-population settings. The closest contribution is generation-relative utilitarianism, suggested by Dasgupta (2005, 2019). Similarly to Dasgupta's proposal, fertilitarianism ranks alternatives recursively from the per-

spective of the existing population (here, the parents). The main difference is that the weight attributed to the potential children is exogenous in generation-relative utilitarianism, while it is here derived endogenously from parents' fertility preferences. Blackorby, Bossert, and Donaldson (1995) characterize the "critical-level generalized utilitarian" criteria by imposing that the utilities of people who lived in the past do not matter. More recently, Asheim and Zuber (2014) and Pivato (2020) propose criteria that aggregate generations in a non-additive fashion. These criteria address instances of the repugnant conclusion by attributing more weight to low-utility generations.

The remainder of the paper is organized as follows. Section 2 presents the welfare criteria and discusses the properties of its functional form. Section 3 contains the main result, that is, the axiomatic characterization. Section 4 illustrates fertilitarianism and how it differs from alternative criteria. Section 5 concludes.

2 The welfare criterion

2.1 The model

Time is discrete, indexed by $t \in T \equiv \{0, 1, 2, ...\}$. Generation t is represented by a single agent who experiences **(personal) utility** $u_t \in \mathbb{R}_+$, that is, a measure of the generation's own well-being, independent of the welfare of others.[4] Generation t has fertility $n_t \geq 0$, that is, the number of children of generation t. An **allocation** is a sequence $\mathfrak{a} \equiv (\mathfrak{u}; \mathfrak{n})$, where $\mathfrak{u} = (u_0, u_1, ...)$ is the utility stream and $\mathfrak{n} = (n_0, n_1, ...)$ is the fertility stream. The set of all conceivable allocations is denoted $\boldsymbol{A} \equiv {\mathbb{R}_+}^{\infty} \times \mathbb{R}_+^{\infty}$. A constant allocation is one in which $(u_t, n_t) = (u, n)$ for each $t \in T$. It is strictly

[4] One can interpret such utility as a money-metric or consumption-equivalent measure of individual welfare. The assumption that each generation's utility is non-negative allows us to obtain a simple parametric representation of social preferences. Moreover, if utility is bounded from below, this restriction entails no loss of generality for the representation result. Importantly, in contrast to the population ethics literature, I do not use any information about the level of utility corresponding to a life just worth living. The social welfare criterion ranks allocations based on parental preferences over such utilities, not on absolute exogenous judgments about existence.

positive if $u, n > 0$, and it has unit fertility if $n = 1$.

Each generation t has preferences over the stream of utility and fertility pairs from period t onward, denoted $\mathfrak{a}_t = (\mathfrak{u}_t; \mathfrak{n}_t)$, where $\mathfrak{u}_t = (u_t, u_{t+1}, ...)$ and $\mathfrak{n}_t = (n_t, n_{t+1}, ...)$. These **dynastic preferences** are represented by a continuous function $U_t(\mathfrak{a}_t)$.[5] Dynastic preferences are assumed to satisfy two assumptions. *Monotonicity* imposes that a strictly larger constant utility is preferred. That is, for any constant streams $\mathfrak{u}_t \gg \mathfrak{u}'_t$, $U_t(\mathfrak{u}_t; \mathfrak{n}_t) > U_t(\mathfrak{u}'_t; \mathfrak{n}_t)$. *Finite substitutability* requires that for any constant strictly positive allocation $\mathfrak{a}_t$, there exist constant allocations $\mathfrak{a}'_t, \mathfrak{a}''_t$ with unit fertility such that $U_t(\mathfrak{a}'_t) < U_t(\mathfrak{a}_t) < U_t(\mathfrak{a}''_t)$.[6] These assumptions are standard in economic models of endogenous fertility and allow both pure (direct or indirect) and paternalistic altruism.

Social preferences, denoted $\succsim$, are defined as a complete and transitive binary relation over the domain $\boldsymbol{D} \subseteq \boldsymbol{A}$. For allocations $\mathfrak{a}, \mathfrak{a}' \in \boldsymbol{D}$, $\mathfrak{a} \succsim \mathfrak{a}'$ indicates that $\mathfrak{a}$ is socially at least as desirable as $\mathfrak{a}'$. The strict and indifferent counterparts, $\succ$ and $\sim$, are defined as usual.

Following standard practice in welfare economics, I define social preferences over allocations independently of feasibility constraints. That is, the desirability of a given path of utilities and population dynamics is derived exclusively from the ethical evaluation of the lives lived in the allocations being compared. Externalities and resource constraints enter downstream—at the policy stage—through feasibility, not

[5]The distinction between personal utility and dynastic preferences is similar to Dworkin's personal vs. external preferences (Dworkin, 1977). Dworkin holds that equal concern requires bracketing external preferences. In the present endogenous-population setting, dynastic preferences can be ethically salient without inducing unequal treatment of personal utilities. The framework allows for different personal and dynastic preferences across generations. Differences within generations can be introduced by interpreting the representative agent t as some aggregation of all individuals belonging to generation t, along the lines of Piacquadio (2017). Piacquadio (2025) discusses a welfare criterion with intragenerational heterogeneity, but does not provide an axiomatic characterization.

[6]Finite substitutability bounds the trade-off between (personal) utility and fertility. It excludes preference orderings that rank positive constant allocations either always higher or always lower than all constant allocations with unit fertility. In particular, this assumption implies that the constant allocation with zero utility (the worst admissible) and unit fertility is ranked lower than any strictly positive constant allocation. The characterization can be extended to allow such preferences by restricting the axioms to the subdomain of allocations that each generation ranks above the unit-fertility allocation with zero utility.

the definition of social preference.[7]

2.2 Fertilitarian social preferences

Fertilitarian social preferences evaluate social welfare recursively. At each t, social welfare W_t depends on the utility u_t of generation t, its number of children n_t, and the continuation welfare W_{t+1}. The central innovation is to derive the social value of population change from the dynastic preferences of parents over the quantity and quality of lives of their dynasties.

Let U_t^f denote the *dynastic-equivalent utility* of generation t. It captures the constant level of utility that all generations would need to achieve—under unit fertility—to make generation t indifferent to having fertility 1 instead of n_t. The function is derived from the representation of dynastic preferences U_t, restricted to allocations with no intergenerational inequalities, that is, at which all generations enjoy the same level of utility. Formally,

$$U_t^f(u, n_t) = k \quad \text{if and only if}$$
$$U_t\left((u, u, \dots), (n_t, 1, 1, \dots)\right) = U_t\left((k, k, \dots), (1, 1, 1, \dots)\right).$$

To evaluate intergenerational inequalities, compute the equally-distributed equivalent utility e_t. This quantity combines *(i)* the utility u_t enjoyed by generation t, *(ii)* its fertility choice n_t, and *(iii)* the continuation welfare W_{t+1} of its descendants. Intuitively, e_t is the single utility level which, if given equally to generation t and to all later generations, would make society indifferent between that hypothetical situation and the actual profile. Formally,

$$e_t \equiv v^{-1}\left(\frac{v(u_t) + n_t B \cdot v(W_{t+1})}{1 + n_t B}\right),$$

[7]For example, larger populations may place more pressure on finite resources or may foster innovation. On this trade-off, see Eden and Kuruc (2023). In this framework, such effects influence optimal policy through constraints, not through changes in the social ranking itself.

where $v : \mathbb{R}_+ \to \mathbb{R}$ is a continuous, strictly increasing, and concave function capturing inequality aversion, and $B > 0$ is a time preference parameter. This expression defines a Kolmogorov–Nagumo mean (generalized mean) of u_t and W_{t+1}, with relative weights 1 and $n_t B$. The term $n_t B$ reflects the weight attributed to the next generation, scaled by its population. Note that B can be recovered from the standard discount factor $\beta \in [0, 1)$ by setting $B = \beta/(1-\beta)$.[8]

Thus, **fertilitarian** social welfare is defined recursively as:

$$W_t = U_t^f(e_t, n_t),$$

where U_t^f is the dynastic-equivalent utility and e_t is the equally-distributed equivalent utility.

In the special case where v has constant elasticity of substitution, it takes the form

$$v(x) = \begin{cases} x^{1-\rho}/(1-\rho) & \text{if } \rho \neq 1, \\ \log x & \text{if } \rho = 1, \end{cases}$$

where $\rho \geq 0$ captures the degree of inequality aversion. I refer to this special case as *CES-fertilitarianism*.

2.3 Graphical Illustration

The recursive nature of the fertilitarian criterion is best illustrated through a stylized allocation. By construction, social welfare is expressed in the same units as utility. In particular, the welfare of a constant allocation with unit fertility is given by that constant level of utility enjoyed by all generations. Thus, for any allocation

[8] This aggregation structure mirrors classical social welfare formulations in exogenous-population settings. In the special case of unit fertility, this reduces to the exponentially-discounted utilitarian criterion, see Lemma 1.

$\mathfrak{a}_{t+1} \in \boldsymbol{D}$, there exists $k \geq 0$ such that

$$W_{t+1}(\mathfrak{a}_{t+1}) = W_{t+1}((k, k, \dots); (1, 1, \dots)) = k.$$

Now consider the allocation in which generation t enjoys utility u_t, while all future generations receive utility k, and fertility is unitary from $t+1$ onward, that is, $\mathfrak{a}_t = ((u_t, k, k, \dots); (n_t, 1, 1, \dots))$.

The right panel of Figure 1 illustrates how the equally-distributed equivalent utility e_t is determined. The value e_t aggregates the utility u_t of generation t, the continuation value k, and the number of children n_t. Graphically, e_t is the utility level corresponding to the point on the 45–degree line that lies on the planner's indifference curve through (u_t, k).

At the point (e_t, e_t), the slope of the indifference curve reflects the relative weight assigned to future generations, captured by $n_t B$. When fertility n_t increases, greater weight is placed on future generations, changing the indifference curve accordingly.

This construction yields an allocation $\mathfrak{a}'_t = ((e_t, e_t, \dots); (n_t, 1, 1, \dots))$ with constant utilities that is socially indifferent to the original allocation: $W_t(\mathfrak{a}_t) = W_t(\mathfrak{a}'_t)$.

The left panel of Figure 1 represents the final step. Given e_t and n_t, the level of welfare $W_t(\mathfrak{a}_t)$ is defined by the dynastic-equivalent utility of t. It is the level of utility—at unit fertility—that makes generation t indifferent to the allocation $\mathfrak{a}'_t$. By construction, $W_t(\mathfrak{a}_t)$ lies on the vertical line at $n_t = 1$, where the indifference curve through (n_t, e_t) intersects.

2.4 Exponentially discounted utilitarianism

The fertilitarian criterion nests the exponentially discounted utilitarianism as a special case under constant population. This result clarifies how the proposed framework generalizes the workhorse welfare criterion for dynamic models to endogenous fertility.

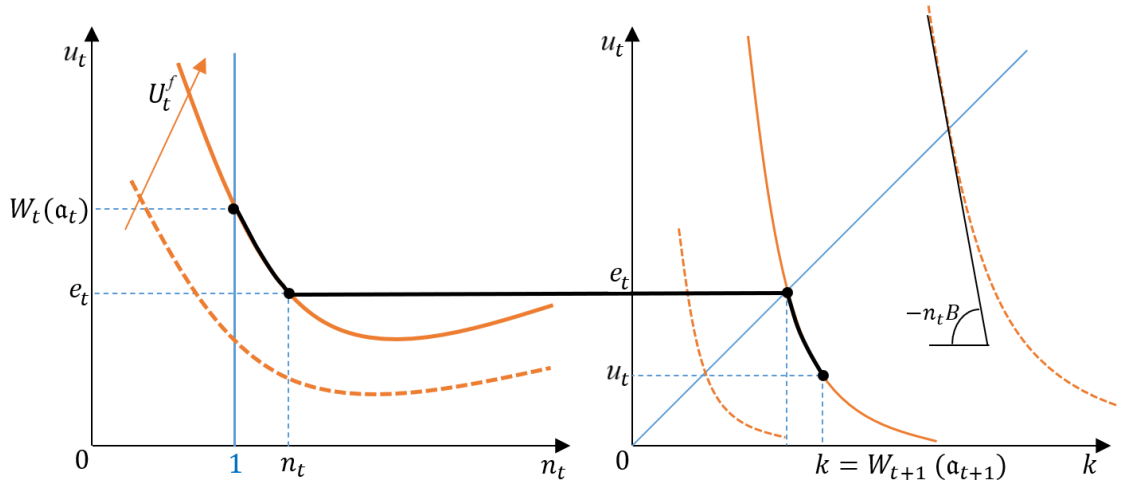


Figure 1: Recursive construction of fertilitarian social welfare. Right: Determination of the equally-distributed equivalent utility e_t. Left: Evaluation of social welfare W_t using the dynastic-equivalent utility U_t^f.

Lemma 1. *Let social preferences $\succsim$ be fertilitarian on $\boldsymbol{D}$, and suppose $u_t = c_t$ for all $t \in T$. Then, for each pair of allocations $\mathfrak{a}, \mathfrak{a}' \in \boldsymbol{D}$ with constant unit fertility,*

$$\mathfrak{a} \succsim \mathfrak{a}' \quad \textit{if and only if} \quad \sum_t \beta^t v(c_t) \geq \sum_t \beta^t v(c'_t),$$

where $\beta \equiv B/(1+B) \in (0,1)$.

Proof. Let $\mathfrak{a} \in \boldsymbol{D}$ and suppose $n_t = 1$ for all t. Since $U_t^f(k,1) = k$ for each k, the recursive structure of the fertilitarian criterion reduces to:

$$W_t(\mathfrak{a}_t) = v^{-1}\left(\frac{v(c_t) + Bv(W_{t+1}(\mathfrak{a}_{t+1}))}{1+B}\right).$$

Define $\bar{W}_t \equiv (1+B)v(W_t)$. Substituting recursively implies:

$$\bar{W}_0 = \sum_t \beta^t v(c_t).$$

Thus, W_0 is ordinally equivalent to $\sum_t \beta^t v(c_t)$. □

This result shows that when population is fixed, fertilitarian social preferences coincide with the classical discounted utilitarian ranking. Differences emerge only

when fertility varies. In this sense, the criterion preserves standard ethical benchmarks in fixed-population environments, while extending them to settings with endogenous population dynamics.

Importantly, the fertilitarian criterion inherits time consistency from exponentially-discounted utilitarianism. In other words, an allocation that is ranked higher today will be preferred also in the future.

2.5 A guide for applications

I next provide a short guide for how to apply the fertilitarian criterion. The following steps are involved:

1. **Specify social preference parameters.** Choose discounting and inequality aversion of society. Discounting is governed by the weight B assigned to future generations. Since the criterion is discounted utilitarian for constant population, B can be recovered from the standard discount factor β, using the fact that $\beta = B/(1+B)$. When assuming homotheticity, inequality aversion is governed by the parameter $\rho \geq 0$, as standard for Atkinson type of aggregators:

$$v(x) = \begin{cases} x^{1-\rho}/(1-\rho) & \text{if } \rho \neq 1, \\ \log x & \text{if } \rho = 1. \end{cases}$$

2. **Construct the dynastic-equivalent utility U_t^f.** The estimated or calibrated function U_t represents the dynastic preferences of generation t over their own utility (or consumption or income), and over the number and utility of their descendants. The dynastic-equivalent utility function $U_t^f(u,n)$ gives the level of constant utility under unit fertility that generation t considers equally desirable. This is defined implicitly by

$$U_t\left((u,u,\dots),(n,1,1,\dots)\right) = U_t\left((k,k,\dots),(1,1,1,\dots)\right),$$

where $k = U_t^f(u, n)$.

3. **Compute social welfare recursively.** Given any allocation $\mathfrak{a} = (\mathfrak{u}, \mathfrak{n})$, compute social welfare recursively:[9]

   - For each generation t, aggregate own utility and social welfare at $t+1$ into the equally-distributed equivalent:

$$e_t = v^{-1}\left(\frac{v(u_t) + n_t B \cdot v(W_{t+1})}{1 + n_t B}\right).$$

   - Then apply the dynastic-equivalent utility:

$$W_t = U_t^f(e_t, n_t).$$

4. **Use for allocation or policy evaluation.** The criterion is then ready to use for the assessment of allocations or policies. One can maximize the fertilitarian social welfare W_0 subject to feasibility constraints. Note that the criterion is *time consistent* even when the dynastic preferences vary over time. Thus, the maximization problem can be solved with standard tools of optimal control/dynamic programming.

2.6 Replacing fertility preferences with ethical choices

The benchmark construction uses U_t^f to align the social value of fertility with parental preferences. If society judges these preferences to be misguided or normatively irrelevant, it can replace them with an explicit *social fertility evaluation*. Concretely, for each t, introduce a function $F_t : \mathbb{R}_+ \times \mathbb{R}_+ \to \mathbb{R}$ to replace the role played by $U_t^f(\cdot, n_t)$ in the fertilitarian criterion. The result is a social welfare function

[9] Because the recursion is increasing in W_{t+1}, existence of a well-defined sequence $\{W_t\}$ follows from the Tarski–Knaster fixed-point theorem.

defined recursively by:

$$W_t \;=\; F_t(e_t, n_t), \qquad e_t \equiv v^{-1}\left(\frac{v(u_t) + n_t B \cdot v(W_{t+1})}{1 + n_t B}\right),$$

where we assume $F_t(u, 1) = u$ for each u to preserve the normalization of social welfare (a unit-fertility constant allocation at level u is assigned social welfare u).

$F(u, n)$ tells how society ranks the *quantity–quality* pairs (u, n). The first is the constant (personal) utility equally-experienced by all generations. The second is the fertility level of generation t, while keeping the population size from period $t + 1$ constant. Thus, it captures the tradeoff between quality of life and quantity of life.

A simple social fertility evaluation function is

$$F(u, n) = u + \phi(n - 1)(u - \hat{u}).$$

Here, $\phi \geq 0$ and $\hat{u} \geq 0$ govern the quality-quantity tradeoff. The parameter ϕ measures the relative value associated with increasing the size of a future population with constant utility $u - \hat{u}$. The parameter $\hat{u}$ is the critical level of utility corresponding to a life worth adding: the above social fertility evaluation function values a larger population only if the utility exceeds this threshold.[10] The marginal rate of substitution between quality of life u and quantity of lives n is given by:

$$MRS_{u,n} \equiv -\frac{\partial F/\partial n}{\partial F/\partial u} = -\frac{\phi(u - \hat{u})}{1 + \phi(n - 1)}.$$

Although the choice of F provides considerable degrees of freedom in the quality-quantity tradeoff, standard criteria do not emerge as special cases (with one exception highlighted below).

The reason is structural. The recursion $W_t = F(e_t, n_t)$ forces welfare at t to assess

[10] The above specification requires a domain restriction to keep $W_t \geq 0$. When the future population is large and the quality of life is very small $F(u, n)$ can be negative. A hinged variant $F(u, n) = u + \phi(n - 1)\max\{u - \hat{u}, 0\}$ keeps $W_t \geq 0$ under the current domain.

the future through W_{t+1} and the fixed averaging rule e_t. Standard benchmarks don't fit this restriction. Average utilitarianism rescales the entire sum by total population and, thus, does not admit a recursive representation.

Total utilitarianism (with respect to personal utilities) admits a recursive formulation.[11] However, it is consistent with our recursive formulation only if: i) e_t is the simple average (v linear); ii) the discount factor is $B = \beta/(1-\beta)$; and iii) the social fertility evaluation function F assumes $\phi = \beta$ and sets the critical level $\hat{u}$ appropriately.

Interestingly, total utilitarianism is equivalent to the fertilitarian social welfare function when individuals' utilities are

$$U_t = \frac{1}{N_t} \sum_{\tau \geq 0} \beta^\tau N_{t+\tau} \left(u_{t+\tau} - \hat{u}\right),$$

v is linear, and the social discount factor is $B = \beta/(1-\beta)$.

3 Axiomatic characterization

3.1 Respecting fertility preferences

The core normative innovation is to treat parents' fertility preferences as ethically relevant. The guiding idea is that a parent's willingness to have children reveals ethically significant judgments about the trade-off between the quantity and quality of lives in the absence of inequalities. Rather than imposing an external criterion, this approach defers—under specific conditions—to the revealed preferences of parents.

This idea aligns with the seminal economic formulation of endogenous fertility in Barro and Becker (1989), but it also addresses conceptual challenges. Parents may discount the future more steeply than society deems appropriate, exhibit lower aversion to intergenerational inequality, or hold views about future fertility decisions

[11] If standard criteria are expressed in terms of dynastic utilities U_t, these generally cannot be expressed recursively.

that their descendants would not endorse. To avoid such normative tensions, deference is limited to situations in which all generations experience the same level of utility.

Axiom (Respect of fertility preferences). *For each pair of allocations* $\mathfrak{a}, \mathfrak{a}' \in \boldsymbol{A}$ *with* $\mathfrak{a} = ((u, u, \dots), (n_0, 1, 1, \dots))$*, and* $\mathfrak{a}' = ((u', u', \dots); (n'_0, 1, 1, \dots))$*, if* $U_0(\mathfrak{a}) \geq U_0(\mathfrak{a}')$*, then* $\mathfrak{a} \succsim \mathfrak{a}'$*.*

This axiom is closely related to the "A-dominance criterion" by Golosov, Jones, and Tertilt (2007), which ranks allocations by exclusively comparing the utility of individuals who exist in both, such as earlier generations. Among the differences, *respect of fertility preferences* embraces the asymmetry between generations and focuses exclusively on parents, while disregarding existing future generations. Moreover, it imposes that existing and potential generations achieve the same utility, as assumed for "representative-agent efficiency" (Michel & Wigniolle, 2007; Conde-Ruiz, Giménez, & Pérez-Nievas, 2010).

A different approach, the "P-dominance criterion," assigns a reference utility to all potential individuals and includes them in the welfare comparison (Golosov et al., 2007).[12] As shown by Pérez-Nievas, Conde-Ruiz, and Giménez (2019) and Pérez-Nievas (2023), A-dominance, P-dominance, and representative-agent efficiency raise analytical and normative issues, particularly in dynamic models with endogenous fertility.

3.2 Toward a Utilitarian Benchmark with Fixed Population

I next introduce axioms that restrict social preferences in the absence of population change. Together with more technical requirements introduced later, these characterize the benchmark case of additive inequality-averse aggregation of utilities.

[12] A generalization of efficiency concepts for population changes is recently proposed by Collignon (2025), which allows for a rich structure of social rules when comparing situations with different populations and studies their consistency with completeness and transitivity.

The first axiom captures the standard Pareto principle: utility improvements are socially desirable if they accrue to generations that are alive.

Axiom (Pareto). *For any* $\mathfrak{a}, \mathfrak{a}' \in \boldsymbol{A}$ *with* $\mathfrak{u} \geq \mathfrak{u}'$ *and* $\mathfrak{n} = \mathfrak{n}'$*:*

(i) If $u_0 > u'_0$*, then* $\mathfrak{a} \succ \mathfrak{a}'$*;*

(ii) If $u_t > u'_t$ *for some* $t > 0$*, and* $(n_0, \ldots, n_{t-1}) \gg 0$*, then* $\mathfrak{a} \succ \mathfrak{a}'$*;*

(iii) Otherwise, $\mathfrak{a} \sim \mathfrak{a}'$*.*

Next, to assess social attitudes toward intergenerational inequality, I introduce a variant of the Pigou–Dalton transfer principle. Rather than comparing arbitrary profiles, I restrict attention to allocations where the first generation has distinct utility, and all later generations enjoy a constant utility. Then, I discipline redistributive transfers between the present and the future.

Axiom (Inequality Aversion). *There exists* $\zeta > 0$ *such that for each pair of allocations* $\mathfrak{a}, \mathfrak{a}' \in \boldsymbol{A}$ *with* $\mathfrak{u} = (u_0, u, u, ...)$*,* $\mathfrak{u}' = (u'_0, u', u', ...)$*,* $\mathfrak{n} = \mathfrak{n}' = (n_0, 1, 1, ...)$ *with* $n_0 > 0$*, and each* $\epsilon > 0$*,*

- *if* $u'_0 - \epsilon = u_0 > u = u' + n_0\zeta\epsilon$*, then* $\mathfrak{a} \succsim \mathfrak{a}'$*;*
- *if* $u'_0 + \epsilon = u_0 < u = u' - n_0\zeta\epsilon$*, then* $\mathfrak{a} \precsim \mathfrak{a}'$*.*

This formulation is similar to Piacquadio (2020) and allows a different weight between the donors and recipients of the Pigou-Dalton transfer. This flexibility is needed to account for the different population sizes of the present generation and all future generations, and for the social attitudes toward discounting.

Finally, I impose a consistency requirement inspired by Thomsen (1927) that ensures additivity in the evaluation of the utility streams.[13]

[13] A more recent application is Karni and Safra (2002). In most intergenerational axiomatic works, it is standard to obtain additive separability over time with stronger axioms, such as "initial trade-off independence," demanding that the trade-off between the consumption of the first two generations is independent of the later consumption stream (see Koopmans, 1960, and Bleichrodt, Rohde, & Wakker, 2008), or by directly imposing "time separability" (see Piacquadio, 2020).

Axiom (Thomsen Condition). *Let* $\mathfrak{a}, \mathfrak{a}', \mathfrak{a}'', \mathfrak{a}''' \in \boldsymbol{A}$ *with*

$$\mathfrak{u} = (u_0, u, u, \dots), \quad \mathfrak{u}' = (u'_0, u', u', \dots),$$
$$\mathfrak{u}'' = (u''_0, u, u, \dots), \quad \mathfrak{u}''' = (u'_0, u'', u'', \dots),$$
$$\mathfrak{n} = \mathfrak{n}' = \mathfrak{n}'' = \mathfrak{n}''' = (n_0, 1, 1, \dots), \quad n_0 > 0.$$

If $\mathfrak{a} \succsim \mathfrak{a}'$ *and* $\mathfrak{a}' \succsim \mathfrak{a}''$*, then* $\mathfrak{a} \succsim \mathfrak{a}'''$*.*

3.3 Infinite-Horizon Structure and Domain Restrictions

This section adapts key axioms from Bleichrodt et al. (2008)'s characterization of discounted utilitarianism to settings with endogenous population. These axioms ensure that social preferences over infinite-horizon allocations admit a tractable and numerically representable structure. They restrict the domain and structure of preferences without committing to specific population assumptions.

Let an allocation $\mathfrak{a} \in \boldsymbol{A}$ be *ultimately constant* if there exists $t \in T$ such that $\mathfrak{a}_t$ is constant. Denote the set of constant and ultimately constant allocations by $\boldsymbol{A}^c$ and $\boldsymbol{A}^{uc}$ respectively.

Axiom (Ultimate Continuity). *For each* $\mathfrak{a}_t \in \boldsymbol{A}^{uc}$*, the upper and lower contour sets*

$$\{\mathfrak{a}'_t \in \boldsymbol{A}^{uc} \mid \mathfrak{a}'_t \succsim \mathfrak{a}_t\} \quad \textit{and} \quad \{\mathfrak{a}'_t \in \boldsymbol{A}^{uc} \mid \mathfrak{a}_t \succsim \mathfrak{a}'_t\}$$

are closed.

This continuity axiom ensures that social preferences over ultimately constant allocations are well-behaved. Then, there exists a numerical representation of such preferences.

Axiom (Stationarity). *For all* $\mathfrak{a}, \mathfrak{a}' \in \boldsymbol{A}$ *such that* $u_0 = u'_0$ *and* $n_0 = n'_0$*,*

$$\mathfrak{a} \succsim \mathfrak{a}' \quad \textit{if and only if} \quad \mathfrak{a}_1 \succsim \mathfrak{a}'_1.$$

Stationarity ensures that the social ranking of allocations is invariant to time shifts: removing a common initial period leaves the relative evaluation unchanged.

Axiom (Constant Equivalence). *For every* $\mathfrak{a} \in \boldsymbol{D}$*, there exists a constant allocation* $\mathfrak{a}' \in \boldsymbol{A}^c$ *such that* $\mathfrak{a} \sim \mathfrak{a}'$*.*

This domain restriction ensures that every admissible allocation is comparable to a constant one. It eliminates divergent paths that would otherwise prevent the representation of social preferences.

Axiom (Tail Robustness). *Let* $\mathfrak{a}, \mathfrak{a}' \in \boldsymbol{A}^c$ *with* $\mathfrak{n} = (1, 1, \dots)$*. If* $\mathfrak{a} \succ \mathfrak{a}'$ *(resp.* $\mathfrak{a}' \succ \mathfrak{a}$*), then there exists* $t \in T$ *such that for all* $t' \geq t$*,*

$$(u_0, \dots, u_{t'-1}, \mathfrak{u}'_{t'}; 1, 1, \dots) \succ \mathfrak{a}' \quad \textit{(resp. } \prec \textit{)}.$$

This axiom limits the sensitivity of social preferences to remote changes in the allocation. It formalizes the intuition that the social ranking of alternatives should not be overturned by small changes in the distant future. It disciplines allocations that are not ultimately constant.

Finally, I state the requirement that the social ranking is invariant to the rescaling of utilities. This assumption simplifies the functional form and justifies the use of homothetic aggregators, as typically done in economic applications. For each $\gamma > 0$ and each $\mathfrak{u}$, let the γ-scaled utility stream be $\gamma\mathfrak{u} \equiv (\gamma u_0, \gamma u_1, ...)$.

Axiom (Scale Invariance). *For each* $\gamma > 0$ *and each pair* $\mathfrak{a} = (\mathfrak{u}; \mathfrak{n}), \mathfrak{a}' = (\mathfrak{u}'; \mathfrak{n}') \in \boldsymbol{A}^{uc}$*, if* $\mathfrak{a} \sim \mathfrak{a}'$ *and* $\mathfrak{n} = \mathfrak{n}' = (1, 1, \dots)$ *then* $(\gamma\mathfrak{u}; \mathfrak{n}) \sim (\gamma\mathfrak{u}'; \mathfrak{n}')$*.*

3.4 Characterization Result

The main result establishes that the fertilitarian criterion is uniquely characterized by the axioms introduced above.

Theorem 1. *Let $\boldsymbol{A}^{uc} \subseteq \boldsymbol{D}$. Then, social preferences $\succsim$ on $\boldsymbol{D}$ are fertilitarian if and only if these satisfy: respect of fertility preferences, Pareto, inequality aversion, Thomsen condition, ultimate continuity, stationarity, constant equivalence, and tail robustness.*

Pareto, ultimate continuity, constant equivalence, and tail robustness ensure that social preferences can be represented by a well-behaved social welfare function. Stationarity implies a recursive structure, while the Thomsen condition delivers an additive representation over generations in the constant-population case. Inequality aversion introduces ethical concern for intergenerational disparities, and respect of fertility preferences pins down the trade-off between population size and individual utility using parents' preferences. The proof is in the appendix.

Two corollaries clarify how the representation simplifies with scale invariance and shed light on the role of the domain assumption.

Corollary 1 (Scale-Invariant Representation)**.** *If social preferences $\succsim$ also satisfy scale invariance, then $\succsim$ admits a CES-fertilitarian representation.*

Corollary 2 (Restricted Domain)**.** *If the domain is restricted to ultimately constant allocations, that is, $\boldsymbol{D} = \boldsymbol{A}^{uc}$, then constant equivalence and tail robustness are no longer required for the characterization.*

4 Fertilitarianism vs. Utilitarian Benchmarks

This section illustrates how fertilitarianism diverges from standard welfare criteria.

4.1 Finite-horizon examples

This subsection deals with finite horizon examples. I first consider allocations involving three generations—grandparents, parents, and children—which differ only in the fertility choices of the parent generation. Comparisons emphasize how these fertility choices affect the trade-off between the number and quality of lives.

Each generation $t \in \{0, 1, 2\}$ has a utility u_t and fertility n_t. Dynastic preferences are represented by:

$$U_t = u_t + \delta n_t u_{t+1},$$

where $\delta = 0.2$ is a fixed weight on child utility. Population size at time t is denoted N_t; its dynamics is given by $N_{t+1} = n_t N_t$, with $N_0 = 1$ normalized.

Table 1 presents three allocations. The boldface entries indicate the key variables that differ across the cases.

Table 1: Three Intergenerational Allocations

Allocation	**Generation**	u_t	n_t	N_t	U_t
$\mathfrak{a}$	Grandparents	30	2	1	54
	Parents	**60**	**2**	2	**84**
	Children	**60**	0	4	60
	Total ut.			7	**462**
	Average ut.				**66**
$\mathfrak{a}'$	Grandparents	30	2	1	50
	Parents	**50**	**3**	2	**80**
	Children	**50**	0	6	50
	Total ut.			9	**510**
	Average ut.				**56.67**
$\mathfrak{a}''$	Grandparents	30	2	1	56
	Parents	**65**	**1**	2	**78**
	Children	**65**	0	2	65
	Total ut.			5	**342**
	Average ut.				**68.4**

Total utilitarianism ranks $\mathfrak{a}'$ as best, since it maximizes the sum of the well-beings: while individuals are worse off, the quantity effect outweighs the reduction in quality of life. Average utilitarianism ranks $\mathfrak{a}''$ highest, favoring instead quality over quantity. Both approaches override the preferences of the parent generation: in both $\mathfrak{a}'$ and $\mathfrak{a}''$, parents are worse off compared to $\mathfrak{a}$. In contrast to these criteria, fertilitarianism respects parental fertility preferences and ranks $\mathfrak{a}$ highest.

Fertilitarianism thus avoids both the repugnant conclusion of total utilitarianism

and the anti-natalist bias of average utilitarianism. This divergence stems from the axiom of *respect of fertility preferences*, which defers to individual fertility preferences, under the equality proviso.

I now contrast this with standard efficiency axioms and highlight the importance of the equality proviso. Consider A-dominance: it supports $\mathfrak{a}$ over $\mathfrak{a}'$ since the parents—those existing in both allocations—prefer $\mathfrak{a}$. P-dominance supports $\mathfrak{a}$ over $\mathfrak{a}'$ only if the well-being of non-existence is larger than 50. Neither criterion conclusively ranks $\mathfrak{a}$ and $\mathfrak{a}''$, where some generations gain and other lose. In contrast, *respect of fertility preferences* assigns full priority to the preferences of parents, capturing the idea that parents have a say on how many children they should have, provided they treat them well enough (i.e., the equality proviso).

A fourth allocation $\bar{\mathfrak{a}}$ highlights the importance of the equality proviso. It differs from $\mathfrak{a}$ by letting children have children of their own, who experience very low utility:

Table 2: Adding low-utility grandchildren.

Allocation	**Generation**	u_t	n_t	N_t	U_t
$\bar{\mathfrak{a}}$	Grandparents	30	2	1	54
	Parents	60	2	2	84
	Children	60	**2**	4	64
	Grandchildren	**10**	0	8	10
	Total ut.				**558**
	Average ut.				**37.2**
$\bar{\bar{\mathfrak{a}}}$	Grandparents	30	2	1	54
	Parents	60	2	2	84
	Children	$\mathbf{e_{t+1}}$	2	4	$1.4e_{t+1}$
	Grandchildren	$\mathbf{e_{t+1}}$	0	8	e_{t+1}
	Total ut.				$\mathbf{222 + 13.6e_{t+1}}$
	Average ut.				$\mathbf{14.8 + .9e_{t+1}}$

Total utilitarianism prefers $\bar{\mathfrak{a}}$ over $\mathfrak{a}$ because the additional lives increase total welfare, despite their low utility. Both A-dominance and P-dominance support $\bar{\mathfrak{a}}$ if the well-being of the grandchildren is above a life worth-living. In contrast, fertilitarianism remains silent: since utility equality across generations fails, the axiom of *respect*

of fertility preferences does not apply, and the judgment depends on discounting and inequality aversion.

In fact, the fertilitarian criterion first transforms $\bar{\mathfrak{a}}$ into an equally-desirable allocation $\bar{\bar{\mathfrak{a}}}$ where children and grandchildren are both assigned the same utility e_{t+1}. For memory, e_{t+1} is the equally-distributed equivalent utility and depends on the social attitudes to discounting and inequality.

Fertilitarianism then asks the generations of children to assess whether the increase in fertility from 0 to 2 is worth the change in utility from 60 to e_{t+1}. This is equivalent to comparing $U_{t+1}(\bar{\bar{a}}) = 1.4e_{t+1}$ with $U_{t+1}(a) = 84$. For this example, any level of discounting and inequality aversion leads to an equally-distributed equivalent utility e_{t+1} that is too low for children to prefer $\bar{\bar{\mathfrak{a}}}$ to $\mathfrak{a}$. Thus, fertilitarianism assesses $\mathfrak{a}$ as socially more desirable than $\bar{\mathfrak{a}}$ (which by construction is equally desirable as $\bar{\bar{\mathfrak{a}}}$).

4.2 A stylized Barro–Becker Model

This subsection applies the fertilitarian criterion to a dynamic macroeconomic environment and compares its policy implications with standard population-sensitive criteria. The model is inspired by the seminal contribution of Barro and Becker (1989).[14]

Each generation t has population size N_t, consumption c_t, and fertility choice n_t, with population evolving as $N_{t+1} = n_t N_t$. Dynastic preferences are represented by:

$$U_t = \frac{1}{N_t} \sum_{\tau=0}^{\infty} \delta^\tau \frac{\left(N_{t+\tau}^{\epsilon} c_{t+\tau}\right)^{1-\theta}}{1-\theta},$$

where $\theta \geq 0$ captures risk aversion, $\epsilon \in (0,1)$ measures population sensitivity, and $\delta \in (0,1)$ is the discount factor.[15]

[14] See also Doepke (2005), Golosov et al. (2007), L. E. Jones and Schoonbroodt (2010), De Silva and Tenreyro (2020), C. I. Jones (2022).

[15] For simplicity, I assume that $u_t = c_t$, in line with the literature. The account of decreasing marginal utilities only matters for the interpretation of the inequality attitudes of the fertilitarian

I study deviations from a steady-state path where $c_t = k$, $n_t = 1$, and $N_t = 1$ for all $t \neq \bar{t}$. Generation $\bar{t}$ chooses fertility $n_{\bar{t}}$ and consumption $c_{\bar{t}}$, subject to the time budget constraint:

$$c_{\bar{t}} \leq (1 - bn_{\bar{t}})w,$$

where w denotes potential earnings and $b \in (0, 1)$ is the time cost per child.

The comparison involves four criteria:

- **Self-interest of generation** $\bar{t}$, that is, $U_{\bar{t}}$.

- **Critical-Level Total Discounted Utilitarianism (CLTDU)**:

$$W^{\text{CLTDU}} = \sum_{\tau=0}^{\infty} \beta^{\tau} N_{\tau}(U_{\tau} - \bar{U}),$$

  where β is the discount factor and $\bar{U}$ is the critical level.[16]

- **Critical-Level Average Discounted Utilitarianism (CLADU)**:

$$W^{\text{CLADU}} = \frac{\sum_{\tau=0}^{\infty} \beta^{\tau} N_{\tau}(U_{\tau} - \bar{U})}{\sum_{\tau=0}^{\infty} \beta^{\tau} N_{\tau}},$$

  where β is the discount factor and $\bar{U}$ is the critical level.

- **CES-Fertilitarianism**, as characterized in Section 3.

The exercise is numerical. The dynastic preference parameters are such that each generation, including $\bar{t}$, would optimally select the steady state. Social preference parameters are set on the basis of the authors' ethical intuition. The following table summarizes these parameters.

Figure 2 shows indifference curves for each criterion in the consumption-fertility space $(n_{\bar{t}}, c_{\bar{t}})$. The steady-state allocation $(1, 90)$ is optimal for generation $\bar{t}$ by

social welfare function: with the adopted formulation, these are to be interpreted as attitudes towards inequality in consumption, rather than inequality in utilties.

[16]Note that the critical-level criteria characterized by Blackorby et al. (1995) do not allow for discounting the well-being of future generations.

Table 3: Parameterization for Numerical Simulation

Variable	=	Value	Justification
N_0	=	1	Normalization
w	=	100	Normalization
b	=	0.1	10% time cost per child
k	=	90	Steady-state consumption
δ	=	0.55	2% annual discount rate (35 year period)
θ	=	1	Log preferences (see C. I. Jones (2022))
ϵ	=	0.091	Implied by steady-state optimality
β	=	0.55	Equal social and individual discounting
$\bar{U}$	=	*well-being at* $c = k/3$	Life barely worth living
ρ	=	3	CES inequality aversion

construction. Since the average coincides with generation $\bar{t}$'s preferences, the steady state is also optimal for CLADU.

The steady state is also top-ranked by fertilitarianism. Since the steady-state involves no intergenerational inequality, the axiom of *respect of fertility preferences* applies and forces social welfare to agree with the preferences of generation $\bar{t}$. However, fertilitarianism has flatter indifference curves, due to role of social inequality attitudes.

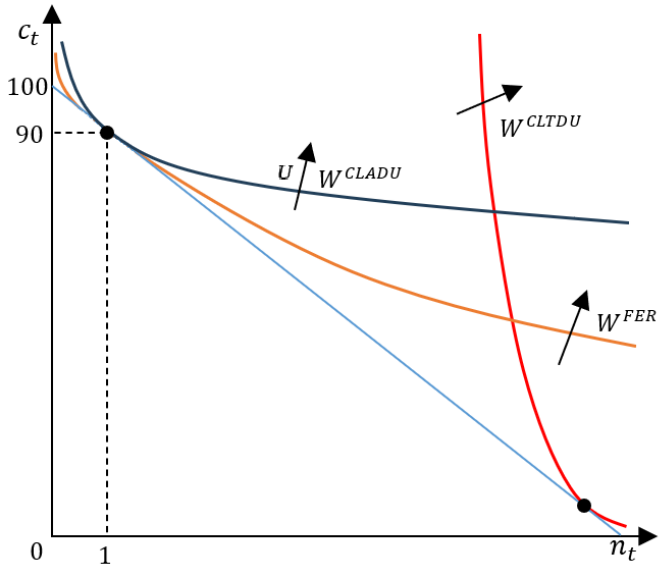


Figure 2: Indifference curves for different criteria in $(n_{\bar{t}}, c_{\bar{t}})$ space.

In contrast, CLTDU recommends very high fertility ($n_{\bar{t}}^{\text{CLTDU}} \approx 9.26$) and low

consumption ($c_{\bar{t}} \approx 7.4$), maximizing population size, despite the large well-being loss of each individual. This is yet another instance of the repugnant conclusion.

Importantly, both CLTDU and CLADU fail a basic consistency requirement. Their recommendation for $\bar{t}$ depends on the path of utility and fertility before $\bar{t}$. The "Egyptology objection" (see Parfit, 1984) is well-known for average utilitarianism: if previous generations had a lower well-being, a larger fertility for $\bar{t}$ would increase average well-being. More surprisingly, CLTDU also fails this consistency due to intergenerational altruism: earlier generations' concern for future generations causes the ranking of $\bar{t}$'s fertility to depend on past allocations.[17]

Fertilitarianism avoids this critique and remains time consistent. As in discounted utilitarianism, the recursive formulation guarantees time consistency for fixed-population allocations. More importantly, potential conflicts between the dynastic preferences of grandparents and those of parents are eliminated by the equality proviso respecting fertility preferences: at time t, only generation t's preferences determine the quality–quantity trade-off.

5 Conclusion

This paper re-examines population ethics by placing parental fertility preferences at the core of intergenerational evaluation. The central normative claim is that, under conditions of equality, society may defer to parents' willingness to have children. This principle—and the separation between aggregation over existing generations and the ethical choice of when to add people—distinguish the present proposal from standard alternatives such as total and average utilitarianism.

The view developed here is related to the neutrality principle in population ethics (Narveson, 1973)—the idea that a person's existence is in itself neutral, provided

[17]If CLTDU and CLADU are evaluated on personal utilities u_τ (rather than overall dynastic utilities U_τ), the qualitative contrasts are unchanged. CLTDU still exhibits a strong incentive for very high fertility. While CLADU still displays the "Egyptology" dependence on past allocations, bracketing out altruism avoids this drawback for CLTDU.

it does not affect others. However, neutrality as commonly formalized (e.g., via A-dominance, Golosov et al., 2007) grants full priority to existing individuals and fails to justify any concern for the welfare of added children. The challenge, then, is to respect fertility preferences while still treating generations equitably.

The key innovation is to leverage the dynastic structure of overlapping-generations models, standard in macroeconomics. This allows the criterion to defer to parental preferences, while keeping aggregation and population choice on separate tracks. The resulting criterion—*fertilitarianism*—has a tractable, recursive structure: it aggregates well-being over existing generations with transparent discounting and inequality attitudes, and it governs the quality-quantity decision via ethically relevant fertility preferences. In the constant-population case it coincides with discounted utilitarianism; in variable-population settings it avoids both the repugnant conclusion and the anti-natalist bias, while remaining time-consistent. Illustrative examples and simulations highlight its normative advantages over standard criteria.

Fertilitarianism opens new avenues for applied normative analysis where population dynamics are central—such as growth, education, and climate change. By combining ethical clarity (disentanglement and deference) with policy tractability (simple recursion and time consistency), it offers a viable alternative to the contested criteria that dominate population ethics today.

References

Arrhenius, G. (2000). An impossibility theorem for welfarist axiologies. *Economics & Philosophy*, *16*(2), 247–266.

Asheim, G. B., & Zuber, S. (2014). Escaping the repugnant conclusion: Rank-discounted utilitarianism with variable population. *Theoretical Economics*, *9*(3), 629–650.

Barro, R. J., & Becker, G. S. (1989). Fertility choice in a model of economic growth. *Econometrica: Journal of the Econometric Society*, *57*(2), 481–501.

Blackorby, C., Bossert, W., & Donaldson, D. (1995). Intertemporal population ethics: critical-level utilitarian principles. *Econometrica: Journal of the Econometric Society*, 1303–1320.

Blackorby, C., Bossert, W., & Donaldson, D. J. (2005). *Population issues in social choice theory, welfare economics, and ethics* (No. 39). Cambridge University Press.

Bleichrodt, H., Rohde, K. I., & Wakker, P. P. (2008). Combining additive representations on subsets into an overall representation. *Journal of Mathematical Psychology*, *52*(5), 304–310.

Broome, J. (2004). *Weighing lives*. Oxford University Press.

Collignon, P.-E. (2025). Generalizing the pareto principle with endogenous populations. *SSRN working paper*.

Conde-Ruiz, J. I., Giménez, E. L., & Pérez-Nievas, M. (2010). Millian efficiency with endogenous fertility. *The Review of Economic Studies*, *77*(1), 154–187.

Dasgupta, P. (2005). Regarding optimum population. *Journal of Political Philosophy*, *13*(4), 414–442.

Dasgupta, P. (2019). *Time and the generations: population ethics for a diminishing planet*. Columbia University Press.

Debreu, G. (1960). Topological methods in cardinal utility theory. In K. Arrow, S. Karlin, & P. Suppes (Eds.), *Mathematical methods in the social sciences*. Stanford University Press.

De Silva, T., & Tenreyro, S. (2020). The fall in global fertility: A quantitative model. *American Economic Journal: Macroeconomics*, *12*(3), 77–109.

Doepke, M. (2005). Child mortality and fertility decline: Does the Barro-Becker model fit the facts? *Journal of population Economics*, *18*(2), 337–366.

Dworkin, R. (1977). *Taking rights seriously*. Harvard University Press, Cambridge,

MA.

Eden, M., & Kuruc, K. (2023). The long-run relationship between per capita incomes and population size. *CEPR working paper DP18353*.

Golosov, M., Jones, L. E., & Tertilt, M. (2007). Efficiency with endogenous population growth. *Econometrica*, *75*(4), 1039–1071.

Greaves, H. (2017). Population axiology. *Philosophy Compass*, *12*(11), e12442.

Jones, C. I. (2022). The end of economic growth? Unintended consequences of a declining population. *American Economic Review*, *112*(11), 3489–3527.

Jones, L. E., & Schoonbroodt, A. (2010). Complements versus substitutes and trends in fertility choice in dynastic models. *International Economic Review*, *51*(3), 671–699.

Karni, E., & Safra, Z. (2002). Individual sense of justice: a utility representation. *Econometrica*, *70*(1), 263–284.

Koopmans, T. C. (1960). Stationary ordinal utility and impatience. *Econometrica: Journal of the Econometric Society*, 287–309.

McMahan, J. (1981). Problems of population theory. *Ethics*, *92*(1), 96–127.

Michel, P., & Wigniolle, B. (2007). On efficient child making. *Economic Theory*, *31*, 307–326.

Narveson, J. (1973). Moral problems of population. *The Monist*, 62–86.

Parfit, D. (1984). *Reasons and persons*. Oxford University Press, Oxford, UK.

Pérez-Nievas, M. (2023). Efficiency with endogenous population growth. do children have too many rights? *Theoretical Economics*, *18*(4), 1375–1406.

Pérez-Nievas, M., Conde-Ruiz, J. I., & Giménez, E. L. (2019). Efficiency and endogenous fertility. *Theoretical Economics*, *14*(2), 475–512.

Piacquadio, P. G. (2017). A fairness justification of utilitarianism. *Econometrica*, *85*(4), 1261–1276.

Piacquadio, P. G. (2020). The ethics of intergenerational risk. *Journal of Economic Theory*, *186*, 104999.

Piacquadio, P. G. (2025). The tyranny of parenthood, dynastic worlds, and fertilitarianism. *mimeo*.

Pivato, M. (2020). Rank-additive population ethics. *Economic Theory*, *69*(4), 861–918.

Thomsen, G. (1927). Un teorema topologico sulle schiere di curve e una caratterizzione geometrica dellesuperficie isometro-asinttotiche. *Bolletino dell'Unione Matematica Italiana*, *6*, 80-85.

A Proofs

A.1 Theorem 1: The fertilitarian criterion satisfies the axioms

Pareto. Let fertility be fixed. By assumption on preferences, the utility-equivalent function $U_0^f(e_0, n_0)$ is strictly increasing in the first argument. Moreover, the equally-distributed equivalent e_t is strictly increasing in u_t and, when $n_t > 0$, strictly increasing in W_{t+1}. Thus, W_0 is strictly increasing in u_t if and only if the population at t is positive.

Respect of fertility preferences. By construction, the social welfare at any allocation $\mathfrak{a} = ((u, u, \dots); (n_t, 1, 1, \dots))$ is given by $W_0 = U_0^f(u, n_t)$. Since U_0^f is a representation of the dynastic preferences of 0 on a subdomain of allocations, the axioms holds.

Inequality aversion. Consider a pair of allocations

$$\mathfrak{a} \equiv ((u_0, x, x, \dots); (n_0, 1, 1, \dots)), \mathfrak{a}' \equiv ((u_0', x', x', \dots); (n_0, 1, 1, \dots)) \in \boldsymbol{A}^{uc}$$

with $u_0, u_0', x, x' \geq 0$ and $n_0 > 0$ and some $\epsilon > 0$, for which $u_0' - \epsilon = u_0 > x = x' + \zeta\epsilon$ with $\zeta = (Bn_0)^{-1}$. By construction, $W_1({}_1\mathfrak{a}) = x$ and $W_1({}_1\mathfrak{a}') = x'$. The equally-

distributed equivalents are

$$e_0(\mathfrak{a}) = v^{-1}\left[\frac{v(u_0) + Bn_0 v(x)}{1 + Bn_0}\right]$$

and

$$e_0(\mathfrak{a}') = v^{-1}\left[\frac{v(u_0') + Bn_0 v(x')}{1 + Bn_0}\right].$$

To prove that $e_0(\mathfrak{a}) \geq e_0(\mathfrak{a}')$ and establish that the axiom holds, take the first-order Taylor expansion of $v(u_0 + \epsilon) + Bn_0 v\left(x - \epsilon(Bn_0)^{-1}\right)$ and, rearranging,

$$[v(u_0') - Bn_0 v(x')] = [v(u_0) + Bn_0 v(x)] + \epsilon[v'(u_0) - v'(x)] + ...$$

Since $u_0 > x$ and v is concave, $[v'(u_0) - v'(x)] \leq 0$ and, thus, $e_0(\mathfrak{a}) \geq e_0(\mathfrak{a}')$. Since the utility-equivalent function U_0^f is strictly increasing in the first argument,

$$W_0(\mathfrak{a}) = U_0^f(e_0(\mathfrak{a}), n_0) \geq U_0^f(e_0(\mathfrak{a}'), n_0) = W_0(\mathfrak{a}')$$

and, thus, $\mathfrak{a} \succsim \mathfrak{a}'$. The case where $u_0' + \epsilon = u_0 < x = x' - \zeta\epsilon$ is symmetric and is omitted.

Thomsen condition. By construction,

$$((u_0, x', x', \dots); (n_0, 1, 1, \dots)) \succsim ((u_0', x, x, \dots); (n_0, 1, 1, \dots))$$

if and only if

$$v(u_0) + n_0 Bv(x') \geq v(u_0') + n_0 Bv(x).$$

Similarly, $((u_0', x'', x'', \dots); (n_0, 1, 1, \dots)) \succsim (u_0'', x', x', \dots); (n_0, 1, 1, \dots))$ if and only if

$$v(u_0') + n_0 Bv(x'') \geq v(u_0'') + n_0 Bv(x').$$

Combining these inequalities gives

$$[v(u_0') + n_0 Bv(x'')] - [v(u_0') + n_0 Bv(x)] \geq [v(u_0'') + n_0 Bv(x')] - [v(u_0) + n_0 Bv(x')],$$

and, simplifying,

$$v(u_0) + n_0 Bv(x'') \geq v(u_0'') + n_0 Bv(x).$$

Since this is equivalent to

$$((u_0, x'', x'', \dots); (n_0, 1, 1, \dots)) \succsim ((u_0'', x, x, \dots); (n_0, 1, 1, \dots)),$$

the axiom holds.

Ultimate continuity. Ultimately constant allocations reduce to a finite-dimensional object. Moreover, both U_0^f and e_0 are continuous in their arguments. Thus, social welfare is continuous over A^{uc}.

Stationarity. This is a direct implication of the time-independent recursive structure of the fertilitarian criterion.

Constant equivalence. For any allocation $\mathfrak{a} \in \boldsymbol{D}$, the recursive structure defines a unique equally desirable allocation $((u, u, \dots), (1, 1, \dots))$ with $u = W_0(\mathfrak{a})$. Thus, the axiom holds.

Tail robustness. This follows from the fact that discounted utility converges under constant population, ensuring that welfare is unaffected by tail changes.

A.2 Theorem 1: The axioms single out the fertilitarian criterion

The proof is split in two. First, it shows how to represent social preferences over ultimately constant allocations by constructing an equally desirable constant allocation. The argument involves recursive steps, using the axioms respect of fertility preferences, Pareto, inequality aversion, Thomsen condition, ultimate continuity,

and stationarity. Second, the proof extends the representation result to the entire domain by imposing constant equivalence and tail robustness.

Part 1. The comparison of ultimately constant allocations

I begin by considering the subset of ultimately constant allocations that are constant from $t = 1$ onward, and show that each such allocation is equivalent to a constant allocation with unit fertility. I then extend the argument recursively to any $t \geq 0$. This yields a unique representation of social welfare for all ultimately constant allocations.

Let $\mathfrak{a} \in A^{uc}$ be of the form:

$$\mathfrak{a} = ((u_0, u, u, \dots), (n_0, 1, 1, \dots)) .$$

I shall show that $\mathfrak{a} \sim \big((u^f, u^f, \dots), (1, 1, \dots)\big)$ for a uniquely defined utility level u^f, which I denote by $W_0(\mathfrak{a})$.

By *ultimate continuity*, there exists a functional representation $F : \mathbb{R}^2_+ \to \mathbb{R}$ of the social preferences $\succsim$ restricted to pairs of allocations that are constant from $t = 1$ onward, that is, $((x, x', x', \dots), (n_0, 1, 1, \dots))$ for some $x, x' \geq 0$.

By Debreu (1960)'s theorem, any continuous preference satisfying Thomsen's condition admits an additive representation. That is,

$$((x, x', x', \dots), (n_0, 1, 1, \dots)) \succsim ((\bar{x}, \bar{x}', \bar{x}', \dots), (n_0, 1, 1, \dots))$$

if and only if

$$f_{n_0}(x) + g_{n_0}(x') \geq f_{n_0}(\bar{x}) + g_{n_0}(\bar{x}') ,$$

for some real-valued functions f_{n_0} and g_{n_0}.

By *Pareto*, f_{n_0} is strictly increasing and, whenever $n_0 > 0$, so is g_{n_0}. By *inequality*

aversion, $g_{n_0} = (1 + Bn_0) f_{n_0}$ for some $B > 0$ and f_{n_0} is concave.[18] Moreover, f_{n_0} is independent of n_0; let v be such a function. Then,

$$((x, x', x', \dots), (n_0, 1, 1, \dots)) \succsim ((\bar{x}, \bar{x}', \bar{x}', \dots), (n_0, 1, 1, \dots))$$

if and only if

$$v(x) + n_0 B v(x') \geq v(\bar{x}) + n_0 B v(\bar{x}').$$

It follows that

$$((u_0, u, u, \dots), (n_0, 1, 1, \dots)) \sim ((e_0, e_0, \dots), (n_0, 1, 1, \dots)),$$

where

$$e_0 \equiv v^{-1}\left(\frac{v(u_0) + Bn_0 v(u)}{1 + Bn_0}\right).$$

Note that e_0 is the equally-distributed equivalent utility defined in Section 2.

Next, by *respect of fertility preferences*,

$$((e_0, e_0, \dots), (n_0, 1, 1, \dots)) \sim ((u_f, u_f, \dots), (1, 1, \dots))$$

if and only if

$$U_0((e_0, e_0, \dots), (n_0, 1, 1, \dots)) = U_0((u_f, u_f, \dots), (1, 1, \dots)).$$

By definition of utility equivalent function, this equality holds only if

$$u_f = U_0^f(e_0, n_0),$$

which exists for any pair of e_0 and n_0 by finite substitutability of preferences. By setting $W_0(\mathfrak{a}) = u_f$, W_0 represents social preferences $\succsim$ on the ultimately constant

[18] See the derivation in Piacquadio (2020).

allocations that are constant from $t = 1$.

To generalize, consider any ultimately constant allocation that becomes constant after period t. Using repeated applications of the above construction and *stationarity*, one can iteratively fold the allocation back in time to obtain a unique equivalent constant allocation with unit fertility. This yields the recursive representation:

$$v(e_t) \equiv \frac{v(u_t) + Bn_t v(e_{t+1})}{1 + Bn_t},$$

and

$$W_t(\mathfrak{a}_t) \equiv U_t^f(e_t, n_t),$$

where e_t is the equally-distributed equivalent utility and U_t^f is the utility equivalent function of generation t. Thus, social preferences over ultimately constant allocations are represented by W_0.

Part 2. Extending the criterion to all allocations

I now extend the representation from ultimately constant allocations to the entire domain D. The key idea is to use the *constant equivalence* axiom, which ensures that any allocation can be assigned the same welfare as a constant one. I then argue that this constant welfare level can be obtained as the limit of welfare levels of ultimately constant allocations that approximate the original allocation.

Let $\mathfrak{a} = ((u_0, u_1, \dots), (n_0, n_1, \dots)) \in D$ be any allocation. By *constant equivalence*, there exists a constant allocation $\mathfrak{a}^* = ((u^*, u^*, \dots), (1, 1, \dots)) \in A^c$ such that $\mathfrak{a} \sim \mathfrak{a}^*$. I next show that the welfare level of $\mathfrak{a}$ can be expressed as the limit of the welfare levels of a sequence of ultimately constant allocations $\{\mathfrak{a}^{(t)}\}$ that approximate $\mathfrak{a}$.

Define the sequence of ultimately constant allocations:

$$\mathfrak{a}^{(t)} \equiv \left((u_0, u_1, \dots, u_t, u_{t+1}^*, u_{t+1}^*, \dots), (n_0, n_1, \dots, n_t, 1, 1, \dots)\right),$$

where $u^*_{t+1} = \min_{s \leq t} u_s$. This construction ensures that the tail of the allocation is constant and bounded by the earlier utility levels. Each $\mathfrak{a}^{(t)} \in A^{uc}$, so by Part 1, the social welfare $W_0(\mathfrak{a}^{(t)})$ is well-defined and recursively computable.

Since $\mathfrak{a}^{(t)} \to \mathfrak{a}$ pointwise and changes after period t become irrelevant as $t \to \infty$, we can invoke the *tail robustness* and *continuity* of the social welfare function to conclude that:

$$\lim_{t \to \infty} W_0(\mathfrak{a}^{(t)}) = W_0(\mathfrak{a}).$$

On the other hand, since $\mathfrak{a} \sim \mathfrak{a}^*$, $W_0(\mathfrak{a}) = W_0(\mathfrak{a}^*) = u^*$. Hence:

$$\lim_{t \to \infty} W_0(\mathfrak{a}^{(t)}) = u^*.$$

It follows that the welfare level of any allocation $\mathfrak{a} \in D$ coincides with that of its constant equivalent $\mathfrak{a}^*$, and that this value is the limit of welfare levels of ultimately constant approximations to $\mathfrak{a}$. Therefore, the fertilitarian criterion extends uniquely and consistently to the full domain D.

A.3 Corollary 1

Proof. When population is constant, the level of social welfare at an allocation $\mathfrak{a} \equiv ((u_0, u, u, \dots), (1, 1, \dots)) \in D$ is measured by

$$W_0(\mathfrak{a}) = v^{-1}\left[\frac{v(u_0) + Bv(u)}{1 + B}\right].$$

This criterion satisfies scale invariance if and only if v is constant elasticity of substitution, that is, v takes the form

$$v(x) = \frac{x^{1-\rho}}{1-\rho} \text{ for } \rho \geq 0, \rho \neq 1,$$

$$v(x) = \ln x \text{ for } \rho = 1.$$

□

A.4 Corollary 2

Proof. This result follows from Theorem 2 of Bleichrodt et al. (2008). In particular, see Observation 3. □