The Irrelevance of Intergenerational Altruism for Social Discounting

Maya Eden*

University of Zurich, Brandeis University, GPI and CEPR

July 5, 2024

Abstract

Should people's concern for the wellbeing of their descendants affect policy decisions? I consider a model in which people's dynastic utilities depend on the consumption of their descendants. The social welfare function is a discounted sum of past, present and future dynastic utilities. I establish that, when the social welfare function places sufficient weight on the dynastic utilities of future generations, intergenerational altruism has only negligible effects on the social ranking, and can be mostly ignored for the purpose of policy decisions. The reason is that, given persistant population growth, each generation's concern for its children roughly cancels out with their parents' concern for them.

JEL Classification: E61, D64

Keywords: Social discounting, intergenerational altruism, dynamic consistency

---

*Email: maya.eden@econ.uzh.ch. I am grateful to Geir Asheim, Luis Mota Freitas, Thomas Houlden, Kevin Kuruc, Frikk Nesje, Florian Scheuer, Itai Sher and Ivan Werning for helpful comments, as well as for comments made may seminar participants at the Normative Economics and Economic Policy Webinar, the Center for Population Level Bioethics at Rutgers University and Oxford University.

1 Introduction

The normative assessment of intergenerational tradeoffs is complicated by the presence of intergenerational altruism. For example, an argument can be made that, in addition to hurting future generations, climate change also hurts people today, because they worry about its negative effects on the lives of their children. In principle, a complete cost-benefit analysis would have to take into account how each generation's consumption affects the welfare of all other generations.

This paper asks, how do these indirect effects change the normative assessment of intergenerational policy tradeoffs? The answer turns out to be, "not much". I consider a discounted-utilitarian social objective, in which each generation's dynastic utility depends not only on their own consumption but also on the consumptions of their ancestors and their descendants. In this framework, I establish that, when the social ordering places sufficient weight on the dynastic utilities of future generations and population growth is constant, intergenerational altruism doesn't matter for the normative assessment of intergenerational tradeoffs in consumption.

The intuition for this result is that, when population growth is constant, each person matters to his family in exactly the same way: the degree to which I matter to my parents is the same as the degree to which my children matter to me; the degree to which I matter to my grandparents is the same as the degree to which my children matter to my parents, and so on and so forth. Regardless of how much people care about their descendants (and ancestors), there is a sense in which all generations are *equally* loved. Consequently, the presence of intergenerational altruism does not create an asymmetry between present and future generations.

Except in special cases, this theoretical result is limited to situations in which population growth is constant. However, human population growth has not been constant throughout history, nor is it expected to remain constant in the future. This implies that, in principle, intergenerational altruism may play a role in welfare assessments. In the second part of the paper, I conduct a quantitative assessment of the effects of intergenerational altruism on the social welfare function. It turns out that, because population growth is highly persistent, the theoretical insight obtained in the constant-growth case provides a good approximation for the empirically-relevant case. I find that, even for extreme parameterizations, intergenerational altruism has only negligible effects on the social welfare function.

This paper is related to a large literature on social discounting (see Millner and Heal [2023] for a recent review). Much of the disagreement in this literature reflects a philosophical disagreement about the reasons for acting on behalf of future generations. The market-based approach (which, in the context of climate change, is identified with Nordhaus [2007]) focuses on reasons of intergenerational altruism. According to this view, the social welfare function should reflect a concern for future generations *because* people today care about them.[1][2]

In contrast, the approach taken by Stern [2008], Greaves and MacAskill [2024] and many others focuses on reasons of independent moral significance. According to this approach, the reasons for acting on behalf of future generations are qualitatively similar to the reasons for acting on behalf of people today; future generations are deserving of consideration, regardless of how much their ancestors care about them.

This paper continues in the tradition of Farhi and Werning [2007], Farhi and Werning [2013] and Ray [2018], who consider both of these reasons jointly. Farhi and Werning [2007] present a model in which these two types of reasons complement each other. In their framework, considering reasons of independent moral significance on top of reasons of intergenerational altruism implies a social objective that cares more about future generations. Here, I consider a similar framework that also takes into account the preferences of past generations, and find that reasons of intergenerational altruism no longer play an important role.

It is important to note that, within the utilitarian framework that I consider here, it is not obvious whether it is appropriate to incorporate people's altruistic preferences. Dworkin [1990], Milgrom [1993] and Sher [2020] argue that, for the purpose of welfare analysis, other-regarding preferences should be ignored, especially if they represent a moral concern towards others. In this case, counting the utility that an individual derives through his altruism leads to "double counting" of the moral value of others' consumption. However, as Ray [2018] notes, it is not obvious that intergenerational altruism represents a moralistic concern of this type. Rather, it seems that people genuinely care for their descendants, in the sense that they derive

[1] Similar arguments can be found in Barro [1974] and Galperti and Strulovici [2017] as well as in Scheffler [2013] and Scheffler [2018].

[2] In recent work, Nesje [2024] shows that when intergenerational altruism extends across dynasties – that is, when people care about each other's children and not just about their own children – then the social discount rate is lower than the market interest rate. This illustrates that the market-based approach is not necessarily implied by a normative approach that is based only on reasons of intergenerational altruism.

true pleasure from their happiness and true pain from their suffering. As a result, the utility derived through intergenerational altruism is indistinguishable from other forms of utility. While, in this paper, I have nothing new to contribute to this debate, my results suggest that perhaps it is not a very important one: whether or not we choose to incorporate intergenerational altruism into our welfare analysis has only negligible effects on the social ranking.

2 Illustration

The main result builds on the symmetry of adjacent generations in situations in which population growth is constant. This symmetry has been overlooked by models that ignore the preferences of past generations. To illustrate, consider the following two examples.

Example 1. The first example is a two-period model, in the spirit of Farhi and Werning [2007] and Ray [2018]. In this model, there are only two consecutive generations of equal size, alive in periods $t = 0$ and $t = 1$. The *dynastic utility* of the first generation (which I refer to as generation 0) is

$$U_0(c_0, c_1) = u(c_0) + \beta u(c_1) \tag{1}$$

where c_0 is generation 0's own consumption and c_1 is the consumption of its children, who will be alive in period 1. The parameter β captures intergenerational altruism of parents towards their children. I refer to $u(c_t)$ as the *consumption utility* of generation t.

For the purpose of this example, I follow Farhi and Werning [2007] and Ray [2018] and ignore children's altruism towards their parents (the case of backward looking altruism will be covered in the more general model in the next section).[3] Since generation 1 has no children, its dynastic utility is simply

$$U_1(c_1) = u(c_1) \tag{2}$$

[3] The broader implications of backward looking altruism have been studied recently in Ray et al. [2024].

The social objective takes the discounted utilitarian form,

$$W(c_0, c_1) = U_0(c_0, c_1) + \frac{U_1(c_1)}{1+\rho}$$

Here, $\rho > -1$ captures the social rate of time preference. It is the rate at which the social objective discounts the dynastic utility of the children's generation relative to the dynastic utility of their parents.

When we substitute in the expressions for the dynastic utilities (equations 1 and 2), we obtain

$$W(c_0, c_1) = u(c_0) + \left(\beta + \frac{1}{1+\rho}\right) u(c_1) \tag{3}$$

Note that the weight on $u(c_1)$ is higher than $1/(1+\rho)$. This is because the social welfare function takes into account that increasing the consumption of generation 1 also improves the welfare of generation 0. This effect is stronger when β is higher; so, in this example, the presence of intergenerational altruism makes future consumption relatively more valuable from a social perspective.

Example 2. In the above example, there were no past generations to consider at time 0. Now let us consider a modification in which there is an additional generation alive in period (-1). The dynastic utility of generation (-1) is

$$U_{-1}(c_{-1}, c_0) = u(c_{-1}) + \beta u(c_0)$$

Note that the extent to which generation (-1) cares about its children is the same as the extent to which generation 0 cares about theirs. For the purpose of this example, I have assumed that people care about their children's consumption utility ($u(c_0)$), and not about their dynastic utility ($U(c_0) = u(c_0) + \beta u(c_1)$). This assumption is useful for developing intuition for the main result; however, as I will establish in the following sections, it is not essential for it.

With this modification, the discounted-utilitarian objective becomes

$$W(c_{-1}, c_0, c_1) = (1+\rho)U_{-1}(c_{-1}, c_0) + U_0(c_0, c_1) + \frac{U_1(c_1)}{1+\rho}$$

$$= (1+\rho)u(c_{-1}) + ((1+\rho)\beta + 1)\, u(c_0) + \left(\beta + \frac{1}{1+\rho}\right) u(c_1)$$

Dividing through by $(1+\rho)\beta+1$, these preferences are also represented by the social welfare function

$$\frac{W(c_{-1},c_0,c_1)}{(1+\rho)\beta+1} = \left(\frac{1+\rho}{(1+\rho)\beta+1}\right)u(c_{-1}) + \left(\frac{(1+\rho)\beta+1}{(1+\rho)\beta+1}\right)u(c_0) + \left(\frac{\beta+\frac{1}{1+\rho}}{(1+\rho)\beta+1}\right)u(c_1) =$$

$$\left(\frac{1+\rho}{(1+\rho)\beta+1}\right)u(c_{-1}) + u(c_0) + \frac{\left(\beta+\frac{1}{1+\rho}\right)u(c_1)}{\left(\beta+\frac{1}{1+\rho}\right)(1+\rho)} =$$

$$\left(\frac{1+\rho}{(1+\rho)\beta+1}\right)u(c_{-1}) + u(c_0) + \frac{u(c_1)}{1+\rho}$$

As in Example 1, intergenerational altruism affects the rate at which $u(c_0)$ is discounted relative to $u(c_{-1})$. However, it does not affect the rate at which $u(c_1)$ is discounted relative to $u(c_0)$.

Intuitively, once we take the preferences of generation (-1) into account, generations 0 and 1 become completely symmetric. The extra boost that the welfare weight of generation 1 gets on account of the altruism of generation 0 is offset by the boost that generation 0 gets on account of the altruism of generation (-1). Consequently, the relative weights on $u(c_1)$ and $u(c_0)$ are completely independent of the intergenerational altruism parameter, β, and depend only on the social rate of time preference, ρ.

This example suggests two conclusions:

- Once we take into account the preferences of past generations, intergenerational altruism doesn't matter as much for social discounting.

- If we ignore the preferences of past generations (which, in this example, amounts to considering period (-1) as the current period), then intergenerational altruism matters for short-run social discounting (in this example, discounting between periods (-1) and 0), but not for the long-run social discount rates (in this example, the discount rate between periods 0 and 1).

2.1 Reasons for acting on behalf of past generations

The above examples illustrate that, within the discounted-utilitarian framework, whether or not we take the altruistic preferences of past generations into account matters for the purpose of making plans for the future, especially over shorter time horizons. This begs the question: should policymakers take the preferences of past generations into account and, if so, why?

The answer to this normative question depends somewhat on how we interpret the dynastic utility functions, $\{U_t\}$, and the social welfare function, W. One possibility is to interpret $\{U_t\}$ as representations of individual preferences, and W as an aggregation of these preferences. In this framework, any differential treatment of individuals' preferences over their own consumption and their preferences over the consumption of others would lead to violations of the Pareto condition. In this sense, taking the preferences of past generations into account is unavoidable.

Things are somewhat more nuanced when we consider U_t to be a measure of experienced utility (as in, for example, Ray [2018]). In this case, U_t cannot be affected by events that take place after an individual's death. However, it may depend on the individual's beliefs about the future, during his own lifetime. For example, a father may experience joy from the belief that his children will continue to have good lives after he passes. At the same time, whether or not they actually continue to have good lives will not affect his experienced utility.

Even in such a framework, the preferences of past generations may play a role in shaping the social objective. To illustrate this (well-understood[4]) point, consider a planner whose objective is to maximize the discounted sum of dynastic utilities, as in Example 2. The planner is able to commit at time (-1) to the entire trajectory of consumption. By doing so, he is able to set the expectations of generation (-1) regarding the future consumption trajectory. In period (-1), he faces the optimization problem

$$\max_{c_{-1},c_0,c_1} (1+\rho)U_{-1}(c_{-1},c_0) + U_0(c_0,c_1) + \frac{U_1(c_1)}{1+\rho} \text{ s.t. } (c_{-1},c_0,c_1) \in B_{-1}$$

where B_{-1} is the budget constraint at time (-1). Importantly, this planner considers how his choices of c_0 and c_1 will affect the experienced utility of generation (-1), who cares about the future trajectory of consumption.

At time 0, the planner simply carries out the plans to which he had committed in period (-1). However, note that, by the principle of optimality, his choice of (c_0, c_1)

[4] See, for example, Kydland and Prescott [1977], Barro and Gordon [1983] and Golosov et al. [2007], and particularly Ahlvik [2022].

also solves the optimization problem

$$\max_{c_0,c_1}(1+\rho)U_{-1}(c_{-1},c_0)+U_0(c_0,c_1)+\frac{U_1(c_1)}{1+\rho}\text{ s.t. }(c_{-1},c_0,c_1)\in B_0(c_{-1})$$

where $B_0(c_{-1})$ is the budget constraint at time 0 (which may depend on the choice of c_{-1} in the previous period). That is, at time 0, the planner behaves *as if* his actions can affect the expectations of generation (-1) posthumously. The way that he evaluates tradeoffs between c_0 and c_1 is exactly the same as the way that he evaluated them in period (-1). This illustrates that the ability to commit implies a norm of taking the preferences of past generations into account.

It is important to clarify that the framework that I consider here is an axiological framework, which assumes a time-invariant social objective. There is no sense in which "welfare from the perspective of time (-1)" can be different from "welfare from the perspective of time 0" – there is only one level of social welfare, that does not depend on when it happens to be evaluated. In Appendix C, I explore a more relativist approach, in which the social objective itself changes over time. There, I assume that people's utilities matter before and during their lifetimes, but cease to matter after their deaths. As illustrated in Bernheim [1989], Galperti and Strulovici [2017], Asheim et al. [2021] and Nesje [2024], this environment often gives rise to dynamic inconsistency, in the sense that, in the future, it might not be optimal to carryout plans that are optimal from today's perspective. I show that, even in such a framework, it may be beneficial to uphold a norm in which policymakers respect the wishes of past generations. However, it is also possible that such a norm increases welfare in certain periods but not in others.

3 An infinite-horizon model

Time periods are discrete, and indexed $t \in \mathbb{Z}$. This formulation assumes that, in each period, both the future and the past are infinite. Let $N = (..., N_t, ...) \in \mathbb{R}^{\mathbb{Z}}_{++}$ denote a vector of population levels, with N_t denoting the size of the cohort born in period t. Similarly, let $c = (..., c_t, ...) \in \mathbb{R}^{\mathbb{Z}}_{++}$ denote a vector of consumption levels, with c_t denoting the per-capita consumption level of generation t. Throughout, I assume that members of the same generation have the same consumption level and the same number of children. This ignores within-generation consumption inequality or heterogneity in the number of children.

For the purpose of deriving a first theoretical benchmark, I assume that, conditional on N, generation t's dynastic utility is given by

$$U_t(c) = \frac{1}{N_t} \sum_{\tau=-\infty}^{\infty} \psi_\tau N_{\tau+t} u(c_{\tau+t}) \tag{4}$$

where $\psi_0 = 1$, $\psi_\tau \geq 0$ for all τ, and u is strictly increasing. It is worth noting that this specification can nest "non-paternalistic" models in which generations care about one another's dynastic utilities, rather than directly about each other's consumption utilities. For example, it is straightforward to show that dynastic preferences of the form $U_t = u(c_t) + \beta(N_{t+1}/N_t)U_{t+1}$ (as in Barro [1974]) are consistent with the above specification with $\psi_t = \beta^t$ for $t \geq 0$ (and $\psi_t = 0$ for $t < 0$).[5]

To guarantee that dynastic utilities are always well-defined, I restrict attention to consumption streams for which $u(c_t)$ is nonzero only in a finite number of periods. This restriction guarantees that the infinite sum in (4) always converges, regardless of the parameters ψ_τ. While this assumption excludes the most realistic scenarios, such as those involving increasing consumption sequences, it also allows me to sidestep the well-known technical issues of convergence that come up in infinite horizon models. As my purpose here is to learn about the properties of the social ranking, this restriction is innocuous, and will play no role other than guaranteeing convergence.

The functional form in (4) assumes that each person is equally-loved by his ancestors, regardless of the dynastic structure. To see this, note that N_{t+1}/N_t is the number of children of each member of generation t; N_{t+2}/N_t is the number of grandchildren of each member of generation t; etc. The parameter ψ_1 governs the extent to which each person cares about each of his children's consumption utility; ψ_2 parameterizes the extent to which each person cares about each of his grandchildren's consumption utility; etc. According to this functional form, a person's willingness to tradeoff his own consumption with the consumption of one of his children does not depend on

[5]See also Saez-Marti and Weibull [2005].

the number of children he has.[6] This property, which has some normative appeal, is empirically controversial. I will return to this point in the following section.

The above functional form also assumes that the extent to which each person's children *collectively* care about them is independent of the number of children that they have. It implies, for example, that a person with only a few children can expect the same level of care in old age as a person with many children. How realistic this assumption is is difficult to assess. However, it is useful for establishing an interesting theoretical benchmark, and it will be relaxed in the following section.

Social preferences. Given a population stream, N, the social ranking of consumption streams is assumed to be represented by the discounted-utilitarian form,

$$W(c) = \sum_{t=-\infty}^{\infty} \frac{N_t U_t(c)}{(1+\rho)^t} \tag{5}$$

where $\rho > -1$ is the social rate of time preference. Note that, here, I am assuming a time-invariant social objective, in which the altruistic preferences of past generations continue to matter after their deaths.

Given our restriction to consumption streams for which $u(c_t)$ is nonzero only in a finite number of periods, it turns out that this infinite sum converges when

$$\sum_{t=-\infty}^{\infty} (1+\rho)^t \psi_t < \infty \tag{6}$$

The term on the left hand side is the discounted sum of the altruistic weights that are placed on each person by his ancestors and descendants. A person places a weight of ψ_0 on themselves. Their parents place weight τ_1 on them: this weight is inflated by $(1+\rho)$, because the dynastic utility of the parents is $(1+\rho)$ times more important than the dynastic utility of their kids. Conversely, one's children place a weight of ψ_{-1} on them. This weight is discounted by a factor of $1/(1+\rho)$.

[6] To see this, recall that c_{t+1} represents the *per-capita* consumption of generation $t+1$. Thus, the marginal rate of substitution between $u(c_t)$ and $u(c_{t+1})$ captures an individual's willingness to sacrifice his own consumption utility in order to simultaneously increase the consumption utility of all of his children. The marginal rate of substitution between one's own consumption utility and the consumption utility of one of their children is given by $\left(\frac{\frac{\partial U_t(c)}{\partial u(c_{t+1})}}{\frac{\partial U_t(c)}{\partial u(c_t)}}\right)\left(\frac{N_t}{N_{t+1}}\right) = \psi_1$. Similarly, an individual's willingness to substitute between his own consumption utility and the consumption utility of one of his τ-th descendants is given by ψ_τ.

For a given sequence $\{\psi_t\}_{t=-\infty}^{\infty}$, whether or not this condition holds may depend on the social rate of time preference, ρ. For example, consider the standard, forward-looking exponential discounting model, in which $\psi_t = \beta^t$ for $t \geq 0$ (and $\psi_t = 0$ for all $t < 0$). In this model, $1/\beta - 1$ is the individual's "generational discount rate", which captures the extent to which an individual discounts the consumption utilities of his descendants relative to his own. It is straightforward to verify that, in this case, the inequality in (6) holds if and only if $\beta(1+\rho) < 1$ — that is, if and only if the social rate of time preference is lower than the individual's generational discount rate. This condition means that the social preference relation must place sufficient weight on the dynastic utilities of future generations (relative to the dynastic utilities of past and present generations).

The first theoretical result is stated in the following proposition.

Proposition 1. *Assume that dynastic utilities are given by (4), and that the inequality in (6) holds. Then, the social ranking represented by (5) is also represented by*

$$W^u(c) = \sum_{t=-\infty}^{\infty} \left(\frac{1}{1+\rho}\right)^t N_t u(c_t) \tag{7}$$

The proof of this proposition is detailed in the appendix, along with other omitted proofs. It follows simply from substituting in the relevant expressions and applying a strictly monotone transformation to the social welfare function.

The social welfare function in (7) is a discounted sum of consumption utilities, whereas the social welfare function in (5) is a discounted sum of dynastic utilities. The social rate of time preference, ρ, is the same in both. Importantly, (7) does not depend on the coefficients of intergenerational altruism, $\{\psi_t\}$. This result therefore illustrates that, under some assumptions, intergenerational altruism is completely irrelevant for the social ranking of consumption streams. If we are willing to accept these assumptions, then, for the purpose of welfare analysis, we can ignore any indirect effects that the consumption of one generation may have on the dynastic utilities of all other generations.

4 Robustness around the theoretical benchmark

The functional form in (4) assumes that the degree to which each person is loved by his family is independent of the structure of the dynasty. This section relaxes this assumption.

For each $N \in \mathbb{R}^{\mathbb{Z}}_{++}$ define the shifted vector, ${}_tN$, as ${}_tN_\tau = N_{\tau+t}$. Consider the class of utility functions,

$$U_t(c) = \sum_{\tau=-\infty}^{\infty} f_\tau \left(\frac{{}_tN}{N_t} \right) u(c_\tau) \tag{8}$$

where $\{f_\tau \geq 0\}_{\tau=-\infty}^{\infty}$ and $f_0 > 0$. The function f_τ captures the altruistic weight that each generation places on the consumption utility of each of their τ-th descendants (for a negative τ, this is the weight placed on the $(-\tau)$-th ancestors). This weight may depend on the number of descendants, and even on the entire structure of the dynasty. Note that (4) is a special case of this family of dynastic utility functions, with $f_\tau(N) := \psi_\tau N_\tau / N_0$. This more-general formulation allows for the possibility that the extent of intergenerational altruism varies with family size. For example, it is possible that parents' willingness to sacrifice on behalf of each of their children is decreasing in the number of children that they have.

It turns out that, when population growth is constant, the conclusion of Proposition (1) can be extended to this family of dynastic utility functions.

Proposition 2. *Assume that dynastic utilities are given by (8), and that* $N_t/N_0 = (1+n)^t$ *for some* $n > -1$*. Further, assume that the inequality in (6) holds for* $\psi_t := f_t(N/N_0)N_0/N_t$*. Then, the social ranking represented by 5 is also represented by 7.*

Intuitively, when population growth is constant, all individuals are perfectly symmetric, regardless of $\{f_\tau\}$. Consider, for example, the case in which each person has two children. In this scenario, each person has two children, four grandchildren, eight great-grandchildren etc. Similarly, each person, regardless of when he is born, is one of his parents' two children, one of his grandparents' four grandchildren, and one of his great-grandparents' eight grandchildren. So, everybody is equally loved by their respective ancestors and descendants. The relative importance of our love for our children and our parents' love for us is discounted at the social rate of time preference, ρ.

5 When is intergenerational altruism relevant?

In this section, I discuss the possible ways in which intergenerational altruism may affect social discounting, away from the theoretical benchmarks covered in Propositions 1 and 2.

A high social rate of time preference, ρ. Note that both propositions assume that the inequality in (6) holds. Intuitively, this condition guarantees that the altruistic preferences of our distant ancestors have only negligible contributions to the weight that the social welfare function places on our consumption utility today. To see this, note that the social welfare function can be written as a linear combination of consumption utilities,

$$W = \sum_{t=-\infty}^{\infty} a_t N_t u(c_t) \text{ where } a_t = \sum_{\tau=-\infty}^{\infty} \frac{\psi_{t-\tau}}{(1+\rho)^{\tau}}$$

The term $\psi_{t-\tau}/(1+\rho)^{\tau}$ is the contribution of generation τ to the coefficient on $N_t u(c_t)$. When people care very little about their distant relatives, and ρ is not too far from 0, then this term will go to zero as $\tau \to -\infty$. However, if ρ is large, then the negligible effects that our consumption today has on the dynastic utilities of our distant ancestors may come to dominate the welfare calculations.

In this case, intergenerational altruism may matter for social discounting – in fact, it may be the *only* thing that matters. To illustrate, consider an example in which the altruism parameters are given by

$$\psi_t = \begin{cases} \beta^t & \text{if } t \geq 0 \\ 0 & \text{Otherwise.} \end{cases}$$

where $\beta(1+\rho) > 1$. Given this parametric restriction, it turns out that the infinite sum in the social objective (5) converges only if there is a "first" generation, that is, only if $N_\tau = 0$ for all τ sufficiently small. So, for the purpose of this example, we must depart slightly from our framework and assume that there is a first generation, born at $t = 0$.

In this case, the coefficient on $N_t u(c_t)$ in the social welfare function is given by

$$a_t = \sum_{\tau=0}^{t} \frac{\beta^{t-\tau}}{(1+\rho)^\tau} = \beta^t \sum_{\tau=0}^{t} \left(\frac{1}{\beta(1+\rho)} \right)^\tau$$

It is straightforward to show that, if $\beta(1 + \rho) > 1$, then

$$\lim_{t\to\infty} \frac{a_{t+1}}{a_t} = \beta$$

That is, the social discount rate between two generations that are sufficiently far from the first humans is approximately the same as the "generational discount rate" that is implied by the parameters of intergenerational altruism. The normative parameter, ρ, no longer plays any role.

This example illustrates that intergenerational altruism may be relevant for social discounting if the social objective discounts the dynastic utilities of future generations at a sufficiently high rate. In the exponential discounting example, there is an interesting dichotomy: when $\beta(1 + \rho) < 1$, then the social social discount rates depend only on ρ, and not on β; but, when $\beta(1 + \rho) > 1$, then the social discount rates depend only on β, and not on ρ.

Time-varying population growth. For the purpose of this discussion, consider a stylized model in which

$$U_t(c) = u(c_t) + \beta \left(\frac{N_{t+1}}{N_t} \right)^{1-\epsilon} u(c_{t+1})$$

where $\epsilon \in [0, 1]$ and $\beta > 0$. If $\epsilon = 0$, then this functional form an instance of (4). In this case, by Proposition 1, intergenerational altruism plays no role regardless of the population dynamics.

Thus, for intergenerational altruism to play a role, it must hold that $\epsilon > 0$. In this case, the assumptions of Proposition 2 are satisfied; consequently, for intergenerational altruism to be relevant, we must consider a scenario in which population growth varies across time. To illustrate the role of intergenerational altruism in this scenario, assume that generation (-1) is the first generation, which has N_{-1} members. Generations 0 and 1 have 1 member each; after that, there are no more people.[7] Under these assumptions, the discounted utilitarian social welfare function is

$$W(c_{-1}, c_0, c_1) = (1+\rho)N_{-1}U_{-1}(c) + U_0(c) + \frac{U_1(c)}{1+\rho} =$$

$$= (1+\rho)N_{-1}u(c_{-1}) + (1 + (1+\rho)N_{-1}\beta\left(\frac{1}{N_{-1}}\right)^{1-\epsilon})u(c_0) + (\beta + \frac{1}{1+\rho})u(c_1)$$

$$= (1+\rho)N_{-1}u(c_{-1}) + (1 + (1+\rho)N_{-1}^{\epsilon}\beta)u(c_0) + (\beta + \frac{1}{1+\rho})u(c_1)$$

and thus the social ranking of (c_0, c_1) conditional on a predetermined c_{-1} is represented by

$$W(c_0, c_1|c_{-1}) = (1 + (1+\rho)N_{-1}^{\epsilon}\beta)u(c_0) + (\beta + \frac{1}{1+\rho})u(c_1)$$

Note that, as $N_{-1} \to 0$, the term N_{-1}^{ϵ} converges to 0, and so the coefficient on $u(c_0)$ converges to 1. In this limit, the above social welfare function is the same as (3). In this case, intergenerational altruism matters because generation 1 is much more loved than generation 0: each member of generation 0 had many siblings which diluted their parents' love towards them. In contrast, each member of generation 1 has only 1 sibling (in a 2-parent household), which, according to this model, makes them more loved by their parents. This example suggests that, when population growth is declining, intergenerational altruism makes future generations relatively more important. Indeed, as we will see in the next section, this seems to be the empirically-relevant case.

At the other extreme, as $N_{-1} \to \infty$, each child of generation 0 is loved more by their parents than each child of generation 1. In this case, population growth is increasing, because it is negative between periods (-1) and 0 and zero between periods 0 and 1. This example therefore suggests that, when population growth is increasing, intergenerational altruism decreases the relative importance of future generations.

[7] For $U_1(c)$ to be well-defined, let us assume that $\epsilon < 1$ and $u(0) = 0$, so that $\beta\left(\frac{0}{1}\right)^{1-\epsilon}u(0) = 0$ and hence $U_1(c) = u(c_1)$.

6 Quantitative results

In this section, I quantify the potential effects of intergenerational altruism in practice, given past and future population growth trajectories, and given realistic parameterizations of intergenerational altruism. It turns out that, although human population growth is not constant, it is highly persistent. Consequently, the intuition we get from considering the constant population growth model provides a good benchmark for the empirically-relevant case.

For the purpose of this discussion, I restrict attention to the standard functional form

$$f_t(N) = \begin{cases} \beta^t \left(\frac{N_t}{N_0}\right)^{1-\epsilon} & \text{if } t \geq 0 \\ 0 & \text{otherwise.} \end{cases}$$

where $\beta \in [0, 1)$ and $\epsilon \in (0, 1)$. This functional form follows the standard specification in Barro and Becker [1988] and Barro and Becker [1989]. In addition, I follow the standard approach and assume that people do not care at all about their ancestors. Including child-to-parent altruism counters some of the effects of parent-to-child altruism; thus, I expect intergenerational altruism to have even smaller effects under more realistic parameterizations.

I set $\epsilon = 0.288$, as in Cordoba [2015] and Cordoba and Ripoll [2019], who calibrate this parameter based on the relationship between health expenditure per child and family size, as documented in Dickie and Messman [2004]. This parameter choice is also broadly in line with Adhami et al. [2024], who calibrate $\epsilon \approx 0.2$ based on analysis of household expenditure patterns documented in Lino [2011].

It is worth noting that the choice of $\epsilon > 0$ is at odds with the work of Angrist et al. [2010] who, using a variety of natural experiments, find no evidence of a quality-quantity tradeoff in the number of children. Instead, they find that, consistent with (4), the amount that parents invest in each of their children does not depend on the number of children that they have. This finding suggests a parameterization of $\epsilon = 0$, which, by Proposition 1, implies no role for intergenerational altruism.

In my calibration, a period corresponds to 25 years, which is roughly the age difference between parents and their children. Note that the parameter β represents the extent to which people care about the *lifetime* consumption utility of their children, relative to their own *lifetime* consumption utility. As Cordoba [2015] notes, this

Table 1: Calibration Parameters

Variable	Parameter value
Length of a period	25 years
β	0.067
ϵ	0.288
ρ (annualized)	0; 0.01; 0.05; 0.1

parameter can be calibrated as

$$\beta = \delta^F \alpha$$

where δ is the individual's annual time discount factor, F is the age of fertility (in this model, $F = 25$), and α captures the extent to which people discount the consumption utility of their children relative to their own consumption utility at the age of parenthood.

To calibrate δ, Cordoba [2015] uses a subjective rate of time preference of 2%, so that $\delta = 1/1.02$. The parameter α is calibrated based on parents' fertility choices. The idea is that, if parents can choose how many children to have, then they will set the number of children so that their marginal utility from an additional child is equal to the cost of raising the child. Given ϵ, an estimate of the cost of raising children can be used to calibrate α. Using this approach, Cordoba [2015] calibrates $\alpha = 0.11$.

This calibration procedure implies $\beta = 0.067$. At first glance, this parameter may seem low. However, it is worth remembering that a period in the model corresponds to 25 years; hence, this implies an annualized discount rate of about 10% (or an annualized β of about 0.9).

The parameter ρ is a normative choice. The purpose of this exercise is not to argue for a particular choice of ρ — rather, it is to show that, for plausible values of ρ, intergenerational altruism parameters have only negligible effects on the social ranking of consumption streams. I therefore consider a range of values for ρ. The lowest value that I consider is $\rho = 0$, which reflects the stance that all generations matter equally for social welfare. On the other end, I consider a value of ρ that discounts future generations by 10% *annually*. This implies $1/(1+\rho) \approx 0.1$, that is, the dynastic utility of the current generation matters roughly 10 times as much as the dynastic utility of the next generation, and 100 times as much as the dynastic utility of the generation after that.

It is worth noting that all of the values of ρ that I consider in this calibration satisfy $\beta(1+\rho) < 1$, which guarantees that the inequality in (6) holds. As discussed in the previous section, higher values of ρ may yield very different results.

Finally, for the purpose of this exercise, I must consider both past population growth rates and the future growth trajectory. Based on Roser and Ritchie [2023], I assume that between the dawn of humanity and the year 1700, population growth was 0.04% annually. Global population levels between 1700 and 2100 are taken from Roser and Ritchie [2023], who base the growth rates between 2023 and 2100 on UN projections. The UN projections end in the year 2100, with a population decline of 0.1%. I assume, perhaps conservatively, that −0.1% remains humanity's population growth rate until its eventual extinction.[8]

Figure 1 presents the population levels that I assume for the sake of this calibration. Because population growth rates are not constant, Proposition 2 does not apply, making it possible for intergenerational altruism to play a role in social welfare analysis.

The purpose of this calibration is to quantify the plausible effects of intergenerational altruism on the social welfare function. For this purpose, it is useful to observe that the social welfare function in (5) can be written as a linear combination of the consumption utilities, $\{u(c_t)\}_{t=-\infty}^{\infty}$, with strictly positive coefficients.[9] Consequently, there exist some $\{\rho_t^*(N) > -1\}_{t=-\infty}^{\infty}$ such that the social welfare function in (5) can be written as

$$aW(c) = \sum_{t=-\infty}^{\infty} \left(\frac{1}{1+\rho_t^*(N)} \right)^t N_t u(c_t)$$

where $a > 0$ is some positive constant (that can be ignored, as it affects only the level of the social welfare function and not the ranking that it represents). Note that $\rho_t^*(N)$ is the average rate at which the social welfare function discounts the consumption utility of generation t relative to the consumption utility of generation

[8]This may be conservative, given that this would constitute a break in the trend of declining growth rates between the years 1963-2100.

[9]This follows from substituting in the relevant expressions for the dynastic utilities.

Figure 1: Population path

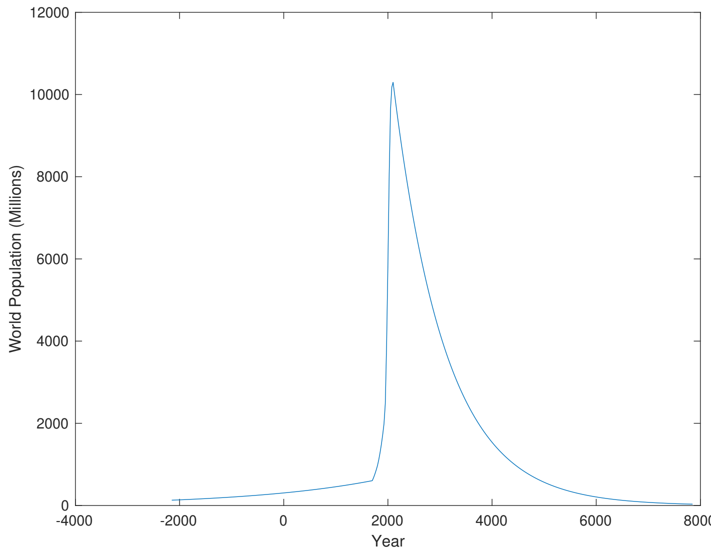

Note: Past population levels and growth trajectories until the year 2100 are based on Roser and Ritchie [2023]. After the year 200, I assume an annual population growth rate of -0.1%. The figure depicts population levels around the current time; however, in the calibration, I consider population levels between the years -47,000 and 24000. The starting point is calibrated so that, given 0.04% population growth in the pre-industrial era, humanity starts with two people. The end point is calibrated so that, given a projected population growth rate of -0.1%, humanity ends with two people.

0.

To express the extent to which this social objective differs from the social objective represented by (7), I report the difference between ρ and $\rho_t^*(N)$. For ease of interpretation, I express this difference in discount rates in annual percentage points. Formally, this means that I report the difference

$$\text{Annualized error (percentage points)} = 100\left((1+\rho)^{1/25} - (1+\rho_t^*(N))^{1/25}\right) \tag{9}$$

(where 25 years is the length of a period in this model). The annualized error is the difference between the naive social rate of time preference, ρ, which we get from using the social welfare function in (7), and the correct social rate of time preference that we would obtain if we were to use the social welfare function in (5). For example, a positive annualized error means that ignoring intergenerational altruism causes us to discount the future at a rate that is too high.

Figure 2 presents the difference between ρ and $\rho_t^*(N)$, between now ($t = 0$) and the next 1,000 years. As population growth is declining throughout this period, the difference tends to be positive, indicating that ignoring intergenerational altruism would lead one to discount the future at a rate that is too high (that is, to underestimate the extent to which future consumption matters for social welfare). However, the most important thing to observe about this graph is the range of the y-axis. For $\rho = 0$, the annualized error is at most 1 basis point (that is, one hundredth of one percent). Even for a value of ρ that discounts future generations at a rate of 10% annually, the annualized error is at most fifteen basis points. This suggests that, as long as ρ is in this range, we can mostly ignore intergenerational altruism for the purpose of computing social discount rates.

It may also be useful to illustrate what happens when we choose not to take the preferences of past generations into account. Technically, this amounts to setting $N_t = 0$ for all $t < 0$ (that is, assuming that the current generation is the first one that matters for the purpose of welfare analysis going forward). Figure 3 illustrates that, in this case, ignoring intergenerational altruism results larger errors, particularly in the medium run (25-100 years). When $\rho = 0$, then ignoring intergenerational altruism leads to a 25-year social discount rate that is about 25 basis points too high, and a 100-year social discount rate that is about 7 basis points too high. The errors are larger for higher values of ρ: for a ρ that discounts the dynastic utilities of future

Figure 2: The annualized error

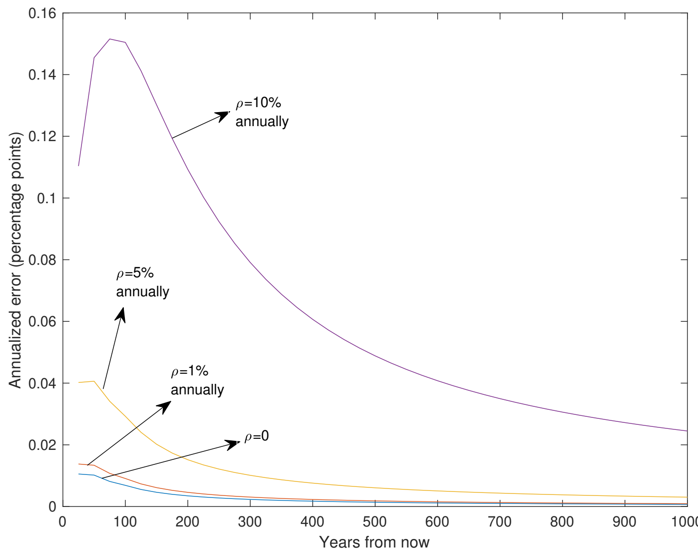

Note: The annualized error at time t is given by the formula in (9). Different lines represent values of ρ.

Figure 3: The annualized error: ignoring past generations

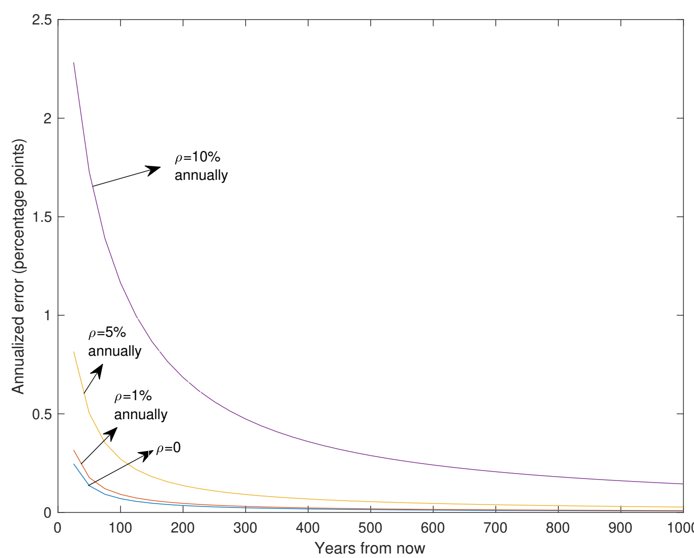

Note: The annualized error at time t is given by the formula in (9); however, $\rho_t^*(N)$ is calculated in a way that ignores the dynastic utilities of past generations.

generations at 10% annually, the annualized error is 2 percentage points for a 25 year horizon, and 1 percentage point for a 100 year horizon. However, even here, the long-run social discount rates are unaffected by intergenerational altruism. This is because the omission of past generations becomes less important the farther out we look into the future.

7 Conclusion

The daunting task of social discounting just became a little bit easier. In principle, this task is complicated by the presence of intergenerational altruism: when evaluating intergenerational tradeoffs, it may be necessary to take into account how each generation's consumption affects the dynastic utility of every other generation. Because people care deeply about their families, these indirect effects may be large, and may add up across multiple generations.

Luckily, it turns out that, for a broad range of parameters, intergenerational altruism is quantitatively irrelevant for social discounting. Importantly, this is not because people "don't care" about their families; rather, it is because each person is, approximately, *equally* loved by their family. Consequently, intergenerational tradeoffs are tradeoffs among people that are equally loved.

Three caveats are in order. The first is that, while intergenerational altruism is irrelevant for the purpose of evaluating intergenerational tradeoffs in consumption, it may not be irrelevant for the purpose of comparing different population trajectories. For example, the extent to which people care about having children may affect the importance of population growth relative to consumption growth.

The second caveat is that the results here rely on the additive-separability of the social welfare function, as well as on the additive separability of the dynastic utility function. When these assumptions are violated (as in, for example, Galperti and Strulovici [2017]), intergenerational altruism may be more relevant. However, the direction of the effect is not obvious: intergenerational altruism may either increase or decrease the social discount rate.

Finally, the analysis here is based on the premise that the preferences of past generations should affect policy decisions today. This premise arises naturally in normative frameworks that are based on preference satisfaction, but is more nuanced in ones that are based on the aggregation of experienced utilities. In the latter case, it is generally beneficial to establish and uphold a norm of honoring the wishes of the deceased. However, in the absence of such a norm, there may be no reason to take the preferences of past generations into consideration. In this case, intergenerational altruism may affect the social discount rate in the short run, but will have only negligible effects on the long run social discount rates.

References

Mohamad Adhami, Mark Bils, Charles I. Jones, and Peter J. Klenow. Population and welfare: The greatest good for the greatest number. 2024.

Lassi Ahlvik. Will of the living dead — the case for a backward-looking welfare function. *Economics Letters*, 214:110418, May 2022.

Joshua Angrist, Victor Lavy, and Analia Schlosser. Multiple experiments for the causal link between the quantity and quality of children. *Journal of Labor Economics*, 28(4):773–824, 2010.

Geir B. Asheim, Kuntal Banerjee, and Tapan Mitra. How stationarity contradicts intergenerational equity. *Economic Theory*, 72:423–444, 2021.

Robert J. Barro. Are government bonds net wealth? *Journal of political economy*, 82(6), 1974.

Robert J. Barro and Gary S. Becker. A reformulation of the economic theory of fertility. *The Quarterly Journal of Economics*, 103(1):1–25, 1988.

Robert J. Barro and Gary S. Becker. Fertility choice in a model of economic growth. *Econometrica*, 57(2):481–501, 1989.

Robert J. Barro and David B. Gordon. Rules, discretion and reputation in a model of monetary policy. *Journal of monetary economics*, 12(1):101–121, 1983.

B. Douglas Bernheim. Intergenerational altruism, dynastic equilibria and social welfare. *The review of economic studies*, 56(1):119–128, 1989.

Juan Carlos Cordoba. Children, dynastic altruism and the wealth of nations. *Review of Economic Dynamics*, 18(4), 2015.

Juan Carlos Cordoba and Marla Ripoll. The elasticity of intergenerational substitution, parental altruism, and fertility choice. *The Review of Economic Studies*, 86 (5), 2019.

Mark Dickie and Victoria L. Messman. Parental altruism and the value of avoiding acute illness: are kids worth more than parents? *Journal of Environmental Economics and Management*, 48(3):1146–1174, 2004.

Ronald Dworkin. The 'double counting' objection. In Jonathan Glover, editor, *Utilitarianism and Its Critics*, pages 103–106. MacMillan, 1990.

Emmanuel Farhi and Ivan Werning. Inequality and social discounting. *Journal of political economy*, 115(3):365–402, 2007.

Emmanuel Farhi and Ivan Werning. Estate taxation with altruism heterogeneity. *American Economic Review*, 103(3):489–495, 2013.

Simone Galperti and Bruno Strulovici. A theory of intergenerational altruism. *Econometrica*, 85(4):1175–1218, 2017.

Mikhail Golosov, Aleh Tsyvinski, and Ivan Werning. New dynamic public finance: A user's guide. In Kenneth Rogoff Daron Acemoglu and Michael Woodford, editors, *NBER Macroeconomics Annual 2006, Volume 21*, pages 317–388. The MIT Press, 2007.

Hilary Greaves and William MacAskill. The case for strong longtermism. In Greaves Barrett and Thorstad, editors, *Essays on Longtermism*. Oxford, UK: Oxford University Press, 2024.

Finn E. Kydland and Edward C. Prescott. Rules rather than discretion: The inconsistency of optimal plans. *Journal of political economy*, 85(3):473–491, 1977.

Mark Lino. Expenditures on children by families. Technical Report, U.S. Department of Agriculture, Center for Nutrition Policy and Promotion, Miscellaneous Publication No. 1528-2011, 2011.

Paul Milgrom. Is sympathy an economic value? philosophy, economics, and the contingent valuation method. In J. A. Hausman, editor, *Contingent valuation: A critical Assessment*, pages 417–441. Elsevier Science Publishers, 1993.

Antony Millner and Geoffrey Heal. Choosing the future: Markets, ethics, and rapprochement in social discounting. *Journal of Economic Literature*, 61(3):1037–87, September 2023.

Frikk Nesje. Cross-dynastic intergenerational altruism. 2024.

William D. Nordhaus. A review of the stern review on the economics of climate change. *Journal of Economic Literature*, 45(3):686–702, 2007.

Debraj Ray. Hedonistic altruism and welfare economics. 2018.

Debraj Ray, Nikhil Vellodi, and Ruqu Wang. Past and future: Backward and forward discounting. *Journal of the European Economic Association*, 22(2):837–875, April 2024. doi: https://doi.org/10.1093/jeea/jvad036.

Kevin Roberts. The theoretical limits to redistribution. *The review of economic studies*, 51(2):177–195, 1984.

Max Roser and Hannah Ritchie. How has world population growth changed over time? *Our World in Data*, 2023. https://ourworldindata.org/population-growth-over-time.

Maria Saez-Marti and Jörgen W. Weibull. Discounting and altruism to future decision-makers. *Journal of Economic Theory*, 122(2):254–266, 2005. ISSN 0022-0531. doi: https://doi.org/10.1016/j.jet.2004.06.003. URL `https://www.sciencedirect.com/science/article/pii/S0022053104001358`.

Samuel Scheffler. *Death and the Afterlife*. Berkeley Tanner Lectures, 2013.

Samuel Scheffler. *Why worry about future generations?* Oxford University Press, 2018.

Itai Sher. How perspective-based aggregation undermines the pareto principle. *Politics, Philosophy & Economics*, 19(2):182–205, 2020.

Nicholas Stern. The economics of climate change. *American Economic Review*, 98 (2):1–37, 2008.

A Proof of Proposition 1

$$
\begin{aligned}
W(c) &= \sum_{t=-\infty}^{\infty} \left(\frac{1}{1+\rho}\right)^t N_t \left(\sum_{\tau=-\infty}^{\infty} \psi_\tau \frac{N_{t+\tau}}{N_t} u(c_{t+\tau})\right) \\
&= \sum_{t=-\infty}^{\infty} \left(\frac{1}{1+\rho}\right)^t \left(\sum_{\tau=-\infty}^{\infty} \psi_\tau N_{t+\tau} u(c_{t+\tau})\right) \\
&= \sum_{t=-\infty}^{\infty} N_t u(c_t) \sum_{\tau=-\infty}^{\infty} \frac{\psi_{t-\tau}}{(1+\rho)^\tau} \\
&= \sum_{t=-\infty}^{\infty} N_t u(c_t) \left(\frac{1}{(1+\rho)^t} \sum_{\tau=-\infty}^{\infty} (1+\rho)^{t-\tau} \psi_{t-\tau}\right) \\
&= \sum_{t=-\infty}^{\infty} N_t u(c_t) \left(\frac{1}{(1+\rho)^t} \sum_{\tau=-\infty}^{\infty} (1+\rho)^{-\tau} \psi_{-\tau}\right) \\
&= \left(\sum_{\tau=-\infty}^{\infty} (1+\rho)^{-\tau} \psi_{-\tau}\right) \sum_{t=-\infty}^{\infty} \frac{N_t}{(1+\rho)^t} u(c_t)
\end{aligned}
$$

Because $\psi_\tau \geq 0$ for all τ and strictly positive for $\tau = 0$, the sum in the brackets is strictly positive (and, by (6), finite). This social welfare function is therefore a multiplication of the social welfare function in (7) by a positive constant, which is a strictly monotone transformation. It follows that they represent the same social preference relation.

B Proof of Proposition 2

By definition of ${}_tN$,

$$
{}_tN_\tau = N_{\tau+t} = (1+n)^{\tau+t} \text{ and } N_t = (1+n)^t
$$

It follows that the vector ${}_tN/N_t$ is the same as the vector N/N_0; to see this, note that it's τ-th entry is given by

$$
\frac{{}_tN_\tau}{N_t} = \frac{(1+n)^{\tau+t}}{(1+n)^t} = (1+n)^\tau = \frac{N_\tau}{N_0}
$$

Thus, for every t and τ, it holds that

$$f_\tau\left(\frac{{}_tN_\tau}{N_t}\right) = f_\tau\left(\frac{N}{N_0}\right)$$

Define

$$\psi_\tau := f_\tau\left(\frac{N}{N_0}\right)\frac{N_0}{N_\tau}$$

Note that, given this definition ψ_τ, dynastic preferences are represented by (4), and, by the assumption of the proposition, (6) holds. Thus, Proposition 2 follows from Proposition 1.

C Time-varying social objectives

In this section, I consider an alternative normative framework in which the social welfare function itself changes over time, and disregards the preferences of past generations. Assume that the social preferences *at time t* are represented by

$$W_t = \sum_{\tau=t}^{\infty}\left(\frac{1}{1+\rho}\right)^\tau U_\tau$$

Note that this social objective disregards the preferences of past generations. This reflects the view that, after their death, people's utilities no longer matter for social welfare.

Even in this framework, there can be an argument that policymakers should act *as if* they cared about the preferences of past generations. This is because, in some situations, such a norm leads to higher welfare in every period.[10]

To illustrate, consider an example that builds on the three-period model in Example 2, and, for simplicity, set $u(c) = c$. Assume that generation (-1) is endowed with one storable consumption good. When policymakers in period 0 commit to carrying out the plans of policymakers at time (-1), then the policymakers at time (-1) face the optimization problem

$$\max_{c_{-1},c_0,c_1}(1+\rho)c_{-1} + (1+(1+\rho)\beta)c_0 + \frac{1}{1+\rho}(1+(1+\rho)\beta)c_1 \text{ s.t. } c_{-1}+c_0+c_1 = 1$$

[10]This idea appears also in Roberts [1984] in the context of dynamic public finance.

If ρ is very large and $\beta > 1$, the solution to this optimization problem is to set $c_0 = 1$ and $c_{-1} = c_1 = 0$. In this case, we have

$$W_{-1} = (1+\rho)\beta + 1 > (1+\rho)$$

$$W_0 = 1$$

$$W_1 = 0$$

Next, consider the case in which the policymakers in period 0 do not commit to respecting the wishes of the policymakers at time (-1). Instead, they forget about the preferences of past generations and pursue policies that maximize welfare from the perspective of time 0. Given c_{-1}, their optimization problem is

$$\max_{c_0, c_1} c_0 + (\beta + \frac{1}{1+\rho})c_1 \text{ s.t. } c_0 + c_1 = 1 - c_{-1}$$

Because $\beta > 1$, the optimal solution sets $c_1 = 1 - c_{-1}$.

Of course, in the absence of a norm for respecting the preferences of past generations, policymakers in period (-1) would take into account that policymakers in period 0 will behave in this way. Their optimization problem becomes

$$\max_{c_{-1}, c_0, c_1} (1+\rho)c_{-1} + \frac{1}{1+\rho}(1 + (1+\rho)\beta)c_1 \text{ s.t. } c_{-1} + c_1 = 1$$

If ρ is sufficiently large, the solution to this optimization problem is to set $c_{-1} = 1$. So, without a norm of respecting the preferences of past generations, we have

$$W_{-1} = (1+\rho)$$

$$W_0 = 0$$

$$W_1 = 0$$

So, in this example, we have that a norm of respecting the preferences of past generations increases welfare in every period. Of course, it is also possible to construct examples in which this norm leads to lower welfare in some periods. For example, if the endowment were given in period 0 rather than in period (-1), policymakers in period 0 would achieve higher welfare in the case in which they are not committed to respecting the preferences of past generations. However, this norm always leads to higher welfare in period (-1). Consequently, we may have something of a tradeoff: the norm of honoring the wishes of past generations may lead to higher welfare in some periods, but not in others. Is this norm to be considered "good"? Without a time-invariant axiology, we cannot say.