The Cross-Sectional Implications of the Social Discount Rate

Maya Eden*

August 26, 2023

Abstract

In this paper, I consider two normative questions: (1) how should policymakers approach tradeoffs that involve different age groups, and (2) at what rate should policymakers discount the consumption of future generations? I demonstrate that, under standard assumptions, these two questions are equivalent: caring more about the future means caring less about the elderly. Even small differences between the social discount rate and the market interest rate can have significant quantitative implications for the relative value placed on the consumption of different age groups.

JEL Classification: D63, E61
Keywords: Overlapping generations, utilitarianism, consequentialism, age-based inequality

*Affiliations: Brandeis University, CEPR, and the Global Priorities Institute, Oxford. Contact address: Maya Eden, Economics Department, Brandeis University, 415 South Street, Waltham, MA 02453; Email: meden@brandeis.edu. I am grateful to Thomas Houlden for outstanding research assistance. I thank Geir Asheim, Matthias Doepke, Ben Eden, Luis Mota Freitas, Mikhail Golosov, Chad Jones, Paul Kelleher, Pete Klenow, Kevin Kuruc, Antony Millner, John Quiggin, Itai Sher and three anonymous referees for helpful comments. This paper has also benefited from comments by seminar participants at the Chicago Fed, Iowa State University, the CEPR Macroeconomics and Growth Programme Online Meeting, the Oxford Workshop on Global Priorities Research, the Atlanta Fed, the University of Zurich, the University of Connecticut and the workshop on Fairness and Redistribution at St. Gallen.

1 Introduction

The social discount rate is a crucial parameter in policy decision-making, as it reflects how society values the interests of future generations. However, there is an ongoing debate about how to set it.[1] Nordhaus [2007] proposes an annual social discount rate of 6%, based on long-run estimates of the return to capital. At the same time, Stern [2008] suggests a rate of 1.5%, based on considerations of intergenerational equity. These differences have significant implications for the social cost of carbon emissions, among other policy areas.

Another normative question that has received less attention is how policymakers should approach tradeoffs involving different age groups. This question became particularly pressing during the Covid-19 pandemic, which posed different risks for older and younger people. In this context, the literature mostly adopts a utilitarian approach, in which the flow utility of all individuals is given equal weight in social welfare calculations (e.g. Hall et al. [2020] and Giagheddu and Papetti [2023]).

In this paper, I demonstrate that these two normative questions are closely related. Specifically, I show that, under standard assumptions, concern for the future implies concern for the young. A social welfare criterion that values future generations as much as current generations must also prioritize the interests of the young over the interests of the elderly.

I begin by examining the standard discounted-utilitarian framework, in which both individual and social preferences take the discounted utility form. For reasonable calibrations, assigning equal weight to future generations implies that it is optimal to set the consumption of 80-year-olds to be equal to one-tenth of the consumption of 20-year-olds. An equitable distribution of consumption across age groups is optimal only if the social welfare function discounts the lifetime utilities of future generations at the same rate that individuals discount their own future flow utilities. This reveals a tension between concerns for intergenerational equity and concerns for consumption equality across age groups in a given period.

This tension persists under weaker assumptions about individual and social preferences. I consider a setup with arbitrary individual preferences and arbitrary social preferences that satisfy just one condition: the social preferences are *non-paternalistic*,

[1] See Greaves [2017] and Millner and Heal [2023] for recent reviews. See also Baumol [1968], Arrow et al. [2013], Gollier and Hammitt [2014], Kelleher [2017] and Drupp et al. [2018].

in the sense that they respect each individual's choice of how to allocate consumption over their lifecycle. I consider the social desirability of small perturbations along an arbitrary growth trajectory. In this general setup, the rate at which society discounts the consumption of future generations places an upper bound on the relative marginal social welfare weights of older and younger people. Even small differences between the social discount rate and the market interest rate can have significant quantitative implications. For example, if the market interest rate is 6% (as in Nordhaus [2007]), then a 1.5% social discount rate (as in Stern [2008]) implies that it is better to increase the consumption of a 20-year-old by \$6 than to increase the consumption of an 80-year-old by \$100.

This paper adds to the literature on social discounting (see Millner and Heal [2023] for a recent review). The social discount rate is typically analyzed in the context of non-overlapping generations models, with notable exceptions being Calvo and Obstfeld [1988], Quiggin [2012], Schneider et al. [2012], and Fleurbaey and Zuber [2015]. Calvo and Obstfeld [1988] and Schneider et al. [2012] consider a discounted-utilitarian social objective in an overlapping generations economy, and demonstrate that transfers between age groups may be necessary for implementing the social optimum. This paper expands on this literature by studying the quantitative implications of a low social discount rate for the optimal distribution of consumption across age groups.

The theoretical aspects of this paper are also connected to Quiggin [2012], who uses an overlapping generations framework to examine the relationship between intertemporal social preferences and social preferences regarding the distribution of consumption across age groups. Quiggin finds that placing equal weight on the flow utility of each age group implies a social objective of maximizing the undiscounted sum of all current and future flow utilities. However, as I illustrate here, this result relies heavily on the assumption that individuals have a zero subjective discount rate (see also Schneider et al. [2012]).

Similar to this paper, Fleurbaey and Zuber [2015] (section 6) decompose the social discount rate into a within-person component and an across-people component. In Fleurbaey and Zuber, the across-people component reflects society's willingness to transfer resources between two individuals at their respective births, rather than the relative desirability of increasing the consumption of different age groups (as in this paper). Based on their decomposition, Fleurbaey and Zuber argue that the market interest rate is largely irrelevant for social discounting over long time horizons. This paper reaches the opposite conclusion: at any horizon, there is a relationship between the social discount rate, the market interest rate, and the cross-sectional distributional weights of different age groups.

Finally, this paper is related to Farhi and Werning [2007] and Barrage [2018], who study the policy implications of a social discount rate that is lower than the market interest rate. These papers do not derive any implications for the relative marginal social welfare weights of different age groups, which is the focus of this paper.[2] However, similar to this paper, they illustrate that deviations of the social discount rate from the market interest rate have far-reaching implications for optimal policy design.

2 Setup

Time is infinite and discrete, and indexed by $t \in \mathbb{Z}$. The current period is denoted by $t = 0$. In each period a new generation is born. In their first period of life, people are considered to be of age 0. People die at age $T > 0$. Note that the assumption that $T > 0$ implies that generations are overlapping. For example, if $T = 1$, then people live for two periods: in the first period they are aged 0 and in the second period they are aged 1. To simplify the analysis, I assume that all cohorts are of the same size.

Given the overlapping generations structure, it is convenient to introduce two different ways to describe a consumption allocation. In some cases, it is useful to think of a consumption allocation as a list of how much each age group consumes in each time period. To denote the cross-sectional distribution of consumption at time t, I use $\{c_{a,t}\}_{a=0}^{T}$, where $c_{a,t}$ is the consumption of a person aged a at time t. Note that I assume that all people from the same generation have the same consumption sequence, abstracting away from any consumption inequality within generations. The upper-case variable C_t denotes aggregate consumption at time t. Generally, I use t as a subscript to indicate time-t variables.

In other cases, it is more useful to describe a consumption allocation by the consumption stream of each generation. To indicate variables belonging to the cohort born at date t, I use the superscript t. The vector $\mathbf{c}^t = (c_0^t, ..., c_T^t)$ is the lifetime con-

[2] Farhi and Werning [2007] and Barrage [2018] consider a model of consecutive generations, in which only one generation is alive in each period. Consequently, there is no cross-sectional variation in age.

sumption stream of each person in generation t. It should be noted that this notation is somewhat redundant, as it must hold that $c_a^t = c_{a,t+a}$. However, it can be helpful for simplifying the analysis, as some identities are clearer when adopting a lifetime perspective while others are clearer when adopting a cross-sectional perspective.

I use $\mathbf{c} = \{\{c_{a,t}\}_{a=0}^{T}\}_{t=-\infty}^{\infty}$ to denote a complete description of how consumption is allocated across time and across generations. The table below summarizes the notation used for describing consumption allocations in this paper.

Notation	Object	Description
$\mathbf{c}$	Matrix	Consumption allocations across time and people
$\mathbf{c}^t$	Vector	The lifetime consumption stream of people born in period t
c_a^t	Scalar	The consumption of people born in period t at age a
C_t	Scalar	Aggregate consumption in period t
$c_{a,t}$	Scalar	The consumption of people aged a in period t

3 The discounted utilitarian model

As a starting point, it is useful to consider the standard framework in which both individuals and the planner are discounted-utility maximizers. Calvo and Obstfeld [1988] and Schneider et al. [2012] illustrate that, under these assumptions, the social rate of pure time preference not only determines the optimal path of aggregate consumption, but also the optimal distribution of consumption across age groups. This section discusses the quantitative implications of this result.

In the standard discounted-utilitarian framework, individual preferences over lifetime consumption streams, $(c_0, ..., c_T)$, are represented by a time-separable form,

$$u(c_0, ..., c_T) = \sum_{a=0}^{T} \left(\frac{1}{1+\rho}\right)^a \frac{c_a^{1-\gamma}}{1-\gamma} \tag{1}$$

where the parameter $\gamma > 0$ is the inverse of the intertemporal elasticity of substitution (which is assumed to be constant), and ρ is the subjective discount rate.

In addition, social preferences are assumed to be represented by a social welfare function of the discounted-utilitarian (DU) form,

$$W^{DU}(\mathbf{c}) = \sum_{t=-\infty}^{\infty} \left(\frac{1}{1+\rho^s} \right)^t u(c_0^t, ..., c_T^t) \tag{2}$$

where $(c_0^t, ..., c_T^t)$ is the lifetime consumption stream of the people born in period t. Here, ρ^s (with a superscript s) is the rate at which the planner discounts the lifetime utilities of future generations.[3] Stern [2008] refers to this rate as the "pure time discount rate". Here, I will refer to it simply as the "generational discount rate".

Of course, the infinite sum in (2) converges only when the sequence of utilities of past generations converges to 0 sufficiently fast and the sequence of utilities of future generations grows sufficiently slowly. However, we can assume this representation for the limited set of consumption sequences for which we obtain convergence.

The set of feasible consumption allocations is denoted $\mathcal{F}$. A typical element of $\mathcal{F}$ is a matrix $\mathbf{c}$ that satisfies the economy's feasibility constraints. I assume throughout that past consumption is immutable, so that different elements of $\mathcal{F}$ differ only in the consumption allocations in periods 0 onwards. In addition, I assume that feasibility depends only on aggregate consumption, and not on how consumption is distributed across age groups.[4] For example, $\mathcal{F}$ may contain all of the consumption allocations that satisfy the economy's aggregate resource constraints. It is convenient to denote by $\mathcal{F}^* \subset \mathbb{R}_+^\infty$ the set of feasible aggregate consumption sequences from period 0 onwards. A typical element of $\mathcal{F}^*$ is a sequence of aggregate consumption levels, $\{C_t\}_{t=0}^\infty$.

[3] The assumption here is that the utility given by expression 1 is the welfare-relevant measure of utility for social welfare purposes (and not, say, the exponent of (1) which represents the same individual preferences).

[4] This is a simplifying assumption; of course, if redistribution requires distortionary taxation, then some cross-sectional distributions of consumption may not be feasible. Furthermore, this assumption rules out situations in which the planner is bound to honor commitments that were previously made to older people. For example, if older people were promised a certain retirement income, then this constrains the set of consumption distributions that are feasible. My point here is not to argue that such constraints are irrelevant. Rather, it is to learn about the properties of the social objective by answering the question: if the planner were not bound by any such obligations and was able to institute lump-sum taxes and transfers, what would be the consumption allocation that he would choose?

Consider the following proposition.

Proposition 1. *Assume that* $\mathbf{c}^*$ *is a unique optimal policy,*

$$\mathbf{c}^* = \arg\max W^{DU}(\mathbf{c}) \ \textit{s.t.} \ \mathbf{c} \in \mathcal{F}$$

and let $\{C_t^*\}_{t=0}^{\infty}$ *denote the sequence of aggregate consumption implied by this optimal policy.*

(a) The sequence of aggregate consumption levels, $\{C_t^*\}_{t=0}^{\infty}$*, solves*

$$\{C_t^*\}_{t=0}^{\infty} = \arg \max_{\{C_t\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \left(\frac{1}{1+\rho^s}\right)^t \frac{C_t^{1-\gamma}}{1-\gamma} \ \textit{s.t.} \ \{C_t\}_{t=0}^{\infty} \in \mathcal{F}^*$$

(b) In each period, for every a, a'*, it holds that*

$$\frac{c_{a',t}^*}{c_{a,t}^*} = \left(\frac{1+\rho^s}{1+\rho}\right)^{\frac{a'-a}{\gamma}}$$

The proof of this proposition can be found in the appendix, together with other omitted proofs. The proposition establishes that the optimal policy can be solved in two steps: the first is to solve for the optimal aggregate consumption path in a standard representative agent framework, and the second is to divide aggregate consumption across age groups in fixed proportions.

This proposition establishes that the generational discount rate, ρ^s, simultaneously affects both the optimal path of aggregate consumption, and the optimal distribution of consumption across age groups. In what follows, I consider two normative benchmarks:

- Setting $\rho^s = \rho$, so that the generational discount rate is equal to the individuals' rate of pure time preference (as in Nordhaus [2007]).
- Setting $\rho^s \approx 0$, to reflect equal concern for all generations (as in Stern [2008]).

In the first normative benchmark, for each two ages, a and a', it holds that:

$$\frac{c_{a',t}^*}{c_{a,t}^*} = \left(\frac{1+\rho^s}{1+\rho}\right)^{\frac{a'-a}{\gamma}} = \left(\frac{1+\rho}{1+\rho}\right)^{\frac{a'-a}{\gamma}} = 1$$

This indicates that, under this benchmark, an equitable allocation of consumption is optimal in each period.

In the second normative benchmark, for every two ages, a' and a, it holds that:

$$\frac{c^*_{a',t}}{c^*_{a,t}} = \left(\frac{1+\rho^s}{1+\rho}\right)^{\frac{a'-a}{\gamma}} \approx \left(\frac{1+0}{1+\rho}\right)^{\frac{a'-a}{\gamma}} = \left(\frac{1}{1+\rho}\right)^{\frac{a'-a}{\gamma}}$$

Note that ρ and γ are not normative choices; their appropriate values depend on individual preferences. Empirical estimates for these parameters can be found in the literature; see Table 1 for some examples.

Table 1: Empirical estimates of individual time preferences

	$\gamma \leq 0.5$	$\gamma \in (1,2)$	$\gamma \geq 10$
$\rho \leq 0$			Hall [1988]
$\rho = 0.04$	Gourinchas and Parker [2001]	Kydland and Prescott [1982]	Best et al. [2020]
$\rho \approx 0.1$		Lawrance [1991]	
$\rho > 0.2$	Andreoni and Sprenger [2012]	Shapiro [2005]	

Table 2: The optimal consumption ratio, $c^*_{80,t}/c^*_{20,t}$, when $\rho^s = 0$

	$\gamma = 0.5$	$\gamma = 1$	$\gamma = 10$
$\rho = 0$	1	1	1
$\rho = 0.04$	0.01	0.1	0.79
$\rho = 0.1$	0.00	0.00	0.56
$\rho = 0.2$	0.00	0.00	0.33

Note: Gray cells correspond to the estimates in the literature referenced in Table 1.

Table 2 presents the optimal consumption ratio of an 80 year old and a 20 year old when $\rho^s = 0$, for different combinations of ρ and γ. If $\rho = 0$, then the second normative benchmark, $\rho^s = 0$, coincides with the first normative benchmark, $\rho^s = \rho$. In this case, an equitable allocation of consumption is optimal in every period. Estimates of $\rho \approx 0$ are broadly consistent with the literature that finds an intertemporal elasticity of substitution close to 0.[5] However, most direct empirical estimates suggest values of ρ that are significantly above 0. For example, if $\rho = 0.04$ and $\gamma = 1$, then it is optimal to set the consumption of an 80 year old to 10% of the consumption of a 20 year old. This means that, at the *optimal* allocation, the consumption of a 20 year old should be about 10 times that of an 80 year old. Micro-based estimates of ρ and γ result in even more extreme levels of inequality, with the optimal consumption of 20 year olds over 100 times that of 80 year olds.

The intuition for these results is as follows. When individuals discount their own future consumption, it is possible to increase their lifetime utilities by having them consume more earlier in life and less later in life. A policy that tilts consumption towards the young in each period reduces the lifetime utility of the current old, but increases the lifetime utilities of everyone else. This tradeoff is acceptable to a planner with a zero generational discount rate, who places equal weight on the well-being of future generations (of which there are infinitely many). However, even in the long-run, this policy results in extreme consumption inequality between older and younger people.

Of course, rejecting a zero generational discount rate has other problematic implications. It means that, as a society, we place less weight on the well-being of future generations. This implies that it is socially desirable to set policies in a way that neglects their long-term implications, even when those implications are dire.

To illustrate this tension, I consider three examples:

- **Endowment economy with constant consumption**. The set of feasible per-capita consumption sequences, $\mathcal{F}^*$, contains a single element, $\{C_t\}_{t=0}^{\infty}$, where $C_t = C_0 = C$ for all t.

[5] The seminal paper by Hall [1988] estimates an elasticity of substitution that is close to 0 ($\gamma \geq 10$). While the paper does not provide an estimate of ρ, by the Ramsey equation (Ramsey [1928]), $r \approx \rho + \gamma g$, where r is the market interest rate and g is the growth rate of consumption. Hence, $\rho \approx r - \gamma g$. For a growth rate of $g = 0.02$ and an interest rate $r \leq 20\%$, it follows that $\rho \leq 0.2 - 10 * 0.02 = 0$.

- **The neoclassical growth model**. The economy's aggregate resource constraints are given by a standard neoclassical growth model. The initial capital stock is chosen to match an average rate of return to capital of 6%. For details of this model and its calibration, see Appendix D.

- **A climate-economy model**. The economy's aggregate resources are given by the model in Golosov et al. [2014]. The calibration matches the benchmark calibration in that paper.[6] The model augments a standard neoclassical growth model by explicitly modeling energy inputs. There are three possible sources of energy for production: oil, which is a non-renewable resource that is cheap to extract; coal, which is plentiful but more expensive to extract; and green energy, which is expensive to produce. Production using oil or coal results in carbon emissions that reduce future productivity through a climate externality. In this model, the energy mix used today affects the production possibilities of future generations.

For the purpose of these calibrations, I assume that $\gamma = 1$ and $\rho = 4\%$.[7] Given these parameters, the first normative benchmark corresponds to $\rho^s = 4\%$. For the second normative benchmark, I follow Stern [2008] and set $\rho^s = 0.1\%$.[8][9]

Figure 1 plots the optimal consumption paths for 20 year-olds and 80 year-olds in the endowment economy, neoclassical growth model, and climate-economy model. In the endowment economy, the generational discount rate determines only the optimal distribution of consumption across age groups, and not the path of aggregate consumption, which is fixed by assumption. While a generational discount rate of 4% implies that an equitable distribution is optimal, a value of 0.1% prescribes extreme inequality between age groups in each period.

[6] In particular, the calibration assumes 2% productivity growth; full depreciation of capital over a decade; and $\gamma = 1$ (log utility). For details about the other parameters, such as those governing the accumulation of atmospheric CO2, see Golosov et al. [2014]. I compute the optimal path of consumption using the supplemental code provided by Barrage [2014].

[7] The specification $\gamma = 1$ is consistent with the calibration of Golosov et al. [2014]. Because the model in Golosov et al. [2014] is one of an infinitely-lived representative agent, there is no counterpart to ρ (but only a counterpart to the social rate of pure time preference, ρ^s, which they vary).

[8] Stern [2008] sets a small positive social rate of pure time preference to capture a small probability of extinction.

[9] Given $\rho > 0$ and $\gamma = 1$, the discounted-utilitarian objective is well-defined for any constant growth rate of aggregate consumption. In the problems that I consider, long-run consumption growth is bounded, and hence the discounted-utilitarian objective is well-defined on the feasible set.

Figure 1: Optimal consumption paths

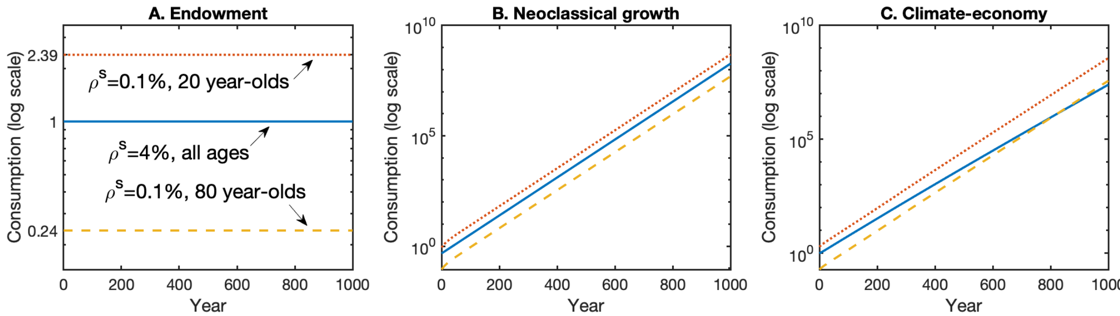

Note: The individual preference parameters are $\rho = 0.04$ and $\gamma = 1$. The solid blue lines represent the optimal consumption paths of a social planner with $\rho^s = \rho = 0.04$. These consumption paths are equitable across age groups. The dashed yellow lines and the dotted red lines assume $\rho^s = 0.001$. The dotted red lines represent the optimal consumption of 20 year-olds and the dashed yellow lines represent the optimal consumption of 80 year-olds. When $\rho^s = 0.001$, the optimal consumption of 80 year-olds is always one-tenth of the optimal consumption of 20 year-olds (see Table 2). The three examples differ in their assumptions about the set of feasible alternatives.

In the neoclassical growth model, the generational discount rate affects both the optimal consumption path and the distribution of consumption across age groups. A more patient planner will choose a higher capital-output ratio along the balanced growth path, financed by a short-term drop in the consumption of older people (but an increase in the consumption of younger people). Along the balanced growth path, aggregate consumption is higher, but consumption inequality is higher as well. Even in the long-term, the consumption of 80 year-olds is lower compared to the optimal policy of the impatient planner.

In the climate-economy model, the generational discount rate not only affects long-term levels of consumption, but also the long-term growth rate. This is because it affects the rate at which natural resources are depleted. As a result, a lower generational discount rate leads to higher consumption for each age group in the distant future. Starting around 800 years from now, the optimal policy of the patient planner features higher consumption for each age group, with the gap widening exponentially over time. For example, 2000 years from now, the consumption of 80 year-olds would be 13 times higher if we adopt a low discount rate than if we adopt a high discount rate. Average consumption would be 56 times higher.

One might be tempted to choose different values of ρ^s in each of these scenarios. For example, the choice of $\rho^s = 4\%$ in the endowment economy seems appealing, as it implies that an equitable distribution is optimal in each period. At the same time, in the climate-economy model, there is a real tradeoff between inequality and long-run prosperity; somebody who is particularly concerned about the well-beings of future generations might opt for $\rho^s = 0.1\%$. However, making the choice of ρ^s dependent on the feasible set of alternatives goes against the fundamental principle of the consequentialist, discounted-utilitarian framework. In this framework, the social objective is always to maximize "welfare", regardless of the feasible set of alternatives. In other words, if you support an equal distribution of consumption in the endowment economy, then you must also accept a faster depletion of natural resources. Similarly, if you advocate for more sustainable climate policies, then you must also support extreme inequality across age groups.

4 Perturbations along a given growth trajectory

The previous section illustrates that, in the standard discounted-utilitarian framework, a tension arises between intergenerational equity considerations and concerns for equality across age groups. This finding prompts the question: is this tension an artifact of the particular assumptions of the discounted-utilitarian framework? The assumption of discounted-utilitarian preferences is unrealistic as it ignores non-homotheticities, time-inseparabilities, and age-dependent consumption needs. Additionally, there are alternative social welfare functions to consider, such as the Prioritarian social welfare function which is averse to inequality in lifetime utilities.[10] Can some combination of more realistic individual preferences and more sophisticated social preferences reverse these uncomfortable conclusions?

While finding the optimal policy requires making assumptions about the functional forms of individual and social preferences, more-general results can be obtained by considering the social desirability of small perturbations along a given growth trajectory. These results suggest that the tension highlighted by the discounted-utilitarian case extends to a large class of models.

To illustrate, I start with a simple example in which there are only two generations, A and B, and three time periods, -1, 0, and 1. Generation A is born in period -1, and generation B is born in period 0. Both generations live for two periods. Suppose we have a benchmark allocation in which consumption is constant over the lifecycle

[10] See, for example, Adler [2019].

and across generations (similar to the optimal policy of the endowment economy when $\rho^s = \rho$). The social discount rate, r^s, is defined based on the indifference condition:

$$\left(\begin{array}{c|ccc} & t=-1 & t=0 & t=1 \\ \hline A & c & c & \\ B & & c & c{+}1{+}r^s \end{array}\right) \sim \left(\begin{array}{c|ccc} & t=-1 & t=0 & t=1 \\ \hline A & c & c+1 & \\ B & & c & c \end{array}\right) \tag{3}$$

where $\sim$ denotes the social indifference relation. The shaded cells highlight the differences between the left hand side and the right hand side. Here, r^s is defined such that increasing the consumption of old people from generation B by $1 + r^s$ units is welfare-equivalent to increasing the consumption of old people from generation A by 1 unit.

It may be useful to point out that, in the discounted-utilitarian framework, $r^s \approx \rho^s$ (provided that one unit is a sufficiently small quantity). As, in this example, consumption is constant across time, discounting the consumption of the future old relative to the consumption of the current old reflects the social discounting of the utilities of future generations. However, note that the social discount rate, r^s, is defined for a much broader class of social preference relations, including preferences that do not take the discounted-utilitarian form.

It turns out that r^s is fully determined by the social preferences over allocations of consumption across age groups at $t = 0$, and the social preferences over the consumption streams of generation B. To see this, let M be such that[11]

$$\left(\begin{array}{c|ccc} & t=-1 & t=0 & t=1 \\ \hline A & c & c+1 & \\ B & & c & c \end{array}\right) \sim \left(\begin{array}{c|ccc} & t=-1 & t=0 & t=1 \\ \hline A & c & c & \\ B & & c+M & c \end{array}\right) \tag{4}$$

The parameter M is defined based on the condition that, in period 0, a one-unit increase in the consumption of the old is welfare equivalent to an M-unit increase in the consumption of the young. If $M = 1$, then the social preference relation is "age-blind": increasing the consumption of the old is just as valuable as increasing the consumption of the young. In contrast, if $M < 1$, then it is better to increase the consumption of the young by one unit than to increase the consumption of the old by one unit.

Next, define R so that[12]

$$\left(\begin{array}{c|ccc} & t=-1 & t=0 & t=1 \\ \hline A & c & c & \\ B & & c+M & c \end{array}\right) \sim \left(\begin{array}{c|ccc} & t=-1 & t=0 & t=1 \\ \hline A & c & c & \\ B & & c & c+MR \end{array}\right) \tag{5}$$

The value of R depends on the social ranking of the consumption sequences for generation B. Usually, these preferences correspond to generation B's own preferences over its consumption stream. In this case, R is determined by the condition that generation B is indifferent between consuming an additional M units in period 0 and consuming an additional MR units in period 1. If members of generation B discount their own future consumption, then it must hold that $R > 1$.

[11] The existence of an M that satisfies this condition is generally not guaranteed. However, it holds for M sufficiently small, under the assumption that the social preference relation can be represented by a differentiable social welfare function that is strictly increasing in all consumption levels. The same is true for R, which is introduced later.

[12] See footnote 4.

By the transitivity of the indifference relation, (3) and (4) imply that

$$\left(\begin{array}{c|ccc} & t=-1 & t=0 & t=1 \\ \hline A & c & c & \\ B & & c & c{+}1{+}r^s \end{array}\right) \sim \left(\begin{array}{c|ccc} & t=-1 & t=0 & t=1 \\ \hline A & c & c & \\ B & & c+M & c \end{array}\right)$$

and, combining the above condition with (5) yields

$$\left(\begin{array}{c|ccc} & t=-1 & t=0 & t=1 \\ \hline A & c & c & \\ B & & c & c{+}1{+}r^s \end{array}\right) \sim \left(\begin{array}{c|ccc} & t=-1 & t=0 & t=1 \\ \hline A & c & c & \\ B & & c & c+MR \end{array}\right)$$

This establishes that increasing the consumption of old people in period 1 by $1+r^s$ units is welfare-equivalent to increasing their consumption by RM units. Under standard monotonicity assumptions, this implies that

$$1 + r^s = RM \tag{6}$$

It is worth highlighting two special cases.

- **Age-blind preferences:** $M = 1$**.** This normative choice means that, in each period, the consumption of the old matters just as much as the consumption of the young.

- **Birthdate-blind preferences:** $r^s = 0$**.** This normative choice means that the timing of birth is irrelevant; the consumption of an old person from generation B matters exactly as much as the consumption of an old person from generation A.

A social preference relation that is both age-blind and birthdate-blind must satisfy

$$1 + r^s = RM \Rightarrow 1 + 0 = R \cdot 1 \Rightarrow R = 1$$

This implies that the social preference relation must be paternalistic, in the sense it disregards people's preferences over their own consumption streams. To illustrate, consider the following two options: (a) increasing the consumption of generation B by one unit when young or (b) increasing its consumption by $1+\epsilon$ units when old. If $R = 1$ then, for any $\epsilon > 0$, the welfare gains from option (b) exceed the welfare gains from option (a). However, empirically, people tend to require some compensation for delaying consumption; this suggests that, for some $\epsilon > 0$ sufficiently small, generation B prefers option (a) to option (b). The result is a violation of the Pareto condition: the planner prefers option (b), even though everyone prefers option (a).

The implication of (6) is that, if we reject paternalism, then we cannot have a social preference relation that is both age-blind and birthdate-blind.

5 A decomposition of social discount rates

As in the example in the previous section, the question of how to make tradeoffs across generations and across time can be decomposed into two simpler sub-questions: first, how to make tradeoffs across age groups within a time period ("M"), and, second, how to make tradeoffs within a generation across time ("R").

In what follows, I generalize this result to settings with arbitrary consumption paths, multiple age groups and multiple generations. To discuss the social desirability of small perturbations to the consumption allocation, it is convenient to fix a consumption allocation, $\bar{\mathbf{c}}$, and assume that, in an open neighborhood of $\bar{\mathbf{c}}$, social preferences are representable by a differentiable social welfare function, W. Note that, in what follows, I do not assume that W takes the discounted-utilitarian form, or any other functional form. At this point, I assume only that W is differentiable (later, I will add an assumption that W is non-paternalistic).

Let $W_{a,t}$ denote the increase in welfare resulting from an incremental increase in the consumption of a person aged a, at time t, given the consumption allocation $\bar{\mathbf{c}}$.

$$W_{a,t} := \frac{\partial W(\bar{\mathbf{c}})}{\partial c_{a,t}} \tag{7}$$

Note that $W_{a,t}$ is the social marginal welfare weight of a person aged a at time t.

5.1 Generalizations of r^s, M and R

As in (3), the social discount rate, $r^s_{a,t}$, is defined based on the social indifference between two perturbations, which are detailed in Table 3. In the first perturbation, we increase the consumption of people aged a by $(1 + r^s_{a,t})^t$ units at time t. In the second perturbation, we increase the consumption of people of the same age by 1 unit today.

Table 3: Welfare-equivalent perturbations involving $r^s_{a,t}$

	Added consumption	Age	Time period	Generation
Perturbation 1	$(1 + r^s_{a,t})^t$	a	t	$t - a$
Perturbation 2	1	a	0	$-a$

The welfare gains from the first perturbation are approximately $(1 + r^s_{a,t})^t W_{a,t}$, and the welfare gains from the second perturbation are approximately $W_{a,t}$. For these welfare gains to be the same, $r^s_{a,t}$ must be defined as

$$r^s_{a,t} := \left(\frac{W_{a,0}}{W_{a,t}}\right)^{\frac{1}{t}} - 1$$

Note that the social discount rate, $r^s_{a,t}$, expresses the social marginal rate of substitution between $c_{a,0}$ and $c_{a,t}$ as an average per period rate of return.

Next, I construct a set of analogs for "M". As in (4), $M_{a,a',t}$ is defined based on the social indifference between two perturbations that affect two different age groups in a given time period (Table 4). In the first perturbation, we add one unit of consumption to people aged a' at time t. The welfare gains from this perturbation are approximately $W_{a',t}$. In the second perturbation, we add $M_{a,a',t}$ units to the people aged a at time t. The welfare gains from this perturbation are $W_{a,t}M_{a,a',t}$. For the welfare gains from these two perturbations to be the same, $M_{a,a',t}$ must be defined as

$$M_{a,a',t} := \frac{W_{a',t}}{W_{a,t}}$$

The value of $M_{a,a',t}$ is the ratio of the social marginal welfare weight of a person aged a' and the social marginal welfare weight of a person aged a, at time t.

Table 4: Welfare-equivalent perturbations involving $M_{a,a',t}$

	Added consumption	Age	Time period	Generation
Perturbation 1	1	a'	t	$t-a'$
Perturbation 2	$M_{a,a',t}$	a	t	$t-a$

Finally, I construct a set of analogs for "R". Generalizing (5), $R_{a,a',t}$ is defined based on the social indifference between two perturbations that affect the same generation in different periods of life (Table 5). In the first perturbation, at time t, we add a unit of consumption to the people aged a. The welfare gains from this perturbation are $W_{a,t}$. In the second perturbation, we wait until the people are aged a' (which happens in period $t+(a'-a)$), and then increase their consumption by $R_{a,a',t}$ units. The welfare gains from this perturbation are $R_{a,a',t}W_{a',t+(a'-a)}$. Social indifference between these two perturbations implies that $R_{a,a',t}$ must be defined as

$$R_{a,a',t} := \frac{W_{a,t}}{W_{a',t+(a'-a)}}$$

Table 5: Welfare-equivalent perturbations involving $R_{a,a',t}$

	Added consumption	Age	Time period	Generation
Perturbation 1	1	a	t	$t-a$
Perturbation 2	$R_{a,a',t}$	a'	$t+(a'-a)$	$t-a$

5.1.1 Examples

Before stating the main result, it is useful to illustrate the concepts of $r^s_{a,t}$, $M_{a,a',t}$ and $R_{a,a',t}$ using the three examples from section 3. Note that all three variables depend on the benchmark consumption allocation, $\bar{\mathbf{c}}$ (which is used in the definition of the social marginal welfare weights in (7)). For the purpose of these illustrations, I assume that the benchmark consumption paths are the ones chosen by a discounted-utilitarian policymaker for whom $\rho^s = \rho$:[13]

$$\bar{\mathbf{c}} = \arg\max_{\mathbf{c}} \sum_{t=0}^{\infty} \left(\frac{1}{1+\rho}\right)^t u(\mathbf{c}^t) \text{ s.t. } \mathbf{c} \in \mathcal{F} \tag{8}$$

Figure 2: $r^s_{a,t}$ for the three examples in section 3

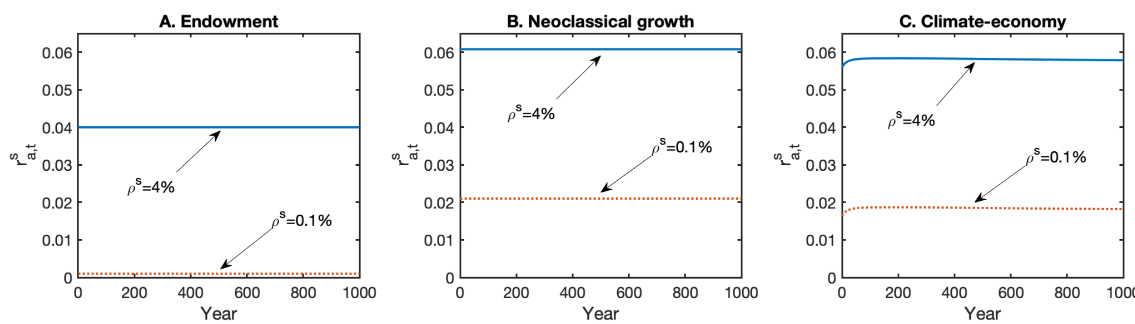

Note: For each of the three examples, the benchmark consumption path, $\bar{\mathbf{c}}$, corresponds to the optimal policy of a discounted utilitarian planner with $\rho^s = 4\%$ (as in equation 8). The social discount rates, $r^s_{a,t}$, are independent of the age, a, along the consumption paths $\bar{\mathbf{c}}$.

Figure 2 plots the social discount rates in the examples in section 3. In these examples, a higher generational discount rate implies a higher social discount rate. While the generational discount rate is a fixed parameter of the social welfare function, the social discount rate depends on the path of aggregate consumption.

Figure 3 illustrates the values of $M_{20,80,t}$ for the three examples in section 3. Note that, as the benchmark allocation $\bar{\mathbf{c}}$ is chosen by a planner for whom $\rho^s = \rho$, by Proposition 1, it must feature an equitable distribution of consumption across age groups. If the generational discount rate is $\rho^s = \rho$, then this distribution is socially

[13] This would be the growth trajectory that is implemented by a policymaker that adopts his constituents' time preferences. Note, however, that this allocation is not necessarily attainable as a *laissez-faire* equilibrium. For example, in the climate-economy model, implementing the policy-maker's optimal policy requires setting carbon taxes to correct for climate externalities.

Figure 3: $M_{20,80,t}$ for the three examples in section 3

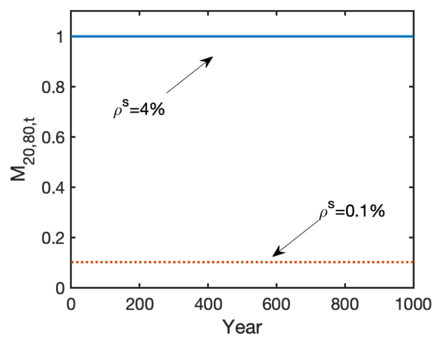

Note: For each of the three examples, the benchmark consumption path, $\bar{\mathbf{c}}$, corresponds to the optimal policy of a discounted utilitarian planner with $\rho^s = 4\%$ (as in equation 8). In these examples, the value of $M_{20,80,t}$ does not depend on the path of aggregate consumption, but only on the distribution of consumption across age groups. Consequently, $M_{20,80,t}$ is constant across time and the same for all three examples.

optimal. In this case, a marginal increase in the consumption of an 80 year-old is just as valuable as a marginal increase in the consumption of a 20 year-old. However, if $\rho^s = 0.1\%$, then it is better to increase the consumption of a 20 year-old by \$10 than to increase the consumption of an 80 year-old by \$100.

Figure 4: $R_{20,80,t}$ for the three examples in section 3

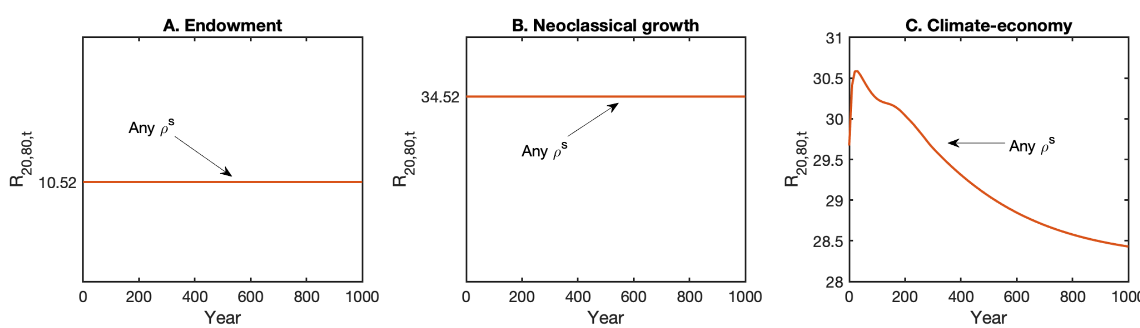

Note: For each of the three examples, the benchmark consumption path, $\bar{\mathbf{c}}$, corresponds to the optimal policy of a discounted utilitarian planner with $\rho^s = 4\%$ (as in equation 8). In these examples, $R_{20,80,t}$ does not depend on the generational discount rate, ρ^s.

Finally, figure 4 illustrates that, in the discounted-utilitarian framework, the values of $R_{a,a',t}$ are independent of the generational discount rate. Regardless of ρ^s, the social preference relation evaluates changes to an individual's lifetime consumption in the same way that the individual would.

5.2 Decomposition

The following proposition generalizes equation (6).

Proposition 2. *Let* $0 \leq a < a' \leq T$ *and* $n > 0$*. For* $t = (a' - a)n$*, it holds that*

$$1 + r^s_{a',t} = \left(\prod_{\tau=0}^{n-1} M_{a,a',(a'-a)\tau} \right)^{\frac{1}{t}} \left(\prod_{\tau=0}^{n-1} R_{a,a',(a'-a)\tau} \right)^{\frac{1}{t}} \tag{9}$$

Proof. By definition,

$$(1 + r^s_{a',t})^t = \frac{W_{a',0}}{W_{a',t}} = \frac{W_{a',0}}{W_{a',(a'-a)n}} \tag{10}$$

The right-hand side can be rewritten as

$$\frac{W_{a',0}}{W_{a',(a'-a)n}} = \underbrace{\left(\frac{W_{a',0}}{W_{a,0}}\right)}_{M_{a,a',0}} \underbrace{\left(\frac{W_{a,0}}{W_{a',a'-a}}\right)}_{R_{a,a',0}} \underbrace{\left(\frac{W_{a',a'-a}}{W_{a,a'-a}}\right)}_{M_{a,a',a'-a}} \cdots \underbrace{\left(\frac{W_{a',(a'-a)(n-1)}}{W_{a,(a'-a)(n-1)}}\right)}_{M_{a,a',(a'-a)(n-1)}} \underbrace{\left(\frac{W_{a,(a'-a)(n-1)}}{W_{a',(a'-a)n}}\right)}_{R_{a,a',(a'-a)(n-1)}}$$

To see this, note that the denominator of $M_{a,a',\tau}$ cancels out with the numerator of $R_{a,a',\tau}$, and the denominator of $R_{a,a',\tau}$ cancels out with the numerator of $M_{a,a',\tau+(a'-a)}$. Writing the above identity in product form and combining with (10) implies that

$$(1 + r^s_{a',t})^t = \left(\prod_{\tau=0}^{n-1} M_{a,a',\tau(a'-a)} \right) \left(\prod_{\tau=0}^{n-1} R_{a,a',\tau(a'-a)} \right)$$

The proposition follows from raising both sides of this equation to a power of $1/t$. □

The proposition generalizes equation 6 in that it decomposes the social discount rate into an "M term" and an "R term" (note that equation 6 can be obtained as a special case of equation 9 by imposing $a' = a + 1$ and $n = 1$). The M term depends on how society values increasing the consumption of older people relative to increasing the consumption of younger people. The R term depends on how society ranks different consumption streams for a single individual over their lifetime.

5.3 The R term

Assume that individuals' preferences over their own consumption streams are represented by utility functions, $\{u^t(\cdot)\}_{t=-\infty}^{\infty}$, which are strictly increasing, concave, and satisfy the Inada conditions.[14] Note that, for the following results, it is not necessary to assume that individual preferences take the discounted-utilitarian form in equation 1 (although of course this is allowable as a special case).

Social preferences are said to be *non-paternalistic* if they align with individual preferences when evaluating changes in consumption over the lifecycle. The discounted utilitarian preferences discussed in section 3 are examples of non-paternalistic social preferences. When social preferences are non-paternalistic, then the social marginal rate of substitution between the consumption of an individual at age a and their consumption at age a' is determined by the rate at which the individual themselves is willing to substitute consumption between these two ages:

$$R_{a,a',t} = \frac{\frac{\partial u^{t-a}(\bar{\mathbf{c}}^{t-a})}{\partial c_a^{t-a}}}{\frac{\partial u^{t-a}(\bar{\mathbf{c}}^{t-a})}{\partial c_{a'}^{t-a}}} \tag{11}$$

In what follows, I restrict attention to economies in which people can trade single period bonds, at the interest rates $\{r_t\}_{t=0}^{\infty}$. Given these interest rates, the consumption stream of generation t, $\bar{\mathbf{c}}^t$, solves the following optimization problem:

$$\bar{\mathbf{c}}^t = \arg\max_{\mathbf{c}^t, \mathbf{b}^t} u^t(\mathbf{c}^t) \text{ subject to} \tag{12}$$

$$c_a^t + b_a^t \leq y_a^t + (1 + r_{t+a}) b_{a-1}^t \tag{13}$$

$$b_a^t \geq \underline{b}_a^t \tag{14}$$

In this optimization problem, y_a^t is generation t's income at age a, b_a^t is the amount of one-period bonds that it holds at age a and $(-\underline{b}_a^t)$ is its debt limit at age a. The first constraint is the flow budget constraint, which states that, in each period, consumption and savings cannot exceed the sum of income and financial wealth. The second constraint is the borrowing constraint, which places a limit on the amount of debt that can be held at each age.[15]

Given the short-term market interest rates, $\{r_\tau\}_{\tau=0}^{t}$, it is useful to define the long-

[14]This assumption allows for intergenerational altruism, as people may care about the consumption streams of other generations in addition to their own. However, it is assumed that such altruistic preferences are separable from individuals' preferences over their own consumption streams, so the utility function u^t does not depend on the consumption of other generations.

[15]For example, it is natural to assume that $\underline{b}_T^t = 0$, so that people cannot die with debt. However, it is worth noting that none of the results that follow assume that borrowing constraints are binding.

term market interest rate, r_t^m, as the average rate of return on a one-period bond that is issued in period 0 and refinanced for t periods:

$$r_t^m := \left(\prod_{\tau=0}^{t-1}(1+r_\tau)\right)^{\frac{1}{t}} - 1$$

Consider the following proposition.

Proposition 3. *Assume that social preferences are non-paternalistic and that, for every t, the consumption sequence, $\bar{\mathbf{c}}^t$, solves problem (12). Then, for every $a' > a$, $n \geq 1$ and $t = n(a' - a)$, it holds that*

$$\left(\prod_{\tau=0}^{n-1} R_{a,a',(a'-a)\tau}\right)^{\frac{1}{t}} \geq 1 + r_t^m$$

To develop intuition, it is useful to focus on the case $n = 1$. Note that, from (11), the social marginal rate of substitution $R_{a,a',t}$ is given by the individual's marginal rate of substitution between consumption at age a and consumption at age a'. If borrowing constraints are not binding then, by the standard Euler condition, this marginal rate of substitution is equal to the market return on saving between these two periods. If borrowing constraints are binding at any point, then the marginal rate of substitution will be greater than the market returns.

This proposition is useful because it sets a lower-bound on the R term in the decomposition result of Proposition 2. When social preferences are non-paternalistic and $\bar{\mathbf{c}}$ is an equilibrium consumption path, then this term must be greater than or equal to the long-term market interest rate.

It is worth noting that the assumption of non-paternalistic social preferences may be controversial (see Caplin and Leahy [2004] for further discussion). Indeed, there are policies in place that can be interpreted as being paternalistic towards individuals' time preferences: for example, subsidized retirement savings capture the sentiment that people should be nudged towards saving more for their own good. However, the focus of the debate surrounding the social discount rate is not on whether individuals save enough for their own benefit, but rather on how to make intertemporal tradeoffs that involve different generations. As such, the non-paternalistic benchmark serves as a logical starting point for this analysis.

5.4 The M term

By combining Proposition 2 and Proposition 3, I obtain the following result.

Proposition 4. *Assume that social preferences are non-paternalistic and that, for every t, the consumption sequence, $\bar{\mathbf{c}}^t$, solves problem (12). Then, for any t that is divisible by 60, it holds that*

$$\left(\prod_{\tau=0}^{t/60-1} M_{20,80,60\tau}\right)^{\frac{1}{(t/60)}} \leq \left(\frac{1+r^s_{80,t}}{1+r^m_t}\right)^{60} \tag{15}$$

This result is immediate from substituting $a = 20$, $a' = 80$ and $n = t/60$ into Propositions 2 and 3. Note that the inequality will be strict if and only if borrowing constraints are binding sometime between the ages of 20 and 79.

The left-hand side of this inequality represents the geometric mean of the social marginal rates of substitution between 20 year-olds and 80 year-olds, across $t/60$ equally spaced periods between 0 and t. The right-hand side is the ratio of the gross social discount rate and the gross market interest rate, raised to the power of 60. Using the standard log approximation, for $r^s_{80,t}$ and r^m_t that are close to 0, it holds that:[16]

$$\left(\frac{1+r^s_{80,t}}{1+r^m_t}\right)^{60} \approx \left(1-(r^m_t - r^s_{80,t})\right)^{60} \tag{16}$$

Figure 5 illustrates the quantitative implications of Proposition 4. The x-axis represents the difference between the market interest rate and the social discount rate, while the y-axis represents the geometric mean of the social marginal rate of substitution between 20 year-olds and 80 year-olds (the M term). The shaded area shows the possible combinations of these two variables that are consistent with the relationship shown in equation 15, using the approximation in equation 16. This figure illustrates that, as the social discount rate falls below the market interest rate, the maximal marginal rate of substitution across age groups becomes very small very quickly.

[16] To see this, note that $\log(((1+r^s_{80,t})/(1+r^m_t))^{60}) = 60(\log(1+r^s_{80,t}) - \log(1+r^m_t)) \approx 60(r^s_{80,t} - r^m_t)$. It thus follows that $((1+r^s_{80,t})/(1+r^m_t))^{60} \approx \exp(60(r^s_{80,t} - r^m_t)) = \exp(r^s_{80,t} - r^m_t)^{60}$. Using the Taylor approximation around $x = 0$, $\exp(x) \approx 1 + x$, and hence $((1+r^s_{80,t})/(1+r^m_t))^{60} \approx \exp(r^s_{80,t} - r^m_t)^{60} \approx (1 + r^s_{80,t} - r^m_t)^{60}$.

Figure 5: Possible values of the M term

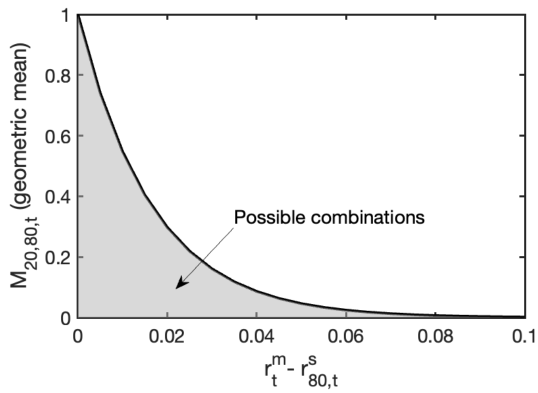

Note: The x-axis represents the difference between the market interest rate and the social discount rate. The y-axis represents the geometric mean of the social marginal rate of substitution between 20 year-olds and 80 year-olds. The shaded area represents the combinations that are consistent with Proposition 4, using the approximation in expression 16.

It may be useful to spell out the quantitative implications of imposing a social discount rate of 1.5%, as in Stern [2008], given a market interest rate of 6%, as in Nordhaus [2007]. In this case, the difference between the social discount rate and the market interest rate is 4.5 percentage points. By expression 15, for a 60-year horizon ($t = 60$), it holds that

$$M_{20,80,0} \leq \left(\frac{1 + r^s_{80,60}}{1 + r^m_{60}}\right)^{80-20} = \left(\frac{1.015}{1.06}\right)^{60} \approx 0.06 \tag{17}$$

This upper-bound means that it is better to increase the consumption of a 20-year old by 6 dollars than to increase the consumption of an 80 year-old by a 100 dollars.

Note that the difference between the 60-year market interest rate and the 60-year social discount rate places an upper-bound on the social marginal rate of substitution between age groups *today.* The above analysis demonstrates that if the social discount rate is 4.5 percentage points below the market interest rate, then society must consider it more beneficial to increase the consumption of a present-day 20 year-old by \$6 than to increase the consumption of a present-day 80 year-old by \$100.

At longer horizons, there is more flexibility. In principle, specifying the 120-year social discount rate to be 1.5% is consistent with the stance that increasing the consumption of the current elderly is just as valuable as increasing in the consumption of the current young (i.e., $M_{20,80,0} = 1$). However, if the social marginal rate of substitution today is higher than its long-run geometric mean, then it must be lower than its long-run geometric mean in the future. Given a market interest rate of 6%, it follows by Proposition 4 that

$$(M_{20,80,0}M_{20,80,60})^{\frac{1}{2}} = (1 \cdot M_{20,80,60})^{\frac{1}{2}} \leq \left(\frac{1+r^s_{80,120}}{1+r^m_{120}}\right)^{80-20} = \left(\frac{1.015}{1.06}\right)^{60} \approx 0.06$$

$$\Rightarrow M_{20,80,60} \leq 0.06^2 = 0.0036$$

which means that, 60 years from now, we must accept that it is more beneficial to increase the consumption of a 20 year old by 36 cents than to increase the consumption of an 80 year-old by 100 dollars.

6 Discussion

This paper demonstrates a tension between concerns for intergenerational equity and the desire to ensure equal consideration for the consumption of people of all ages.

There are several ways forward. One is to simply accept these implications. For instance, perhaps it *is* desirable to favor the consumption of younger people over that of older people. Alternatively, perhaps policymakers *should* discount the utilities of future generations.

If we reject both of these implications, we are left with two options. The first is to adopt a paternalistic social preference relation, which reflects the view that people save too little for their own good. However, such paternalism may be problematic, as it supports policy interference in individual saving decisions over the lifecycle – something that many people would consider objectionable.

The final option is to reject the consequentialist framework altogether. The consequentialist framework assumes that the social objective should be independent of the feasible set of alternatives. The results in this paper highlight the potential limitations of this framework. When designing redistributive policies, we may want to express equal concern for people of all ages; when confronted with intergenerational tradeoffs, we may want to give serious consideration to their long-term implications; and, when deciding on an individual's retirement savings, we may want to defer to their preferences. Unfortunately, it is not possible to formulate a single social objective that satisfies all of these demands. These three demands can only be met in a non-consequentialist framework, in which the social objective depends on the context.

References

Matthew D. Adler. *Measuring Social Welfare: An Introduction.* Oxford University Press, 2019.

James Andreoni and Charles Sprenger. Estimating time preferences from convex budgets. *American Economic Review*, 102(7):3333–3356, 2012.

Kenneth Arrow, Maureen Cropper, Christian Gollier, Ben Groom, Geoffrey Heal, Richard Newell, and William Nordhaus et al. Determining benefits and costs for future generations. *Science*, 341(6144):349–350, 2013.

Lint Barrage. Sensitivity analysis for Golosov, Hassler, Krusell, and Tsyvinski (2014): 'Optimal taxes on fossil fuel in general equilibrium'. *Supplemental Material, Econometrica*, 82(1):41–88, 2014.

Lint Barrage. Be careful what you calibrate for: Social discounting in general equilibrium. *Journal of Public Economics*, 160:33–49, 2018.

William J. Baumol. On the social rate of discount. *American Economic Review*, 58 (4):788–802, 1968.

Michael Carlos Best, James S. Cloyne, Ethan Ilzetzki, and Henrik J. Kleven. Estimating the elasticity of intertemporal substitution using mortgage notches. *The Review of Economic Studies*, 87(2):656–690, 2020.

Guillermo A. Calvo and Maurice Obstfeld. Optimal time-consistent fiscal policy with finite lifetimes. *Econometrica*, 56(2):411–432, 1988.

Andrew Caplin and John Leahy. The social discount rate. *Journal of political economy*, 112(6):1257–1268, 2004.

Moritz A. Drupp, Mark C. Freeman, Ben Groom, and Frikk Nesje. Discounting disentangled. *American Economic Journal: Economic Policy*, 10(4):109–134, 2018.

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

Marc Fleurbaey and Stéphane Zuber. Discounting, risk and inequality: A general approach. *Journal of Public Economics*, 128:34–49, 2015.

Marta Giagheddu and Andrea Papetti. The macroeconomics of age-varying epidemics. *European Economic Review*, 151:104346, January 2023.

Christian Gollier and James K. Hammitt. The long-run discount rate controversy. *Annual Review of Resource Economics*, 6:273–295, 2014.

Mikhail Golosov, John Hassler, Per Krusell, and Aleh Tsyvinski. Optimal taxes on fossil fuel in general equilibrium. *Econometrica*, 82(1):41–88, 2014.

Pierre-Olivier Gourinchas and Jonathan A. Parker. The empirical importance of precautionary saving. *American Economic Review*, 91(2):406–412, 2001.

Hilary Greaves. Discounting for public policy: A survey. *Economics & Philosophy*, 33(3):391–439, 2017.

Robert E. Hall. Intertemporal substitution in consumption. *Journal of political economy*, 96(2):339–357, 1988.

Robert E. Hall, Charles I. Jones, and Peter J. Klenow. Trading off consumption and COVID-19 deaths. *Federal Reserve Bank of Minneapolis Quarterly Review*, 42(1): 1–14, 2020.

J. Paul Kelleher. Pure time preference in intertemporal welfare economics. *Economics & Philosophy*, 33(3):441–473, 2017.

Finn E. Kydland and Edward C. Prescott. Time to build and aggregate fluctuations. *Econometrica*, 50(6):1345–1370, 1982.

Emily C. Lawrance. Poverty and the rate of time preference: Evidence from panel data. *Journal of Political Economy*, 99(1):54–77, February 1991.

Antony Millner and Geoffrey Heal. Choosing the future: Markets, ethics, and rapprochement in social discounting. *Journal of Economic Literature*, 2023. Forthcoming.

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

John Quiggin. Equity between overlapping generations. *Journal of Public Economic Theory*, 14(2):273–283, 2012.

Frank Plumpton Ramsey. A mathematical theory of saving. *The Economic Journal*, 38(152):543–559, 1928.

Maik T. Schneider, Christian P. Traeger, and Ralph Winkler. Trading off generations: Equity, discounting, and climate change. *European Economic Review*, 56(8):1621–1644, 2012.

Jesse M. Shapiro. Is there a daily discount rate? Evidence from the food stamp nutrition cycle. *Journal of public Economics*, 89(2-3):303–325, 2005.

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

A Proof of Proposition 1

I start by proving the second clause. First, note that $\mathbf{c}^*$ is feasible. Because feasibility depends only on the sequence of aggregate consumption, if $\mathbf{c}$ differs from $\mathbf{c}^*$ only in the allocation of consumption across age groups, but not in the sequence of aggregate consumption, then $\mathbf{c}$ is feasible as well.

It is convenient to denote "flow welfare" by

$$w(c_0, ..., c_T) = \sum_{a=0}^{T} \left(\frac{1+\rho^s}{1+\rho} \right)^a \frac{c_a^{1-\gamma}}{1-\gamma}$$

Note that social preferences are represented by

$$W(\mathbf{c}) = \sum_{t=-\infty}^{\infty} \left(\frac{1}{1+\rho^s} \right)^t w(c_{0,t}, ..., c_{T,t})$$

Let $\mathbf{c}^*$ denote the optimal feasible consumption allocation, and let $\mathbf{c}^{**} = \{c_{a,t}^{**}\}_{a,t}$ denote the solutions to the optimization problems:

$$\{c_{a,t}^{**}\}_{a=0}^{T} = \arg \max_{\{c_a\}_{a=0}^{T}} w_t(c_{0,t}, ..., c_{T,t}) \text{ s.t. } \sum_{a=0}^{T} c_{a,t} = C_t^*$$

By construction, $\mathbf{c}^{**}$ maximizes "flow welfare" in each period. If, for some period t, it holds that $w_t(c_{0,t}^{**}, ..., c_{T,t}^{**}) > w_t(c_{0,t}^{*}, ..., c_{T,t}^{*})$, then social welfare, W, would increase if we were to replace $\mathbf{c}^*$ with $\mathbf{c}^{**}$. Given that $\mathbf{c}^*$ is optimal and $\mathbf{c}^{**}$ is feasible (as it has the same sequence of aggregate consumption), it must hold that, for every t

$$w_t(c_{0,t}^{**}, ..., c_{T,t}^{**}) = w_t(c_{0,t}^{*}, ..., c_{T,t}^{*})$$

It is left to show that the solution to the within-period optimization problem delivers the expression in Proposition 1. The first-order conditions of the within-period optimization problem are

$$\left(\frac{1+\rho^s}{1+\rho}\right)^a c_a^{-\gamma} = \lambda$$

where $\lambda > 0$ is the Lagrange multiplier on the constraint. Thus, for every a and a', it holds that

$$\left(\frac{1+\rho^s}{1+\rho}\right)^a c_a^{-\gamma} = \left(\frac{1+\rho^s}{1+\rho}\right)^{a'} c_{a'}^{-\gamma} \Rightarrow \frac{c_{a'}}{c_a} = \left(\frac{1+\rho^s}{1+\rho}\right)^{\frac{a'-a}{\gamma}}$$

To prove the first clause of the proposition, note that, using the above expressions for the optimal allocation of consumption across age groups, in the optimal allocation, the "flow welfare", w, is given by

$$w(c_{0,t}^{*}, ..., c_{T,t}^{*}) = \sum_{a=0}^{T} \left(\frac{1+\rho^s}{1+\rho}\right)^a \frac{(\zeta_a(C_t^*))^{1-\gamma}}{1-\gamma} = \left(\sum_{a=0}^{T} \left(\frac{1+\rho^s}{1+\rho}\zeta_a^{1-\gamma}\right)^a\right) \frac{(C_t^*)^{1-\gamma}}{1-\gamma}$$

where ζ_a is the constant consumption share of people aged a. It follows that

$$W^{DU}(\mathbf{c}^*) = \sum_{t=-\infty}^{\infty} \left(\frac{1}{1+\rho^s}\right)^t \left(\left(\sum_{a=0}^{T} \left(\frac{1+\rho^s}{1+\rho}\zeta_a^{1-\gamma}\right)^a\right) \frac{(C_t^*)^{1-\gamma}}{1-\gamma}\right) =$$

$$\left(\sum_{a=0}^{T} \left(\frac{1+\rho^s}{1+\rho}\zeta_a^{1-\gamma}\right)^a\right) \left(\sum_{t=-\infty}^{\infty} \left(\frac{1}{1+\rho^s}\right)^t \frac{(C_t^*)^{1-\gamma}}{1-\gamma}\right)$$

Note that the first term in brackets is a positive constant, and hence maximizing W^{DU} is equivalent to maximizing the second term. Because $\mathbf{c}^*$ is optimal, it follows that the sequence of aggregate consumption, $\{C_t^*\}_{t=0}^{\infty}$, solves the optimization problem in the first clause of the proposition.

B Proof of Proposition 3

As, for every i, the social preferences over the consumption stream $\mathbf{c}^i$ are represented by u^i, it holds that

$$R_{a,a',t} = \frac{W_{a,t}}{W_{a',t+(a'-a)}} = \frac{\frac{\partial u^i(\mathbf{c}^i_e)}{\partial c_a}}{\frac{\partial u^i(\mathbf{c}^i_e)}{\partial c_{a'}}}$$

where $i = t - a$ is the generation that is aged a in period t.

The first-order condition of individual i's optimization problem with respect to c^i_a yields

$$\frac{\partial u^i(\bar{\mathbf{c}}^i)}{\partial c^i_a} = \lambda^i_a \tag{18}$$

where $\lambda^i_a > 0$ is the Lagrange multiplier on the age-a budget constraint.[17]

For $a < T$, the first-order condition with respect to b^i_a is

$$\lambda^i_a = \zeta^i_a + (1 + r_{i+a})\lambda^i_{a+1}$$

where $\zeta^i_a \geq 0$ is the Lagrange multiplier on the borrowing constraint at age a. Substituting expression 18 yields

$$\frac{\partial u^i(\bar{\mathbf{c}}^i)}{\partial c^i_a} = \zeta^i_a + (1 + r_{i+a})\frac{\partial u^i(\bar{\mathbf{c}}^i)}{\partial c^i_{a+1}}$$

$$\Rightarrow \frac{\frac{\partial u^i(\bar{\mathbf{c}}^i)}{\partial c^i_a}}{\frac{\partial u^i(\bar{\mathbf{c}}^i)}{\partial c^i_{a+1}}} = \frac{\zeta^i_a}{\frac{\partial u^i(\bar{\mathbf{c}}^i)}{\partial c^i_{a+1}}} + (1 + r_{i+a}) \geq 1 + r_{i+a}$$

Thus,

$$R_{a,a',t} = \frac{\frac{\partial u^i(\bar{\mathbf{c}}^i)}{\partial c^i_a}}{\frac{\partial u^i(\bar{\mathbf{c}}^i)}{\partial c^i_{a'}}} = \left(\frac{\frac{\partial u^i(\bar{\mathbf{c}}^i)}{\partial c^i_a}}{\frac{\partial u^i(\bar{\mathbf{c}}^i)}{\partial c^i_{a+1}}}\right)\left(\frac{\frac{\partial u^i(\bar{\mathbf{c}}^i)}{\partial c^i_{a+1}}}{\frac{\partial u^i(\bar{\mathbf{c}}^i)}{\partial c^i_{a+2}}}\right)\cdots\left(\frac{\frac{\partial u^i(\bar{\mathbf{c}}^i)}{\partial c^i_{a'-1}}}{\frac{\partial u^i(\bar{\mathbf{c}}^i)}{\partial c^i_{a'}}}\right) \geq$$
$$(1 + r_{i+a})(1 + r_{i+a+1})\cdots(1 + r_{i+a'-1})$$

[17] Here, I am using the assumption that u is strictly increasing, concave and satisfies the Inada conditions to conclude that the optimal consumption plan satisfies this first-order condition.

As $i = t - a$, it follows that

$$R_{a,a',t} \geq \prod_{\tau=0}^{a'-a-1} (1 + r_{t+\tau})$$

Thus, for $t = (a' - a)n$,

$$\prod_{\tau=0}^{n-1} R_{a,a',\tau(a'-a)} \geq \prod_{\tau=0}^{n-1} \left(\prod_{\tau'=0}^{a'-a-1} (1 + r_{\tau(a'-a)+\tau'}) \right) =$$

$$((1 + r_0) \cdots (1 + r_{a'-a-1})) \cdots \big((1 + r_{(n-1)(a'-a)}) \cdots (1 + r_{(n(a'-a)-1)}\big)$$

$$= \prod_{\tau=0}^{t-1} (1 + r_\tau) = (1 + r_t^m)^t$$

(where the last equality follows from the definition of the long-term market interest rate, r_t^m). Hence,

$$\left(\prod_{\tau=0}^{n-1} R_{a,a',\tau(a'-a)} \right)^{\frac{1}{t}} \geq 1 + r_t^m$$

C Proof of Proposition 4

By Proposition 2 and Proposition 3, when preferences are non-paternalistic then, along an equilibrium path, it holds that

$$1 + r_{80,t}^s \geq \left(\prod_{\tau=0}^{t/60-1} M_{20,80,60\tau} \right)^{\frac{1}{t}} (1 + r_t^m)$$

for a time horizon t that is divisible by 60. Dividing both sides by the gross market interest rate, $1 + r_t^m$, I obtain

$$\frac{1 + r_{80,t}^s}{1 + r_t^m} \geq \left(\prod_{\tau=0}^{t/60-1} M_{20,80,60\tau} \right)^{\frac{1}{t}} = \left(\left(\prod_{\tau=0}^{t/60-1} M_{20,80,60\tau} \right)^{\frac{1}{(t/60)}} \right)^{\frac{1}{60}}$$

Raising both sides by a power of 60 yields the inequality

$$\left(\frac{1+r^s_{80,t}}{1+r^m_t}\right)^{60} \geq \left(\prod_{\tau=0}^{t/60-1} M_{20,80,60\tau}\right)^{\frac{1}{(t/60)}} \tag{19}$$

D The neoclassical growth model

This section presents the feasibility constraints and the calibration of the neoclassical growth model.

I begin by specifying the economy's resource constraints (which determine the set of feasible aggregate consumption sequences, $\mathcal{F}^*$). In each period, aggregate output, Y_t, is produced according to a standard Cobb-Douglas production function,

$$Y_t = K_t^{\alpha}(A_t L_t)^{1-\alpha}$$

where $L_t > 0$ is aggregate labor inputs, $K_t > 0$ is aggregate capital inputs, A is labor-augmenting productivity and $\alpha \in (0,1)$ is the capital intensity of production. Productivity growth is constant at the rate g:

$$A_t = (1+g)^t A_0$$

Labor is supplied inelastically; each person supplies $L_a \geq 0$ units of labor at age a. As the size of each cohort is fixed, labor supply at time t is given by

$$L_t = L = \sum_{a=0}^{T} L_a \tag{20}$$

The initial capital stock, K_0, is given. For $t > 0$, K_t is given by the standard capital accumulation equation,

$$K_t = (1-\delta)K_{t-1} + I_{t-1}$$

where I_t is investment and $\delta \in (0,1)$ is the capital depreciation rate.

The economy's aggregate resource constraint at time t is

$$C_t + I_t = Y_t$$

where C_t is aggregate consumption. Note that, in this model, the set of feasible aggregate consumption sequences depends on the initial capital stock, K_0.

D.1 Calibration

I calibrate the model to the United States. Table 6 summarizes the model's calibrated parameters. Most of the parameters are standard. The capital intensity parameter, α, was chosen to roughly match the capital income share, and the parameter g was chosen to roughly correspond to the average annual growth rate of GDP per-capita in the US. The parameter γ was chosen as $\gamma = 1$ corresponding to log-utility. The capital depreciation rate, δ, was chosen to roughly match the average depreciation rate of capital in the US.

Table 6: Calibration parameters

	Value	Target
α	0.3	Capital income share
δ (annual)	0.05	Capital depreciation rate
$1/\gamma$	1	Intertemporal elasticity of substitution
g (annual)	0.02	GDP per-capita growth rate
A_0	1	(Normalization)
L	1	(Normalization)
K_0		Rate of return of 6%

The initial capital stock, K_0, was calibrated under the assumption that the current rate of return on capital is 6%. The rate of return on capital, r, is given by

$$r = \frac{\partial Y_t}{\partial K_t} - \delta = \alpha K_0^{\alpha-1}$$

(using the normalization $A_0 = L = 1$). Thus, given α and r, it is possible to back out K_0.