

The cross-sectional implications of the social discount rate

Maya Eden

Global Priorities Institute | September 2021

GPI Working Paper No. 12-2021



The Cross-Sectional Implications of the Social Discount Rate

Maya Eden*

Brandeis University and CEPR

September 6, 2021

Abstract

How should policy discount future returns? The standard approach to this normative question is to ask how much society should care about future generations relative to people alive today. This paper establishes an alternative approach, based on the social desirability of redistributing from the current old to the current young. Along the balanced growth path, bounds on the welfare gains from age-based redistribution imply bounds on the social discount rate. A calibration shows that an objective of maximizing the sum of utilities in each period implies social discount rates that are within a percentage point of the market interest rate.

JEL Classification: D63, E61

Keywords: Overlapping generations, utilitarianism, prioritarianism, age-based inequality

*Contact address: Maya Eden, Economics Department, Brandeis University, 415 South Street, Waltham, MA 02453; Email: meden@brandeis.edu. I thank Geir Asheim, Matthias Doepke, Ben Eden, Paul Kelleher, Pete Klenow, Kevin Kuruc and Antony Millner for helpful comments. This paper has also benefited from comments by seminar participants at the macroeconomics seminar at the Chicago Fed, and the CEPR Macroeconomics and Growth Programme Online Meeting, and the Oxford Workshop on Global Priorities Research.

1 Introduction

The social discount rate is the rate at which policy should discount future returns. There is substantial controversy about what is the appropriate benchmark for quantifying it.¹ In Nordhaus [2007], the annual social discount rate is 6%, consistent with long-run estimates of the real interest rate.² In contrast, Stern [2008] calibrates a social discount rate based on concerns for intergenerational equity, and obtains a social discount rate of 1.5%.

This paper provides another perspective on this debate by exploring the cross-sectional implications of the social discount rate. When generations are overlapping, the extent to which social preferences care about future generations simultaneously determines the social discount rate, and the marginal welfare weights of younger people compared to older ones. Along the balanced growth path, the distribution of consumption across age groups is socially optimal if and only if the social discount rate is equal to the market interest rate.

This result is useful because it allows us to check which social discount rates are consistent with our moral views about redistributing consumption across age groups. It turns out that even plausible deviations of the social discount rate from the market interest rate generate uneasy implications. For example, assuming a market interest rate of 6% (as in Nordhaus [2007]), the 1.5% social discount rate proposed by Stern [2008] implies that it must be socially desirable to reduce the consumption of a 70-year-old by \$1 in order to increase the consumption of a 20 year-old by 10 cents. Such extreme ageist implications can be avoided only by choosing a social discount rate that is closer to the market interest rate. A calibration suggests that a social objective of maximizing the sum of utilities in each period is inconsistent with social discount rates that are much below the market interest rate.

¹See Greaves [2017] and Millner and Heal [2021] for recent reviews. See also Arrow et al. [2013], Gollier and Hammitt [2014] and Kelleher [2017].

²It is important to note that there is a separate debate about what is the appropriate benchmark for quantifying the long-run interest rate, especially given the lack of assets with very long maturities. See Stern [2008] and Millner and Heal [2021] for further details.

This paper is related to a large literature on social discounting (see Millner and Heal [2021] for a recent review). The social discount rate is usually discussed in the context of non-overlapping generations models. Notable exceptions are Calvo and Obstfeld [1988] and Schneider et al. [2012], who solve for the optimal policy given a discounted-utilitarian objective and an overlapping generations economy. Calvo and Obstfeld [1988] establish that, in order to implement the social optimum, transfers between age groups may be necessary. This paper adds to this literature by quantitatively exploring the relationship between social discounting and age-based transfers away from the optimal policy. The key insight here is that small deviations of the social discount rate from the market interest rate imply large differences in the marginal social welfare weights of different age groups. This quantitative insight is absent from previous work.

2 A decomposition of the social discount rates

Time is discrete and indexed $t = 0, \dots$. In each time period, one generation is born, and lives for $1 < T < \infty$ periods. The assumption that $T > 1$ implies that generations are overlapping. For simplicity, I assume that the size of each cohort is fixed.³

Let c_a^t denote generation t 's consumption at age a . Let $\mathbf{c}^t = (c_1^t, \dots, c_T^t)$ denote the consumption sequence of generation t , and let $\mathbf{c} = (\mathbf{c}^0, \dots)$ denote the intergenerational consumption allocation.

The social preference relation is represented by a differentiable social welfare function, $W(\mathbf{c})$.⁴ Given an allocation of consumption, \mathbf{c} , define the social marginal rate of substitution between c_a^t and $c_{a'}^{t'}$ as

$$MRS((a, t), (a', t')) = \left(\frac{\partial W(\mathbf{c})}{\partial c_a^t} \right) / \left(\frac{\partial W(\mathbf{c})}{\partial c_{a'}^{t'}} \right) \quad (1)$$

³Allowing for population growth does not change the results, but complicates notation.

⁴This assumption rules out the maximin welfare criterion advocated by Rawls [1974] (see also Asheim and Zuber [2013] in an intergenerational context).

Note that c_a^t is the consumption of generation t that takes place in period $t+a-1$, when that generation is aged a . If $t'-a' > t-a$, then the consumption $c_{a'}^{t'}$ takes place in a later date; in this case, the *social discount rate* between generation t in period $t+a-1$ and generation t' in period $t'+a'-1$ is

$$r_{a,t,a',t'}^s = MRS((a,t), (a',t'))^{\frac{1}{t'+a'-t-a}} - 1 \quad (2)$$

The social discount rate is the required average rate of return on a small investment at time $t+a-1$ which is financed by generation t , and benefits generation t' in period $t'+a'-1$.

There is generally no single social discount rate between periods. Rather, the social marginal rate of substitution between the consumptions of two individuals may depend on their characteristics. For example, the required social rate of return on a project that benefits some individuals in the far future may depend on whether it is financed by current taxes levied on the old or on the young.⁵

Define the *social-individual discount rates* as

$$r_{a,a',t}^{si} = r_{a,t,a',t}^s \quad (3)$$

The social-individual discount rate, $r_{a,a',t}^{si}$, is defined so that society is roughly indifferent with respect to taking ϵ units of consumption away from generation t at age a , and compensating it with $(1 + r_{a,a',t}^{si})^{a'-a}\epsilon$ additional units of consumption at age a' . This is the *social* rate of discount for an individual's own consumption; in principle, it may be different from the rate at which the individual discount his own future consumption.

If, at some period τ , generation t is aged a and generation t' is aged a' , then $MRS((a,t), (a',t'))$ captures the relative *distributional weights* of generations t and t' in period τ . In period τ , it is socially desirable to take one small unit of consumption from generation t and give it to generation t' if and only if $MRS((a,t), (a',t')) < 1$. It is useful to denote these cross-sectional relative

⁵This is also discussed in Fleurbaey and Zuber [2015].

distributional weights by

$$\mu_{a,a'}^\tau = MRS((a, \tau - a + 1), (a', \tau - a' + 1)) \quad (4)$$


The results in this paper build on the insight that, because generations are overlapping, the relative distributional weights, $\{\mu_{a,a'}^\tau\}_{a,a',\tau}$, and the social-individual discount rates, $\{r_{a,a',t}^{si}\}_{a,a',t}$, are sufficient for recovering the entire set of social discount rates, $\{r_{a,t,a',t'}^s\}_{a,a',t,t'}$.

Proposition 1. *For every $t' > t$, the social discount rate $r_{a,t,a',t'}^s$ is given by*

$$(1 + r_{a,t,a',t'}^s)^{t'+a'-(t+a)} = \left(\prod_{\tau=t+a-1}^{t'+a'-2} (1 + r_{1,2,\tau}^{si}) \right) \left(\mu_{a,1}^{t+a-1} \mu_{2,a'}^{t'+a'-1} \prod_{\tau=t+a}^{t'+a'-2} \mu_{2,1}^\tau \right)$$

The proof of this proposition is in the appendix, together with other omitted proofs. This result allows for the decomposition of the social discount rates into two terms. The first depends only on the social-individual discount rates. The second depends only on the relative distributional weights of different age groups.

Figure 1: Decompositions of the social discount rate



Note: The solid back arrow represents a transfer between two 40-year-olds, one in 2020 and the other in 2080. The gray arrows represent the decomposition in Proposition 1, which is based on the insight that this transfer can be done through a sequence of transfers within-people and across-time, and transfers between people in a given time. The dashed arrows correspond to the decomposition in Fleurbaey and Zuber [2015].

Figure 1 illustrates the decomposition. The solid black line represents a dollar transfer between a 40 year-old in 2020, and a 40 year-old in 2080. Proposition 1 is based on the observation that this transfer can be implemented through a sequence of transfers that involve either transferring within-people, across time, or between people, within-time. This sequence is illustrated with the solid grey lines. In 2020, the dollar is transferred from the 40 year-old to a 20 year-old. Then, it is transferred from the 20 year-old in 2020 to the same person 20 years later, when he is 40. At that point, the dollar is transferred to a contemporaneous 20 year-old, and so on and so forth. This sequence of transfers is composed only of transfers between people in the same time period (vertical lines), and transfers within people across time (horizontal lines). The social desirability of the former depends on the distributional weights of different age groups, and the social desirability of latter depends on the social-individual discount rates.

The figure also illustrates the difference between the decomposition here and the decomposition in Fleurbaey and Zuber [2015] (section 6). Their decomposition builds on the observation that a transfer between two individuals can be implemented as a transfer between them at their respective births, plus transfers within their lifetimes. Based on their decomposition, Fleurbaey and Zuber [2015] conclude that the long-run social discount rate is determined by the marginal welfare gains from reallocating resources between two people at their respective births. The decomposition here shows that these welfare gains are determined by the social desirability of transferring resources across age groups: how much we care about the far future is related to how much we care about the current young.

Proposition 1 is useful because it reduces the question of social discounting across generations to two sub-questions. The first is, “how should society discount an individual’s own future consumption?” Here, the common approach is to evoke the Pareto principle: if people discount their own future consumption at a certain rate, then society should respect their preferences and discount their future consumption at that rate as well. The second question is, “how should consumption be distributed across people alive today?” Here,

the most common approach is utilitarian: a transfer between two individuals is desirable provided that it increases the total sum of their flow utilities.

3 An illustrative example

The usefulness of this decomposition can be illustrated with the following simple example. In this example, each generation is alive for two periods. The market interest rate is constant and equal to r . The equilibrium allocation is Allocation A.

Allocation A

Generation	$t = 0$	$t = 1$	$t = 2$	\dots	$t = n$	$t = n + 1$
0	c_1^0	c_2^0				
1		c_1^1	c_2^1			
\vdots				\ddots		
n					c_1^n	c_2^n

Allocation B

Generation	$t = 0$	$t = 1$	$t = 2$	\dots	$t = n$	$t = n + 1$
0	$c_1^0 - 1$	$c_2^0 + (1 + r)$				
1		c_1^1	c_2^1			
\vdots				\ddots		
n					c_1^n	c_2^n

Allocation C

Generation	$t = 0$	$t = 1$	$t = 2$	\dots	$t = n$	$t = n + 1$
0	$c_1^0 - 1$	c_2^0				
1		$c_1^1 + (1 + r)\mu$	c_2^1			
\vdots				\ddots		
n					c_1^n	c_2^n

Allocation D

Generation	$t = 0$	$t = 1$	$t = 2$	\dots	$t = n$	$t = n + 1$
0	$c_1^0 - 1$	c_2^0				
1		c_1^1	$c_2^1 + (1 + r)^2\mu$			
\vdots				\ddots		
n					c_1^n	c_2^n

If individuals from generation 0 can borrow and save at the market interest rate, then they choose their consumption sequences so that they are exactly indifferent with respect to saving an additional unit. It follows that individuals from generation 0 are indifferent between the equilibrium allocation, A , and an alternative allocation, B , in which their consumption is reduced by one (small) unit period 0, and increased by $1 + r$ units in period 1.

If the social preference relation satisfies the standard Pareto condition, then it must be consistent with the Pareto indifference condition:⁶ if all generations are indifferent between two allocations, then society must be indifferent

⁶The social preference relation satisfies the Pareto condition if it can be written as $W(\mathbf{c}) = \tilde{W}(\{U_t(\mathbf{c}^t)\}_{t=0}^\infty)$, where \tilde{W} is strictly increasing.

Allocation E

Generation	$t = 0$	$t = 1$	$t = 2$	\dots	$t = n$	$t = n + 1$
0	$c_1^0 - 1$	c_2^0				
1		c_1^1	c_2^1			
\vdots				\ddots		
n					$c_1^n + ((1+r)\mu)^n$	c_2^n

between them as well. When there are no externalities, other generations care only about their own consumption sequences, which are the same in A and B. In this case, all generations are indifferent between the two allocations, and the Pareto indifference condition requires society to be indifferent as well. Thus,

$$\text{Allocation A} \sim \text{Allocation B}$$

where \sim denotes the social indifference relation.⁷

This reasoning shows that, when people can frictionlessly borrow and save at the market interest rate and there are no externalities, then the social-individual discount rates are pinned down by the Pareto principle, and it holds that $r_{a,a',t}^{si} = r$ for all a, a' and t .

If, instead, people from generation 0 are borrowing constrained, then they strictly prefer allocation A over allocation B. In this case, the social-individual discount rates would be higher than the market interest rate ($r_{a,a',t}^{si} > r$). Saving constraints would have the opposite implication.

The presence of externalities may also affect the social-individual discount rates. If the consumption of generation 0 imposes some externalities on other generations, then other generations may not be indifferent between A and B, even though their consumption sequences are the same in both. For exam-

⁷ $W(\mathbf{c}_A) = W(\mathbf{c}_B)$, where \mathbf{c}_A and \mathbf{c}_B are the consumption allocations in A and B, respectively.

ple, consider the case in which consumption is associated with driving, which creates traffic, noise, and air pollution. In this case, generation 1 may have a strict preference for A over B , because it experiences the negative externalities from generation 0's consumption in period 1, but not in period 0. In this case (assuming that all other generations are indifferent), Pareto requires society to have a strict preference for A as well, implying that $r_{1,2,0}^{si} > r$.⁸ How negative consumption externalities affect social-individual discount rates in periods $t > 0$ is less clear: generation $t - 1$ prefers that generation t saves more (and consumes less in period t), while generation $t + 1$ prefers that generation t saves less (and consumes less in period $t + 1$). Of course, positive consumption externalities would have the opposite implications.

Before proceeding, it is worth noting that, while commonly assumed, the Pareto condition is somewhat controversial in an inter-temporal context.⁹ Because of various forms of present-bias and dynamic inconsistency, it is possible to argue that people save too little for their own good. This implies that, while individuals may be indifferent between A and B , society should strictly prefer B . These paternalistic concerns imply social-individual discount rates that are lower than the market interest rate ($r_{a,a',t}^{si} < r$).

Next, define the scalar $\mu = \mu_{2,1}^1$ so that

$$\text{Allocation B} \sim \text{Allocation C}$$

This parameter captures the desirability of redistributing consumption from older people to younger ones. For example, $\mu < 1$ implies that it is desirable to redistribute towards younger people; in this case, society is willing to take away $(1 + r)$ units from an old person in order to give a younger person the smaller amount of $(1 + r)\mu$.

The next step mimics the first step: if individuals from generation 1 can

⁸Note that $r_{1,2,0}^{si} < r$ is inconsistent with the Pareto condition, because both generation 0 and generation 1 strictly prefer A over an allocation in which generation 0 consumes 1 unit less in period 0, and an additional positive amount in period 1 that is less than $1 + r$. Hence, Pareto requires that generation 0 strictly prefers to save at the social discount rate, and thus $r_{1,2,0}^{si} > r$.

⁹See, for example, Caplin and Leahy [2004].

borrow and lend at the market interest rate, then they choose their consumption sequences so that they are indifferent with respect to saving another $(1+r)\mu$ small units at the market interest rate. This means that they are indifferent with respect to giving up $(1+r)\mu$ units in period 1 in exchange for an additional $(1+r)((1+r)\mu) = (1+r)^2\mu$ units in period 2. Consequently, they are indifferent between allocations C and D. Once again, the Pareto principle implies that

$$\text{Allocation C} \sim \text{Allocation D}$$

If the value of μ is constant throughout time, then repeating the same argument and using the transitivity of the social indifference relation implies that

$$\text{Allocation A} \sim \text{Allocation B} \sim \dots \sim \text{Allocation E}$$

This example establishes that, if the relative distributional weights of different age groups are time invariant, and if the social preference relation is consistent with the Pareto principle, then $1 + r_{1,0,1,n}^s = (1+r)\mu$. This suggests a straightforward mapping between the social discount rates and the cross-sectional distributional weights.

4 Balanced growth path

This section explores the implications of this decomposition along the balanced growth path. Let y_a^t denote the income of generation t at age a . I assume that $y_a^t = (1+g)^t y_a^0$ for some $g > 0$. This allows for life-cycle variation in earnings, but assumes that the income of each age group grows at the same rate. Further, I assume that all generations have the same preferences over consumption sequences, which can be represented by a utility function, $U(\cdot)$. To be consistent with a balanced-growth-path equilibrium, I assume that U is

homogeneous.¹⁰

Appendix C discusses how the market interest rate, r , is determined in the closed-economy equilibrium of a simple over-lapping generations model. However, for the results that follow, it is sufficient to assume that r is exogenously fixed and time-invariant. This can be interpreted as a small-open-economy assumption, or as a partial equilibrium model of a balanced growth path.

The consumption sequence of generation t solves the following constrained optimization problem:

$$\mathbf{c}^t = \arg \max_{\mathbf{c}^t} U(\mathbf{c}^t) \text{ s.t. } \sum_{a=1}^T \frac{c_a^t}{(1+r)^a} = \sum_{a=1}^T \frac{y_a^t}{(1+r)^a} \quad (5)$$

Each generation chooses its consumption stream optimally, subject to the inter-temporal budget constraint. This optimization problem abstracts from any credit or saving constraints, and assumes that people can borrow and save at the market interest rate.

I restrict attention to a particular class of social welfare functions which imply time-invariant distributional weights along the balanced growth path. The social welfare function, W , is of the form

$$W(\mathbf{c}) = \sum_{t=0}^{\infty} \left(\frac{1}{1+\rho^s} \right)^t \phi(U(\mathbf{c}^t)) \quad (6)$$

where ϕ is some strictly increasing and homogeneous function, and $\rho^s > -1$.¹¹

This functional form nests two important special cases. If $\phi(U(\mathbf{c}^t))$ is a measure of generation t 's lifetime utility, then the social welfare function is a discounted sum of lifetime utilities (note that, as ϕ is strictly increasing,

¹⁰In a production economy, there is an equilibrium balanced growth path if preferences are homothetic, total factor productivity grows at a constant rate, and the production function has constant returns to scale in capital and labor. Appendix C discusses the (closed-economy) equilibrium balanced growth path in a simple model with $T = 2$ and no capital.

¹¹Given the assumptions that $g > 0$ and ϕ is strictly increasing, the infinite sum in equation 6 does not converge for $\rho \leq 0$ (and, depending on ϕ , may also not converge for some $\rho^s > 0$). For these parameters, the social welfare function can only be defined over an arbitrarily large, but finite, number of cohorts.

$\phi(U(\cdot))$ is a representation of individual preferences, which can be interpreted as a utility function).¹² If, instead, $U(\mathbf{c}^t)$ is a measure of generation t 's lifetime utility and ϕ is concave, then the social welfare function is of the discounted-prioritarian form. The prioritarian objective differs from the utilitarian objective in its aversion to inequality in lifetime utilities.¹³ In both cases, ρ^s is the social rate of pure time preference, or the rate at which the social objective discounts the well-beings of future generations.

The following proposition establishes that, along the balanced growth path, the social discount rate and the cross-sectional distributional weights are closely related.

Proposition 2. *1. There exists $\mu > 0$ such that $\mu_{a',a}^t = \mu^{a'-a}$ for all $a, a' \leq T$ and t .*

2. For every age, a , and time periods, $t' > t$, it holds that $r_{a,t,a,t'}^s = r^s$, where

$$1 + r^s = (1 + r)\mu$$

The rate r^s is the social discount rate across two people in the same age, alive in different periods. This proposition establishes that whenever $r^s \neq r$, then the cross-sectional distribution of consumption across age groups is suboptimal.


To fully appreciate the ethical implications of this result, consider the case in which $r^s < r$. By the above proposition, in this case, it must hold that $\mu < 1$. Hence, it is socially desirable to redistribute from old people to younger people. By the first clause of the proposition, it is socially desirable to reduce the consumption of someone aged a by 1 unit in order to increase the consumption

¹²In this case, the social welfare function can also be interpreted as rank-discounted utilitarian, as in Zuber and Asheim [2012]. According to this approach, the rate of pure time preference, ρ^s , depends on whether future generations are likely to be better or worse off than current generations. As I am restricting attention to a balanced growth path in which $g > 0$, a specification in which $\rho^s > 0$ can be viewed as contingent on the assumption that lifetime utilities grow at a positive rate. This criterion allows for the rate of pure time preference to change if the rate of growth becomes negative.

¹³For a discussion of the prioritarian welfare criterion, see, for example, Adler [2019], Chapter 3.1.

of a newborn by μ^a units. In the limit $T \rightarrow \infty$, the age a can be chosen to be arbitrarily large; consequently, μ^a can be made arbitrarily small. In this case, it is socially desirable to take away \$1 from an extremely elderly person in order to give a young person basically nothing. This is almost a violation of the Pareto condition, as social welfare is “improved” by making an elderly person worse-off without any detectible gain to anyone else.

Figure 2: The relative distributional weights of 70-year olds and 20-year olds implied by different deviations of the social discount rate from the market interest rate



The x-axis corresponds to $100 * (r - r^s)$. Relative distributional weights are computed based on Proposition 2. By the first clause of the proposition, relative distributional weights are given by $\mu^{70-20} = \mu^{50}$. By the second clause of the proposition, $\mu \approx 1 + r^s - r$. The relative distributional weights are therefore computed as $(1 + r^s - r)^{50}$.

One might accept this problematic implication on the grounds that, in practice, the age distribution is bounded; consequently, the relative marginal social welfare weights of old and young people are bounded by μ^T . However, the quantitative implications are uneasy even for a realistic age distribution. For example, assume that the interest rate is 6% (as in Nordhaus [2007]) and that the social discount rate is 1.5% (as in Stern [2008]). These numbers imply that $\mu \approx 0.955$. Under these assumptions, it is welfare-improving to reduce

the consumption of a 70 year-old by \$1 in order to increase the consumption of a 20 year-old by 10 cents.

Figure 2 plots the implied relative distributional weights of 70 year-olds and 20 year-olds, for a range of social discount rates that are below the market interest rate. The figure illustrates that, unless the social discount rate is very close to the market interest rate, there are substantial gains from redistributing across age groups. For example, a social discount rate that is 1.5 percentage points below the market interest rate implies that it is socially desirable to reduce the consumption of a 70 year-old by \$1 in order to increase the consumption of a 20 year-old by 50 cents.

5 A utilitarian perspective

Proposition 2 suggests a new way to think about the relationship between the market interest rate and the social discount rate. Along the balanced growth path, the difference between them depends on the desirability of redistributing consumption from the current old to the current young.

The most common approach for assessing the desirability of redistribution is utilitarian. According to the utilitarian view, a transfer between two individuals is socially desirable if and only if it increases the sum of their utilities.

Different variants of the utilitarian approach adopt different interpretations of utility. Here, I follow most of the economics profession and identify utility with pleasure or lack of suffering.¹⁴ This notion of utility is inherently time-separable: pleasure and suffering are experienced in time, so it is possible to say, for example, how much pleasure each person experiences in a given unit of time and how much suffering there is in the world in a given time period.

I consider a social objective that aims to maximize the sum of utilities in each period. For clarity, I refer to this objective as the within-period utilitarian objective. As I show below, this objective does not necessarily imply an objective of maximizing the sum of utilities across time.

¹⁴This is the classical, hedonistic utilitarian view, put forth by Jeremy Bentham and John Stuart Mill. See <https://plato.stanford.edu/entries/utilitarianism-history>.

I assume that people have discounted-utilitarian preferences, which can be represented by a utility function of the form:

$$U(\mathbf{c}^t) = \sum_{a=1}^T \beta^a u_a(c_a^t) \quad (7)$$

where $\beta > 0$ is the subjective discount factor. Here, $u_a(c_a^t)$ is interpreted as generation t 's utility experienced at age a .¹⁵

It is worth highlighting the impossibility of inferring the age-specific utility functions, $\{u_a\}_{a=1}^T$, from consumer choice data. In the discounted-utility framework, inter-temporal choices depend both on the age-specific utility functions, and on the discount rate, β . Because of this, they are not separately identified. For example, a consumer with a discount factor of $\beta < 1$ and age-specific utility functions of $\{u_a\}_{a=1}^T$ is observationally-equivalent to a consumer with a discount factor of $\tilde{\beta} = 1$ and age-specific utility functions of $\{\tilde{u}_a = \beta^a u_a\}_{a=1}^T$. However, from a normative-utilitarian perspective, the two are not equivalent because they imply different mappings between consumption and utility at each age.

The Pareto principle requires that social preferences over each individual's consumption stream coincide with that individual's intertemporal preferences. Thus, Pareto requires that equation 7 also represents social preferences with respect to each generation's consumption stream. This means that society must discount future utilities at the subjective discount rate. An implication is that, if $\beta < 1$, then Pareto is inconsistent with an objective of maximizing the sum of utilities across time. It is worth highlighting that this result is not coming from any assumptions about how much society should care about the current generation relative to future generations; rather, it is a necessary implication of the Pareto condition.

In each period, the within-period utilitarian objective is to maximize the

¹⁵This assumes that utility from consumption is time-separable: the pleasures and discomforts experienced in this period depend only on the consumption in this period and on age, but not on how much was consumed in any previous period, or how much is expected to be consumed in the future.

sum of flow utilities, $\sum_{a=1}^T u_a(c_a^{\tau-a+1})$. As discussed in Schneider et al. [2012], this objective is inconsistent with a Paretian social preference relation that cares about all generations equally. To see this, note that, to be consistent with the within-period utilitarian objective, the social preference relation represented by equation 6 must be time-separable. Hence, ϕ must be linear. Consistency with the Pareto principle requires that $1/(1 + \rho^s) = \beta$: otherwise, society’s preferences over each individual’s consumption streams are different from the individual’s preferences. Consequently, the social welfare function in equation 6 must take the form

$$W(\mathbf{c}) = \sum_{t=0}^{\infty} \beta^t U(\mathbf{c}^t) \quad (8)$$

If $\beta < 1$, this social objective discounts the lifetime utilities of future generations. While many would find this ethically unacceptable (Ramsey [1928]), it is an unavoidable implication of the Pareto principle, which prevents placing equal values on the utilities on the same person experienced in different time periods. In a utilitarian framework, the only way to avoid this implication while maintaining consistency with the Pareto principle is to assume that $\beta = 1$.

In non-overlapping generation models with intergenerational altruism, social preferences that are represented by equation 8 always imply social discount rates that are equal to the market interest rate.¹⁶ This is not the case in overlapping generation models. Proposition 2 focuses on the social discount rate

¹⁶Intergenerational altruism has been central to the discussion of the social discount rate. Barro [1974], Bernheim [1989] and Arrow [1999] point out that, if each generation cares about the well-beings of its descendants, then the market interest rate reflects the concern of the present generation for the well-beings of future generations. A policy objective of maximizing the well-being of the present generation implies a social discount rate that is equal to the market discount rate. However, it is possible to argue that policy has an ethical obligation towards future generations that goes beyond the extent to which their ancestors care about them. As Dasgupta [1974] and Farhi and Werning [2007] show, this implies a social discount rate that is lower than the market interest rate. Because of this, in non-overlapping generations models with intergenerational altruism, the “normative” rate of discount is typically perceived as being lower than the market rate (see also Ramsey [1928], Arrow [1999] and Stern [2008]).

between two individuals who are the same age but alive in different periods. Appendix C illustrates that, in an overlapping generations model, this social discount rate may be different from the equilibrium market interest rate, even when the social welfare function is given by equation 8. Consequently, assuming this social welfare function is not equivalent to assuming that the social discount rate is equal to the market interest rate.

Given a within-period utilitarian objective, the redistributive weights, $\mu_{a,a'}^\tau$, are given by the ratio of marginal utilities of consumption:

$$\mu_{a,a'}^\tau = \frac{u'_a(c_a^{\tau-a+1})}{u'_{a'}(c_{a'}^{\tau-a'+1})} \quad (9)$$

Redistributive weights depend on two things: the equilibrium distribution of consumption across age groups, and the age-specific mappings between consumption and utility.

5.1 The distribution of consumption across age groups


I estimate the distribution of consumption across age groups using data from the US Consumer Expenditure Survey, 2019 (henceforth, CEX). The CEX contains household-level expenditure data for a representative sample of the US population. I define consumption expenditure as total household expenditure net of cash contributions and expenditures on education, personal insurance and pensions.¹⁷

Figures 3 and 4 plot the variation in average consumption across age groups, and Table 1 reports the slope of the relationship between consumption

¹⁷Equivalently, consumption expenditure is defined as the sum of the following expenditure categories: food, alcoholic beverages, housing, apparel and services, transportation, healthcare, entertainment, personal care products and services, reading, tobacco products and smoking supplies, and miscellaneous. The reason for excluding expenditure on education is that it is usually considered a form of investment in human capital, rather than consumption (although some people enjoy learning new things). The reason for excluding cash contributions is that this is money devoted to primarily the consumption of others (although people may derive some utility from their altruism). The reason for excluding personal insurance and pensions is that this is spending that does not affect current experienced utility, but rather utility at different states or time periods (although the peace of mind from having insurance and savings may increase utility today).

and age. Absent any controls, the relationship is slightly negative, suggesting that old people tend to consume less than younger ones. However, as consumption is measured at the household level, it is affected by predictable variation in household size over the lifecycle. Adding controls for household composition removes this life-cycle variation, producing a 0 slope between consumption and age.

Figure 3: Household consumption expenditure by age group




Calculations based on CEX. The chart shows average household consumption levels, grouped according to the age of the reference person. Consumption levels correspond to consumptions in the previous quarter, for all quarters in 2019. Error bars correspond to 95% confidence intervals. The sample consists of 21,570 households with a reference person between the ages of 20 and 90.

5.2 The utility functions

It is usually assumed that the mapping from consumption to utility does not depend on age. This assumption is explicit, for example, in Hall et al. [2020], who assume that $u_a(c) = \bar{u} + c^{1-\gamma}/(1-\gamma)$. In this case,

$$\mu_{a+1,a}^\tau = \left(\frac{c_a^{\tau-a+1}}{c_{a+1}^{\tau-a}} \right)^\gamma \quad (10)$$

Figure 4: Household consumption expenditure by age group: married couples with no kids



See notes from Figure 3. The sample here is restricted to households of married couples with no kids. This reduces the sample size to 5,036 households.

Table 1: The relationship between consumption and age

	(1)	(2)
Age	-0.003 (-0.003, -0.002)	0 (-0.001, 0.001)
Family type controls	No	Yes

Note: The dependent variable is log consumption expenditure, where consumption expenditure is defined as total expenditure in the last quarter minus expenditures on education, cash contributions and personal insurance. Age corresponds to the age of the reference person, and family type controls are based on the family type variables in the CEX. Brackets correspond to 95% confidence intervals. The regressions are based on 21,570 observations, and weighted by the CEX sampling weights.

Taking logs, it follows that

$$\ln(\mu) = \ln(\mu_{a+1,a}^\tau) = -\gamma (\ln(c_{a+1}^{\tau-a}) - \ln(c_a^{\tau-a+1})) \quad (11)$$

The expression on the right hand side is a product of two terms: $(-\gamma)$, which is the rate at which the marginal utility of consumption declines, and $(\ln(c_{a+1}^{\tau-a}) - \ln(c_a^{\tau-a+1}))$, which is the log-difference in the consumptions of two people born one year apart. In the expected-discounted utility framework, the first term can be calibrated based on the elasticity of intertemporal substitution, or the coefficient of relative risk aversion. The second term is estimated in Table 1. Using the upper- and lower-bounds of the 95% confidence intervals, and specifying $\gamma = 2$ (as in Hall et al. [2020])¹⁸ we have that (with 95% confidence),

$$-2 * (0.001) = -0.002 \leq \ln(\mu) \leq 0.006 \leq -2 * (-0.003) \quad (12)$$

By Proposition 2, $\ln(\mu) = \ln(1 + r^s) - \ln(1 + r) \approx r^s - r$. Thus, these estimates suggest a social discount rate that is at most 60 basis points above the market interest rate, and at most 20 basis points below the market interest rate:

$$-0.002 \leq r^s - r \leq 0.006 \quad (13)$$

This calibration illustrates that, when flow utility functions are independent of age, then there are limited gains from redistributing from old to young. An implication of Proposition 2 is that, in this case, the within-period utilitarian objective is inconsistent with the combination of the Pareto principle and social discount rates that are significantly below the market interest rate.

5.3 The assumption of age-independent utility functions

The calibration relies crucially on the assumption that utility functions are independent of age. Below I discuss three approaches for evaluating this assumption.

¹⁸The choice of $\gamma = 2$ is based on standard estimates of the coefficient of relative risk aversion. As this model abstracts from risk, one might argue that a calibration of γ based on the elasticity of intertemporal substitution is more appropriate. Estimates of the elasticity of intertemporal substitution range from 0.1–1 (see, for example, Guvenen [2006]), implying $\gamma \in [1, 10]$. This range roughly corresponds to the range of estimates of the coefficient of relative risk aversion.

To frame the discussion, it is useful to allow for two types of age-variation in flow utility functions: one that depends on consumption, and one that does not. Generalizing Hall et al. [2020], I assume that age-specific utility functions are given by:

$$u_a(c) = \bar{u}_a + A^a \frac{c^{1-\gamma}}{1-\gamma} \quad (14)$$

The constant, \bar{u}_a , is a term that depends on age but not on consumption. The term A^a is the age-specific productivity of transforming consumption to utility, which is assumed to be exponential in age.¹⁹ Given this generalized functional form, equation 10 becomes

$$\mu = \mu_{a+1,a}^\tau = A \left(\frac{c_a^{\tau-a+1}}{c_{a+1}^{\tau-a}} \right)^\gamma \quad (15)$$

Given that consumption is similar across age groups, it follows that

$$\mu = A \quad (16)$$

The above expression illustrates that, in order to compute the distributional weights of different age groups, it is necessary to understand how age affects people's ability to transform consumption into utility.

The first way to assess this is using measures of subjective well-being. Because consumption expenditure is uncorrelated with age, the hypothesis that utility functions are independent of age implies that utility should be uncorrelated with age. Subjective well-being measures proxy for utility by estimating people's overall happiness. If \bar{u}_a is independent of age and $A = 1$, then happiness should be independent of age.

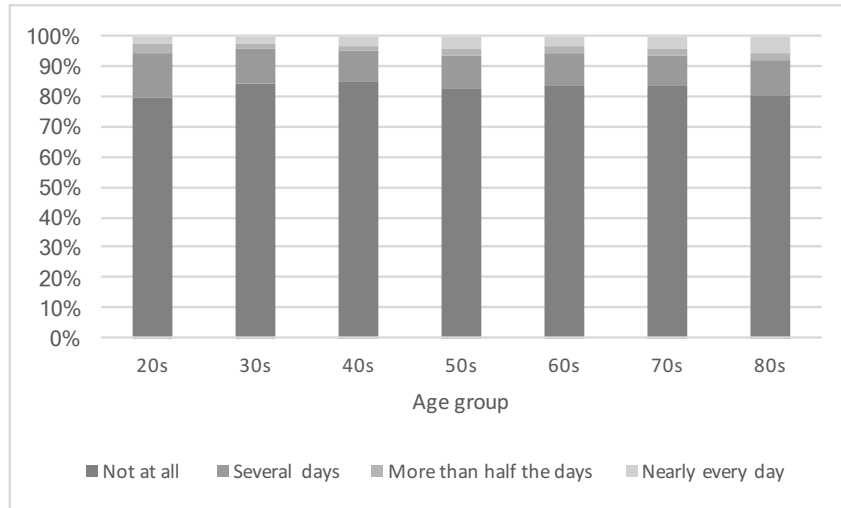
Unfortunately, the empirical relationship between subjective well-being and age is inconclusive (see Ulloa et al. [2013]). Some studies document no relationship, while others find increasing, decreasing, or non-monotone relationships. Many studies document a U-shaped relationship between well-being and age (with the lowest level of well-being usually experienced around age 40 or 50).

¹⁹The assumption that productivity is exponential in age is necessary because, by Proposition 2, $\mu_{a+1,a}^\tau$ must be independent of a .

This rejects the joint hypothesis that \bar{u}_a is independent of age, *and* that $A = 1$. At the same time, it is not consistent with a framework in which \bar{u}_a is independent of age and $A \neq 1$, because this would imply a monotone relationship between happiness and age.

A more direct way of testing the hypothesis that $A = 1$ is to use measures of subjective well-being that relate directly to people’s ability to transform consumption to enjoyment. The National Health Interview Survey (henceforth NHIS) uses various questions to evaluate the mental health of the US population. One elicits the frequency in which respondents have “little interest or pleasure in doing things”. This is a subjective measure of how well people are able to convert “doing things” into pleasure, which is a reasonable proxy for A . As illustrated in Figure 5, the responses are very similar across age groups, consistent with the hypothesis that $A \approx 1$.

Figure 5: Frequency of little interest or pleasure in doing things, by age



The chart shows the age-specific distributions of responses to the question: “Over the last two weeks, how often have you been bothered by little interest or pleasure in doing things?” (NHIS, 2019). Proportions are estimated using the NHIS sample weights. Sample: 31,997 adults.

A second approach for evaluating the assumption of age-independent utility functions is to look at expenditure categories. Intuitively, the marginal utility

of consumption is higher when the marginal dollar is spent on necessities. For example, if a transfer from person A to person B reduces the size of A’s yacht but increases the quality of B’s medical care, then it is reasonable to conjecture that this transfer is welfare-improving. On an intuitive level, the marginal utility of consumption is related to the propensity to spend the marginal dollar on necessities.

Figure 6 documents expenditure shares by age. Older people spend more on medical care, while expenditures on other necessities, such as food and housing, are roughly the same across age groups. Taken together, this suggests that, if anything, older people devote relatively more of their consumption expenditure to meeting basic needs, and less to luxury goods. This suggests that $A \geq 1$.


Another way to get a sense of how different age groups spend their marginal dollars is to document how consumption expenditure varies with income at different age groups. The results are summarized in Figure 7. The results show that the marginal dollar is spent similarly across age groups, except that the marginal propensity to spend on healthcare is increasing with age. This is consistent with the hypothesis that $A \geq 1$.

A final approach for evaluating the assumption of age-independent utility functions is based on sensory sensitivities.²⁰ At least some forms of consumption are converted to utility through the five senses: vision, hearing, taste, smell, and touch. For example, the utility from listening to music depends on the ability to hear it. One concern is that the process of aging reduces sensory sensitivities, making it harder for older people to convert consumption expenditure into enjoyment. This would imply that $A < 1$.

There is ample evidence that people’s abilities to see, hear, taste and smell deteriorate with age (the effects of aging on sensitivity to touch or to pain are more mixed; see Schieber [1992]). Figures 8 and 9 document self-reported dif-

²⁰This approach is inspired by Argenziano and Gilboa [2019], who propose measuring marginal utilities based on “just noticeable differences” - the minimal increase in consumption that each person can notice. Intuitively, if a transfer from A to B is not noticed by A but is noticed by B, then B is better off and A is just-as-well off, so we have a Pareto improvement. The ability to detect changes in consumption levels is related to sensory sensitivity; for example, people with a better sense of taste are more likely to notice a small improvement in the flavor of their food.

Figure 6: Expenditure shares by age




Expenditure shares by age group, based on data from the CEX. Darker shades correspond to expenditure categories that represent higher shares of the consumption expenditures of lower-income households (goods that are more “necessary” in the consumer-theory sense). Proportions are estimated using the CEX sample weights. The sample consists of 21,570 households with a reference person between the ages of 20 and 90.

faculties in hearing and seeing by age, based on the NHIS. The figures broadly confirm that sensory sensitivities deteriorate with age. However, they also reveal that the deterioration takes place very gradually.

The extent to which declining sensory sensitivities can be used to support the view that $A < 1$ depends on the importance of sensory sensitivities for transforming consumption to utility. For example, the enjoyment of music not only depends on the ability to hear it well, but also on the emotional ability to relate to it – which may improve with age, as people have more life experiences to draw from.

To conclude, I have discussed three approaches for evaluating the hypothesis that $A \approx 1$. The first is based on subjective well-being measures, which are broadly consistent with this hypothesis. The second is based on the propensity to spend on necessities, which is somewhat increasing with age, suggesting that $A \geq 1$. The third is based on the relationship between age and sensory sensitivities. Because some sensory sensitivities are decreasing with age, this

Figure 7: Marginal expenditures by age

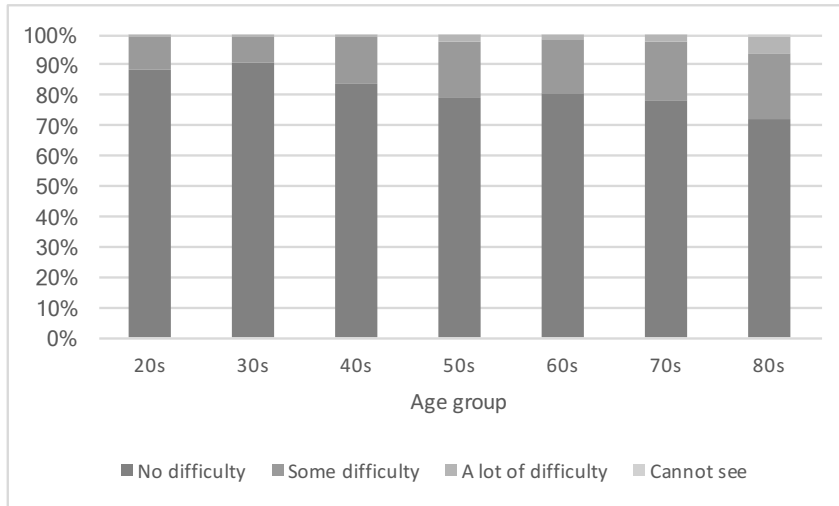


The chart illustrates the age-specific relationships between expenditure categories and total expenditure amounts. For each category, the bar size corresponds to the regression coefficient from regressing expenditure on that category on income. Darker shades correspond to expenditure categories that represent higher shares of the consumption expenditures of lower-income households (goods that are more “necessary” in the consumer-theory sense). Proportions are estimated using the CEX sample weights. The sample consists of 21,570 households with a reference person between the ages of 20 and 90.

approach is consistent with the view that $A < 1$. However, more work is necessary in order to understand the mapping between sensory sensitivities and A .

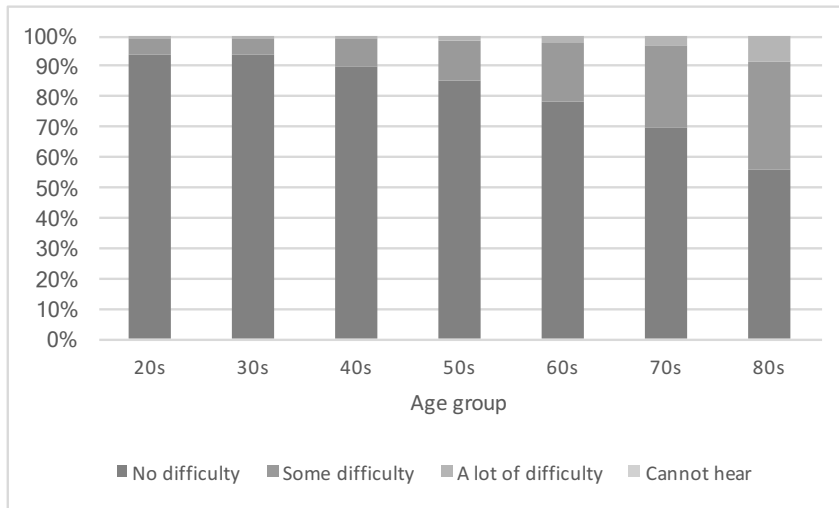
In practice, utility comparisons are often made based on introspection (see De la Croix and Doepke [2021]). We can reflect on whether our ability to extract enjoyment out of consumption expenditure changes as we age. If, upon reflection, we can agree that changes are very gradual, then $A \approx 1$. In this case, the within-period utilitarian objective is inconsistent with social discount rates that are substantially below the market interest rate.

Figure 8: Difficulty seeing by age



The chart shows the age-specific distributions of responses to the question: “Do you have difficulty seeing (even when wearing glasses or contact lenses/seeing, if applicable)?” (NHIS, 2019). Proportions are estimated using the NHIS sample weights. Sample: 31,997 adults.

Figure 9: Difficulty hearing by age



The chart shows the age-specific distributions of responses to the question: “Do you have difficulty hearing (even when using your hearing aid, if applicable)?” (NHIS, 2019). Proportions are estimated using the NHIS sample weights. Sample: 31,997 adults.

6 Conclusion

This paper establishes an equivalence between two normative questions. The first is, how should policy discount future returns? This question is relevant, for example, for evaluating the social gains from climate change mitigation. In a discounted-utilitarian framework, the answer to this question depends on the social rate of pure time preference – the rate at which society should discount the utilities of future generations.

The second question is, how should society distribute resources across people of different age groups? This question became particularly contentious during the COVID-19 pandemic (see, for example, Hall et al. [2020]). Containment measures disproportionately benefitted the elderly, who were at higher risk from the virus. However, some of the costs were born by children and working-age adults, who suffered serious disruptions. This raised the question of how to tradeoff benefits to the elderly with costs to younger people. This question also comes up during normal times, when policy makers face budgetary tradeoffs between programs that benefit the elderly (such as social security and medicaid) and programs that benefit younger people (such as childcare subsidies and playgrounds).

The key result is that, if the social discount rate is lower than the market interest rate, then it is socially desirable to increase the consumption of the young at the expense of the elderly. Because, empirically, consumption tends to be uncorrelated with age, this normative prescription is difficult to square with a social objective of maximizing the sum of utilities in each period.

There are several possibilities forward. First, we can maintain that the social discount rate is lower than the market interest rate, and accept that it is socially desirable to transfer resources from the current old to the current young. Quantitatively, even a small deviation of the social discount rate from the market interest rate implies large welfare gains from redistribution across age groups. An implication is that the government should substantially cut social security benefits and medicare in order to increase spending on schools, child care, and playgrounds.

A second possibility is to maintain that the current (equal) distribution of resources across age groups is close to optimal, and accept that the social discount rate is close to the market interest rate. In a discounted-utilitarian framework, this requires discounting the lifetime utilities of future generations. However, a more flexible social welfare function may deliver this implication without discounting the utilities of future generations. For example, a prioritarian objective that is averse to inequality in lifetime utilities may have a zero rate of pure time preference, and still imply social discount rates that are equal to the market interest rate.

A third possibility is to reject the Pareto principle. Crucially, if people discount their own future utilities, then a (non-discounted) utilitarian objective is inconsistent with the Pareto principle: Pareto requires society to have the same preferences as individuals with respect to their own consumption streams. If people discount their own future utilities, society must do so as well. Maintaining that society should care equally about utilities experienced in different times requires a rejection of the Pareto condition. A social objective of maximizing the sum of all current and future utilities violates the Pareto condition, but implies both that the current distribution of resources across age groups is roughly optimal, and that the social discount rate is lower than the market interest rate.

Finally, it is worth cautioning that the quantitative relationship between social discounting and the desirability of age-based redistribution changes as we move away from the balanced growth path. A growth slowdown, as suggested by Jones [2021], may imply long-run social discount rates that are below the market interest rate, even if there are no welfare gains from current redistribution between old and young. Furthermore, as pointed out by Stern [2008], the social discount rate is a useful statistic only for evaluating the social desirability of marginal inter-temporal changes along a given growth trajectory. In some instances, current saving may affect the growth path itself; for example, uncontrolled climate change may set humanity off on a different growth path. In this case, social discount rates along the previous balanced-growth path are of limited practical use.

References

- Matthew D. Adler. *Measuring Social Welfare: an Introduction*. Oxford University Press, 2019.
- Rossella Argenziano and Itzhak Gilboa. Perception-theoretic foundations of weighted utilitarianism. *The Economic Journal*, 129(620):1511–1528, 2019.
- 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.
- Kenneth J. Arrow. Discounting, morality, and gaming. *Discounting and intergenerational equity*, pages 13–21, 1999.
- Geir B. Asheim and Stéphane Zuber. A complete and strongly anonymous leximin relation on infinite streams. *Social Choice and Welfare*, 41(4):819–834, 2013.
- Robert J. Barro. Are government bonds net wealth? *Journal of political economy*, 82(6):1095–1117, 1974.
- B. Douglas Bernheim. Intergenerational altruism, dynastic equilibria and social welfare. *The review of economic studies*, 56(1):119–128, 1989.
- Guillermo A. Calvo and Maurice Obstfeld. Optimal time-consistent fiscal policy with finite lifetimes. *Econometrica*, 56(2), March 1988.
- Andrew Caplin and John Leahy. The social discount rate. *Journal of political economy*, 112(6):1257–1268, 2004.
- Partha Dasgupta. On some alternative criteria for justice between generations. *Journal of Public Economics*, 3(4):405–423, 1974.
- David De la Croix and Matthias Doepke. A soul’s view of the optimal population problem. *Mathematical Social Sciences*, 112:98–108, 2021.

- 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.
- Christian Gollier and James K. Hammitt. The long-run discount rate controversy. *Annual Review of Resource Economics*, 6:273–295, 2014.
- Hilary Greaves. Discounting for public policy: A survey. *Economics & Philosophy*, 33(3):391–439, 2017.
- Fatih Guvenen. Reconciling conflicting evidence on the elasticity of intertemporal substitution: A macroeconomic perspective. *Journal of Monetary Economics*, 53(7):1451–1472, 2006.
- 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.
- Charles I. Jones. The past and future of economic growth: A semi-endogenous perspective. *Annual Review of Economics*, 2021. Forthcoming.
- J. Paul Kelleher. Pure time preference in intertemporal welfare economics. *Economics & Philosophy*, 33(3):441–473, 2017.
- Antony Millner and Geoffrey Heal. Choosing the future: Markets, ethics, and rapprochement in social discounting. 2021.
- William D. Nordhaus. A review of the stern review on the economics of climate change. *Journal of economic literature*, 45(3):686–702, 2007.
- Frank Plumpton Ramsey. A mathematical theory of saving. *The Economic Journal*, 38(152):543–559, 1928.
- John Rawls. Some reasons for the maximin criterion. *The American Economic Review*, 64(2):141–146, 1974.

Frank Schieber. Aging and the senses. In *Handbook of mental health and aging*, pages 251–306. Academic Press, 1992.

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.

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

Beatriz Fabiola López Ulloa, Valerie Møller, and Alfonso Sousa-Poza. How does subjective well-being evolve with age? a literature review. *Journal of Population Ageing*, 6(3):227–246, 2013.

Stéphane Zuber and Geir B. Asheim. Justifying social discounting: the rank-discounted utilitarian approach. *Journal of Economic Theory*, 147(4):1572–1601, 2012.

A Proof of Proposition 1

$$MRS((a, t), (a', t')) = \frac{\frac{\partial W(\mathbf{c})}{\partial c_a^t}}{\frac{\partial W(\mathbf{c})}{\partial c_{a'}^{t'}}} = \frac{\frac{\partial W(\mathbf{c})}{\partial c_a^t}}{\frac{\partial W(\mathbf{c})}{\partial c_1^{t+a-1}}} \frac{\frac{\partial W(\mathbf{c})}{\partial c_1^{t+a-1}}}{\frac{\partial W(\mathbf{c})}{\partial c_2^{t+a-1}}} \frac{\frac{\partial W(\mathbf{c})}{\partial c_2^{t+a-1}}}{\frac{\partial W(\mathbf{c})}{\partial c_1^{t+a}}} \cdots \frac{\frac{\partial W(\mathbf{c})}{\partial c_1^{t'-a'-2}}}{\frac{\partial W(\mathbf{c})}{\partial c_2^{t'-a'-2}}} \frac{\frac{\partial W(\mathbf{c})}{\partial c_2^{t'-a'-2}}}{\frac{\partial W(\mathbf{c})}{\partial c_{a'}^{t'}}} =$$

$$\mu_{a,1}^{t+a-1} (1 + r_{1,2,t+a-1}^{si}) \mu_{2,1}^{t+a} \cdots (1 + r_{1,2,t'+a'-2}^{si}) \mu_{2,a'}^{t'+a'-1}$$

Rearranging the terms yields the desired expression.

B Proof of Proposition 2

The derivative of equation 6 with respect to c_a^t is

$$\frac{\partial W(\mathbf{c})}{\partial c_a^t} = \left(\frac{1}{1 + \rho^s} \right)^t \phi'(U(\mathbf{c}^t)) \frac{\partial U(\mathbf{c}^t)}{\partial c_a^t} \quad (17)$$

Along the balanced growth path, $c_a^t = (1+g)^t c_a^0$. As U is homogeneous, there exists some η such that $U(\mathbf{c}^t) = U((1+g)^t \mathbf{c}^0) = (1+g)^{\eta t} U(\mathbf{c}^0)$. For this η , it holds that

$$\frac{\partial U(\mathbf{c}^t)}{\partial c_a^t} = (1+g)^{(\eta-1)t} \frac{\partial U(\mathbf{c}^0)}{\partial c_a^0} \quad (18)$$

Because ϕ is homogeneous, so is ϕ' , and hence there exists ζ such that $\phi'(U(\mathbf{c}^t)) = \phi'((1+g)^{\eta t} U(\mathbf{c}^0)) = (1+g)^{\eta \zeta t} \phi'(U(\mathbf{c}^0))$. Substituting yields

$$\frac{\partial W(\mathbf{c})}{\partial c_a^t} = \left(\frac{1}{1+\rho^s} \right)^t (1+g)^{\eta \zeta t} \phi'(U(\mathbf{c}^0)) (1+g)^{(\eta-1)t} \frac{\partial U(\mathbf{c}^0)}{\partial c_a^0} \quad (19)$$

By the Euler equation (for generation 0),

$$\frac{\partial U(\mathbf{c}^0)}{\partial c_a^0} (1+r)^{a-1} = \frac{\partial U(\mathbf{c}^0)}{\partial c_1^0} \quad (20)$$

Substituting yields

$$\frac{\partial W(\mathbf{c})}{\partial c_a^t} = \left(\frac{1}{1+\rho^s} \right)^t (1+g)^{\eta \zeta t} \phi'(U(\mathbf{c}^0)) (1+g)^{(\eta-1)t} \left(\frac{1}{1+r} \right)^{a-1} \frac{\partial U(\mathbf{c}^0)}{\partial c_1^0} \quad (21)$$

Generation t is aged a in period $\tau = t + a - 1$. Substituting yields

$$\frac{\partial W(\mathbf{c})}{\partial c_a^t} = \left(\frac{1}{1+\rho^s} \right)^{\tau-a+1} (1+g)^{(\tau-a+1)\eta \zeta} \phi'(U(\mathbf{c}^0)) (1+g)^{(\eta-1)(\tau-a+1)} \left(\frac{1}{1+r} \right)^{a-1} \frac{\partial U(\mathbf{c}^0)}{\partial c_1^0}$$

In this period, the people aged \tilde{a} are from generation \tilde{t} , where $\tilde{t} = \tau - \tilde{a} + 1$.

Substituting into the above expression implies that

$$\frac{\partial W(\mathbf{c})}{\partial c_a^{\tilde{t}}} = \left(\frac{1}{1+\rho^s} \right)^{\tau-\tilde{a}+1} (1+g)^{(\tau-\tilde{a}+1)\eta \zeta} \phi'(U(\mathbf{c}^0)) (1+g)^{(\eta-1)(\tau-a+1)} \left(\frac{1}{1+r} \right)^{\tilde{a}-1} \frac{\partial U(\mathbf{c}^0)}{\partial c_1^0}$$

and hence

$$\mu_{a,\tilde{a}}^t = \frac{\frac{\partial W(\mathbf{c})}{\partial c_a^t}}{\frac{\partial W(\mathbf{c})}{\partial c_a^{\tilde{t}}}} = \left(\frac{1}{1+\rho^s} \right)^{\tilde{a}-a} (1+g)^{(\tilde{a}-a)\eta \zeta} (1+g)^{(\eta-1)(\tilde{a}-a)} \left(\frac{1}{1+r} \right)^{a-\tilde{a}} \quad (22)$$

$$= \left(\left(\left(\frac{1+r}{1+\rho^s} \right) (1+g)^{\eta-1+\eta\zeta} \right)^{-1} \right)^{a-\bar{a}}$$

setting

$$\mu = \left(\left(\frac{1+r}{1+\rho^s} \right) (1+g)^{\eta-1+\eta\zeta} \right)^{-1}$$

concludes the proof of the first part of the proposition.

To prove the second part of the proposition, note that, by equation 21,

$$\frac{\frac{\partial W(\mathbf{c})}{\partial c_a^t}}{\frac{\partial W(\mathbf{c})}{\partial c_a^{t'}}} = \frac{\left(\frac{1}{1+\rho^s} \right)^t (1+g)^{t\eta\zeta} (1+g)^{(\eta-1)t}}{\left(\frac{1}{1+\rho^s} \right)^{t'} (1+g)^{t'\eta\zeta} (1+g)^{(\eta-1)t'}} = \frac{(\mu(1+r))^{-t}}{(\mu(1+r))^{-t'}} = (\mu(1+r))^{t'-t} \quad (23)$$

concluding the proof.

C The equilibrium of the overlapping generations model

This section solves for the equilibrium balanced growth path in a simple overlapping generations model. The purpose of this illustration is to show that, even when the social objective is given by equation 8, the social discount rate may be different from the equilibrium market interest rate.

To simplify, assume that each generation lives for two periods only ($a = 1, 2$). Labor is the only factor of production. Labor endowments vary with age, and are given by l_1 and l_2 . The aggregate production at time t is given by:

$$Y_t = A_t L_t \quad (24)$$

where $L_t > 0$ is aggregate labor and $A_t > 0$ is aggregate productivity.

Firms are competitive so wages, w_t , are given by the marginal products of labor:

$$w_t = A_t \quad (25)$$

Let r_t denote the interest rate at time t . Each generation maximizes its discounted utility, subject to its budget constraints:

$$\max_{c_1^t, c_2^t, s^t} u_1(c_1^t) + \beta u_2(c_2^t) \text{ s.t. } c_1^t + s^t = w_t l_1 \text{ and } c_2^t = w_{t+1} l_2 + (1 + r_t) s^t$$

Here, s^t denotes household saving. Note that the old cannot borrow from the young, because they won't be around to repay. The young can't borrow from the old, because the old have no interest in future repayment. Consequently, market clearing requires that, for each t , $s^t = 0$, yielding

$$c_1^t = w_t l_1 \text{ and } c_2^t = w_{t+1} l_2 \quad (26)$$

The market interest rate is then determined by the Euler condition:

$$u'_1(c_1^t) = \beta u'_2(c_2^t)(1 + r_t) \quad (27)$$

Assuming that $u'_a(c) = \zeta_a c^{-\gamma}$ (to maintain the assumption of homothetic preferences), substituting yields

$$\begin{aligned} \zeta_1 (c_1^t)^{-\gamma} &= \beta \zeta_2 (c_2^t)^{-\gamma} (1 + r_t) \Rightarrow \zeta_1 (w_t l_1)^{-\gamma} = \beta \zeta_2 (w_{t+1} l_2)^{-\gamma} (1 + r_t) \quad (28) \\ \Rightarrow (1 + r_t) &= \frac{\zeta_1 (w_t l_1)^{-\gamma}}{\beta \zeta_2 (w_{t+1} l_2)^{-\gamma}} = \frac{\zeta_1 (A_t l_1)^{-\gamma}}{\beta \zeta_2 (A_{t+1} l_2)^{-\gamma}} \end{aligned}$$

Assuming that productivity grows at a rate g_A , this yields

$$(1 + r_t) = \frac{\zeta_1 l_1^{-\gamma}}{\beta \zeta_2 l_2^{-\gamma}} (1 + g_A)^\gamma$$

In particular, r_t is constant over time.

By Proposition 2, the social discount rate is equal the market interest rate if and only if the distribution of consumption across age groups is socially optimal. Given a within-period utilitarian social objective (as in equation 8), this holds if and only if the distribution of consumption across age groups is socially optimal. This holds if and only if the marginal utilities of consumption

are equalized across old and young:

$$r^s = r \Leftrightarrow u'_1(c_1^t) = u'_2(c_2^{t-1}) \Leftrightarrow \zeta_1(w_t l_1)^{-\gamma} = \zeta_2(w_t l_2)^{-\gamma} \Leftrightarrow \zeta_1 l_1^{-\gamma} = \zeta_2 l_2^{-\gamma} \quad (29)$$

For example, if $\zeta_1 = \zeta_2$ (flow utility functions are independent of age), then this condition is satisfied if and only if labor endowments are the same in both periods of life. Otherwise, the equilibrium distribution of consumption across age groups is not socially optimal, and the social discount rate is different from the market interest rate.