ANNUAL REVIEWS

Annual Review of Financial Economics

Missing Participants, Missing Markets, and the Social Discount Rate: Borrowing Constraints, Intergenerational Transfers, Altruism, and the Desire for Legacy

Andrew Caplin[1] and John Leahy[2,3]

[1]Department of Economics and Center for Experimental Social Science, New York University, New York, NY, USA; email: andrew.caplin@nyu.edu

[2]Department of Economics and Gerald R. Ford School of Public Policy, University of Michigan, Ann Arbor, Michigan, USA

[3]National Bureau of Economic Research, Cambridge, Massachusetts, USA

Annu. Rev. Financ. Econ. 2023. 15:165–84

First published as a Review in Advance on August 16, 2023

The *Annual Review of Financial Economics* is online at financial.annualreviews.org

https://doi.org/10.1146/annurev-financial-101822-112948

JEL codes: E7, H43

Keywords

social discount rate, incomplete markets, altruism, legacy

Abstract

Market interest rates reflect the preferences of market participants. When market participants are missing, the average discount rate in the population may therefore differ from the market rate. Missing current market participants, such as constrained borrowers, tends to imply an average discount rate that is above the market rate, whereas missing future market participants, such as future generations, tends to imply an average rate below the market rate. Nonetheless, a government with the ability to transfer wealth intratemporally across agents will generally wish to use the market interest rate as a guide to policy. One robust argument for the use of a lower social discount rate is intrapersonal: Future selves discount the past, whereas current selves discount the future. Legacy utility may also justify a low social discount rate.

1. INTRODUCTION

The social discount rate is the rate at which society compares costs and benefits that occur in the future to costs and benefits that occur today. The appropriate discount rate is a very important component in policy evaluation, and given the power of compound interest, this is especially the case when projects or policies have long horizons. Prominent examples of projects or policies with long horizons include climate change mitigation, infrastructure such as roads and airports, education reform, and social security. In each of these cases, the relative weights placed on costs and benefits at different points in time may drastically affect the desirability of any given project or policy proposal. For example, the conclusion in the Stern Review that the failure to address climate change could cost upwards of 5% of global GDP rests crucially on the assumed discount rate of 0.1% (Nordhaus 2006).

In this article, we consider a government that is contemplating a project that has immediate costs and future benefits. We assume that the government seeks to maximize some weighted average of the utilities of the agents in the economy. We evaluate the desirability of relying on market prices to measure the social trade-off between resources today and resources in the future. We highlight the role that missing market participants play in determining the usefulness of market prices. The reason to focus on market participation is obvious: Market prices can only reflect the preferences of market participants. It would appear natural that market prices might be a poor guide for policy when the preferences of excluded groups systematically differ from those that participate. To keep the focus on missing market participants, we keep other aspects of the environment as simple as possible. For example, we assume that the project is small, so that we may ignore general equilibrium feedback from the project to the economy as a whole, and we assume that the project has a certain payoff, so that we may abstract from issues of risk and uncertainty.

We focus, in particular, on two groups that are generally excluded from participation in asset markets: constrained potential borrowers and future generations. The intertemporal preferences of each group differ systematically from market interest rates and differ in a different way. Constrained borrowers tend to care more about the present than the future, whereas future generations tend to care more about the future than the present. Missing market participants in the form of excluded current borrowers therefore points to an average discount rate that is larger than the market rate. This would seem to imply that society as a whole should discount future benefits at a rate above the market rate. Missing market participants in the form of missing future generations, on the other hand, would seem to imply that society as a whole should discount future benefits at a rate lower than the market rate. Whether the social discount rate should be above or below the market rate would then seem to depend on which form of market imperfection is more prevalent.

The case for deviating from market rates, however, is not so simple. The argument of the previous paragraph holds if the costs and benefits of the government's project are borne by all agents in the economy so that the social discount rate reflects some average of individual discount rates. Governments, however, have tools that allow them to concentrate the benefits on those who value them most and to concentrate the costs on those most able to bear these costs. Taxation is one such tool, but borrowing in asset markets is also a way to focus the costs and benefits on asset market participants.

We show that a government with sufficient tools may still want to benchmark its projects against market rates of return, proceeding with projects that exceed market returns and refraining from projects that earn returns less than the market. The presence of constrained borrowers who value the present over the future does not necessarily mean that the government should forgo a project that earns a market return. If the government can raise funding in asset markets from agents who value the future at market rates and can redistribute the benefits from the project through lump sum transfers, then any project that earns above-market returns improves welfare, while any project that earns below-market returns destroys resources. The absence of future generations who value the future over the present does not necessarily mean that the government should undertake projects that earn below-market rates of return. Following Calvo & Obstfeld (1988), we show that a government with access to lump sum transfers can always do better transferring wealth from the old to the young within periods, rather than engaging in low-return projects. The reason is that current generations value future wealth at the market rate of return, so that shifting a transfer to the future effectively earns this market return. Adding altruism in the form of a concern for future generations (Barro 1974, Farhi & Werning 2007) does not alter this conclusion.

One argument for a social discount rate below the market rate survives this critique. Future selves, like future generations, do not participate in current market choices. Caplin & Leahy (2004) argue that this may justify the use of a social discount rate that is below the market rate of interest. In their model, individuals discount past utility in much the same way that they discount future utility. For this reason, tastes change over time (Strotz 1956) as the future becomes the present and the present becomes the past. Future selves place a greater weight on future consumption than do current selves, in much the same way that future generations do. The difference is that, whereas it is possible for a government to make intratemporal transfers between overlapping generations, there is no way to make intratemporal transfers between intertemporal selves. Since the market rate of interest reflects the views of current selves and future selves value current consumption less than current selves do, a government that places weight on future perspectives will discount the future at a rate below the market rate of interest. Such a government will willingly undertake projects with below-market rates of return. This argument holds even if an agent's preferences are time consistent. If discount rates are hyperbolic, so that preferences are time inconsistent, as in the model by Phelps & Pollak (1968), then even the current generation may discount the far future at less than the market rate of interest (Barro 1999). In such cases, even a government that considers only the views of current market participants may find illiquid projects that lock in future returns desirable even if they earn returns less than the market rate of return (Laibson 1997).

The idea that how people think about the future and the past matters can be taken a bit further. One motivation that is familiar to psychology but has not received much attention in economics is the desire for legacy (Waggoner, Bering & Halberstadt 2023). If agents care about how they are remembered, then they may be willing to undertake certain projects with below-market rates of return if these projects provide sufficient memory utility in the future—whether from future generations believing that their parents did the right thing or older selves being satisfied with their lives. Alumni who donate buildings to universities and wealthy individuals who set up foundations are potential examples. This desire for legacy does not in itself motivate government intervention. Individuals would be willing to undertake such projects on their own if they had sufficient resources. Many potential projects, however, require coordinated effort or resources beyond the means of most individuals. Climate change might be one example. The desire for legacy might also have a generational component. In such cases, there might be a role for a government to undertake such a project.

We focus our survey on missing markets and missing market participants. For a general survey of the literature, the reader should see the excellent recent surveys by Groom et al. (2022) and Millner & Heal (2023) and the references therein. To simplify the presentation and focus on market participation, we abstract from many other interesting and important issues in the economics of the social discount rate. We abstract from well-known reasons to be skeptical of market interest rates, such as externalities and imperfect competition. We abstract from the general equilibrium impact of large projects. We abstract from issues related to the cost of taxation. When considering transfers between generations, we generally assume that the government has access to lump sum transfers that it can use to address concerns regarding the distribution of wealth at any given point in time. This separates intertemporal issues from intratemporal ones and keeps the focus on the intertemporal. We abstract from issues of risk and liquidity. We study situations in which there is one risk-free rate between any two periods and in which the term structure hypothesis holds when thinking long-term. Finally, we abstract from issues of political feasibility. We consider a social planner that maximizes a weighted average of the utilities of all agents: market participants and nonmarket participants, current agents and future agents. One might argue that market access is likely correlated with political access and that, in practice, governments likely reflect the views of current generations and wealthy individuals.[1] The interaction between market participation and any of these concerns is an interesting and important topic for future research.

We begin in Section 2 with the classic case for the use of market interest rates to evaluate public projects. This case rests on the fact that, in a world with all participants present and complete markets, all agents view intertemporal trade-offs in exactly the same way. If they did not, then there would be an opportunity for trade: Agents who value future consumption would lend to those who want to consume today. The market interest rate reflects this shared trade-off and therefore correctly values resources in the future relative to resources today. The simple conclusion is that a government should undertake a project if and only if its return exceeds the market interest rate. Projects that meet this criterion expand the set of economic possibilities. Projects that do not destroy resources.

In Section 3, we consider an economy in which some agents do not participate fully in asset markets and markets are incomplete. We consider a standard macroeconomic setting in which agents only trade a risk-free bond and so are unable to insure fluctuations in their income. In addition, some agents find it difficult to borrow against their human capital. Models of this type have been considered by Aiyagari (1994), Hugget (1996), and, more recently, Kaplan & Violante (2014). In these models, market interest rates are determined in large part by the preferences of rich savers and therefore reflect savers' preferences for future consumption. Borrowing-constrained households who are excluded from the market tend to value current consumption more than do asset market participants. A government that cannot target the costs and benefits of a project should therefore discount the future at a rate above the market rate in order to incorporate the preferences of these excluded market participants. A government that can target the costs and benefits, however, will generally find it in its interest to pursue only projects with returns above the market rate. In our simple model, the government can target the costs of a project by funding the project in asset markets and targeting the benefits through lump sum transfers from the agents who benefit from the project to those who financed the project. This strategy is a Pareto improvement if and only if the project earns a return that exceeds the market rate of return.

In Section 4, we consider another class of missing market participants: future generations. Discomfort with discounting the utilities of future generations lies at the heart of many arguments for low social discount rates, starting with Ramsey (1928), who, while acknowledging that private agents discount the future, argued that it was "ethically indefensible" for the government to do so (p. 543). Pigou (1929) and Allais (1947) have expressed similar views. We present a simple overlapping generations model in which future generations value future consumption more than current generations do. As in the case with excluded borrowers, such a project is a Pareto improvement if and only if the return to the project exceeds the market rate of interest. The reason is that the market interest rate between any two adjacent periods reflects the intertemporal preferences

---

[1]It is possible to design models in which governments place greater weight on the utilities of nonmarket participants. For example, Sen (1961) and Marglin (1963) have shown how altruism may lead agents to vote for policies that are more forward-looking than the choices they would make for themselves and their children.

of agents whose lives span those periods. The government can finance the project by taxing the young. Then, in each subsequent period, it can compensate those who were taxed in the past with a transfer from the new young generation until it taxes those who benefit when the return to the project is realized.

We can take this argument a step further. A government that places greater weight than the market on the utilities of future generations may be tempted to proceed with projects with returns below the market rate. Such projects are not Pareto improvements, but they may improve the government's objective if the gain to future generations outweighs the cost to current generations. Following Calvo & Obstfeld (1988) and Eden (2021), however, we show that the government can achieve the same end more efficiently through a sequence of intratemporal transfers between generations. Again, the intuition is that the market interest rate reflects the intertemporal preferences of agents whose lives span those periods, so that taking a given amount from agents today to finance a project with a below-market return is less efficient than simply transferring resources in the future.

Farhi & Werning (2007) have argued that intergenerational altruism provides a rationale for a social discount rate that is below the market interest rate. The argument is similar to that in the simple overlapping generations model. Even if current generations care about the utility of future generations, a government that places weight on the utility of both current and future generations will place even greater weight on future generations' utility. The utility of future generations will matter both directly and indirectly through that of current generations. In Section 5, however, we show that even in the presence of altruism the government can generally achieve its ends more efficiently through intratemporal transfers. As in the overlapping generations model, the intuition is that future transfers are valued by current generations at the market interest rate.

In Section 6, we consider another missing market participant: future selves. We present the model of Caplin & Leahy (2004), in which agents discount both past and future utility. In this model, an agent's utility changes over time. A self in period t values consumption in that period more than a self in earlier or later periods does. Mathematically, the model is similar to a model of intergenerational altruism in which the self is a generation. There is, however, one important difference: While generations overlap, selves are tied to a particular period in time. Intratemporal transfers cannot reallocate resources across intertemporal selves. Hence, even a government that has access to lump sum transfers cannot execute a sequence of transfers that makes all agents better off at the market rate of interest. The model of Caplin & Leahy (2004) therefore presents a robust argument for a social discount rate that is below the market rate of interest. Note, however, that these low-return projects are desirable because the government places greater weight on the perspectives of future selves than the market does. They are not Pareto optimal. Current selves lose, while future selves gain.

In Section 7, we investigate the desire for legacy and show that it too presents a potential argument for a social discount rate that is below the market rate. The desire for legacy is a preference over whether, how, and by whom one is remembered after death. Since one cannot experience how one is remembered after death, legacy utility is entirely anticipatory and belief-based. Legacy utility therefore provides an additional, nonmonetary component to the return on a project that may compensate for any loss relative to the market rate of return. Importantly for the social discount rate, legacy need not be individualistic. We argue that legacy has a generational component that may motivate investment by the government.

In summary, if the government cannot target the costs and benefits of a project, missing current market participants tend to push the social discount rate above the market interest rate, whereas missing future market participants tend to push it down. If the government can target costs and benefits, however, it generally can do better with intratemporal transfers. If the government places weight on the future views of agents currently alive, however, the social discount rate will be below the market rate.

2. THE CASE FOR MARKET DISCOUNT RATES

> It is hard to see why the revealed preference of individuals should be disregarded in the realm of time, where it is accepted, broadly speaking, in evaluating current commodity flows.
>
> —Arrow & Kurz (1970, p. 12)

We begin with the case for using market prices to discount the future. We illustrate the case in an endowment economy with a finite number of infinitely lived agents who trade a complete set of Arrow-Debreu securities in competitive markets at date zero. We add a government that is considering whether or not to enact a policy or project that entails a cost at date zero and a benefit at some future date. The government's goal is to maximize some weighted average of agents' utilities. We show that if markets are complete, all agents evaluate intertemporal trade-offs in the same way and preferences are aligned with the market interest rate. The government may therefore use the market interest rate as a guide for whether or not to implement the policy or project regardless of how it weights the preferences of individual agents.

2.1. An Economy with Complete Markets

Our presentation loosely follows the work by Ljungqvist & Sargent (2018). We allow time preference to differ across agents, and we allow for both idiosyncratic and aggregate income risk.

Time is discrete, with $t = \{0, 1, 2, \ldots\}$. At any point in time, the economy can be in one of a finite set of states S. The state s_t is realized at date t, and $s^t = (s_0, s_1, \ldots s_t)$ denotes the history through t. S^t denotes the set of all histories through t, and $\pi(s^t)$ denotes the probability of the history s^t. For simplicity, we assume that the states follow a Markov chain with transition probabilities $\pi(s'|s)$, so that, given a history, $s^t = (s_0, s_1, \ldots s_t)$ has probability

$$\pi(s^t) = \pi(s_t|s_{t-1})\pi(s_{t-1}|s_{t-2}) \ldots \pi(s_1|s_0).$$

The history s^t is publicly observable, and agents can condition trade on this history. Markets are therefore complete.

There are I agents named $i = 1, \ldots, I$. There is a single, perishable, consumption good. In each period t each agent i receives an endowment of the good, $y_t^i = y^i(s^t)$, which depends on the history until t. All trade takes place in a competitive market at date $t = 0$ after observing s_0.[2] A price system is a sequence of functions $\{q_t(s^t)\}_{t=0}^{\infty}$. Each function $q_t(s^t)$ assigns a price to the consumption good at date t after history s^t. Given the price system, each agent i chooses a complete, state-contingent consumption plan $c^i = \{c_t^i(s^t)\}_{t=0}^{\infty}$, where for each t, $c_t^i : S^t \to R_+$ maps histories through t into realizations of consumption. The agent chooses this plan to maximize a time-separable utility function,

$$U^i(c^i) = \sum_{t=0}^{\infty} \sum_{s^t \in S^t} (\beta^i)^t u\left(c_t^i(s^t)\right) \pi(s^t),$$

subject to the lifetime budget constraint,

$$\sum_{t=0}^{\infty} \sum_{s^t \in S^t} q_t(s^t) c_t^i(s^t) \leq \sum_{t=0}^{\infty} \sum_{s^t \in S^t} q_t(s^t) y^i(s^t).$$

[2] One can easily consider sequential trade. Date zero trading simplifies the analysis, as we do not need to keep track of the evolution of asset positions over time.

We allow each individual to have a different discount factor β^i. We assume that the individual specific discount factors all lie between zero and one, $\beta^i \in (0, 1)$, and that the period utility function $u(\cdot)$ is increasing, differentiable, and concave. We impose the Inada condition $\lim_{c\to 0} u'(c) = \infty$ so that all consumption plans are interior, $c_t^i(s^t) > 0$ for all t and all s^t.

An allocation is a set of consumption plans $c^i = \{c_t^i(s^t)\}_{t=0}^{\infty}$, one for each i. An allocation is feasible if aggregate consumption is less than the aggregate endowment,

$$\sum_i c_t^i(s^t) \leq \sum_i y^i(s^t),$$

for each history s^t that occurs with positive probability. A competitive equilibrium is a feasible allocation and a price system such that the allocation solves each household's problem. This completes the description of the complete markets economy.

2.2. The Optimality of Market Interest Rates

We assume that the economy begins in competitive equilibrium. We consider a government that is considering whether to undertake a project that takes x units of the consumption good at date zero and produces z units at some future date T. We assume that z is independent of the history s^T to focus on issues of discounting rather than attitudes toward risk. To abstract from distributional issues in taxation, we assume that at any point in time the government has the ability to tax any agent or transfer resources freely among agents in a lump sum manner. We assume that the project is small relative to the economy as a whole, so that we do not have to consider the effect that it has on market prices.[3] The government's objective is to maximize a weighted average of the utilities of the agents in the economy. Let μ_i denote the weight on agent i. The government maximizes,

$$W = \sum \mu_i U^i.$$

A key feature of this complete markets setup is that all agents value intertemporal trade-offs in exactly the same way. For each agent i, the trade-off between consumption at date zero and consumption after some history s^T satisfies the first-order condition:

$$q(s_0)u'(c_0^i) = q(s^T)(\beta^i)^T u'(c^i(s^T))\pi(s^T).$$

The left-hand side is the benefit of a marginal unit of consumption at date zero in terms of the (arbitrary) numeraire. The right-hand side is the benefit of a marginal unit of consumption after history s^T times the probability of that history. Taking expectations across histories s^T yields the standard consumption Euler equation,

$$u'(c_0^i) = (1 + r_{0,T})(\beta^i)^T E\{u'(c^i(s^T)|s_0\}, \qquad 1.$$

where $1 + r_{0,T} = \left(\sum_{s^T \in S^T} \frac{q(s_0)}{q(s^T)}\right)^{-1}$ is the risk-free rate between dates zero and T. Since in equilibrium the consumption plan of each agent satisfies Equation 1, all agents are indifferent between one unit of the consumption good at date zero and $1 + r_{0,T}$ units of the consumption good at date T. Allowing for individual specific discount factors does not alter this conclusion, as agents arrange the trajectory of their consumption plans to compensate for such differences.

In this setting, the social discount rate is equal to the market rate $r_{0,T}$. If a project has a return greater than this market return, it creates wealth and the government should undertake it. Given the government's ability to transfer resources among agents, the costs and the benefits of the project can be allocated so that all agents are better off. If the return is less than the market rate,

[3] If the project is large, then using ex ante market prices is equivalent to maximizing the compensating variation.

the project should be shelved. The project destroys resources, and the utility of some agent must fall.[4]

The power of this argument lies in the result that all agents evaluate intertemporal trade-offs in exactly the same manner and would make the same choice if the project were their own. This raises the question as to the difference between a government and a firm in this setting. One possibility is that the return z is in the form of a nontradeable public service such as clean air or public safety. Another possibility is that the costs and benefits are borne by different agents. In both cases, the project may require the government's ability to tax and transfer in order for all to benefit.

2.3. Intertemporal Transfers and Intratemporal Transfers

Since all agents agree on the intertemporal trade-offs, the principal welfare problem faced by the government is distributional. We have assumed that the government has the ability to collect lump sum taxes and make lump sum transfers in order to address how the costs and benefits of the project are distributed across agents. A recurring theme in what follows is that often when a project with a below-market rate of return would appear to improve welfare, the government can achieve the same ends more efficiently through a set of intratemporal transfers.

Consider any competitive equilibrium of our complete markets economy. Let λ_i denote agent i's marginal utility of wealth in equilibrium. It turns out that this competitive equilibrium is a Pareto optimum. It is the Pareto optimum that maximizes the weighted average of individual utilities in which the Pareto weights are equal to the inverse of these marginal utilities of wealth,

$$W = \sum \frac{1}{\lambda_i} U^i.$$

If the government employs its full power to make transfers, the government can reallocate wealth so that the λ_i are proportionate to the government's welfare weights μ_i. Therefore, if the government makes full use of its ability to transfer resources, there will be no distributional motive to undertake a project with a return below the market rate.

Suppose, however, that the government has not made full use of its ability to transfer resources, so that the welfare weights and inverse marginal utilities differ. In this case, taking x units of the consumption good from some agent i with a relatively low marginal utility of wealth, investing in a project with a below-market rate of return, and returning z units of the consumption good in period T to some other agent j with a relatively high marginal utility of wealth may increase the government's objective W. What the government loses in the project's return, it makes up by improving the distribution of income. It goes without saying that such a policy is not a Pareto improvement.

In such cases, however, the government generally can do better by making an intratemporal transfer. Suppose that the return on the project is less than the market rate, so that $z < (1 + r_{0,T})x$. According to Equation 1, the agent i that the government taxes to finance the project places the same value on x units of the consumption good in period zero as they do $(1 + r_{0,T})x$ units of the consumption good in period T. A transfer of $(1 + r_{0,T})x$ units from agent i to agent j in period T therefore has the same present-value cost to agent i and transfers strictly more resources to agent j.

[4]We have considered a policy that transfers resources from date zero to some later date T. We could have just as easily considered a policy that transfers resources at some date t after history s^t to some date $T > t$. In this case, the risk-free rate between history s^t and T would be $1 + r_{s^t,T} = \left(\sum_{s^T|s^t} \frac{q(s_t)}{q(s^T)}\right)^{-1}$, where the summation is over all histories s^T that follow upon s^t. As above, the policy improves welfare if its return exceeds this rate.

One might argue that such intratemporal transfers are not always feasible or costless and that such costs are the primary reason that the inverse marginal utilities of wealth in competitive equilibrium are not proportionate to the government's welfare weights in the first place. In such cases, a project with below-market returns might represent a second-best means of reallocating resources among agents and achieving the government's ends. There is nothing inherently wrong with this argument. It ignores, however, that the project also represents a transfer—a transfer of resources both between agents and across time. In order for the project to be a more efficient means of redistributing resources, the government must make sure to tax the agents that have too much wealth according to its criteria and must make sure that the returns to the project go to the agents that have too little wealth according to its criteria. If intratemporal transfers and taxes are costly, then presumably it is also costly to target the costs and benefits of the project. The questions become "How much does each policy cost?" and, if the project costs less, "Do these cost savings compensate for a return that is below the market rate?" Such questions must be answered on a project-by-project basis.

3. MISSING MARKETS

In this section, we relax the assumption that markets are complete. There are many ways to introduce market incompleteness. We consider one of the most popular: The set of assets is limited, and for this reason, participation in asset markets may be limited as well. Given limited participation, not all consumption plans satisfy the Euler equation (Equation 1). The market interest rate therefore only represents the intertemporal trade-offs of a subset of agents in the economy.

The most popular class of these models assume that agents can only borrow and save in one-period risk-free bonds and that their ability to borrow on these terms is limited by their wealth (Aiyagari 1994, Hugget 1996, Kaplan & Violante 2014). For tractability, these models usually assume a large number of agents with idiosyncratic income fluctuations and consider a steady state in which the distribution of wealth and income is constant. Aggregate consumption and the short-term interest rate are therefore also constant. Agents with low wealth are either unable or unwilling to borrow at this interest rate. The wealthy attempt to save but their ability to save is limited by the ability of the poor to borrow. A general feature of these models is that the real rate of interest in steady state is lower than what we would see if markets were complete. The rich cannot save as much as they would like, so their consumption is expected to decline relative to the complete markets outcome. This downward sloping consumption path implies a lower interest rate.

An example due to Werning (2015) illustrates how these markets work and makes clear their implications for social discounting. In Werning's example, the borrowing constraint is severe: Borrowing is not allowed at all. Without borrowing, there can be no saving in equilibrium and the interest rate adjusts so that savers do not save. Without saving, the asset distribution is degenerate, greatly simplifying the calculation of equilibrium quantities. All agents simply consume their endowment. To simplify further, we assume that income $y^i(s^t)$ can take on one of two values y^H and y^L, with $y^H > y^L$, and that transitions between y^H and y^L are independent across agents and across time and follow a Markov chain, with all transitions having positive probability.

In this economy, agents hold no assets, so agents who receive y^H in period t are temporarily wealthy. They attempt to save, but there is no one to lend to. The market interest rate adjusts so that zero saving is their optimal choice. This equilibrium interest rate satisfies the Euler equation for the high earners:

$$u'(y_t|y_t = y^H) = (1 + r)\beta E\{u'(y_{t+1})|y_t = y^H\}. \quad 2.$$

Agents who receive y^L in period t would like to borrow at this equilibrium interest rate r but cannot. Their Euler equation does not hold with equality,

$$u'(y_t|y_t = y^L) > (1+r)\beta E\{u'(y_{t+1})|y_t = y^L\}. \qquad 3.$$

In this example, the market interest rate reflects the views of potential savers who are temporarily wealthy. Those with temporarily low incomes have a greater preference for current consumption than would be inferred from the market interest rate.

In the complete markets economy, limits on the government's ability to tax and transfer opened the door to the possibility that the government might want to undertake projects with below-market returns if such a project helped the government achieve its distributional objectives. In this incomplete markets economy, the situation is reversed: A government that cannot target the costs and benefits of a project may not wish to pursue a project even if its return is greater than the market rate. The required rate of return implicit in Equation 3 is above the market rate in Equation 2. The average discount rate across agents in the economy is therefore above the market rate. Even projects with above-market returns may reduce the government's objective if they transfer resources away from agents who are currently borrowing constrained. These agents prefer current consumption.

Even though the market interest rate does not reflect the views of all market participants, however, it may still provide a useful guide to policy. Consider again our policy of investing x units of the consumption good today in return for z units at some future date T. A feasible policy for a government with access to lump sum transfers is to tax potential savers in each period by an amount $(1+r)^t x$. It can use the revenue in period zero to finance the project and the revenue in each subsequent period $t \in (1, T-1)$ to compensate those taxed in period $t-1$. In period T, the government receives the return to the project and can compensate those taxed in $T-1$ loans if $z > (1+r)^T x$. This plan represents a Pareto improvement if and only if the project earns a return greater than the market return.

The ability of the government to undertake a complex series of lump sum transfers raises the question, "Why the government does not eliminate the uninsurable labor income risk, which is the reason that the poor are not participating in asset markets?" Note, however, that while eliminating labor income risk requires a complicated set of transfers that affects every agent in the economy, the government can easily implement this transfer policy to finance the project. The government does not need to identify and tax savers. It needs only to float bonds with a market rate of return and roll them over from period to period. Only savers will buy these bonds, so that all of the transfers in periods $t \in (0, T-1)$ will take place through capital market transactions. The government needs only to figure out how to allocate an amount $(1+r)^T x$ of the final return to the project z to bond holders from period $T-1$. This is a much easier task than removing uninsurable labor income risk.

To this point we have only considered a small project that does not alter equilibrium prices. A lesson from this class of models is that the market interest rate is determined in large part by the characteristics of wealthy savers. When projects or policies have a large impact on the economy, as they might in the case of efforts to mitigate climate change, it will be important to understand how they affect wealthy savers and through wealthy savers the market interest rate. There has been a lot of theoretical and empirical work identifying characteristics that promote wealth accumulation. This literature, however, provides few clear lessons for social discounting and very little guidance on how policy may impact market interest rates. Patience plays some role in determining who becomes wealthy. Beyond that, some argue that the rich earn higher returns (Fagereng et al. 2020), while others argue that they take more risk (Bach, Calvet & Sodini 2020). The rich may be more prudent (Kimball 1990) or they may be ambiguity averse (Borovicka 2020).

Some have argued that the rich have more accurate beliefs (Friedman 1953, Blume & Easley 2006), while others have argued that irrational noise traders may earn higher returns (Delong et al. 1990). Bernheim, Garrett & Maki (1997) argue that financial literacy is correlated with wealth, while Ameriks, Caplin & Leahy (2003) find that the propensity to plan is correlated with wealth, and Ameriks et al. (2007) argue that the wealthy exhibit greater self-control. Overall, while much work has been done on the possible forces that may promote wealth accumulation, there is less work on characterizing the actual wealthy and even less work on the implications for the relationship between the social discount rate and the market interest rate. This is a fertile ground for future research.

4. FUTURE GENERATIONS

One common argument against using market discount rates to discount long-horizon projects is that discounting the future in this way is equivalent to discounting the utility of future generations and placing a lower weight on the utility of future generations is unethical. This was the view expressed by Ramsey (1928). In this section, we argue that overlapping generations do not introduce substantively new issues beyond those we encountered in the complete markets economy. A government that can target the costs and benefits of a project through lump sum transfers can make all agents better off if and only if the project has a return greater than the market rate of return. The only difference is that, in the model of Section 2, every agent would be willing to finance the project and receive the project's return, whereas with overlapping generations the government must string together a sequence of transfers, borrowing from one generation and then the next until the project's return is realized. In this section, we show that a government that places greater weight on the utility of future generations than the market does can sometimes improve aggregate welfare by pursuing projects that have returns below the market rate. In such cases, however, the government can generally achieve the same end more efficiently by scrapping the project and instead making a series of intratemporal transfers.

To illustrate these issues, we make several changes to our complete markets model of Section 2. To simplify the analysis and focus on intergenerational equity, we remove all uncertainty. There is only one state in each period so that the income of agent i in period t, y_t^i, is a known quantity. Next, we assume that there are an infinite number of agents and that the set of agents is partitioned into a sequence of overlapping generations. Let j index the generation born at date j and let J represent the set of all generations. For simplicity we assume that each generation lives for exactly $N > 1$ periods. For each agent i in generation j, $y_t^i = 0$ and $c_t^i = 0$ whenever $t < j$ or $t > j + N$. In order to facilitate comparisons of utility across generations, we evaluate the utility of agent i in generation j at birth,

$$U^i(c^i) = \sum_{t=j}^{j+N} (\beta^i)^{t-j} u\left(c_t^i\right),$$

so that all discounting is relative to the date that a generation is born. There are also $N - 1$ generations alive at date zero who live truncated lives. We also discount their utility back to their birth dates.[5] As in our complete markets model, we assume that the government maximizes a weighted average of the individual utilities,

$$W = \sum_J \mu_j N_j U^j,$$

[5] Calvo & Obstfeld (1988) show that this normalization is needed if the government's optimal policy is to be time consistent.

where μ_j is the weight on generation j and N_j is the size of the generation. We assume that the weights are such that W is finite and return to this issue below.

With these changes, all of the analysis of the complete markets model goes through, except that not all agents participate in the market in all periods. The market interest rate between any two periods t and $t+1$, $r_{t,t+1}$ equates the marginal utility in period t to the discounted marginal utility in period $t+1$ of all agents alive in both periods:

$$u'(c_t^j) = (1 + r_{t,t+1})\beta^i u'(c_{t+1}^j).$$

To see that this interest rate remains a sensible criterion for evaluating projects, consider again whether a government should proceed with a public project that takes x units of the consumption good at date zero and produces z units of the consumption good at some future date T. As in our discussion of the last section, the government can borrow x from generation $j = 0$ at rate $r_{0,1}$ and roll this loan over in each subsequent year, borrowing from generation $j = t$ at rate $r_{t,t+1}$. Finally generation j_{T-1} is willing to receive the final payment z if and only if $z \geq (1 + r_{0,T})x$.[6] Hence, overlapping generations do not alter the social trade-off between resources in different periods. Any project that returns more than the market rate of interest expands economic opportunity and makes possible a Pareto improvement relative to the competitive equilibrium. Any project that returns less than the market rate of interest reduces available resources, so that someone must take a loss.

Might the government be tempted to pursue a project that is not a Pareto improvement because this project redistributes resources in a way that improves its objective? As in the case of complete markets, the market equilibrium when it is Pareto efficient (more on this below) represents a specific Pareto optimum with specific Pareto weights.[7] Suppose that these weights differ from the weights in the government's objective. One possibility is that the government places greater weight on current generations, possibly because future generations might be richer and have a lower marginal utility of wealth. In this case, the government still wants to undertake any project with a return above the market rate of interest, since such projects represent Pareto improvements and therefore make everyone (weakly) better off. It may be tempted to reallocate the project's gains to earlier generations by paying above-market interest rates on the bonds that it issues, but this incentive is orthogonal to the decision of whether or not to undertake the project. The government will not undertake projects with below-market returns, as these do not make current generations better off. The government really wants to undertake projects that take future resources and create current resources, but these are not technically feasible.

Consider now a government that places greater weight on future generations, possibly because it places greater weight on the utilities of future generations than does the market. Such a government might be tempted to undertake a project with a below-market rate of return, since it transfers resources from current to future generations. The government could make up what it loses in returns by redistributing wealth toward generations it sees as underserved by the market. As in the case of complete markets, the problem with this argument is that there is a simpler and more efficient way for the planner to transfer resources between generations. Suppose that the government is contemplating a project with a return below the market rate $z < x(1 + r_{0,T})$. Rather than take x from generation $j = 0$, invest in the project, and transfer z to generation

[6] Recall, we are only considering situations in which the term structure hypothesis holds.

[7] In this setting, the government will only be happy with the market allocation if $\mu_j \sim \beta^j/\lambda_j$, where the marginal utility of wealth is equal to the marginal utility of consumption at birth. With these weights, the government discounts the utility of future generations at the same rate that agents discount their own utility within lifetimes.

$j = T$, the government can instead make a sequence of intratemporal transfers. At each date t, beginning with $t = 1$, the government can transfer $x(1 + r_{t-1,t})$ from generation $j = t - 1$ to generation t. The cost of this transfer scheme is the same as the cost of the project, as generation $j = 0$ values a tax of $x(1 + r_{0,1})$ in period 1 the same as a tax of x in period zero. Each generation $j \in (1, T - 1)$ is unaffected by the transfer scheme. They place the same value on the transfer that they receive while young as they do on the transfer they make when old. Finally, generation $j = T$ receives $x(1 + r_{0,T})$, which is the market return on the investment x and is strictly greater than z. Hence, a government with the ability to make intratemporal transfers can always do better than invest in a project with a below-market rate of return.

We can take this argument a step further. Following Calvo & Obstfeld (1988), under fairly general conditions, if the government makes all of the intratemporal transfers it wants, then the market interest is the appropriate benchmark for the government's objective. The government can reallocate resources across generations until the competitive equilibrium aligns with its objective.

There are at least two caveats to this argument. As in the complete markets case, the government may not have access to intratemporal transfers or it may be costly for the government to make intratemporal transfers. As many of these costs also apply to the project, the question becomes how to balance the relative returns and the relative transaction costs. In such cases, projects with below-market returns might be an appropriate substitute for intratemporal transfers, so long as the costs x and the benefits z can be targeted to generations in a way that mimics the intratemporal transfers that the government is not able to make.

The second caveat is that the government may not be able to commit to future intragenerational transfers but may be able to commit to the project. Calvo & Obstfeld (1988) argue that the optimal transfer policy is time consistent so long as all generations are treated symmetrically and the government discounts future generations geometrically. This argument, however, assumes that the tastes and priorities of the government do not change. In such cases, one could imagine a government using an investment project to lock in an intergenerational transfer. At the same time, it is not obvious how to evaluate the welfare impact of such a policy if the government favors it today, but some future government wishes it had never been enacted. Presumably, policies enacted to benefit future generations should be seen by those future generations as beneficial. Time consistency may therefore be a good assumption.

The argument in this section is incomplete in one sense: We have assumed that the market equilibrium is a Pareto optimum. The addition of an infinite number of agents makes new trades possible. A standard result in growth theory establishes that if the market rate of interest is too low, then the present value of endowments may be infinite and the competitive equilibrium may not be Pareto efficient. In such cases, the government can generally improve the allocation by taxing the young and transferring resources to the old. None of this, however, alters the conclusion that the government should invest in projects that pay returns above the market. Hence, overlapping generations in and of themselves do not alter the advisability of using market interest rates.

5. ALTRUISM

Farhi & Werning (2007) argue that intergenerational altruism may lead to a social discount rate that is lower than the market rate. Their argument can be illustrated in a two-period example taken from their paper. Consider a parent who lives in period one and a child who lives in period two. Suppose that c_2 is the consumption of the child and $U^2 = u(c_2)$ is the utility of the child. Suppose that the parent derives utility from their own consumption c_1 as well as the utility of the child. The parent's utility is

$$U^1 = u(c_1) + \delta u(c_2),$$

where δ is the weight that the parent places on their child's utility. The market interest rate then equates

$$u'(c_1) = (1 + r)\delta u'(c_2) \quad 4.$$

and reflects the marginal utilities of the two generations and the parent's weight on the child's utility. Now consider a government that maximizes a weighted average of the utilities of both the parent and the child,

$$W = U^1 + \mu U^2, \quad 5.$$

or equivalently,

$$W = u(c_1) + (\delta + \mu)u(c_2).$$

So long as $\mu > 0$, the government places a greater weight on the utility of the child than does the parent. The child's utility gets counted twice: First, as it enters the utility of the parent, and second, on its own terms. The government is therefore willing to invest x in period 1 in return for z in period 2 so long as the return z/x is greater than

$$\frac{1}{\delta + \mu} \frac{u'(c_1)}{u'(c_2)} < (1 + r). \quad 6.$$

It is easy to extend this model to multiple periods and to situations in which children also care about their parents. Many of these extensions complicate the math and lead to questions such as "Given the feedback between generation's utilities, does there exist a utility function?" or "Is choice time consistent?" In the end, what matters for social discounting is not that children do not care about their parents, but that children care more about themselves than their parents do. If this is the case, then the social discount rate will be less than the market interest rate.

An unrealistic assumption in Farhi & Werning (2007) is that generations do not overlap.[8] This makes it impossible for the government to transfer resources across generations within a time period. The only possible transfers are intertemporal. Once we allow for intratemporal transfers, we see that the introduction of altruism does not alter the conclusions of the overlapping generations model of the last section. There are more efficient ways for the government to transfer resources across generations than by undertaking projects with below-market rates of return.

To see this, we amend Farhi & Werning's (2007) example so that the parent lives for two periods and has utility function

$$U_1 = u(c_1^1) + \beta u(c_2^1) + \delta u(c_2^2),$$

where c_t^i refers to the consumption of generation i in period t. The parent cares about their own consumption in the second period as well as the consumption of the child. β is the parent's discount factor, and as before, δ captures the weight that the parent places on the child's utility in period 2. Suppose that agent i receives income y_t^i in period t. If absent all transfers $\beta u'(y_2^1) < \delta u'(y_2^2)$, then the parent makes a transfer to their child in period 2 so that $\beta u'(c_2^1) = \delta u'(c_2^2)$.[9] In this case, the first-order condition (Equation 4) still captures the parent's trade-off between consumption in period 1 and their child's consumption in period 2. A government that maximizes Equation 5 would still want to invest in the project if Equation 6 holds, so long as the returns z accrue to the child.

[8]Generations also do not overlap in Phelps & Pollak's (1968) model of imperfect altruism.

[9]If $\beta u'(y_2^1) < \delta u'(y_2^2)$, the parent would like to receive a transfer from the child. As the child does not place weight on the parent's utility, no transfer is made in this example.

Given that the parent values their own consumption in period 2 according to the Euler equation $u'(c_1^1) = (1 + r)\beta u'(c_2^1)$, however, the alternative policy of simply transferring z from the parent to the child in period 2 produces the same outcome as the project for the child at lower cost to the parent. The transfer of z reduces the parent's utility by only $z/(1 + r)$, so that this policy has an effective return equal to the market interest rate. As in the standard overlapping generations model, the possibility that the government might wish to pursue a project with a below-market rate of return points to the existence of a superior policy of intratemporal transfers.[10,11]

If the government makes full use of its ability to make transfers, it will set

$$\beta u'(c_2^1) = (1 + r)(\mu + \delta)u'(c_2^2)$$

and the market interest rate will reflect the government's intertemporal preferences.[12]

6. FUTURE SELVES

> The existence of preference for the present over equally certain future pleasures does not imply that any economic dissatisfaction would be suffered if future pleasures were substituted at full value for present ones. The nonsatisfaction this year of a man's preference to consume this year rather than next year is balanced by the satisfaction of his preference next year to consume next year rather than this year.
>
> —Pigou (1929, p. 25)

To this point, it would appear that the market interest rate is an appropriate criterion in intergenerational contexts, particularly when the government has other tools for redistributing resources across agents such as intertemporal transfers. While there may be situations in which the market interest rate places lower weight on the utility of future generations than that preferred by the government, these are situations in which there exist intratemporal transfers that achieve similar ends at lower cost.

A related argument, however, survives this critique. Instead of thinking about different generations at different times, consider the same individual at different points in their life. Much the same way that generations might care more about their own utility than the utility of other generations, selves may care more about their utility in their time than do past or future selves. The main difference between selves and generations is that selves do not overlap. Each self is associated with their fixed moment in time. The fact that selves do not overlap makes intratemporal transfers impossible. Every transfer is from a self at one date to a self at another date.

Caplin & Leahy (2004) develop this idea in a version of the complete markets model of Section 2 without uncertainty. Instead of evaluating consumption plans $c^i = \{c_t^i\}_{t=0}^{\infty}$ from the perspective of time zero, they follow Strotz (1956) and allow each individual's utility to change over

---

[10]The argument in this paragraph shares similarities to Barro's (1974) argument that Ricardian equivalence holds in the presence of intergenerational altruism.

[11]If $\beta u'(y_2^1) > \delta u'(y_2^2)$, then the optimal transfer depends on $\beta u'(y_2^1)$ relative to $(\mu + \delta)u'(y_2^2)$. If the latter is larger, the government transfers resources from the parent to the child. Otherwise, it is optimal to transfer resources from the child to the parent. In either case, it is optimal only to undertake projects with returns greater than the market rate.

[12]Note that if the government makes full use of its ability to make transfers, parents will want to receive transfers from their children. In the example, children will not be willing to make these transfers since they do not place weight on their parent's utility. If children also value their parent's utility, this will limit the government's ability to make transfers, since any attempt by the government to give too much to the children would be undone by the children making transfers to their parents. Presumably, however, the government's relative valuation of parents and children would lie between that of the parents and children, so that the government would not make transfers that might be undone by either party.

time. Specifically, they assume that agent i in period t has a utility function of the form

$$U_t^i(c^i) = \sum_{s=0}^{t-1} \delta^{t-s} u\left(c_s^i\right) + u\left(c_t^i\right) + \sum_{s=t+1}^{\infty} \beta^t u\left(c_s^i\right). \qquad 7.$$

Here, $u(c_t^i)$ is the utility from consumption of agent i in period t. The first term captures the value of past consumption, and the third term the value of future consumption. If both β and δ are positive and less than one, then the agent discounts both past and future utility. If δ equals zero, then the utility function is similar to the utility function in Farhi & Werning (2007), except that in their model U_t represents the utility of a generation rather than the utility of an individual at a specific point in time. For simplicity, we assume that all agents have the same β and δ.

Since the agent in period zero has the same utility function as an agent in Section 2, the equilibrium with time-zero trading is the same as it is in Section 2. Since each agent in each period discounts each future period by the same factor β, period zero plans are time consistent. Even though preferences are changing over time, no agent has any incentive to alter these plans at any future date. Given the additive separability of the utility function, past consumption is irrelevant for decision-making in this model. This is possibly why past consumption is normally ignored in welfare analysis and in economic models in general.[13] We shall see that just because past consumption is irrelevant for decision-making does not mean that it is irrelevant for welfare.

Clearly the same agent at different points in time has different views regarding intertemporal trade-offs. Given any consumption plan c^i, we can calculate the path of interest rates that would make an agent choose that stream at any given date. Consider consumption at two dates t and s with $t > s$ and consider an agent at date v, who is evaluating the utility from that consumption stream. Then for the agent's Euler Equation 1 to hold the rate of return between s and t for an agent living in period v, $r_{s,t,v}$ must satisfy

$$1 + r_{s,t,v} = \begin{cases} \beta^{s-t} u'(c_s)/u'(c_t) & \text{if } v \leq s \\ \delta^{v-s} \beta^{v-t} u'(c_s)/u'(c_t) & \text{if } s < v < t \\ \delta^{t-s} u'(c_s)/u'(c_t) & \text{if } t \leq v \end{cases}.$$

When $v < t$, both t and s are in the future. $r_{s,t,v}$ depends on β and is identical to the market rate of return in the complete markets model. This is because all agents discount the future at rate β. When $v > s$, so that both t and s are in the past, $r_{s,t,v}$ depends on δ. Only if $\delta = \beta^{-1} > 1$ is $r_{s,t,v}$ equal to the return in the complete markets model. In this case, U_t is simply a monotonic transformation of U_0. All agents agree on the ranking of consumption streams, and $r_{s,t,v}$ is independent of v.

The case in which $\delta = \beta^{-1}$, however, is a strange one. It implies that agents place greater weight on past utility than they do on the present. A more natural assumption might be that $\delta < 1 < \beta^{-1}$ so that agents place greater weight on current utility, possibly because imperfect memory makes past experience less salient. In this case, preferences over consumption streams change over time. If $\delta < 1$, then $\delta^{t-s} < \delta^{v-s}\beta^{v-t} < \beta^{s-t}$ and $r_{s,t,v}$ is monotonically declining in v. Future selves place greater weight on future consumption and do not need to be compensated with high returns in order to be satisfied with any given consumption stream.

What are the implications for using market interest rates to discount the costs and benefits of public projects? Consider our investment of x in period zero with a return of z in period T. Since only current selves participate in asset markets, market interest rates reflect the views of agents looking forward. As in Section 2, this project makes period-zero selves better off if and only if

[13] The exception is when preferences over current and future consumption depend directly on past choices, as they might if there were habit formation.

they earn a return above the market rate of interest. Future selves, however, cannot participate in asset markets, and all future selves place a greater relative value on consumption in period T than do period-zero selves. All would gain from the project even if it had a return slightly below that market rate of interest. In particular, the period-T self would benefit if the rate of interest were above $1 + r_{0,T,T}$, which is $(\delta/\beta)^T$ times the market interest rate. If the government places any weight on the perspectives of future selves, some projects with returns below the market rate will increase its objective. These projects will not be Pareto improvements, as time-zero selves will be made worse off, but the gain to future selves will compensate for these losses. Only if the government places zero weight on the perspective of future selves will the market interest rate be the benchmark.[14]

Unlike the overlapping generations models, there is no obvious alternative policy that the government can pursue that achieves the same ends at lower cost. It is not possible to make intratemporal transfers between selves in different periods. The only way to make future selves better off is to save and invest.

These arguments become even more potent if preferences are myopic, as they are in models of hyperbolic discounting (Phelps & Pollak 1968, Laibson 1997). In such models, agents fear that if they save, their future selves will spend any savings. Phelps & Pollak (1968) show that this leads to savings that are low even from the current perspective. This low rate of saving tends to reduce capital accumulation in the long run and to raise market rates of interest (Barro 1999). Laibson (1997) shows that, with hyperbolic discounting, current selves have an incentive to commit resources to long-lived illiquid assets in order to prevent their future selves from consuming out of their savings. Long-horizon government projects may fill this role. Hence, in models with hyperbolic discounting, projects with below-market returns represent Pareto improvements. They may prove desirable even from the perspective of the current generation.

7. DESIRE FOR LEGACY

Desire for legacy is a preference over whether, how, and by whom one is remembered after death. Waggoner, Bering & Halberstadt (2023) point out that "individual human beings have long strived to create and curate an enduring postmortem reputation—their 'legacy'. . . . They do this in a variety of ways, such as via sports achievement, creative works, having children, leaving a financial endowment, philanthropy, passing down family heirlooms, or extreme attention-grabbing acts" (p. 1). The basis for stressing the desire for legacy in this article is our view that it is an extremely important intergenerational motive and plays a key role in how we value different projects of long duration. As parents, we strive to bring children up well in the hope that this will earn us the desired place in their hearts. As academics, we strive to introduce ideas that will stand the test of time. Alumni donate money for university buildings. The wealthy set up foundations.[15] These are all costly actions made in the hopes that these actions are appreciated in the future by others.

Since the desire for legacy is a nonstandard source of utility, we discuss it in some detail. For present purposes, three characteristics stand out. First, the desire for legacy is a form of belief-based anticipatory utility in the sense of Caplin & Leahy (2001). Whereas traditional utility comes from the experience of some outcome such as consumption or leisure, the payoff to legacy is the

[14]Parents often play the role of the government in this model. Parents often urge their children to study and work hard—not to maximize their children's current utility but to maximize their children's future utility as an adult.

[15]Just as intertemporal optimization can be thought of as a form of altruism for our future selves, the desire for legacy has an intrapersonal counterpart. Agents may wish to think well of their own lives. Undertaking projects that they deem virtuous may be a way to build self-esteem.

belief that one will be remembered well by future generations. In many cases one does not ever learn what future generations actually believe since these acts of remembrance, if they occur, often occur after one's lifetime. Whether or not one is actually remembered is irrelevant. Agents derive legacy utility from the belief that they will be remembered.

Second, the desire for legacy operates as a type of second-order belief. Whereas altruism reflects an individual's utility derived from making another party better off, legacy motives represent the desire to have another party view one fondly. Legacy utility depends on a current generation's belief about future generations' beliefs. Because these future beliefs are realized largely after the current generation is passed, there are no guarantees that actions chosen today will in fact be viewed positively. Projects that one individual or generation undertakes to be honored by future generations may at times be viewed very differently than they had hoped.

Third, these beliefs are attention-based and manipulable. Anticipatory feelings can be triggered, changed, and enhanced in a variety of ways. Particularly salient manipulations operate through attentional channels. In addition to bringing the future more to mind, it is important to provide a sense of efficacy. In the case of legacy, this channel is familiar to all who seek donations: A key goal is to bring possible respect for legacy to mind and to tie gifts to particular named outcomes.

Our model of the legacy motive is simple. The desire for legacy provides an additional psychological return to investment that can justify investing in a project with a rate of return below the market. Consider the two-period model of altruism from Section 5, in which there is a parent and a child. Suppose that the parent's utility is

$$U_1 = u(c_1^1) + \beta u(c_2^1) + \delta u(c_2^2) + v(x),$$

where $v(x)$ is the additional legacy utility that the parent receives from investing x in a project and believing that their child will remember this investment. Suppose that the project returns z to the child in terms of the consumption good. Suppose that the parent makes transfers in the second period so that $\beta u'(c_2^1) = \delta u'(c_2^2)$, where again c_t^i is the consumption of generation i in period t. Then the parent's generation invests in the project if

$$u'(c_1^1)x < \beta u'(c_2^1)z + v(x)$$

or

$$z = (1 + r)\,[x - v(x)].$$

The additional legacy utility justifies pursuing the project even if the financial return on the investment x is below the market rate r. The psychological return in the form of legacy utility makes up the difference.

Importantly for the social discount rate, legacy motives can operate at both the level of the individual and the community. Some legacies, such as parenting or endowments, tend to be individualistic. Others might relate to group identities. While our simple model was cast in terms of an individual parent, we could have just as well considered a generation. We define generational legacy projects as projects undertaken by one generation to be better viewed in the collective sense by future generations. Examples might include the national parks in the United States or the European Union. Many of these are consciously undertaken to preserve and protect the future, and while we also remember the individuals who started the movements, there is also a generational aspect to respect for their achievements.

Not all projects are likely to produce legacy, and an interesting topic for future research is to identify precisely which aspects of a project are legacy enhancing. There is obviously a strong link with public goods provision: Parks, libraries, and schools are all common legacy projects. The extent to which current generations sacrifice might be important. Those born between 1900 and 1925 are commonly known as the greatest generation largely because they lived during the Great Depression and World War II. A belief about what future generations will be grateful for and what they will resent about the behavior of past generations is also involved. Preservation of important coral reefs is more likely to produce a positive legacy than grand projects that pave over large stretches of rainforest such as the Amazon. There may also be activities that would earn respect if undertaken on scale. For such projects to work, there may be need for leaders to identify particular generationally supported activities and overcome the collective action and coordination problems.

According to Waggoner, Bering & Halberstadt (2023, p. 6), "given the intergenerational issues currently facing humanity, such as climate change, threats to biodiversity, and rising political strife, research on the legacy drive can not only shed light on the conceptual issues explored here, but also help to mobilize practical efforts toward solving critical problems that do not affect us personally. For example, if psychological distance limits our own generation's capacity to weigh the consequences of our actions on those who will come after us, efforts to drive change using the legacy drive may prove to be a powerful 'selfish' means towards the selfless end of addressing critical global issues. Along these lines, future researchers might investigate whether the pursuit of communal legacies, focusing on community and future generations, . . . or agentic legacies, focusing on individual admiration and fame, . . . are more likely to lead to acts of environmental stewardship."

8. CONCLUSION

In this article, we surveyed some of the main implications of limited market participation for the social discount rate. Market interest rates reflect the preferences of market participants, whereas the social discount rate reflects the views of society as a whole. Constrained borrowers generally value current consumption more than the market; this provides an argument for a social discount rate above the market rate of interest. On the other hand, future generations and future selves generally value future consumption more. This argues for a lower social discount rate.

The ability to make intratemporal transfers often restores market interest rates as a benchmark. In the case of constrained borrowers, such a government can target the costs and benefits of a project to market participants so that all projects with above-market rates of return are welfare improving. In the case of future generations, such a government can use transfers to shift resources between generations at market rates of return.

It is impossible, however, to make transfers from current selves to future selves intratemporally. Hence, if agents discount the past, a government that cares about future selves will entertain a social discount rate below the market rate. It is also possible that certain projects, because of their perceived importance, can generate legacy utility. Both of these mechanisms provide robust arguments for a social discount rate that lies below the market rate of interest.

DISCLOSURE STATEMENT

The authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

We thank Chris House, Pablo Ottonello, and the editor, Deborah Lucas, for helpful comments and suggestions. J.L. thanks the National Science Foundation for their support under grant SES-1919362.

LITERATURE CITED

Aiyagari SR. 1994. Uninsured idiosyncratic risk and aggregate saving. *Q. J. Econ.* 109:659–84

Allais M. 1947. *Economic et Intent*. Paris: Imprimerie Nationale

Ameriks J, Caplin A, Leahy J. 2003. Wealth accumulation and the propensity to plan. *Q. J. Econ.* 118:1007–47

Ameriks J, Caplin A, Leahy J, Tyler T. 2007. Measuring self-control problems. *Am. Econ. Rev.* 97:966–72

Arrow K, Kurz M. 1970. *Public Investment, the Rate of Return, and Optimal Fiscal Policy*. Baltimore, MD: Johns Hopkins Univ. Press

Bach L, Calvet L, Sodini P. 2020. Rich pickings? Risk, return, and skill in household wealth. *Am. Econ. Rev.* 110:2703–47

Barro R. 1974. Are government bonds net wealth? *J. Political Econ.* 82:1095–117

Barro R. 1999. Ramsey meets Laibson in the neoclassical growth model. *Q. J. Econ.* 114:1125–52

Bernheim BD, Garrett D, Maki D. 1997. *Education and saving: the long-term effects of high school curriculum mandates*. NBER Work. Pap. 6085

Blume L, Easley D. 2006. If you're so smart why aren't you rich? *Econometrica* 74:929–66

Borovicka J. 2020. Survival and long-run dynamics with heterogeneous beliefs under recursive preferences. *J. Political Econ.* 128:206–51

Calvo G, Obstfeld M. 1988. Optimal time-consistent fiscal policy with finite lifetimes. *Econometrica* 56(2):411–32

Caplin A, Leahy J. 2001. Psychological expected utility theory and anticipatory feelings. *Q. J. Econ.* 116:55–79

Caplin A, Leahy J. 2004. The social discount rate. *J. Political Econ.* 112:1257–68

Delong JB, Schleifer A, Summers L, Waldmann R. 1990. Noise trader risk in financial markets. *J. Political Econ.* 98:703–38

Eden M. 2021. *The cross-sectional implications of the social discount rate*. CEPR Discuss. Pap. DP16114

Fagereng A, Guiso L, Malacrino D, Pistaferri L. 2020. Heterogeneity and persistence in returns to wealth. *Econometrica* 88:115–70

Farhi E, Werning I. 2007. Inequality and social discounting. *J. Political Econ.* 115:365–402

Friedman M. 1953. The case for flexible exchange rates. In *Essays in Positive Economics*. Chicago: Univ. Chicago Press

Groom B, Drupp MA, Freeman MC, Nesje F. 2022. The future, now: a review of social discounting. *Annu. Rev. Resour. Econ.* 14:467–91

Hugget M. 1996. Wealth distribution in life-cycle economies. *J. Monet. Econ.* 38:469–94

Kaplan G, Violante G. 2014. A model of the consumption response to fiscal stimulus payments. *Econometrica* 82:1199–239

Kimball MS. 1990. Precautionary saving in the small and in the large. *Econometrica* 58:53–73

Laibson D. 1997. Golden eggs and hyperbolic discounting. *Q. J. Econ.* 112:443–78

Ljungqvist L, Sargent T. 2018. *Recursive Macroeconomic Theory*. Cambridge, MA: MIT Press

Marglin S. 1963. The social rate of discount and the optimal rate of investment. *Q. J. Econ.* 77:100–9

Millner A, Heal G. 2023. Choosing the future: markets, ethics and rapprochement in social discounting. *J. Econ. Lit.* Forthcoming

Nordhaus W. 2006. *The Stern Review on the economics of climate change*. NBER Work. Pap. 12741

Phelps E, Pollak RRA. 1968. On second-best national saving and game-equilibrium growth. *Rev. Econ. Stud.* 35:185–99

Pigou A. 1929. *The Economics of Welfare*. London: Macmillan and Co. 3rd ed.

Ramsey F. 1928. A mathematical theory of saving. *Econ. J.* 38:543–59

Sen A. 1961. On optimizing the rate of saving. *Econ. J.* 71:479–96

Strotz R. 1956. Myopia and inconsistency in dynamic utility maximization. *Rev. Econ. Stud.* 23:165–80

Waggoner B, Bering JM, Halberstadt J. 2023. The desire to be remembered: a review and analysis of legacy motivations and behaviors. *New Ideas Psychol.* 69:101005

Werning I. 2015. *Incomplete markets and aggregate demand*. NBER Work. Pap. 21448