Optimal Default Retirement Saving Policies: Theory and Evidence from OregonSaves

Mingli Zhong*

January 19, 2021

Abstract

I study the optimal default savings rate in automatic enrollment retirement saving plans. If individuals tend to procrastinate to make an active decision, the optimal default rate should be high to encourage people to opt out of the default. If individuals tend to actively undersave, the optimal default rate should be low to encourage people to stay at the default. Using exogenous increases in the default savings rate in OregonSaves, the first state-sponsored auto-enrollment plan in U.S., I identify individual adherence to the default rate. Using survey data from OregonSaves-eligible workers, I estimate the degree of undersaving if workers actively switch to a non-default rate. Combining individual-level administrative data with survey data, I suggest that the optimal default savings rate 7%.

JEL Classification: D14, D60, D91, G51, H00
Keywords: state-sponsored retirement plans, automatic enrollment, optimal default savings rate, undersaving, inattention, sufficient statistics

*Email: mzhong@nber.org. Zhong is a postdoctoral fellow at the National Bureau of Economic Research (NBER). I am extremely grateful to Hanming Fang, Ben Lockwood, and Olivia S. Mitchell for their guidance and support, and to Mike Abito, Douglas Bernheim, John Chalmers, Judd Kessler, Ilyana Kuziemko, Corinne Low, Sendhil Mullainathan, James Poterba, Alex Rees-Jones, Jonathan Reuter, Katja Seim, Lin Shen, Jonathan Skinner, and Juuso Toikka for comments and discussions. I also thank participants at the 21st Annual SSA Research Consortium Meeting, Wharton, Vanguard, Federal Reserve Bank of Dallas, and World Risk and Insurance Congress for helpful comments. This research was supported by a grant from the US Social Security Administration (SSA) to the Michigan Retirement Research Center (MRRC) as part of the Retirement Research Consortium (RRC). The research reported herein was also performed pursuant to grant RDR18000003 from the US Social Security Administration (SSA) funded as part of the Retirement and Disability Research Consortium. The support was also provided by the AARP; the Pew Foundation; the Pension Research Council/Boettner Center of the Wharton School at the University of Pennsylvania; and the Quartet program at the University of Pennsylvania. I thank many individuals from the OregonSaves program for numerous discussions and insights into the OS program, and Yong Yu as well as Wenliang Hou for excellent research assistance. The opinions and conclusions expressed are solely those of the author and do not represent the opinions or policy of SSA, any agency of the Federal Government, the MRRC, OregonSaves, NBER, or any other institutions with which I am affiliated. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of the contents of this report. Reference herein to any specific commercial product, process or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply endorsement, recommendation or favoring by the United States Government or any agency thereof. ©Zhong.

1 Introduction

Since the outbreak of the COVID-19 pandemic, millions of newly unemployed workers have lost access to employment-based retirement plans. Moreover, survey data show that about one-third of the U.S. population has used their savings to cover daily spending.[1] This has prompted increasing discussion among policymakers and researchers about mitigating the longer-term impact of the COVID-19 pandemic on retirement security. A new state-level policy that was introduced in several states before the pandemic and is now gaining more attention is state-sponsored automatic-enrollment retirement plans. This is an effort to expand the coverage of automatic-enrollment retirement plans among private-sector workers. Seven states (Oregon, Illinois, California, Maryland, Connecticut, New Jersey, and Colorado) have passed legislation requiring that private-sector employers make state-sponsored auto-enrollment plans available to workers if employers do not offer a plan.

A large strand in the literature has documented that automatically enrolling individuals into a retirement plan significantly increases their savings (Madrian and Shea, 2001; Choi et al., 2004; Beshears et al., 2009; Chetty et al., 2014). Moreover, a substantial fraction of participants passively save at the default savings rate determined by the plan designer. Despite the significant influence of auto-enrollment and the default savings rate on savings behavior, little is known about how to optimally set the default savings rate.

In this paper, I theoretically analyze and empirically identify the determinants of the optimal default savings rate that maximizes individual welfare. I then provide an explicit formula for the optimal default savings rate. A growing body of literature on optimal defaults suggests that, when individuals tend to procrastinate to make an active savings decision, setting the default savings rate at an undesirable level can compel individuals to opt out of the default rate and actively select a non-default preferred rate (Carroll et al., 2009; Goldin

[1] Survey data from Pew Research Center in August 2020: `https://www.pewsocialtrends.org/2020/09/24/economic-fallout-from-covid-19-continues-to-hit-lower-income-americans-the-hardest/`.

and Reck, 2018). In other words, the primary welfare effect of the default savings rate is to protect individuals from *inaction* caused by procrastination, status quo bias, or inattention.

This argument is based on a key assumption that individuals would actively choose a non-default preferred rate that maximizes their lifetime utility if opting out of the default rate. Previous empirical studies have found that individuals might underestimate the amount of retirement savings they need due to financial illiteracy, misinformation, or myopia (Goda et al., 2020).

Motivated by these findings, this paper proposes an additional welfare effect of the default savings rate, namely, to protect individuals from *undersaving*. The distinction between inaction and undersaving is of interest because the two types of behavioral biases have opposite implications for the optimal default savings rate. If individuals tend to procrastinate rather than making an active decision, it might be welfare-improving to set the default rate at an undesirable level to compel individuals to opt out of the default rate and choose a non-default savings rate. By contrast, if the non-default rate that individuals actively choose does not maximize their lifecycle utility, it might be welfare-improving to set the default rate at a desirable level to encourage individuals to stay at the default rate.

To understand how the optimal default savings rate is shaped by its welfare effects, I propose a unified framework that incorporates both types of behavioral biases: inaction and undersaving. The behavioral biases arise because individuals and the policymaker maximize individual utility differently. Individuals rely on their *perception* of how their lifetime utility should be modeled. To maximize their perceived utility, commonly referred to as decision utility, individuals either passively accept the default savings rate or actively select a non-default preferred savings rate. Compared to individuals' binary choices of either staying with or opting out of the default rate, the policymaker with more information—or who is more forward-looking than individuals—maximizes the normative utility, describing the *reality* of how individual choices actually affect individuals' lifetime welfare.

Based on the unified welfare framework, I derive a formula for the optimal default savings rate as a function of reduced-form sufficient statistics that can be empirically identified. The welfare analysis and the sufficient statistic formula do not necessarily rely on specific underlying behavioral models of why the default options affect individual behavior. I directly characterize how individual responses to different default rates impact their welfare. This approach allows my welfare analysis to incorporate a range of underlying behavioral models that may explain the default effect on individual behavior.[2]

I implement the formula for the optimal default savings rate empirically by estimating two sets of key statistics. The first is a semi-elasticity measuring how individuals react to different levels of the default savings rate at an aggregate level. Using the exogenous increase in the default savings rate in the first state-sponsored retirement plan in the U.S., OregonSaves, I find that, on average, about half of passive savers no longer stayed at the default rate when the default rate increased by one percentage point. While extensive previous research has examined the effect of default contributions, the causal effect of the default savings rate on saving behavior remains largely unclear due to data limitations. Previous studies dating back to Madrian and Shea (2001) have relied on data from employer-sponsored retirement plans whereby employers often match employee contributions to encourage employees to save. The presence of employer matching confounds the impact of the default rate on saving behavior, as employees' saving decisions are now influenced by both the default rate and employer matching. Given that the state-sponsored retirement plan, OregonSaves, does not allow employer matching, it provides a unique opportunity to tease out the causal effect of the default savings rate on retirement savings.

The second set of statistics reports revealed time preferences to enable inferences about

[2]The traditional approach to such welfare and optimal policy questions usually requires structural estimation of an underlying model's primitives, followed by a numerical simulation of the effects of policy changes. I instead identify a set of sufficient statistics directly measured from exogenous policy variations and survey data.

whether individuals would actively undersave if they were to opt out of the default rate. These statistics reveal the extent to which the default savings rate improves individual welfare by protecting individuals from actively undersaving. If individuals are unlikely to actively undersave, the welfare effect of correcting undersaving is dominated by the opposite welfare effect of correcting inaction. I conduct an online survey to elicit the time preferences for OregonSaves-eligible workers, and I find that, on average, workers weakly prefer spending most of their income now over dividing the income between now and the future. Plugging the point estimates into the optimal formula, I find the optimal default savings rate to be 7%.

Contributions to the Literature. The welfare analysis of the default savings rate proposed in this paper is related to three strands of literature. First, the optimal design of default retirement saving policies—the default savings rate in particular—has been a focus of recent research. Based on some early discussions about the welfare impact of the default rate (Thaler and Sunstein, 2003; Carroll et al., 2009), Bernheim et al. (2015) provided the first explicit guidance, concluding that the optimal default savings rate should be set at the employer matching cap in an employer-sponsored retirement plan. My analysis complements the findings of Bernheim et al. (2015), as I evaluate the optimal default without employer matching. I also develop the first sufficient statistics formula for the optimal default rate that directly connects the causal effect of the default rate with welfare analysis.

The present paper also contributes to the literature on optimal public policy with behavioral agents, which has been studied extensively in the context of income tax (Saez, 2001, 2002; Farhi and Gabaix, 2020), commodity tax (O'Donoghue and Rabin, 2006; Allcott et al., 2019), unemployment insurance (Chetty, 2008), and energy policy (Allcott et al., 2014; Allcott and Taubinsky, 2015).[3] I extend this literature to "nudge" policies—in particular,

[3]The following studies summarize how to incorporate behavioral biases into the welfare analysis of public policy: Chetty (2009); Mullainathan et al. (2012); Chetty (2015).

the default retirement savings rate. Farhi and Gabaix (2020) discussed optimal nudges in taxation, while I focus on optimal nudges in the context of default saving policies.

Finally, the present paper sheds light on two long-standing questions in household finance. The first is why a substantial fraction of American households save so little. Hubbard et al. (1995) and Scholz et al. (2006) argue that the explanation largely lies in asset-based and means-tested welfare programs and Social Security benefits. Here, I provide an alternative explanation: one reason people do not save is that they have limited access to automatic-enrollment retirement plans. Much research in development economics has studied the impact of expanding access to auto-enrollment retirement plans. Blumenstock et al. (2018) used experimental interventions in Afghanistan to conclude that automatic enrollment not only increases savings but also helps cultivate a saving habit. For instance, Carvalho et al. (2016) found that expanding access to savings accounts improves financial planning.[4] The second question I address pertains to the optimal level of savings. There is little consensus on the optimal level of retirement savings, given the substantial heterogeneity in health, life expectancy, desired retirement lifestyle, and family structure across the population. Instead of thinking about individuals' optimal level of savings, I provide a new perspective stemming from social welfare. Retirees lacking sufficient personal savings have to rely on social safety nets, which increases the fiscal burden imposed on all taxpayers to finance such programs. From a policy perspective, social welfare is maximized when the majority of people who can afford to save do so, leaving means-tested welfare programs to support only those people who cannot afford to save.

The rest of this paper is organized as follows. Section 2 describes the institutional background of the OregonSaves program and provides descriptive statistics for the first two years of the program as of April 2020. Section 3 discusses the welfare impact of the default rate in

[4]Related research has also been conducted in developed countries. For example, Epper et al. (2020) used Danish administrative data and found that more impatient individuals are less likely to save or accumulate wealth than patient individuals, which could drive wealth inequality exceeding income inequality.

a sufficient statistic framework and presents an explicit formula for the optimal default rate. Section 4 describes the identification strategies and estimation results of the key statistics in the optimal default formula using OregonSaves administrative and survey data. Section 5 calculates the optimal default rate using the estimation results from Section 4. Section 6 concludes.

2 An Overview of OregonSaves

In this section, I provide the institutional background and some preliminary empirical evidence on the first state-based mandatory retirement savings program in the United States, called OregonSaves.

2.1 Institutional Background

The 2015 passage of Oregon House Bill 2960 set into motion the creation of the OregonSaves program, the first U.S. state-sponsored retirement savings program. The Oregon Retirement Savings Board was given statutory authority to research and design the plan, with a target launch date of July 2017. OregonSaves requires that all private-sector employers including non-profit organizations either offer their own retirement plans or enroll their employees in OregonSaves. Besides Oregon, nine states have passed the legislation to establish a state-sponsored retirement plan to date. Table 1 provides information on the state-sponsored plans across the states.

OregonSaves is structured as a Roth Individual Retirement Account (IRA), with automatic enrollment, a default (after-tax)savings rate of 5%, and employee-only contributions. Once an employer registers and provides OregonSaves with employees data, employees enter a 30-day enrollment period during which time their identity is verified and employees may

choose to opt out. A Roth account[5] is created at the end of the enrollment period for each employee that has not opted out and whose identity is successfully verified. Enrollment in OregonSaves sets contributions levels at a 5% default rate, though employees can choose to save at different rates (up to 100% of pay), or opt out at any time. By default, the first $1,000 contributed into each participant's OregonSaves account is invested in a money market account. When a saver's account balance reaches $1,000, subsequent contributions default into an age-appropriate target date fund.

OregonSaves differs from conventional employer-sponsored retirement plans such as 401(k) or 403(b) plans in two key ways. First, OregonSaves participants may access their contributions invested in the money market account without penalty. The OregonSaves account is a combination of an emergency savings account (first $\$1,000$ withdrawal without penalty), and a retirement savings account (long-term investment returns from target date funds). Second, OregonSaves allows workers to contribute to the same account via different employers. In other words, workers can accumulate retirement savings in the same account over time. This feature of account-specific contributions can potentially encourage employees to accumulate more personal savings, especially for those working in smaller firms with high job turnover rates.

OregonSaves was rolled out to private-sector workers lacking access to workplace retirement plans in seven waves. A first wave of firms volunteered to be in the pilot program, followed by six compulsory waves. Employer waves were determined by the number of employees at the firm, with larger employers having to register earlier than smaller firms. For example, the largest firms (100+ employees) began a compulsory registration period on October 1, 2017. The smallest firms (4 or fewer employees) was scheduled to start enrolling May 12, 2020, but the compulsory registration period for smaller firms has been postponed

[5]Contributions to a Roth account are not tax free, while qualified withdrawals and earnings in the account are tax-free.

to 2021 due to the disruption of the COVID-19 pandemic. In practice, however, some smaller firms did register earlier than required, and some unknown number of larger firms may not have registered to date. It has been announced that an employer penalty will be levied on companies that do not provide access to their own retirement plans or to OregonSaves for employees, but the date for implementation of the fines has been pushed back.

Once an employer is registered, the firm submits employees' Social Security numbers, dates of birth, and names to OregonSaves, after which a 30-day enrollment period begins. During the enrollment period, employees may opt out of the program. If they do not do so during the first 15 days, OregonSaves then conducts an identity verification check. Employees who are successfully identified are then deemed eligible for enrollment at the end of the 30-day window.

2.2 Data and Descriptive Statistics

Using individual-level administrative data from the first three years of the OregonSaves program, I present empirical evidence on the impact of OregonSaves on expanding access to workplace retirement plans and on savings.[6]

As of April 30, 2020, 11,088 employers had registered their businesses in OregonSaves. This means that they previously did not offer employer-sponsored retirement plans, and all current and future employees would have access to OregonSaves with an option to opt out. About 263,660 individuals had information provided to OregonSaves by an employer, and median firm size was 10. Employees' median age in 2019 was 33. As of April 30, 2020, OregonSaves had accumulated $51 million in total assets.

Panel A of Table 2 itemizes the status of the 263,660 employees who had a chance to enroll in OregonSaves. During the enrollment window, 88,246 (33.5%) employees opted out, while another 23,076 (8.8%) employees opted out after the enrollment window. There were

[6] Additional descriptive results are documented in Chalmers et al. (2019).

30,808 (11.7%) employees awaiting the background check, which, in many cases, extends their pending status. There were 11,760 (4.5%) employees who enrolled and passed their background checks, but their employers had not yet submitted payroll. Finally, 109,770 (41.6%) names were enrolled, where the background check was successfully completed, the employer was submitting deferrals for at least one employee, and the employee had not opted out. In a sense, these are the employees who may now participate in OregonSaves. Nevertheless, the 30,808 pending cases and the 11,760 employees still to contribute are also potential participants for whom I cannot yet observe their choices.

Panel B of Table 2 describes the group I term Eligible Active Workers (EAW): these are employees eligible for an OregonSaves account and who appear to be actively working for at least one employer making payroll contributions for at least one employee. To be more precise, the EAW group includes those who opted out of OregonSaves while still actively working, plus people with a positive account balance in the past but currently a zero balance, plus people with a positive account balance currently, plus people with a positive balance and positive current contribution. This group comprises 152,112 people. In this group, 23,579 individuals made a monthly contribution to their accounts in April 2020, with a median contribution amount of $90. For employees with a positive contribution amount and a positive savings rate, I estimate their median monthly incomes (=contribution/contribution rate) to be $1,800. By way of comparison, the March 2018 Current Population Survey reports average monthly income of $4,843 (and median income of $3,411) for individuals who worked in the previous year. This comparison supports the conclusion that OregonSaves serves a population with low- and mid- income levels.

Panel A of Table 3 presents data for the 67,668 OregonSaves participants having a positive OregonSaves account balance. Given total assets of $51 million, the median balance per account stood at $316 as of April 2020. Panel B of Table 3 shows that 46,179 of the 67,668 with a positive balance are classified as eligible active workers. The median account balance

for EAW is $420.

3 A Sufficient Statistic Framework for the Optimal Default Savings Rate

Optimally designing the default savings rate is one of the key policy considerations for state and municipal governments interested in launching a government-sponsored retirement savings program similar to OregonSaves. In this section, I develop a sufficient statistics framework to derive the optimal default rate depending on statistics that can be directly estimated from the OregonSaves data described in the previous section.

3.1 Setup

In a two-period intertemporal choice model, workers need to divide their labor income Z between consumption and savings for retirement. Each worker has a underlying preferred savings rate, denoted θ. The preferred savings rate is determined by three exogenous parameters, her income Z, her normative time preference δ, and her behavioral time preference γ. The normative time preference δ captures the normative reasons to discount future utility (e.g. non-labor wealth, family structure, health, or bequest motive). The behavioral time preference δ captures the behavioral reasons to underestimate future utility (e.g. time inconsistency or misinformation). Appendix A presents details about the microfoundation of the preferred saving rate θ. Workers with the same preferred saving rate are defined as the same type, $\theta := (Z, \delta, \gamma) \in \Theta$, where the density of each type is $m(\theta)$.

A policymaker launches an automatic enrollment retirement savings program with a default saving rate $r \in (0, 1]$. The default saving amount for a given type-θ of workers with earnings $Z(\theta)$ is $R(\theta) = r \cdot Z(\theta)$. Each type of workers chooses a pension saving amount $P(\theta)$

from two discrete options: the default saving amount $R(\theta)$ or the preferred saving amount $S(\theta)$. The preferred saving amount equals the preferred savings rate θ times their income: $S(\theta) = \theta \cdot Z(\theta)$. Workers allocate their income between consumption C and savings P. The indirect decision utility function for a type-θ worker in the presence of a default rate r is expressed as:

$$U(C(\theta), P(\theta); \theta, r, K) = u(C(\theta)) + \gamma(\theta)\delta(\theta)v(P(\theta)) - K \cdot \mathbf{1}\{P(\theta) = S(\theta)\}, \tag{1}$$

where $C(\theta) + P(\theta) = Z(\theta)$ and $P(\theta) \in \{S(\theta), R(\theta)\}$. The functions $u(\cdot)$ and $v(\cdot)$ are both increasing and concave. The disutility K in the presence of a non-zero default rate represents the perceived costs of actively opting out of the default choice. I will refer to K as the opt-out costs,[7] which include but are not limited to time and psychological costs of switching from the default rate to the worker's preferred rate. I assume that the preferred saving rate θ, which determines each worker's type, is independent of the default rate r.

The policymaker thinks workers should maximize normative utility N, which can differ from decision utility U. The indirect normative utility function N is formally expressed as:

$$\begin{aligned} N(C(\theta), P(\theta); \theta, r, K, \pi) &= u(C(\theta)) + \delta(\theta)v(P(\theta)) - \pi K \mathbf{1}\{P(\theta) = S(\theta)\} \\ &= U + (1 - \gamma(\theta))\delta(\theta)v(P(\theta)) + (1 - \pi)K\mathbf{1}\{P(\theta) = S(\theta)\}, \end{aligned} \tag{2}$$

subject to the same budget constraint $C(\theta) + P(\theta) = Z(\theta)$. Following Goldin and Reck (2018), I define πK as the fraction of the normative opt-out costs: that is the realized cost after workers take action to opt out of the default that reduces their welfare by πK. The remaining fraction $(1 - \pi)K$ is the psychological opt-out costs that opted-out workers perceive ex ante but do not affect their welfare ex post. Similar to K, π is assumed to be

[7] The opt-out costs K specifically mean the costs of opting out of the default option, not opting out of the savings program. Opting out of the program is considered as electing a zero saving rate.

homogeneous across the population.

Equation (2) presents two sources of discrepancy between U and N. First, from the policymaker's perspective, workers might undervalue the utility of savings. The size of the underestimation, $(1-\gamma(\theta))\delta(\theta)v(P(\theta))$, is defined as the *welfare internality of savings*. This is the welfare gain of savings that workers do not consider when making saving decisions. One potential cause of this underestimation is due to the difference in time preferences between workers and the policymaker. Specifically, the policymaker is more forward-looking and discounts the value of future utility less than workers. This hypothesis is related to a large body of literature examining the disagreement in time preferences between the long-run self and the short-run self, where a policymaker can act like the long-run self (Laibson, 1997; O'Donoghue and Rabin, 1999). Moser and Olea de Souza e Silva (2017) and Choukhmane (2018) analyze the welfare consequences of time inconsistency in the context of retirement saving policies. A paper by Ericson and Laibson (2018) uses the term "present-focused" preferences to characterize individuals overestimating immediate utility compared to future utility documented in models such as hyperbolic and quasi-hyperbolic discounting, procrastination, and naivetè. Another potential reason for the underestimation of the utility of savings could be misinformation: that is, the policymaker may have more accurate information than do workers regarding public sources of retirement income such as Social Security and means-tested social transfers. Based on ambiguous or incorrect information, workers could be too optimistic about retirement support from social insurance and undervalue the importance of accumulating personal savings.

A second source of discrepancy between decision utility U and normative utility N could be that workers overlook the benefit from making an active decision.The size of the benefit from taking action, $(1-\pi)K$, is defined as the *welfare internality of action*. This is the welfare gain of taking action because workers perceive the cost before opting out of the default but the cost does not exist after opting out. One potential cause is that workers

overestimate opt-out costs. Such a miscalculation could explain why people stay at the default even though it may not be their preferred choice (Bernheim et al., 2015; Goldin and Reck, 2018; Luco, 2019). Underestimation of the benefit from making an active decision could also be caused by inattention (Caplin and Dean, 2015; Karlan et al., 2016; Gabaix, 2019). In the context of retirement savings, workers may fail to pay attention to planning for retirement or notice any policy changes that could impact their retirement security, so that they remain at the default.

Given worker's type-specific choices of consumption $C(\theta)$ and savings $P(\theta) \in \{S(\theta), R(\theta)\}$, the policymaker will select a default rate r to maximize aggregate normative utility weighted by type-specific Pareto weights $\alpha(\theta)$:

$$W(r) = \max_{r} \int_{\Theta} \alpha(\theta) N(C(\theta), P(\theta); \theta, r, K, \pi) dm(\theta) \tag{3}$$

subject to individual optimization

$$\{C(\theta), P(\theta)\} = \arg\max_{\{C,P\}} U(C, P; \theta, r, K) \tag{4}$$

where

$$C + P = Z(\theta) \text{ for all } \theta. \tag{5}$$

A key goal of the model section is to derive a formula for the optimal default rate from the policymaker's maximization problem presented in Equations (3)-(5). Moreover, the optimal formula can be empirically implemented using statistics observed from the data. I do this in two steps. First, I characterize the key economic forces in the formula that determine the optimal default. Second, I quantify these economic forces using statistics that can be empirically estimated.

3.2 Key Economic Forces Determining the Optimal Default Savings Rate

To understand how different economic forces shape the optimal level of the default savings rate, I use the marginal perturbation approach to derive the optimal default from the policymaker's maximization problem presented in Equations (3) - (5). Marginal perturbation was introduced by Saez (2001, 2002) and has been widely used in the optimal tax literature. I apply the same method to the context of the optimal default. The intuition of the marginal perturbation in the context of the optimal default is the following: When the default is set at an optimal level, the social welfare function $W(r)$ in Equation (3) is maximized. Consequently, a marginal change in the default rate r will not generate any first-order effect on welfare $W(r)$. In math, it means that $\frac{dW(r^*)}{dr} = 0$, where r^* stands for the optimal default savings rate.[8]

Figure 1 displays how marginal perturbation can be used to decompose the economic forces that determine the optimal level of the default. Figure 1a presents the distribution of the preferred rate s before the marginal perturbation. The preferred rate s is an underlying preference that each individual has. The y axis shows the density of individuals having the same preferred rate s, denoted m_s. For any default savings rate r, m_r is the total fraction of individuals passively accepting the default. These individuals are called passive savers. The fraction of passive savers m_r consists of two parts: The first part is individuals with a preferred rate between s_l and r. The fraction of these individuals is m_l. Although their preferred rates are lower than the default, their preferred rates are close enough to the default so that they are willing to stay at the default to avoid making an active decision. The second part of passive savers comes from individuals with a preferred rate between r and s_h. The fraction of these individuals is m_r. Similar to m_l, individuals in m_r have a preference close

[8]The application of marginal perturbation on the optimal default is based on a key assumption that the social welfare function $W(r)$ is smooth.

to the default so that they are willing to accept the default. Individuals with a preferred rate below s_l and above s_h opt out of the default r and elected their preferred rate. This is because their preference greatly deviates from the default so that they are not willing to accept the default.

Figures 1b - 1d present the decomposition of the determinants of the optimal default. The decomposition is based on three types of reaction from individuals affected by a marginal perturbation of the default rate. The marginal perturbation specifically means that instead of setting the initial default rate at r, the policymaker sets an alternative initial default rate which is marginally higher than r. The new default $r' = r + dr$. Figure 1b presents that when the initial default rate is r', the total fraction of passive savers is $m_{r'}$. A subset of passive savers under the old default r remain passive under the new default r'. This group of passive savers, denoted by P, consistently stay at the default because both defaults are close to their preferred rates. Since the default rate that they adhere to changes, it affects their savings level and consequently their normative utility described in Equation (2), which reflects how their savings level influences individual actual lifetime welfare. The change in normative utility for all passive savers in Group P can be aggregated and denoted by dI, which indicates the impact of the default rate on individual utility on the *intensive* margin. The aggregate change in normative utility on the intensive margin dI impacts the social welfare function in Equation (3) that the policymaker maximizes to determine the optimal default. The specifications of dI and other utility components introduced in this section will be explained in detail in the next section and Appendix B.

Figure 1c shows the second type of reaction from individuals affected by a marginal increase in the default rate. Under the default r, individuals with a preferred rate between s_l and s_l' are passive savers because r is close to their preferred rates. When these individuals face an increased default r' that further deviates from their preferences, they become active savers by opting out of r' and switching to their preferred rates. This group of savers on

the margin of opting out of the default is denoted by L, which indicates that they would actively elect a savings rate *lower* than the default after the marginal increase in the default. Their active decisions of changing the savings rate from the default to the preferred rate affect their normative utility in the social welfare function in Equation (3). For all savers in Group L, their aggregate normative utility changes by $dS_l - dR_l$, where dR_l indicates the welfare loss associated with the decision of opting out of the default rate and dS_l indicates the welfare gain associated with the decision of electing their preferred rates. Additionally, the aggregate normative utility for Group L also increases by dK_l, which is the welfare gain for making an active savings decision. As specified in the decision utility in Equation (1), individuals perceive the cost of opting out of the default and switching to their non-default preferred rate as K. When they make an active decision, the real opt-out cost is a portion of the perceived cost, which is specified in the normative utility in Equation (2) as πK. This means that individuals overestimate the opt-out cost by $(1-\pi)K$, which could be perceived as a welfare gain by active savers because their active decisions incur less cost than their estimate. This welfare gain for all savers in Group L is denoted by dK_l.

Figure 1d presents the third type of reaction from savers with a preferred rate between s_h and s_h'. This group of savers is denoted by H, which indicates that their preferred rates are *higher* than the default rate r. When the default rate is r, savers in Group H actively elected their preferred rates. When they face an alternative default r' which is marginally higher than r, savers in Group H become passive savers by accepting the default r' as their savings rate. The reason for this behavioral change is that compared to r, r' is closer to their preference so that staying at r' does not deviate from their preferred rates as much as r. Because Group H change their savings rate after the marginal increase, their normative utility in Equation (2) as a function of the savings rate also changes. At the aggregate level, the normative utility for all individuals in Group H changes by $dR_h - dS_h$, where dR_h indicates the welfare gain associated with the decision of accepting the default after the marginal increase in the

default and dS_h indicates the welfare loss associated with the decision of no longer saving at the preferred rate. Additionally, the aggregate normative utility for Group H also decreases by dK_h. For individuals in Group H who are on the margin of becoming passive savers under the default r', they feel indifferent between passively accepting the default and paying the perceived opt-out cost that they overestimate. In other words, instead of being active savers under r', they are worse off by implicitly paying more opt-out cost than they actually would need to pay.

3.3 Formula for the Optimal Default Savings Rate

Based on all the economic forces described in the previous section, we can derive the condition for the optimal default savings rate. When the default rate reaches the optimum, it will not generate any first-order effect on the sum of all the economic forces. Appendix B provides detailed proof for the following equation:

$$\begin{aligned}\frac{dW(r^*)}{dr} &= \frac{dI}{dr} + \frac{dS_l - dR_l \cdot r^*}{dr} + \frac{dK_l}{dr} + \frac{dR_h \cdot r^* - dS_h}{dr} + \frac{dK_h}{dr} \\ &= 0.\end{aligned}$$

3.4 Key Statistics and Assumptions

The optimal default condition is also based on a few key assumptions:

1. Individuals make their saving decisions once at the beginning of their working lives.[9]

2. The total opt-out costs K and the fraction of the normative opt-out costs π are homogeneous across types.

[9]Most retirement saving plans allow people to adjust their savings rates anytime, although in reality few people do so. Usually plan participants do not make active adjustments after they make their initial saving decisions (accepting the default, switching to a non-default rate, or opting out of the program) unless they face some exogenous shocks (i.e., income or unemployment shocks).

3. Individual preferred rates s are independent of the default rate r.

4. The utility function of savings $P(\theta)$ is linear: $v(P(\theta)) = P(\theta)$.

Proposition 1. *Based on Assumptions 1-4, the default savings rate satisfies the following equation at the optimum:*

$$r^* = \frac{dI + dS_l - dS_h + dK_l - dK_h}{dR_l - dR_h}.$$

Proof. See Appendix B.

3.5 Key Statistics and Assumptions

To derive a formula for the optimal default rate that is empirically implementable from the theoretical welfare framework, I introduce the following sufficient statistics:

- ϵ_l: the (observed) semi-elasticity of the percentage change in the density of l-type passive savers with preferred rates below the default (dm_l) with respect to all passive savers ($m_r = m_l + m_h$), as the default rate increases by 1 percentage point (dr), equal to $\frac{dm_l}{m_r}\frac{1}{dr}$;
- ϵ_h: the (observed) semi-elasticity of the percentage change in the density of h-type passive savers with preferred rates above the default (dm_h) with respect to all passive savers ($m_r = m_l + m_h$), as the default rate increases by 1 percentage point (dr), equal to $\frac{dm_h}{m_r}\frac{1}{dr}$;
- $g(\theta)$: type-specific social marginal welfare weights. This indicates the social marginal value of savings for a given type-s worker relative to the marginal value of public funds (λ) evaluated at the optimal default rate in units of dollars. The social marginal welfare weight measures the social value of each dollar that a type-θ worker saves from the

policymaker's perspective. Specifically, the policymaker values an additional dollar of savings from a type-θ worker as much as $\$g(\theta)$ from public funds. The welfare weights can be formally expressed as:

$$g(\theta) := \frac{\alpha(\theta) v'_{R^*(\theta)}}{\lambda}. \tag{6}$$

The derivative $v'_{R^*(\theta)} := \dfrac{dv(R^*(\theta))}{dR^*(\theta)}$, where $R^*(\theta) = r^* Z(\theta)$, r^* is the optimal default rate, and $Z(\theta)$ is the labor income for type-θ workers.

The welfare analysis is also based on several key assumptions sufficient to derive the optimal default rate:

1. Individuals make their saving decisions once at the beginning of their working lives.[10]
2. The total opt-out costs K and the fraction of the normative opt-out costs π are homogeneous across types.
3. Individual preferred rates s are independent of the default rate r.
4. The utility function of savings $P(\theta)$ is linear: $v(P(\theta)) = P(\theta)$.

4 Estimating Key Parameters for the Optimal Default Savings Rate

In this section, I outline empirical strategies to identify key statistics to calculate the optimal default savings rate in Proposition 1 using OregonSaves data described in Section 2. Table

[10]Most retirement saving plans allow people to adjust their savings rates anytime, although in reality few people do so. Usually plan participants do not make active adjustments after they make their initial saving decisions (accepting the default, switching to a non-default rate, or opting out of the program) unless they face some exogenous shocks (i.e., income or unemployment shocks).

8 lists all the statistics that need to be estimated and their values. Key statistics discussed in this section are:

- ϵ_l: the (observed) semi-elasticity of the percentage change in the density of l-type passive savers with preferred rates below the default (dm_l) with respect to all passive savers ($m_r = m_l + m_h$), as the default rate increases by 1 percentage point (dr), equal to $\frac{dm_l}{m_r}\frac{1}{dr}$;
- ϵ_h: the (observed) semi-elasticity of the percentage change in the density of h-type passive savers with preferred rates above the default (dm_h) with respect to all passive savers ($m_r = m_l + m_h$), as the default rate increases by 1 percentage point (dr), equal to $\frac{dm_h}{m_r}\frac{1}{dr}$;
- δ_l, δ_h: the normative time preference for l- and h-type passive savers; and
- γ_l, γ_h: the behavioral time preference for l- and h-type passive savers

4.1 Semi-Elasticities ϵ_l and ϵ_h

The (observed) semi-elasticity ϵ_l measures how many l-type passive savers stop saving at the default rate as the default marginally increases. L-type passive savers are previously defined as savers accepting the default and having a preferred rate below the default. As illustrated in Section 3.2 and Figure 1, when the default rate marginally increases, it further deviates from the preferred rates of some l-type passive savers. Consequently, in the presence of a higher default, some l-type passive savers under the lower default before the marginal increase now have a strong incentive to opt out of the default and switch to their preferences. The semi-elasticity ϵ_l quantifies the decrease in l-type passive savers. Specifically, it represents the percentage change in the fraction of l-type passive savers (dm_l) with respect to all passive savers ($m_r = m_l + m_h$), as the default rate increases by one percentage point (dr), equal to

$\frac{dm_l}{m_r}\frac{1}{dr}$.

The (observed) semi-elasticity ϵ_h quantifies how many h-type savers become passive savers as the default marginally increases. H-type passive savers are previously defined as savers adhering to the default and having a preferred rate above the default. When the default rate marginally increases, it moves closer to the preferred rates of some h-type savers. Therefore, more h-type savers who actively elected their preferred rates under the lower default are now likely to accept the higher default after the marginal increase. The semi-elasticity ϵ_h quantifies the increase in h-type passive savers. Specifically, it measures the percentage change in the fraction of h-type passive savers (dm_h) with respect to all passive savers ($m_r = m_l + m_h$), as the default rate increases by one percentage point (dr), equal to $\frac{dm_h}{m_r}\frac{1}{dr}$.

In this section, I describe how to use the exogenous increase in the default savings rate in the OregonSaves program from 2018 - 2020 to estimate both semi-elasticities.

4.1.1 Identification of ϵ_l

I exploit the exogenous increase in the default rate resulting from automatic escalation in the OregonSaves program to identify the changes in the fraction of passive savers. Automatic escalation is an exogenous increase in workers' non-zero savings rate that happens annually starting from 2019. Every year on January 1, workers who had opened an OregonSaves account for at least six months will be eligible for automatic escalation. Their savings rate automatically increases by 1 percentage point. It applies to all OregonSaves participants even if they have opted out of the default savings rate. In addition, participants eligible for automatic escalation will experience an increase in their savings rate only if they agree to accept the automatic escalation arrangement. By default, all participants accept the arrangement unless they actively opt out. Section 2 describes more institutional background about automatic escalation.

I use the following sample to identify ϵ_l that quantifies how the fraction of passive savers

changes with the default rate. Individuals who meet the following criteria are included in the sample: (1) Workers had actively worked for at least one employer offering OregonSaves one month before and three months after automatic escalation occurred; (2) Workers were eligible for automatic escalation. That is, workers' OregonSaves accounts had been open for at least six months before automatic escalation; (3) Workers had saved at the *default* rate before automatic escalation; and (4) Workers had accepted the default automatic escalation arrangement one month before it took effect.

The first row of Panel A in Table 4 indicates the estimate of ϵ_l using automatic escalation for passive savers that occurred on January 1, 2019. The default rate automatically increased from 5% to 6%. I find that 13,389 workers staying at the 5% default rate were eligible for auto-escalation in December 2018: $m_r = 13,389$. Three months after their default rate automatically increased to 6%, 6,246 of these workers actively elected a rate lower than 6%: $dm_l = -6,246$. That is the total decrease in l-type passive savers who actively elected a non-default rate lower than the new default after auto-escalation. Based on the definition of ϵ_l:

$$\begin{aligned} \epsilon_l &= \frac{dm_l}{m_r}\frac{1}{dr} \\ &= \frac{-6,246}{13,389}\frac{1}{6-5} \\ &= -0.47. \end{aligned} \tag{7}$$

The estimate of ϵ_l suggests that after the default rate marginally increases by 1 percentage point, 47% of passive savers opted out of the default rate and elected a non-default rate lower than the new default.

Besides passive savers who experienced auto-escalation in 2019, I also use another sample of passive savers who experienced auto-escalation in 2020 to test the robustness of the estimate of ϵ_l. The second row of Table 4 shows the estimate of ϵ_l using automatic esca-

lation that took effect on January 1, 2020. Workers who joined the OregonSaves program between July 1, 2018 and June 30, 2019 were not eligible for auto-escalation in 2019 but eligible for auto-escalation in 2020.[11] Among these workers, 27,846 of them saved at the 5% default rate of in December 2019 and their default rate increased to 6% on January 1, 2020: $m_r = 27,846$. Three months after the increase, 15,747 of these workers actively elected a rate lower than 6% : $dm_l = -15,747$.[12] Using auto-escalation from 5% to 6% in 2020, I find that ϵ_l is estimated to be -0.57.[13]

Taking the average of these two estimates using the 2019 sample and the 2020 sample, I find that as the default savings rate marginally increases by 1 percentage point, the fraction of passive savers reduces by 52%. These are passive savers who no longer saved at the default and switched to a non-default rate lower than the new default after the marginal increase.[14]

4.1.2 Identification of ϵ_h

The semi-elasticity ϵ_h measures the increase in passive savers as the default rate marginally increases. Some active savers under the low default might become passive savers when they face a high default rate. This is because for active savers who wish to save a large fraction of earnings, the high default is closer to their preferred savings rates than the low default. To identify the increase in passive savers, I use auto-escalation for active savers. Although active savers opted out of the initial default savings rate and elected a non-default rate,

[11] Workers' OregonSaves accounts are required to be open for at least six months before each automatic escalation to be eligible for the arrangement. Workers joined OregonSaves before July 1, 2018 were eligible for auto-escalation in 2019. Workers joined OregonSaves between July 1, 2018 and June 30, 2019 were eligible for auto-escalation in 2020.

[12] Based on the employment statistics from the U.S. Bureau of Labor Statistics, the unemployment rate in the state of Oregon caused by the coronavirus pandemic dramatically increased starting from April 2020. The first three months of 2020 were not significantly affected by the pandemic: `https://www.bls.gov/eag/eag.or.htm`. Here is the employment statistics from the U.S. Bureau of Labor Statistics Workers who changed their savings rate in the first three months of 2020 were mostly their responses to automatic escalation instead of to the macroeconomic conditions.

[13] $\epsilon_l = \frac{dm_l}{m_r}\frac{1}{dr} = \frac{-15,747}{27,846}\frac{1}{6-5} = -0.57$.

[14] Some passive savers completely opted out of the OregonSaves program after the default rate increases. They were considered as workers who switched to a zero savings rate.

some of them were still defaulted in the auto-escalation arrangement. Their non-default rate increase by 1 percentage point every year on January 1 until they actively opt out of auto-escalation. I use this exogenous variation to examine how many workers were active savers when the initial default rate was far from their preferred rates but passively accepted the marginal increase in their elected rates.

I use the following sample to study the transition from active savers to passive savers after auto-escalation: (1) Workers had actively worked for at least one employer offering OregonSaves one month before and three months after automatic escalation occurred; (2) Workers were eligible for automatic escalation. That is, workers' OregonSaves accounts had been open for at least six months before automatic escalation; (3) Workers had switched to a *non-default* and non-zero savings rate before automatic escalation; and (4) Workers had accepted the default automatic escalation arrangement one month before it took effect.

The first row in Panel B of Table 4 shows the estimate of ϵ_h using auto-escalation for active savers that occurred on January 1, 2019. Among all active savers who elected 7% as their savings rate before auto-escalation and met the criteria to be included in the sample, 26 of them stayed at 8% three months after auto-escalation. In other words, after auto-escalation, these 26 workers became passive savers by accepting their new default rate. The increase in passive savers is quantified by ϵ_h as follows:

$$\begin{aligned} \epsilon_h &= \frac{dm_h}{m_r}\frac{1}{dr} \\ &= \frac{26}{13,389}\frac{1}{8-5} \\ &= 0.0006. \end{aligned} \tag{8}$$

The estimate of ϵ_h suggests that when the default rate marginally raises by 1 percentage point, the fraction of passive savers increases by 0.06%. Equation (8) indicates that when the default rate is 5%, 13,389 were passive savers (the same denominator as in Equation (7)).

After auto-escalation, 26 workers who were previously active savers and elected 7% became passive savers by accepting the 8% default rate.

The second row in Panel B of Table 4 presents an estimate of ϵ_h using the sample of active savers who elected 8% before auto-escalation in 2019. After auto-escalation, their savings rate increased to 9%. Three months after auto-escalation, 14 workers who actively elected 8% before auto-escalation passively accepted 9%. Based on the definition of ϵ_h as shown in Equation (8), I find ϵ_h to be 0.03%.[15]

The third and fourth rows in Panel B of Table 4 present the estimates of ϵ_h using auto-escalation for active savers on January 1, 2020.[16] The average estimate of ϵ_h across all samples in 2019 and 2020 is 0.04%. This means that as the default rate marginally increases by 1 percentage point, the fraction of passive savers raises by 0.04%. These are active savers who elected a non-default rate higher than the old default before the marginal increase and passively accepted the new default after the marginal increase.

The size of ϵ_l is larger than ϵ_h by three orders of magnitude. This suggests that individuals with a low preferred savings rate are much more sensitive to the default rate than individuals with a high preferred savings rate. For individuals with a low preferred rate, many of them would opt out of the default when the default deviates from their preference. For individuals with a high preferred rate, their savings decisions are not significantly affected by the default rate.

[15] $\epsilon_h = \frac{14}{13{,}389} \frac{1}{9-5} = 0.0003$.

[16] The total number of passive savers before auto-escalation in the 2020 sample is 32,987. These are passive savers who joined OregonSaves before June 30, 2019 and were still active employees in 2020. This number is smaller than the sum of passive savers in 2019 and 2020 in Panel A of Table 4: $32,987 < 13,389 + 27,846 = 41,235$. This is because some passive savers counted in 2019 (13,389) were no longer active employees in 2020.

4.2 Normative and Behavioral Time Preferences δ and γ

The time preference parameters in the optimal default rate formula in Proposition 1 captures how a normative and a present self would discount future utility differently due to reasons including present bias, inattention, and misinformation. This section illustrates how I elicit and estimate the time preferences.

4.2.1 Survey Design and Sample

The survey was done in two stages. The first stage is a baseline survey that was fielded in June 2018 to workers who had access to OregonSaves in the first few months since its launch. Both workers who opted out of OregonSaves and those who were participating were invited to answer the survey. The baseline survey includes questions to collect socioeconomic information and savings and debt outside OregonSaves. The second stage is a follow-up survey that was fielded in June 2019. The follow-up survey includes questions to elicit time preferences. There were 441 workers who completed the baseline survey. Among the 441 workers, 143 completed the follow-up survey (32.4% response rate).

Survey respondents had two weeks to answer both the baseline and the follow-up surveys through an email link. They answered the survey by clicking the customized link in the invitation email. For all respondents who completed the survey, they received a $40 Starbucks electronic gift card that was sent through an email two weeks after the completion of the surveys. In the survey invitation email, workers were informed that both the invitation email and the gift card were sent from the Oregon State Treasury. This ensures the credibility of the payments.

To address the question about whether workers who participated in the survey are representative of the entire OregonSaves population, I compare the age and the savings rate between survey respondents and all OregonSaves workers in Figures 9 and 10 of Appendix D. Both figures present evidence that the age distribution and the savings rate distribution

of survey respondents closely follow the distributions of the overall population.

Although I was unable to provide survey payments that are based on individual survey responses, I show some supportive evidence that we can reasonably believe survey respondents truthfully answered survey questions. Figure 8 in Appendix D displays the difference in age between survey answers and administrative records. Survey respondents reported their age by answering the following question in the survey: "How old were you on your last birthday?" Additionally, all workers were asked to report their birth year when their OregonSaves savings account was open. The birth-year information was used to automatically invest workers' contributions to an age-appropriate fund. Since I do not know workers' exact birth date, there could be a one-year difference when we use survey responses and administrative records to compare age. Figure 8 in Appendix D shows that for more than 95% of all survey respondents, the age information from survey data is within the one-year difference from the age data in the administrative records. This suggests that most workers truthfully answered the survey questions. Additionally, Chapman (1996) has argued that despite the lack of incentives, participants often give truthful answers in hypothetical choices (Also see Frederick et al. (2002); Benhabib et al. (2010)).

4.2.2 Eliciting and Estimating Time Preferences

To elicit time preferences, survey respondents answered 16 hypothetical questions about how to allocate 100 experimental "tokens" to either a "sooner" time t, or a "later" time $t+k$, at different "token exchange rates" r. Participants faced 16 intertemporal decisions involving 16 combinations of $(t, k, 1+r)$, where $t = (0, 1)$, $k = (1, 2)$, and $1 + r = (1, 1.01, 1.02, 1.05)$. Table 5 shows the time periods, token budgets, token unit values, and annual interest rates for all 16 combinations. Appendix C provides the survey questions where four questions with the same set of (t, k) combination are displayed on the same page. Participants could change their answers to questions within the same set, but they could not change answers

after they moved on to the next page with a different (t, k) combination.

For each question, survey respondents moved a slider to choose C tokens to receive at a sooner time and R tokens to receive at a later time continuously along a convex budget set:

$$(1 + r)C + R = 100, \tag{9}$$

where participants had a budget of 100 tokens for each question. Tokens allocated at a sooner time were worth a_t while tokens allocated to a later time were worth a_{t+k}. For example, in the first question, each token was worth \$100 today and \$100 in a year. Participants were asked to move a slider to divide the 100 tokens between two time points as they preferred. In this question, $t = 0$, $k = 1$, and $1 + r = \frac{a_{t+k}}{a_t} = 1$. If one allocated 60 tokens today and 40 tokens to a year away, the survey would show the total dollar amount she would have today, \$6,000 (= \$100 $\times$ 60), and the total dollar amount she would have in a year, \$4,000 (= \$100 $\times$ 40). The total dollar amount allocated to a sooner time was denoted by C and the total dollar amount allocated to a later time was denoted by R in Equation (9).

Since the specific purpose of the survey is to help estimate time preferences for OregonSaves workers, the interest rates in the experiment are designed to be close to the real interest rate in the OregonSaves program. The initial \$1,000 contributions are invested in the money market fund and the vast majority of workers had accumulated savings less than \$1,000 when the survey was fielded. Therefore, the interest rate r in the experiment ranges from 0% to 5% to resemble the interest rate in the money market fund. This design helps to estimate the time preferences for the specific population in this specific OregonSaves setting. It also addresses the concern discussed in the literature that choices in time preference experiments could be influenced by interest rates in their real life (Cubitt and Read, 2007; Epper et al., 2020; Carvalho et al., 2016; Dean and Sautmann, 2020).

Given the intertermporal allocations that survey respondents made, I use the convex time

budget (CTB) approach to identify time preferences. It was first introduced by Andreoni and Sprenger (2012) to simultaneously estimate the time preferences and the curvature of the utility function, as time preferences could be affected by utility function curvature (Anderhub et al., 2001; Frederick et al., 2002; Andersen et al., 2008). The CTB approach has been widely used in monetary discounting experiments. Imai et al. (2019) summarized the recent 28 articles using this approach. Andreoni et al. (2015) compared the CTB approach with other widely used approach to elicit time preferences and specified the advantage of the CTB approach. Cohen et al. (2020) also provided a summary of the long history of experimental tests to elicit time preferences.

I use variations in starting times t to identify respondents' present bias parameter γ. I use variations in delay length k and interest rates $(1+r)$ to identify the annual discount factor δ and utility function curvature. Given consumption at a sooner time C and consumption at a later time R, I express decision utility U in the following three specifications. The first specification is a multi-period time separable CRRA (constant relative risk aversion) utility function subject to budget constraint (9):

$$U(C,R) = \frac{1}{\alpha}(C+W)^{\alpha} + \gamma\delta^{k}\frac{1}{\alpha}(R+W)^{\alpha}.$$

The parameter α is the CRRA curvature parameter, γ is the present bias parameter, δ is the annual discount factor, and k is the delay length between the two time points. The variable W is background consumption. Following Andreoni and Sprenger (2012), I assume that the background consumption level at two time points is the same. When I log-linearize the decision utility function $U(C,R)$, I obtain:

$$\ln\Big(\frac{C+W}{R+W}\Big) = \Big(\frac{\ln\gamma}{\alpha-1}\Big)\mathbf{1}\{t=0\} + \Big(\frac{\ln\delta}{\alpha-1}\Big)k + \Big(\frac{1}{\alpha-1}\Big)\ln(1+r). \tag{10}$$

W is the minimum annual consumption level asked in the survey. C and R are survey

responses to the intertemporal allocation questions described in Appendix C; $\mathbf{1}\{t = 0\}$ is an indicator if the sooner time period is today; k is the delay length between the sooner time and the later time described in Table 5; and $\ln(1 + r)$ is the natural log of the annual interest rate in Table 5. In order to handle corner solutions, I use a two-limit Tobit maximum likelihood regression to estimate parameters γ, δ, and α[17].

The second specification is the constant absolute risk aversion (CARA) utility function. The decision utility U in this formulation subject to budget constraint (9) is expressed as:

$$U(C, R) = -\exp(-\rho C) - \gamma\delta^k \exp(-\rho R),$$

where ρ is the coefficient of absolute risk aversion. The log-linearized utility function is:

$$C - R = \Big(\frac{\ln\gamma}{-\rho}\Big)\mathbf{1}\{t = 0\} + \Big(\frac{\ln\delta}{-\rho}\Big)k + \Big(\frac{1}{-\rho}\Big)\ln(1 + r). \tag{11}$$

Further simplify the tangency condition above, I obtain the following solution function:

$$C = \Big(\frac{\ln\gamma}{-\rho}\Big)\frac{\mathbf{1}\{t = 0\}}{2 + r} + \Big(\frac{\ln\delta}{-\rho}\Big)\frac{k}{2 + r} + \Big(\frac{1}{-\rho}\Big)\frac{\ln(1 + r)}{2 + r} + \frac{m}{2 + r}. \tag{12}$$

The third specification is the linear utility function as in Meier and Sprenger (2015):

$$U(C, R) = (C + W) + \gamma\delta^k(R + W).$$

After log-linearization, we have the following functional form:

$$\ln\frac{C}{R} = \ln(\gamma)\mathbf{1}\{t = 0\} + \ln(\delta)k. \tag{13}$$

[17]See Andreoni and Sprenger (2012); Augenblick et al. (2015) for a detailed explanation of using Tobit to estimate time preference parameters

[17]In estimation, I follow the approach in Andreoni and Sprenger (2012) to set $m/(2 + r)$ to 1.

4.2.3 Results

Table 6 shows baseline estimates of γ and δ taking into account the curvature of utility. There were 143 survey respondents who answered the time preference survey questions, and they made 1,765 intertemporal choices in total. Columns 1-2 show estimates of the CRRA regression (Equation (10)). The annual background consumption $W = 1,040$, equal to the negative of the minimum consumption level among all OregonSaves survey respondents. The average δ is 0.987 (standard deviation 0.005), and the average present bias parameter γ is 0.994 (s.d. 0.006). Columns 3-4 show estimates of the CARA regression (Equation (12)). The average δ is 0.988 (s.d. 0.005) and the average present bias parameter γ is 0.993 (s.d. 0.006).

Tables 9 - 11 in Appendix E present additional estimates of γ and δ and compare my estimates to those in previous studies. Despite some differences in the survey design between my experiment and previous experiments, my estimate of the present bias parameter γ (0.994) is close to the average estimate in all previous studies (0.95), which is discussed in Imai et al. (2019). In summary, studies using the CTB approach find little evidence of present bias. Additionally, Carvalho et al. (2016) further test that there is no differential present bias between liquidity constrained and liquidity unconstrained individuals. Although some papers find present bias close to what Laibson et al. (2015) suggested, these experiments usually ask individuals to allocate real-effort task instead of money (e.g., Augenblick et al. (2015); Carvalho et al. (2016); Dellavigna and Pope (2018); Augenblick and Rabin (2019)). In this OregonSaves setting where workers need to decide how much to save, using the intertemporal time preferences over money is more appropriate.

5 Computing the Optimal Default Savings Rate

For a baseline calculation of the optimal default rate, I use the average value of the point estimates of the semi-elasticities presented in Table 4, where $\epsilon_l = -0.52$ and $\epsilon_h = 0.0004$. I assign the point estimate of the normative time preference presented in Table 6 to both l- and h-type passive savers: $\delta_l = \delta_h = 0.987$. The behavioral time preference for h-type and l-type passive savers are assumed to be the average level under CRRA utility shown in Table 4: $\gamma_l = \gamma_h = 0.994$.

Another key statistics in the formula for the optimal default is the type-specific social marginal welfare weight $g(\theta)$. It generally characterizes how the policymaker evaluates each dollar, or equivalently the social value of each dollar. This concept was introduced by Saez (2001, 2002) and widely used in the tax literature (Lockwood and Weinzierl, 2016; Saez and Stantcheva, 2016; Allcott et al., 2019; Hendren, 2019). In the tax literature, the social marginal welfare weight for each income group usually measures the social value of transferring one dollar to a certain income group. In the context of default savings where there is no transfer across people, it indicates the social marginal value of savings for a given type-s worker relative to the marginal value of public funds (λ) evaluated at the optimal default rate in units of dollars.[18]

In Table 7, I report the estimates of the social marginal welfare weights for l-type passive savers (preferred rates below the default while saving at the default) and for h-type passive savers (preferred rates above the default while saving at the default). The welfare weight for a given type g is the Pareto weight α normalized by the aggregate weighted Pareto

[18]Specifically, the policymaker values an additional dollar of savings from a type-θ worker as much as $\$g(\theta)$ from public funds. The welfare weights can be formally expressed as:

$$g(\theta) := \frac{\alpha(\theta) v'_{R^*(\theta)}}{\lambda}. \tag{14}$$

The derivative $v'_{R^*(\theta)} := \dfrac{dv(R^*(\theta))}{dR^*(\theta)}$, where $R^*(\theta) = r^* Z(\theta)$, r^* is the optimal default rate, and $Z(\theta)$ is the labor income for type-θ workers.

weight $\bar{\alpha}$. The normalization ensures that the welfare weights only depend on the relative difference in income across types but are independent of the absolute size of income within type. I use observed data when the default rate is 5% to estimate the welfare weights at the optimal default. The first row of Table 7 presents the percent of passive savers for each type. Since the underlying preference of passive savers is unobserved, I use the following method to approximate the fraction of passive savers for each type. I use the observed fractions of active savers of each type to estimate the fraction of passive savers of each type. This approximation is based on an assumption that although the fraction of each type (active and passive savers combined) could be different, the ratio of active versus passive savers within type is the same.

The second row presents the estimated income for each type using the OregonSaves savings data in December 2019, which was the last month that all individuals faced 5% default rate. The individual-level monthly income equals the contribution amount divided by the contribution rate. Only individuals with a positive contribution amount and a positive rate are included. Imputed average annual income equals the average monthly income times 12. The third row shows that the primitive Pareto weight α equals the inverse of income. I follow the same method used in Saez (2002) to estimate the Pareto weight α. The fourth row presents that the aggregate weighted Pareto weight is the primitive Pareto weight α weighted by the percent of each type. The last row presents the final social marginal welfare weight for each type, where $g_l = 2.02$ and $g_h = 2.99$.

The optimal default rate is computed by plugging the baseline values listed in Table 8

into Proposition 1. The baseline optimal default rate is expressed as follows:[19]

$$\begin{aligned} r^*\% &= \frac{dI + dS_l - dS_h + dK_l - dK_h}{dR_l - dR_h} \\ &= \frac{588.4 + 1,223.4 - 2.1 + 262.6 - 0.3}{305.8 - 0.2}\% \\ &= 7\%. \end{aligned} \tag{15}$$

To put my calculation into perspective, we need to consider two additional factors. The first is that I use individual responses to auto-escalation of the default rate from 5% to 6% to estimate how two identical groups of workers would respond to two *initial* default rates 5% and 6%. Although using the exogenous auto-escalation could generate reasonable estimates as we compare different default rates in a local region, we cannot completely rule out the possibility that individuals who are presented with an initial 6% default rate could be more responsive than those who are auto-escalated from 5% to 6%. In the next section, I provide sensitivity analysis of how the optimal default rate could change when individuals are more responsive to the default.

The second factor that could impact the level of the optimal default is other retirement income sources such as Social Security. Individuals usually plan for how much to save based on the amount of Social Security benefits they think they would get in retirement. In the presence of a default option, the *perceived* level of Social Security benefits informs the decision on whether to passively accept the default or actively switch to a non-default rate. In my model, this binary decision is captured by the semi-elasticities, which are key determinants of the optimal default rate. In short, the impact of Social Security on the optimal default is embedded in the measure of the semi-elasticities. In the next section, I discuss how policy changes in Social Security or other savings-related policy changes impact individual

[19] An empirical assumption required to calculate dI, the welfare impact on the intensive margin, is that I use an unweighted average welfare component to approximate a weighted average welfare component, as the underlying fractions of h-type and l-type savers are unobserved.

perceptions about their Social Security benefits, and subsequently impact how individuals respond to the default, and then eventually shape the level of the optimal default rate.

5.1 Sensitivity Analysis

Figures 2 - 7 present the optimal default savings rate r^* under alternative values of the key parameters. In Figure 2, the solid curve presents how the optimal default r^* varies with l-type semi-elasticity ϵ_l, where ϵ_l measures the *decrease* in the fraction of passive savers whose preferred rates are below the default (l-type passive savers) as the default marginally increases. The intersection of the dash-dotted line and the solid curve is r^* at the baseline value of ϵ_l. The optimal default r^* exponentially increases with ϵ_l. When $\epsilon_l = -1$, all passive savers with preferred rates below the default would opt out of the default as it marginally increases. The policymaker should set the default around 6% to encourage some individuals to stay at the default. When $\epsilon_l = 0$, no passive savers would ever opt out regardless of the level of the default rate. The policymaker should set the default very high to increase savings which could maximize lifecycle utility subject to the budget constraint.

In Figure 3, the solid curves present how r^* varies with h-type semi-elasticity ϵ_h, where ϵ_h measures the *increase* in the fraction of passive savers whose preferred rates are above the default (h-type passive savers) as the default marginally increases. When ϵ_h is between 0 and 0.5, r^* is bounded below 7% and decreases with ϵ_h. When ϵ_h is between 0.5 and 1, r^* is bounded above 14% and decreases with ϵ_h. When the magnitude of ϵ_h is close to the magnitude of the baseline value of $\epsilon_l (= -0.52)$, the size of the increase in l-type passive savers equals the size of the decrease in h-type passive savers. The changes in the default rate induce little welfare effect as the total number of passive savers remains constant. The optimal default r^* could be either 0 or very large. If $\epsilon_h = 0$, no h-type active savers under a low default would want to become passive savers under a high default even though the high default is close to their preferred rate. The policymaker should set the default around 7%

to maximize the lifecyle utility for existing passive savers. If $\epsilon_l = 1$, all h-type active savers under a low default would save at a high default. The policymaker should set the default relatively high around 14% to minimize the aggregate opt-out cost of h-type workers.

The solid curve in Figure 4 displays how r^* varies with δ_l, which is the average annual discount factor for savers with a preferred rate lower than the default. The dash-dotted line presents how r^* varies with δ_h, which is the average annual discount factor for savers with a preferred rate higher than the default. The dashed vertical line highlights r^* under the baseline values of δ_l and δ_h. The optimal default r^* exponentially decreases with δ_l. This suggests that when δ_l is low, the (positive) welfare internality of saving at the default for savers with a low preferred rate is small, measured by dR_l in Proposition 1. The policymaker should set the default very high to encourage people with a low preferred rate to opt out of the default. When δ_l is high, the welfare internality of saving at the default is large. The policymaker should set the default around 7% to attract savers to stay at the default. Additionally, r^* slightly increases with δ_h. When δ_h is low, the (positive) welfare internality of saving at the default for savers with a high preferred rate is low, measured by dR_h in Proposition 1. The policymaker should set the default relatively low to encourage these savers to opt out of the default. Since a low default deviates from these savers' preferences, they do not have strong incentives to stay at the default. Moreover, since savers with a high preferred rate are not very responsive to the default implied by the small magnitude of ϵ_h, the impact of δ_h on the optimal default r^* is weak.

Figure 5 shows r^* under different values of present bias parameters. The solid curves present how r^* varies with γ_l, which is the average present bias parameter for savers with a preferred rate lower than the default (l-type savers). The dash-dotted line presents how r^* varies with γ_h, which is the average present bias parameter for savers with a preferred rate higher than the default (h-type savers). The dashed vertical line highlights r^* under the baseline values of γ_l and γ_h. When γ_l is close to 1, l-type savers are neither present-biased nor

future-biased. The policymaker should set the default either to a very high level or to zero so that l-type savers are compelled to select their preferred rates. When l-type savers are either present-biased ($\gamma_l < 1$) or future-biased ($\gamma_1 > 1$), the optimal default is bounded between 5% and 8% to correct individuals from either actively under-saving or actively over-saving. The dash-dotted line shows that r^* decreases with γ_h, which is the present bias parameter for h-type savers. If h-type savers are present-biased ($\gamma_h < 1$), the policymaker should set a high default that is above their preferred rate to encourage them to save more. Using a high default to encourage more default savings could be more effective to h-type savers than l-type savers because h-type savers are less likely to opt out of the default than l-type savers. If h-type savers are future-biased ($\gamma_h > 1$), r^* is below 8% to attract these savers to save less than their preferred level by staying at the default.

In Figure 6, I show r^* under a relaxed assumption that the total opt-out cost K perceived by individuals is heterogeneous across different types of savers. The solid line presents how r^* varies with K_l, which is the average perceived opt-out cost for savers with a preferred rate lower than the default. The dash-dotted line presents how r^* varies with K_h, which is the average perceived opt-out cost for savers with a preferred rate higher than the default. The dashed vertical line highlights r^* under the baseline values of K_l and K_h, where $K_l = K_h = K = \$250$. Both K_l and K_h vary between 0 to $\$2,000$, which is the maximum estimate in previous studies (DellaVigna, 2009; Bernheim et al., 2015). The optimal default r^* increases with K_l and is bounded under 13%. When the total opt-out cost perceived by savers with a low preferred rate (K_l) is high, the policymaker should set a relatively high default that significantly deviates from savers' preferences. Consequently, the benefit of opting out of the default outweighs the perceived cost. More l-type savers would opt out of the default despite the high perceived opt-out cost. Additionally, r^* does not change with K_h. This is mainly because savers with a high preferred rate are unlikely to adjust their savings rate based on the level of the default rate, measured by the low semi-elasticity ϵ_h.

Finally, in Figure 7, I present r^* under another relaxed assumption that the fraction of the normative opt-out cost π (with respect to the total perceived opt-out cost) is heterogeneous across different types of savers. The solid line presents how r^* varies with π_l, which is the average fraction of the normative opt-out cost for savers with a preferred rate lower than the default. The dash-dotted line presents how r^* varies with π_h, which is the average fraction of the normative opt-out cost for savers with a preferred rate higher than the default. The dashed vertical line highlights r^* under the baseline values of π_l and π_h, where $\pi_l = \pi_h = \pi = 0$. Both lines show that r^* are bounded between 5% and 7% for any value of π_l and π_h between 0 and 1. This suggests that the extent to which individuals overestimate the opt-out cost does not significantly impact the level of the optimal default. This is partly explained by the relatively small magnitude of the welfare internalities of action dK_l and dK_h relative to other components in Proposition 1. The magnitude of each component in the baseline calculation is displayed in Equation 15.

6 Conclusion

This paper has proposed a tractable framework that directly connects empirical analysis of the causal impact of the default savings rate on individual saving behavior with welfare analysis of the optimal design of the default savings rate. I introduced a novel set of sufficient statistics to capture individual adherence to the default savings rate as the default rate varies. Given individual responsiveness to the default savings rate, I characterized how the level of the default savings rate impacts individual welfare. Specifically, the default savings rate could improve individual welfare by protecting workers from two types of behavioral biases: actively undersaving and inaction. In a unified welfare framework that incorporates these two biases, I showed that if workers tend to procrastinate to make an active decision, it might be welfare-improving to set the default rate at an undesirable level to compel individuals

to opt out of the default rate and choose a non-default savings rate. In contrast, if the non-default preferred rate that individuals actively choose does not maximize their lifecycle utility, it might be welfare-improving to set the default rate at a desirable level to encourage workers to stay at the default rate.

Using individual-level administrative and survey data from the first state-sponsored auto-enrollment plan in the U.S. called OregonSaves, I found that, when the default rate increased by one percentage point, on average about half of workers who had passively stayed at the previous default rate switched to a non-default rate or opted out of the program. I also found that OregonSaves-eligible workers show little evidence of undervaluing the utility of savings. Given these estimates, a baseline calculation suggested that the optimal default savings rate should be set at 7%, somewhat higher than the current 5% default rate.

Analyzing optimal policy with reduced-form empirical identification has been widely adopted in the context of income transfer programs such as tax policy and social insurance. The present paper extends the applicability of this approach to default saving policy, and shows that the same approach to addressing welfare and optimal policy questions based on empirical evidence can be applied to a broader context of economic policies including nudge policy and retirement policy.

References

Allcott, Hunt, Benjamin B. Lockwood, and Dmitry Taubinsky (2019). Regressive Sin Taxes, with an Application to the Optimal Soda Tax. *Quarterly Journal of Economics 134*(3), 1557–1626.

Allcott, Hunt, Sendhil Mullainathan, and Dmitry Taubinsky (2014). Energy Policy with Externalities and Internalities. *Journal of Public Economics 112*, 72–88.

Allcott, Hunt and Dmitry Taubinsky (2015). Evaluating Behaviorally Motivated Policy: Experimental Evidence from the Lightbulb Market. *American Economic Review 105*(8), 2501–2538.

Anderhub, Vital, Uri Gneezy, Werner Guth, and Doron Sonsino (2001). On the Interaction of Risk and Time Preferences: An Experimental Study. *German Economic Review 2*(3), 239–253.

Andersen, Steffen, Glenn W. Harrison, Morten I. Lau, and E. Elisabet Rutström (2008). Eliciting Risk and Time Preferences. *Econometrica 76*(3), 583–618.

Andreoni, James, Michael A. Kuhn, and Charles Sprenger (2015). Measuring Time Preferences: A Comparison of Experimental Methods. *Journal of Economic Behavior and Organization 116*, 451–464.

Andreoni, James and Charles Sprenger (2012). Estimating Time Preferences from Convex Budgets. *American Economic Review 102*(7), 3333–3356.

Augenblick, Ned, Muriel Niederle, and Charles Sprenger (2015). Working Over Time: Dynamic Inconsistency in Real Effort Tasks. *Quarterly Journal of Economics* (2001), 1067–1115.

Augenblick, Ned and Matthew Rabin (2019). An Experiment on Time Preference and Misprediction in Unpleasant Tasks. *Review of Economic Studies 86*(3), 941–975.

Benhabib, Jess, Alberto Bisin, and Andrew Schotter (2010). Present-Bias, Quasi-Hyperbolic Discounting, and Fixed Costs. *Games and Economic Behavior 69*(2), 205–223.

Bernheim, B. Douglas, Andrey Fradkin, and Igor Popov (2015). The Welfare Economics of Default Options in 401 (k) Plans. *American Economic Review 105*(9), 2798–2837.

Beshears, John, James J. Choi, David Laibson, and Brigitte C. Madrian (2009). The Importance of Default Options for Retirement Saving Outcomes: Evidence from the United States. In *Social Security Policy in a Changing Environment*, Number June, pp. 167 – 195. University of Chicago Press.

Blumenstock, Joshua, Michael Callen, and Tarek Ghani (2018). Why Do Defaults Affect Behavior? Experimental Evidence from Afghanistan. *American Economic Review 108*(10), 2868–2901.

Caplin, Andrew and Mark Dean (2015). Revealed Preference, Rational Inattention, and Costly Information Acquisition. *American Economic Rreview 105*(7), 2183–2203.

Carroll, Gabriel D, James J Choi, David Laibson, Brigitte C Madrian, and Andrew Metrick (2009). Optimal Defaults and Active Decisions. *Quarterly Journal of Economics 124*(4), 1639–1674.

Carvalho, Leandro S., Silvia Prina, and Justin Sydnor (2016). The Effect of Saving on Risk Attitudes and Intertemporal Choices. *Journal of Development Economics 120*, 41–52.

Chalmers, John, Olivia S. Mitchell, Jonathan Reuter, Geoffrey T. Sanzenbacher, and Mingli Zhong (2019). Auto-Enrollment Retirement Plans for the People: Choices and Outcomes in OregonSaves. *Working Paper*.

Chapman, Gretchen B. (1996). Temporal Discounting and Utility for Health and Money. *Journal of Experimental Psychology: Learning Memory and Cognition 22*(3), 771–791.

Chetty, R (2008). Moral hazard vs. Liquidity and Optimal Unemployment Insurance. *Journal of Political Economy 116*(2), 173–234.

Chetty, Raj (2009). Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods. *Annual Review of Economics 1*(1), 451–488.

Chetty, Raj (2015). Behavioral Economics and Public Policy: A Pragmatic Perspective. *American Economic Review: Papers and Proceedings 105*(5), 1–33.

Chetty, Raj, John N Friedman, Sore Leth-Petersen, Toben Heien Nielsen, and Tore Olsen (2014). Active vs. Passive Decisions and Crowd-Out in Retirement Savings Accounts. *Quarterly Journal of Economics 129*(3), 1141–1219.

Choi, James, David Laibson, Brigitte Madrian, and Andrew Metrick (2004). For Better or for Worse: Default Effects and 401(k) Savings Behavior. In *Perspectives on the Economics of Aging*, Volume 401, pp. 81–126.

Choukhmane, Taha (2018). Default Options and Retirement Saving Dynamics. *Working Paper*.

Cohen, Jonathan, Keith Marzilli Ericson, David Laibson, and John Myles White (2020). Measuring Time Preferences. *Journal of Economic Literature 58*(2), 299–347.

Cubitt, Robin P. and Daniel Read (2007). Can Intertemporal Choice Experiments Elicit Time Preferences for Consumption? *Experimental Economics 10*(4), 369–389.

Dean, Mark and Anja Sautmann (2020). Credit Constraints and the Measurement of Time Preferences. *Review of Economics and Statistics*, 1–52.

DellaVigna, Stefano (2009). Psychology and Economics: Evidence from the Field. *Journal of Economic Literature 47*(2), 315–372.

Dellavigna, Stefano and Devin Pope (2018). What Motivates Effort? Evidence and Expert Forecasts. *Review of Economic Studies 85*(2), 1029–1069.

Epper, Thomas, Ernst Fehr, Helga Fehr-Duda, Claus Thustrup Kreiner, David Dreyer Lassen, Soren Leth-Petersen, and Gregers Nytoft Rasmussen (2020). Time Discounting and Wealth Inequality. *American Economic Review 110*(4), 1177–1205.

Ericson, Keith Marzilli and David Laibson (2018). Intertemporal Choice: Theory. *NBER Working Paper 25358*.

Farhi, Emmanuel and Xavier Gabaix (2020). Optimal Taxation with Behavioral Agents. *American Economic Review 110*(1), 298–336.

Frederick, Shane, George Loewenstein, and Ted O'Donoghue (2002). Time Discounting and Time Preference: A Critical Review. *Journal of Economic Literature 40*(2), 351–401.

Gabaix, Xavier (2019). Behavioral Inattention. In *Handbook of Behavioral Economics: Applications and Foundations, Vol. 2*, pp. 261–343.

Goda, Gopi Shah, Matthew R. Levy, Colleen Flaherty Manchester, Aaron Sojourner, and Joshua Tasoff (2020). Who Is a Passive Saver under Opt-In and Auto-Enrollment? *Journal of Economic Behavior and Organization 173*, 301–321.

Goldin, Jacob and Daniel H. Reck (2018). Optimal Defaults with Normative Ambiguity. *SSRN Working Paper 2893302*.

Hendren, Nathaniel (2019). Efficient Welfare Weights. *Working Paper*.

Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes (1995). Precautionary Saving and Social Insurance. *Journal of Political Economy 103*(2), 360.

Imai, Taisuke, Tom Rutter, and Colin Camerer (2019). Meta-Analysis of Present-Bias Estimation Using Convex Time Budgets. *CESifo Working Paper*.

Karlan, Dean, Margaret McConnell, Sendhil Mullainathan, and Jonathan Zinman (2016). Getting to the Top of Mind: How Reminders Increase Saving. *Management Science 62*(12), 3393–3411.

Laibson, David (1997). Golden Eggs and Hyperbolic Discounting. *Quarterly Journal of Economics 112*(2), 443–477.

Laibson, David, Andrea Repetto, and Jeremy Tobacman (2015). A Debt Puzzle. In *Household Finance*, Volume 309, pp. 810–861.

Lockwood, Benjamin B. and Matthew Weinzierl (2016). Positive and Normative Judgments Implicit in U.S. Tax Policy, and the Costs of Unequal Growth and Recessions. *Journal of Monetary Economics 77*, 30–47.

Luco, Fernando (2019). Switching Costs and Competition in Retirement Investment. *American Economic Journal: Microeconomics 11*(2), 26–54.

Madrian, Brigitte C. and Dennis F. Shea (2001). The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior. *Quarterly Journal of Economics 116*(4), 1149–1187.

Meier, Stephan and Charles D. Sprenger (2015). Temporal Stability of Time Preferences. *Review of Economics and Statistics 97*(2), 273–286.

Moser, Christian and Pedro Olea de Souza e Silva (2017). Optimal Paternalistic Savings Policies. *Columbia Business School Research Paper No. 17-51*.

Mullainathan, Sendhil, Joshua Schwartzstein, and William J. Congdon (2012). A Reduced-Form Approach to Behavioral Public Finance. *Annual Review of Economics 4*(1), 511–540.

O'Donoghue, Ted and Matthew Rabin (1999). Doing It Now or Later. *American Economic Review 89*(1), 103–124.

O'Donoghue, Ted and Matthew Rabin (2006). Optimal Sin Taxes. *Journal of Public Economics 90*(10-11), 1825–1849.

Saez, Emmanuel (2001). Using Elasticities to Derive Optimal Income Tax Rates. *Review of Economic Studies*, 205–229.

Saez, Emmanuel (2002). Optimal Income Transfer Programs: Intensive versus Extensive Labor Supply Responses. *Quarterly Journal of Economics 117*(3), 1039–1073.

Saez, Emmanuel and Stefanie Stantcheva (2016). Generalized Social Marginal Welfare Weights for Optimal Tax Theory. *American Economic Review 106*(1), 24–45.

Scholz, John Karl, Ananth Seshadri, and Surachai Khitatrakun (2006). Are Americans Saving " Optimally " for Retirement? *Journal of Political Economy 114*(4), 607–643.

Thaler, Richard H. and Cass R. Sunstein (2003). Libertarian Paternalism. *American Economic Review 93*(2), 175–179.

Table 1: State Legislation Establishing a State-Sponsored Retirement Plan, October 2020

State	Type of program	Status	Default rate	Program/Bill website
Oregon	Mandatory auto Roth IRA	Launched in July 2017	5%, auto escalation up to 10%	OregonSaves
Illinois	Mandatory auto Roth IRA	Launched in May 2018	5%, no auto escalation	Illinois Secure Choice
California	Mandatory auto Roth IRA	Launched in July 2019	5%, auto escalation up to 8%	CalSavers
Maryland	Mandatory auto Roth IRA	Scheduled to launch in mid-2020	To be determined	Maryland$aves
Connecticut	Mandatory auto Roth IRA	Bill passed in 2016	To be determined	Connecticut program
New Jersey	Mandatory auto Roth IRA	Bill passed in March 2019	3%	New Jersey Secure Choice Savings Program Act
Colorado	Mandatory auto Roth IRA	Bill passed in June 2020	5%	Colorado Secure Savings Program Act
Vermont	Voluntary to employers; auto Roth IRA to workers	Bill passed in June 2017	To be determined	Green Mountain Secure Retirement Plan
New York	Voluntary to employers; auto Roth IRA to workers	Bill passed in February 2018; scheduled to launch in April 2020	To be determined	New York State Secure Choice Savings Program Act
Washington	Expanding from a voluntary program to a mandatory program to all private-sector businesses	Voluntary program launched in 2015; bill for the mandatory program passed the State Senate in March 2019; waiting for a House floor vote	To be determined	Washington Secure Choice Savings Program Act
Massachusetts	Expanding from a voluntary program only to non-profits to a mandatory program to all private-sector businesses	Voluntary program launched in October 2017; bill for the mandatory program introduced in January 2019	To be determined	Massachusetts Secure Choice Savings Program Act

Note: In a mandatory auto Roth IRA program, private-sector employers are required to provide employees access to either a state-sponsored plan or an employer-sponsored plan such as 401(k). Employees are automatically enrolled in a retirement plan with a default contribution rate. They can always opt out or elect a non-default contribution rate. Roth IRA is an individual retirement account where contributions are not tax-free but qualified withdrawals and earnings in the account are tax-free. Besides these 11 states that have passed the legislation for a voluntary or a mandatory program, about another 20 states have introduced legislation but not yet enacted. AARP summarized the status of these states: `https://www.aarp.org/ppi/state-retirement-plans/savings-plans/`.

Table 2: Summary Statistics for Individuals Ever Had Access to OregonSaves, April 2020

	N	%	$
Panel A: All individuals			
Total unique individuals entered by employers	263,660	100.0	–
Immediate opted-out individuals	88,246	33.5	–
Delayed opted-out individuals	23,076	8.8	–
Pending individuals	30,808	11.7	–
Enrolled individuals w/o payroll info	11,760	4.5	–
Enrolled individuals with payroll info	109,770	41.6	–
Panel B: Eligible active workers (EAW)			
Total EAW	152,112	100.0	–
Immediate opted-out workers	68,320	44.9	–
Delayed opted-out workers	17,228	11.3	–
EAWs with no balance	25,624	16.9	–
Suspended contributors	17,361	11.4	–
Contributors	23,579	15.5	–
Average monthly contributions if > 0, April 2020	–	–	121
Average monthly income	–	–	2,284

Notes: Data from anonymized administrative records as of April 30, 2020. In Panel A, immediate opted-out individuals left the OregonSaves program during the first 30-day enrollment window. Delayed opted-out individuals left the OregonSaves program after the 30-day window. Pending individuals were in the background check, failed the background check, or in the 30-day window (all employers). Enrolled individuals with payroll information passed the background check and the initial 30-day window (at least 1 employer), but program is waiting for payroll information. Enrolled individuals with payroll information passed the background check, passed the initial 30-day window (at least 1 employer), and the same employer(s) provided payroll information. In Panel B, eligible active workers (EAW) are persons eligible for contributions (at least one employer) and inferred to be actively working on April 30, 2020. Individuals eligible for contributions have passed the background check and the 30-day enrollment window (at the same employer(s)), which provided payroll information for at least one employee at the firm. Suspended contributors are EAWs with a positive balance but no monthly contributions in April 2020. Contributors are EAWs with a positive balance and positive monthly contributions in April 2020.

Table 3: Summary Statistics for Individuals with a Positive Account Balance, April 2020

	N	$
Panel A: All individuals with a positive balance		
All individuals with a positive balance	67,668	–
Opted-out individuals with a positive balance	7,347	–
Participating individuals with a positive balance	60,321	–
Median balance if positive	–	316
Total assets	–	51 million
Panel B: Eligible active workers (EAW) with a positive balance		
EAWs with a positive balance	46,179	–
Opted-out EAWs with a positive balance	5,239	–
Participating EAWs with a positive balance	40,940	–
Median balance if positive	–	420

Notes: Data from anonymized administrative records on April 30, 2020. Panel A reports statistics for all individuals ever had access to OregonSaves with a positive balance on April 30, 2020. Opted-out individuals with a positive balance are persons who opted out of the program before April 30, 2020 but had ever contributed and did not withdraw all contributions. Participating individuals with a positive balance are persons who were participating in the program on April 30, 2020, had ever contributed, and did not withdraw all contributions. Panel B presents statistics for eligible active workers (EAW) with a positive balance on April 30, 2020. EAWs are persons eligible for contributions (at least one employer) and inferred to be actively working on April 30, 2020. Individuals eligible for contributions have passed the background check and the 30-day enrollment window (at the same employer(s)), which provided payroll information for at least one employee at the firm.

Table 4: Estimates of Semi-Elasticities

Panel A: ϵ_l					
Auto-escalation year	Sample size (m_r)	Decrease in passive savers (dm_l)	Default rate before increase (%)	Default rate after increase (%)	Value of ϵ_l
2019	13,389	-6,246	5	6	-0.47
2020	27,846	-15,747	5	6	-0.57
Average value					-0.52
Panel B: ϵ_h					
Auto-escalation year	Sample size (m_r)	Increase in passive savers (dm_h)	Default rate before increase (%)	Default rate after increase (%)	Value of ϵ_h
2019	13,389	26	5	8	0.0006
2019	13,389	14	5	9	0.0003
2020	32,987	61	5	8	0.0006
2020	32,987	30	5	9	0.0005
Average value					0.0004

Notes: This table presents the estimates of two semi-elasticities ϵ_l and ϵ_h. The identification strategy and the results are discussed in Section 4.1. Panel A shows the results of ϵ_l. Based on its definition expressed in Equation (7), ϵ_l measures how many passive savers no longer stay at the default rate after an exogenous increase in the default rate. The first row in Panel A shows that before the default rate increased from 5% to 6% on January 1, 2019, 13,389 passive savers stayed at the old default 5%. After the increase, 6,246 passive savers opted out of the 6% new default rate and switched to their preferred non-default rate lower than 6%: $\epsilon_l = -0.47 = \frac{-6,246}{13,389}\frac{1}{6-5}$. Using the samples in 2019 and 2020, I find the average value of ϵ_l is -0.52. Only eligible active workers (EAW) summarized in Tables 2 and 3 are included in the sample. Panel B presents the results of ϵ_h. Based on its definition expressed in Equation (8), ϵ_h measures how many active savers under a lower default rate become passive savers under a higher default as the higher default is closer to their preferred rate. The first row in Panel B shows that when the default rate is 5%, 13,389 were passive savers while some workers actively elected 7% and their elected rate would be automatically increased to 8% on January 1, 2019. Among these workers who elected 7%, 26 of them accepted the 8% new default rate after their elected rate increased from 7% to 8% in 2019. These 26 workers were active savers under 5% default rate but passive savers under 8% default rate. Using this sample, I conclude that when the default rate marginally increases by 1 percentage point, passive savers increase by 0.06%: $\epsilon_h = 0.0006 = \frac{26}{13,389}\frac{1}{8-5}$. The second row of Panel B indicates that among all workers who actively elected 8%, after their elected rate increased to 9%, 14 of them accepted the 9% new default rate. This suggests that passive savers increase by 0.03% for each one percentage point increase in the default rate. Taking the average of all the estimates of ϵ_h using various samples in 2019 and 2020, I conclude the average value of ϵ_h is 0.0004.

Table 5: Choice Sets to Identify Time Preferences from Survey Responses

Start date t (year)	Delay length k (year)	Total # of tokens	Token unit value sooner time a_t	Token unit value later time a_{t+k}	Annual interest rate $(1+r)$
0	1	100	100	100	1
0	1	100	99	100	1.01
0	1	100	98	100	1.02
0	1	100	95	100	1.05
0	2	100	100	100	1
0	2	100	99	100	1.01
0	2	100	98	100	1.02
0	2	100	95	100	1.05
1	1	100	100	100	1
1	1	100	99	100	1.01
1	1	100	98	100	1.02
1	1	100	95	100	1.05
1	2	100	100	100	1
1	2	100	99	100	1.01
1	2	100	98	100	1.02
1	2	100	95	100	1.05

Notes: This table shows variations in starting times t, delay length k, and interest rates $(1+r)$ to identify the key parameters from survey responses (see text). These include the normative time preference δ, the behavioral time preference β, and the utility function curvature. The survey was conducted in June 2019 to participants and opted-out workers ever had access to OregonSaves. Survey questions are provided in Appendix C. Parameters of interest are identified using regression models specified in Equations (10) and (11). Estimation results are presented in Table 6.

Table 6: Parameter Estimates of Time Preferences and Utility Function Curvature

	CRRA		CARA	
	(1)	(2)	(3)	(4)
Annual discount factor δ	0.987	0.987	0.988	0.988
	(0.005)	(0.005)	(0.005)	(0.005)
Present bias parameter γ	0.995	0.994	0.994	0.993
	(0.006)	(0.006)	(0.006)	(0.006)
CRRA curvature α	0.501	0.499		
	(0.089)	(0.089)		
CARA curvature ρ			0.00001	0.00001
			(0.000)	(0.000)
ln(annual income)		-0.00007		13.529
		(0.007)		(129.541)
Age		0.002		28.364
		(0.001)		(18.826)
Number of dependents		0.022		394.440
		(0.016)		(278.043)
Background consumption W ($)	1,040	1,040	–	–
Observations	1,765	1,765	1,765	1,765
N. unique subjects	143	143	143	143
Log likelihood	-70.093	-57.766	-15,075.172	-15,062.091
R^2	0.310	0.432	0.002	0.003

Notes: Data from anonymized survey responses collected in June 2019. An online experimental survey was sent to 440 OregonSaves-eligible workers, including those who opted out and participating as of June 2019. There are 143 survey respondents who answered the time preference survey questions provided in Appendix C, and these respondents made 1,765 intertemporal decisions in total. All columns present estimation results from two-limit Tobit maximum likelihood regressions. Columns 1-2 show estimates of the regression specification in the form of Eq.(10) assuming constant relative risk aversion utility (CRRA). The dependent variable indicates the percentage change of consumption at a sooner time relative to a later time scaled by the background consumption: $\ln\left(\frac{C+W}{R+W}\right)$. The annual background consumption ($W = 1,040$) was set to equal to the negative of the minimum consumption level among all survey respondents. The average annual normative discount factor δ under CRRA is 0.987, and the average annual behavioral discount factor γ under CRRA is 0.995. Columns 3-4 present estimates of the regression specification in the form of Eq.(12) assuming constant absolute risk aversion utility (CARA). The dependent variable is the consumption level at a sooner time: C. The average δ under CARA is 0.988 and the average γ is 0.994. Columns 2 and 4 control for covariates as in Table 2 of Meier and Sprenger (2015). Individual-level data of covariates are collected from survey responses. Standard deviations clustered on the individual level in parentheses. Additional estimation results are presented in Appendix E.

Table 7: Social Marginal Welfare Weight g Calculations

Statistics $s = \{l, h\}$	l-type passive savers	h-type passive savers
Approximated percent of type h_s	45%	3%
Average annual income Z_s	\$49,168	\$33,240
Primitive Pareto weight $\alpha_s = \frac{1}{Z_s}$	0.00002	0.00003
Aggregate weighted Pareto weight $\bar{\alpha} = \sum_{s=\{l,h\}} \alpha_s h_s$	0.00001	0.00001
Social marginal welfare weight $g_s = \frac{\alpha_s}{\bar{\alpha}}$	2.02	2.99

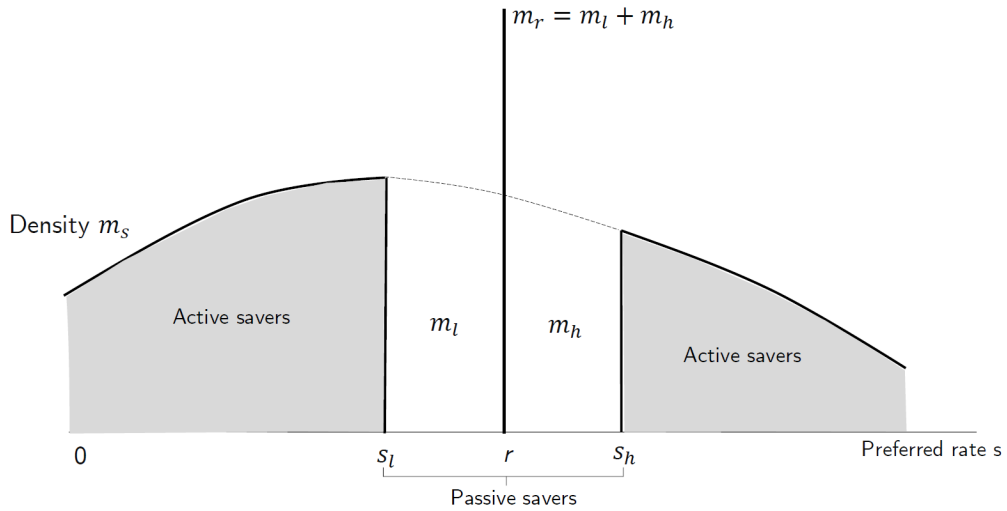
Notes: This table reports estimates of the social marginal welfare weights for l-type passive savers (preferred rates below the default while saving at the default) and for h-type passive savers (preferred rates above the default while saving at the default). The welfare weight for a given type $g_{s=\{l,h\}}$ is the Pareto weight α_s normalized by the aggregate weighted Pareto weight $\bar{\alpha}$. The normalization ensures that the welfare weights g_s only depend on the relative difference in income across types but are independent of the absolute size of income within type. We use observed data when the default rate is 5% to estimate the welfare weights at the optimal default. The first row presents the approximated (unobserved) percent of passive savers for each type. I use the observed fractions of active savers of each type to approximate the fraction of passive savers of each type. This approximation is based on an assumption that although the fraction of each type (active and passive savers combined) could be different, the ratio of active versus passive savers within type is the same. The second row presents the estimated income for each type Z_s using the OregonSaves savings data in December 2019, which was the last month that all individuals faced 5% default rate. The individual-level monthly income equals the contribution amount divided by the contribution rate. Only individuals with a positive contribution amount and a positive rate are included. Imputed average annual income equals the average monthly income times 12. The third row shows that the primitive Pareto weight α_s equals the inverse of income $\frac{1}{Z_s}$ same as in Saez (2002). The fourth row shows that the aggregate weighted Pareto weight is the primitive Pareto weight α_s weighted by the percent of each type h_s. The last row presents the final social marginal welfare weight g_s for each type.

Table 8: Baseline Optimal Default Savings Rate Calculations

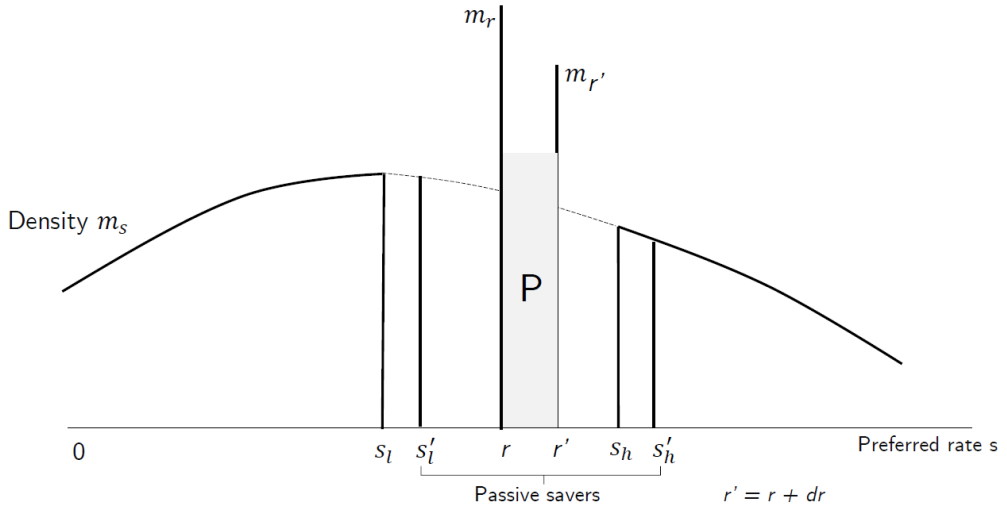
Statistics	Values
Panel A: Statistics for l-type passive savers	
Semi-elasticity ϵ_l	-0.52
Normative time preference δ_l	0.987
Behavioral time preference γ_l	0.994
Annual income Z_l	\$49,168
Social marginal welfare weight g_l	2.02
Preferred rate of passive savers on the margin s_l	0.04
Panel B: Statistics for h-type passive savers	
Semi-elasticity ϵ_h	0.0004
Normative time preference δ_h	0.987
Behavioral time preference γ_h	0.994
Annual income Z_h	\$33,240
Social marginal welfare weight g_h	2.99
Preferred rate of passive savers on the margin s_h	0.09
Panel C: Opt-out costs	
Money-metric cost of opting out of the default rate K	\$250
Fraction of normative opt-out cost π	0
Panel D: Optimal default rate	
Baseline optimal default rate r^*	7%

Notes: Estimates of key statistics used to compute the optimal default savings rate in Proposition 1: All statistics in Panel A and Panel B are estimated from the OregonSaves data (see text). Estimates for δ_l, γ_1, δ_h, and γ_h are identified using survey data collected from OregonSaves-eligible workers in Table 6 (see Section 4.2). Estimation procedures for g_l and g_h are are provided in Table 7. In Panel C, the value of K borrows from Choukhmane (2018). Calculation details for the baseline optimal default rate in Panel D are provided in Section 5.

Figure 1: Impact of a Marginal Perturbation of the Default Savings Rate r

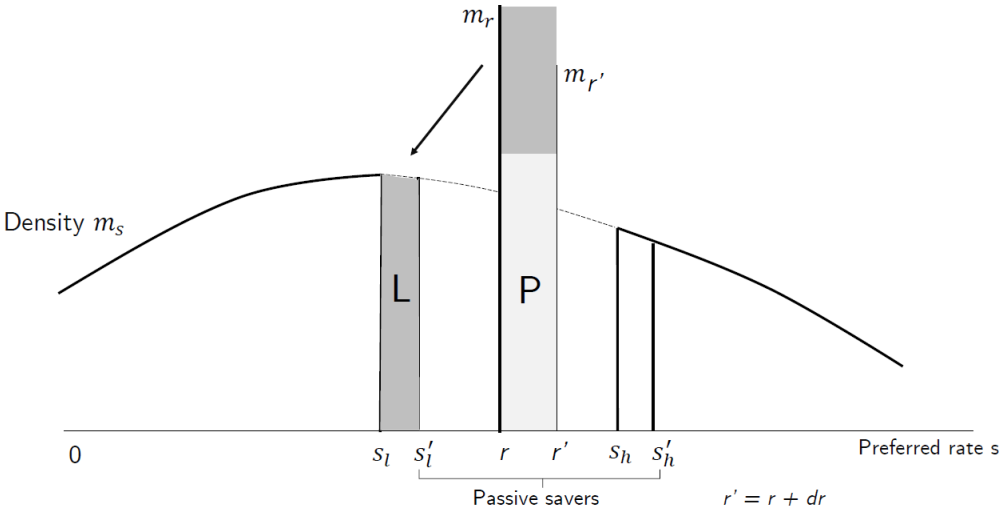


(a) In this figure, r is the default savings rate. Workers with a preferred savings rate between s_l and s_h save at the default rate r because the default is close to their preferred rates. These workers are defined as passive savers, with density $m_r = m_l + m_h$, where m_l are the fraction of passive savers with an underlying preferred rate between s_l and r, and m_h are the fraction of passive savers with an underlying preferred rate between r and s_h. Workers with a preferred rate below s_l or above s_h actively opt out of the default rate. They are defined as active savers.

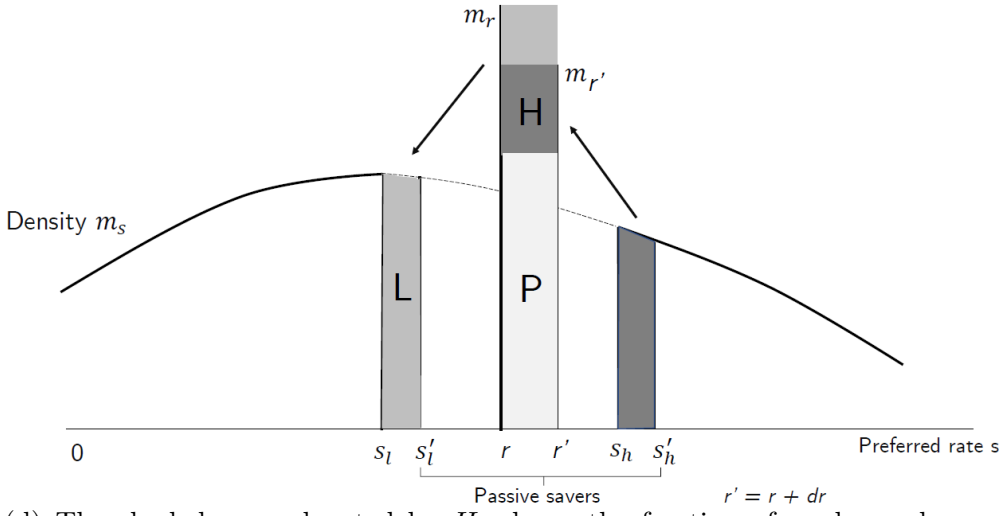


(b) Suppose the policymaker sets a default rate at r' instead of r, where $r' = r + dr$. When the default rate is r', passive savers are workers with a preferred savings rate between $s_l^{'}$ and $s_h^{'}$. Their preferred rates are close to the new default rate r'. The total density of passive savers is denoted by $m_r^{'}$. Active savers are workers with a preferred rate below $s_l^{'}$ or above $s_h^{'}$. The shaded rectangle, denoted by P, shows the fraction of workers who are passive savers both under r and r'. These are workers with a preferred rate between $s_l^{'}$ and s_h.

Figure 1: Continued

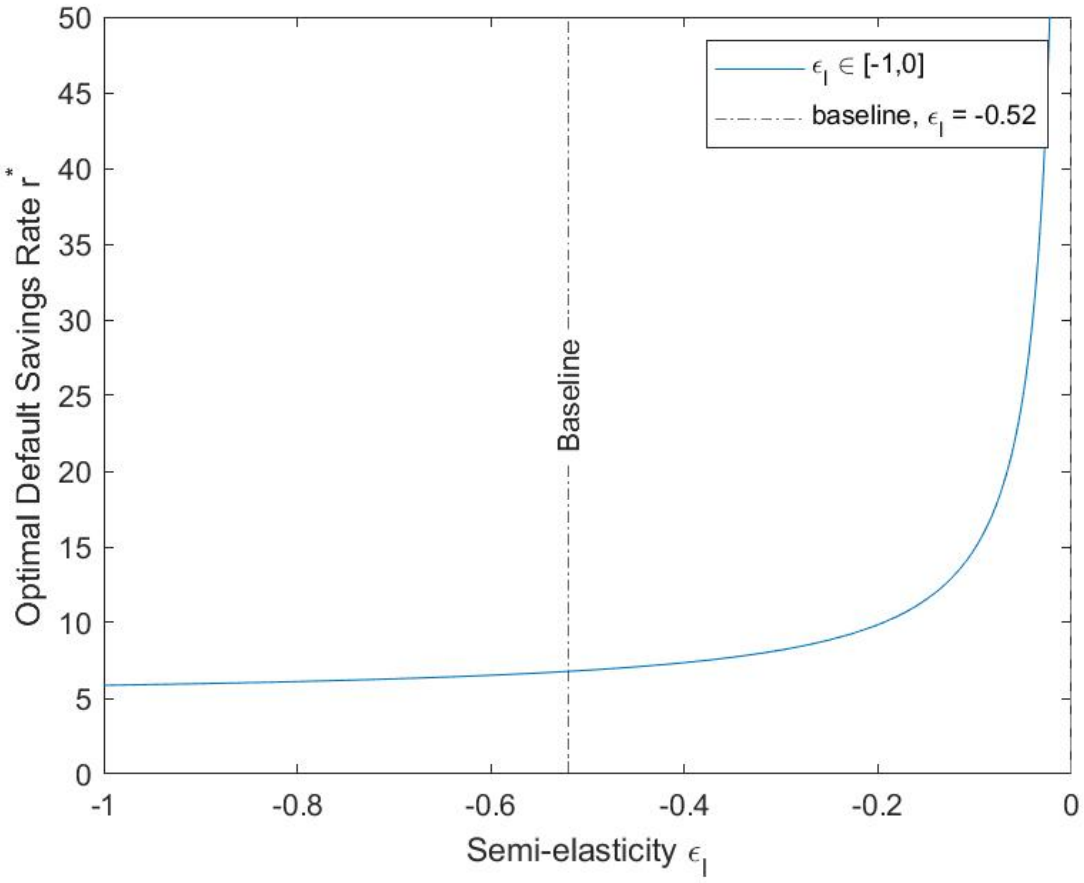


(c) The shaded area, denoted by L, indicates the fraction of workers who are passive savers under the low default rate r but active savers under the high default rate r'. These are workers with a preferred rate between s_l and s_l'. They are passive savers under the low default r because it is close enough to their preferred rates. They opt out of the high default r' and become active savers because the high default r' is far from their preferred rate.



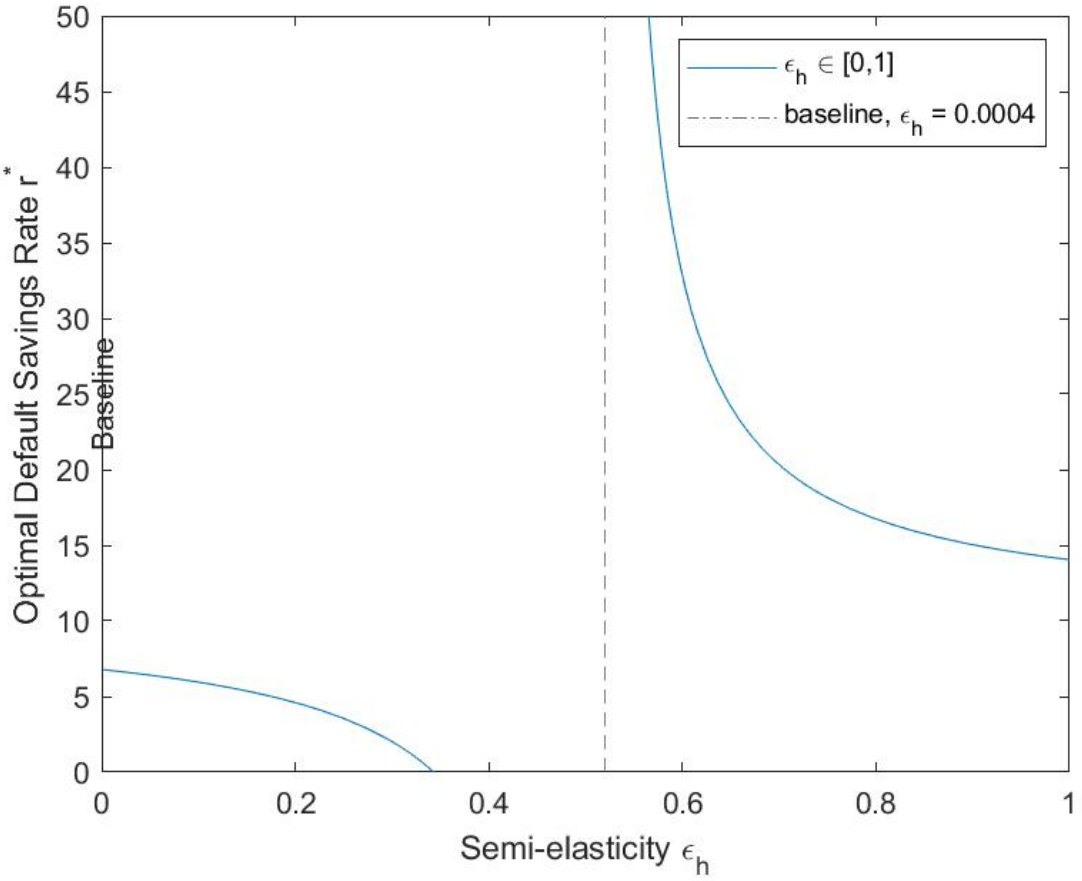
(d) The shaded area, denoted by H, shows the fraction of workers who are active savers under the low default rate r but passive savers under the high default rate r'. These are workers with a preferred rate between s_h and s_h'. They are active savers under the low default r because the low default is far from their preferred rates. They are passive savers under the high default r' because the high default is close enough to their preferred rates.

Figure 2: Optimal Default Savings Rate r^* under Different Values of Semi-Elasticity ϵ_l



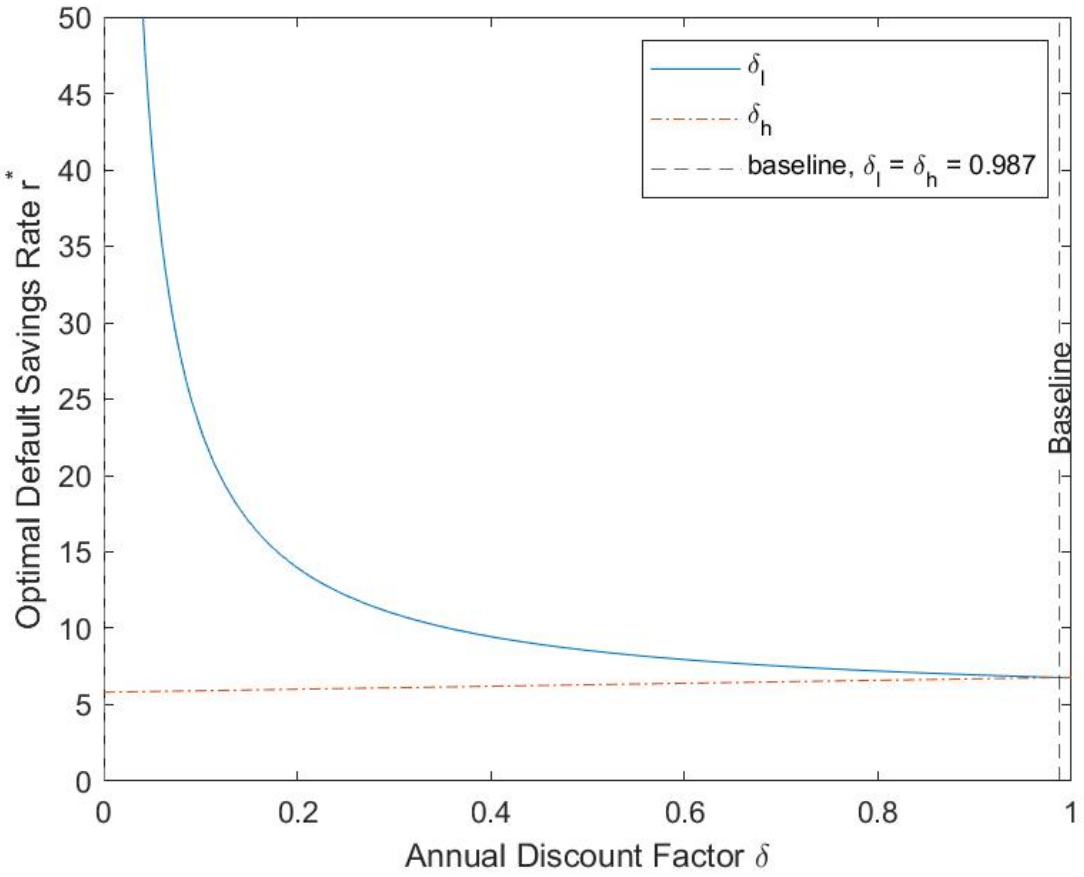
Notes: the solid curve presents how the optimal default r^* varies with l-type semi-elasticity ϵ_l, where ϵ_l measures the *decrease* in the fraction of passive savers whose preferred rates are below the default (l-type passive savers) as the default marginally increases. The intersection of the solid curve and the vertical dash-dotted line is r^* at the baseline value of ϵ_l. The optimal default r^* exponentially increases with ϵ_l. When $\epsilon_l = -1$, all passive savers with preferred rates below the default would opt out of the default as it marginally increases. The policymaker should set the default around 6% to encourage some individuals to stay at the default. When $\epsilon_l = 0$, no passive savers would ever opt out regardless of the level of the default rate. The policymaker should set the default very high to increase savings which could maximize lifecycle utility subject to the budget constraint.

Figure 3: Optimal Default Savings Rate r^* under Different Values of Semi-Elasticity ϵ_h



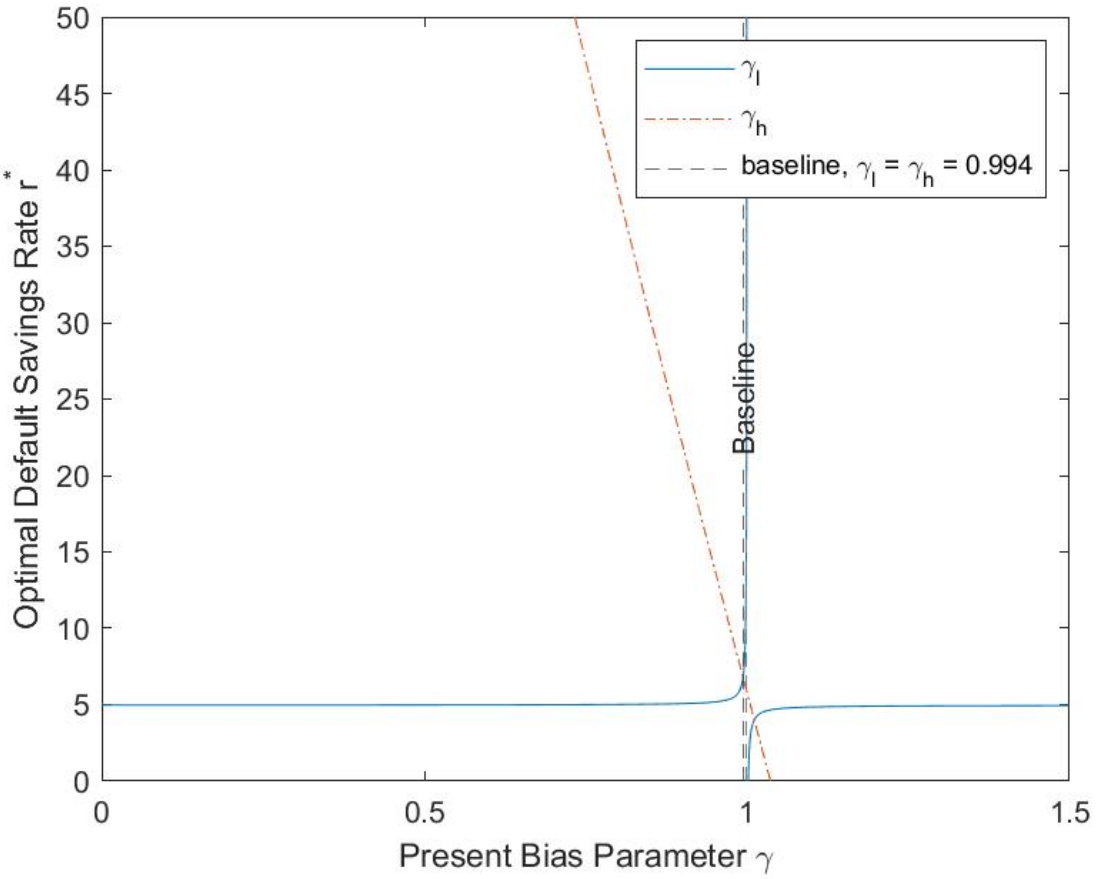
Notes: The solid curves present how r^* varies with h-type semi-elasticity ϵ_h, where ϵ_h measures the *increase* in the fraction of passive savers whose preferred rates are above the default (h-type passive savers) as the default marginally increases. The intersection of the solid curve and the vertical dash-dotted line is $r*$ at the baseline level of ϵ_h. When ϵ_h is between 0 and 0.5, r^* is bounded below 7% and decreases with ϵ_h. When ϵ_h is between 0.5 and 1, r^* is bounded above 14% and decreases with ϵ_h. When the magnitude of ϵ_h is close to the magnitude of the baseline value of $\epsilon_l(= -0.52)$, the size of the increase in l-type passive savers equals the size of the decrease in h-type passive savers. The changes in the default rate induce little welfare effect as the total number of passive savers remains constant. The optimal default r^* could be either 0 or very large. If $\epsilon_h = 0$, no h-type active savers under a low default would want to become a passive saver under a high default even though the high default is close to their preferred rate. The policymaker should set the default around 7% to maximize lifecyle utility for existing passive savers. If $\epsilon_l = 1$, all h-type active savers under a low default would save at a high default. The policymaker should set the default relatively high around 14% to minimize the aggregate opt-out cost of h-type workers.

Figure 4: Optimal Default Savings Rate r^* under Different Values of Annual Discount Factors δ_l and δ_h



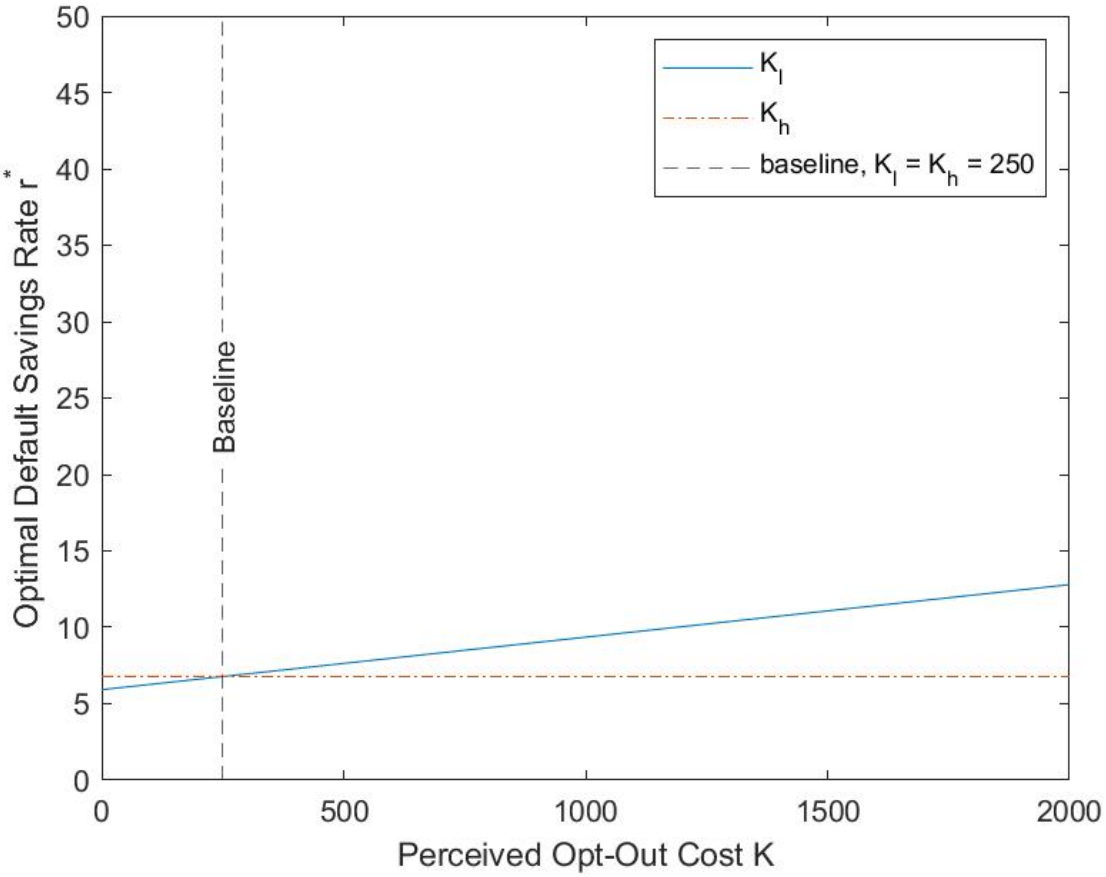
Notes: The solid curve displays how r^* varies with δ_l, which is the average annual discount factor for savers with a preferred rate lower than the default. The dash-dotted line presents how r^* varies with δ_h, which is the average annual discount factor for savers with a preferred rate higher than the default. The dashed vertical line highlights r^* under the baseline values of δ_l and δ_h. The optimal default r^* exponentially decreases with δ_l. This suggests that when δ_l is low, the (positive) welfare internality of saving at the default for savers with a low preferred rate is small, measured by dR_l in Proposition 1. The policymaker should set the default very high to encourage people with a low preferred rate to opt out of the default. When δ_l is high, the welfare internality of saving at the default is large. The policymaker should set the default around 7% to attract savers to stay at the default. Additionally, r^* slightly increases with δ_h. When δ_h is low, the (positive) welfare internality of saving at the default for savers with a high preferred rate is low, measured by dR_h in Proposition 1. The policymaker should set the default relatively low to encourage these savers to opt out of the default. Since a low default deviates from these savers' preferences, they do not have strong incentives to stay at the default. Moreover, since savers with a high preferred rate are not very responsive to the default implied by the small magnitude of ϵ_h, the impact of δ_h on the optimal default r^* is weak.

Figure 5: Optimal Default Savings Rate r^* under Different Values of Present Bias Parameters γ_l and γ_h



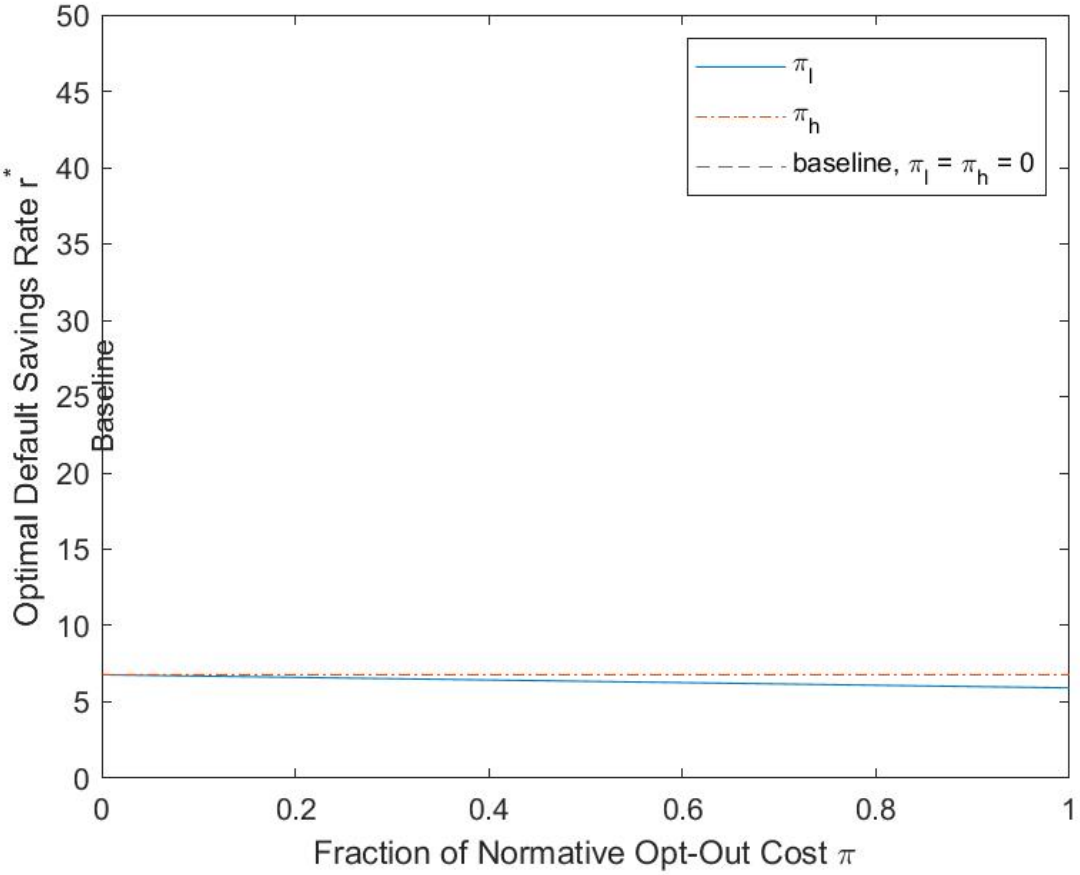
Notes: The solid curves present how r^* varies with γ_l, which is the average present bias parameter for savers with a preferred rate lower than the default (l-type savers). The dash-dotted line presents how r^* varies with γ_h, which is the average present bias parameter for savers with a preferred rate higher than the default (h-type savers). The dashed vertical line highlights r^* under the baseline values of γ_l and γ_h. When γ_l is close to 1, l-type savers are neither present-biased nor future-biased. The policymaker should set the default either to a very high level or to zero so that l-type savers are compelled to select their preferred rates. When l-type savers are either present-biased ($\gamma_l < 1$) or future-biased ($\gamma_1 > 1$), the optimal default is bounded between 5% and 8% to correct individuals from either actively under-saving or actively over-saving. The dash-dotted line shows that r^* decreases with γ_h, which is the present bias parameter for h-type savers. If h-type savers are present-biased ($\gamma_h < 1$), the policymaker should set a high default that is above their preferred rate to encourage them to save more. Using a high default to encourage more default savings could be more effective to h-type savers than l-type savers because h-type savers are less likely to opt out of the default than l-type savers. If h-type savers are future-biased ($\gamma_h > 1$), r^* is below 8% to attract these savers to save less than their preferred level by staying at the default.

Figure 6: Optimal Default Savings Rate r^* under Different Values of the Total Perceived Opt-Out Costs K_l and K_h



Notes: In this figure, I show r^* under a relaxed assumption that the total opt-out cost K perceived by individuals is heterogeneous across different types of savers. The solid line presents how r^* varies with K_l, which is the average perceived opt-out cost for savers with a preferred rate lower than the default. The dash-dotted line presents how r^* varies with K_h, which is the average perceived opt-out cost for savers with a preferred rate higher than the default. The dashed vertical line highlights r^* under the baseline values of K_l and K_h, where $K_l = K_h = K = \$250$. Both K_l and K_h vary between 0 to $\$2,000$, which is the maximum estimate in previous studies (DellaVigna, 2009; Bernheim et al., 2015). The optimal default r^* increases with K_l and is bounded under 13%. When the total opt-out cost perceived by savers with a low preferred rate (K_l) is high, the policymaker should set a relatively high default that significantly deviates from savers' preferences. Consequently, the benefit of opting out of the default outweighs the perceived cost. More l-type savers would opt out of the default despite the high perceived opt-out cost. Additionally, r^* does not change with K_h. This is mainly because savers with a high preferred rate are unlikely to adjust their savings rate based on the level of the default rate, measured by the low semi-elasticity ϵ_h.

Figure 7: Optimal Default Savings Rate r^* under Different Values of the Normative Opt-Out Costs π_l and π_h



Notes: In this figure, I present r^* under another relaxed assumption that the fraction of the normative opt-out cost π (with respect to the total perceived opt-out cost) is heterogeneous across different types of savers. The solid line presents how r^* varies with π_l, which is the average fraction of the normative opt-out cost for savers with a preferred rate lower than the default. The dash-dotted line presents how r^* varies with π_h, which is the average fraction of the normative opt-out cost for savers with a preferred rate higher than the default. The dashed vertical line highlights r^* under the baseline values of π_l and π_h, where $\pi_l = \pi_h = \pi = 0$. Both lines show that r^* are bounded between 5% and 7% for any value of π_l and π_h between 0 and 1. This suggests that the extent to which individuals overestimate the opt-out cost does not significantly impact the level of the optimal default. This is partly explained by the relatively small magnitude of the welfare internalities of action dK_l and dK_h relative to other components in Proposition 1. The magnitude of each component in the baseline calculation is displayed in Equation 15.

Appendix

A Microfoundation of the Preferred Saving Rate θ

The preferred saving rate, denoted θ, is the non-default rate that individuals would choose if they opt out of the default savings rate. The preferred saving rate θ is chosen from a policy space $\tilde{\theta} \in [0, 1]$, where θ maximizes the following utility function:

$$\theta = \arg\max_{\tilde{\theta}} U(\tilde{\theta}) = \arg\max_{\tilde{\theta}} u((1 - \tilde{\theta}) \cdot Z) + \gamma\delta v(\tilde{\theta} \cdot Z) - K, \tag{16}$$

where Z is the labor income, γ is the behavioral time preference, δ is the normative time preference, and K is the perceived cost of making an active decision. See Section 3 for detailed descriptions of these variables. Workers who choose θ as their preferred saving rate are defined as type-θ workers.

For a given type-θ worker, her preferred consumption amount $C(\theta) = (1 - \theta)Z(\theta)$, and preferred savings amount $S(\theta) = \theta Z(\theta)$. When there is a default savings rate r, she decides her pension saving amount $P(\theta)$ between two options: the default savings amount $R(\theta) = r \cdot Z(\theta)$ and her preferred saving amount $S(\theta)$. Her observed choice of saving amount in the presence of a default rate r maximizes Equation (1) in Section 3.

B Proof of Proposition 1

The first-order condition for the social welfare function, Equation (3), equals zero at the optimal default rate r^*:

$$\begin{aligned}\frac{dW(r^*)}{dr} &= \frac{d}{dr}\int_{\theta=\theta_l}^{\theta_h}\alpha(\theta)N(P(\theta))dm(\theta) \\ &\approx \int_{\theta=\theta_l}^{\theta_h}\alpha(\theta)\frac{dN(R^*(\theta))}{dr}dm(\theta) + \frac{dm_h}{dr}\alpha_h(N(R_h)-N(S_h)) - \frac{dm_l}{dr}\alpha_l(N(S_l)-N(R_l)) \qquad (17)\\ &= 0.\end{aligned}$$

The first term in Equation (17) can be decomposed into two terms:

$$\begin{aligned}&\int_{\theta=\theta_l}^{\theta_h}\alpha(\theta)\frac{dN(R^*(\theta))}{dr}dm(\theta) \\ \approx &\int_{\theta_l}^{r^*}\alpha_l\frac{dN(R_l^*)}{dr}dm_l + \int_{r^*}^{\theta_h}\alpha_h\frac{dN(R_h^*)}{dr}dm_h \\ &= \alpha_l\frac{dN(R_l^*)}{dr}m_l + \alpha_h\frac{dN(R_h^*)}{dr}m_h \\ &= \alpha_l\frac{dN}{dR_l^*}\frac{dR_l^*}{dr}m_l + \alpha_h\frac{dN}{dR_h^*}\frac{dR_h^*}{dr}m_h \\ &= \alpha_l\frac{dN}{dR_l^*}Z_l m_l + \alpha_h\frac{dN}{dR_h^*}Z_h m_h, \qquad (18)\end{aligned}$$

where $R_l^* = r^* \cdot Z_l$ so that $\frac{dR_l^*}{dr} = Z_l$. Based on Equation (2) that $N = U + (1-\gamma_l)\delta_l v(R_l^*)$, the partial derivative $\frac{dN}{dR_l^*}$ can be rewritten as:

$$\begin{aligned}\frac{dN}{dR_l^*} =&\frac{d}{dR_l^*}(U + (1-\gamma_l)\delta_l v(R_l^*)) \\ =&(1-\gamma_l)\delta_l v'_{R_l^*} \\ =&(1-\gamma_l)\delta_l\frac{g_l\lambda}{\alpha_l}, \qquad (19)\end{aligned}$$

where $g_l := \frac{\alpha_l v'_{R_l^*}}{\lambda}$ as defined in Equation (14) and $v'_{R_l^*} := \frac{dv(R_l^*)}{dR_l^*}$. Similarly, $\frac{dN}{dR_h^*} = (1 - \gamma_h)\delta_h \frac{g_h \lambda}{\alpha_h}$. Combining Equations (18) and (19), we rewrite the first term in Equation (17) as:

$$\begin{aligned}
&\int_{\theta=\theta_l}^{\theta_h} \alpha(\theta) \frac{dN(R^*(\theta))}{dr} dm(\theta) \\
\approx &\alpha_l \frac{\partial N}{\partial R_l^*} Z_l m_l + \alpha_h \frac{\partial N}{\partial R_h^*} Z_h m_h \\
= &(1 - \gamma_l)\delta_l g_l \lambda Z_l m_l + (1 - \gamma_h)\delta_h g_h \lambda Z_h m_h. \qquad (20)
\end{aligned}$$

Based on Equation (2) that $N = U + (1 - \gamma(\theta))\delta(\theta)v(P(\theta)) + (1 - \pi)K\mathbf{1}\{P(\theta) = S(\theta)\}$, where $P(\theta) \in \{R(\theta), S(\theta)\}$ and $\theta \in \{h, l\}$, the second term in Equation (17) can be rewritten as:

$$\begin{aligned}
&\frac{dm_h}{dr} \alpha_h (N(R_h) - N(S_h)) \\
= &\frac{dm_h}{dr} \alpha_h \Big(U(R_h) + (1 - \gamma_h)\delta_h v(R_h) - U(S_h) - (1 - \gamma_h)\delta_h v(S_h) - (1 - \pi)K \Big) \\
= &\frac{dm_h}{dr} \alpha_h \Big((1 - \gamma_h)\delta_h (v(R_h) - v(S_h)) - (1 - \pi)K \Big). \qquad (21)
\end{aligned}$$

Workers on the margin of switching to their preferred saving amount $S_h (= \theta_h Z_h)$ are indifferent between saving at the default and their preference when they evaluate these two choices using the decision utility U. The same argument was applied in the optimal taxation literature in Saez (2002). Therefore, $U(R_h) = U(S_h)$. Based on Assumption (4) that $v(R_h) = R_h$ and the definition of g_h in Equation (14), $\alpha_h = \frac{g_h \lambda}{v'_{R_h}} = g_h \lambda$. Equation (21) can

be expressed as:

$$\begin{aligned}
&\frac{dm_h}{dr}\alpha_h(N(R_h)-N(S_h))\\
=&\frac{dm_h}{dr}\alpha_h\Big((1-\gamma_h)\delta_h(v(R_h)-v(S_h))-(1-\pi)K\Big)\\
=&\frac{dm_h}{dr}\alpha_h\Big((1-\gamma_h)\delta_h(R_h-S_h)-(1-\pi)K\Big)\\
=&\frac{dm_h}{dr}g_h\lambda\Big((1-\gamma_h)\delta_h(r^*-\theta_h)Z_h-(1-\pi)K\Big).
\end{aligned}\tag{22}$$

Similarly, the third term in Equation (17) can be expressed as:

$$\begin{aligned}
&\frac{dm_l}{dr}\alpha_l(N(S_l)-N(R_l))\\
=&\frac{dm_l}{dr}\alpha_l\Big(U(S_l)+(1-\gamma_l)\delta_l S_l+(1-\pi)K-U(R_l)-(1-\gamma_l)\delta_l R_l\Big)\\
=&\frac{dm_l}{dr}g_l\lambda\Big((1-\gamma_l)\delta_l(\theta_l-r^*)Z_l+(1-\pi)K\Big).
\end{aligned}\tag{23}$$

Combining Equations (20), (22), and (23), we get

$$\begin{aligned}
\frac{dW(r^*)}{dr}=&(1-\gamma_l)\delta_l g_l\lambda Z_l m_l+(1-\gamma_h)\delta_h g_h\lambda Z_h m_h\\
&+\frac{dm_h}{dr}g_h\lambda\Big((1-\gamma_h)\delta_h(r^*-\theta_h)Z_h-(1-\pi)K\Big)\\
&-\frac{dm_l}{dr}g_l\lambda\Big((1-\gamma_l)\delta_l(\theta_l-r^*)Z_l+(1-\pi)K\Big)\\
=&0.
\end{aligned}\tag{24}$$

We rearrange Equation (24) and plug in semi-elasticities $\epsilon_l = \frac{dm_l}{dr}\frac{1}{m_r}$ and $\epsilon_h = \frac{dm_h}{dr}\frac{1}{m_r}$:

$$
\begin{aligned}
&\frac{dW(r^*)}{dr} \\
&= (1-\gamma_l)\delta_l g_l Z_l m_l - \frac{dm_l}{dr} g_l (1-\gamma_l)\delta_l(\theta_l - r^*)Z_l - \frac{dm_l}{dr} g_l(1-\pi)K \\
&+ (1-\gamma_h)\delta_h g_h Z_h m_h + \frac{dm_h}{dr} g_h (1-\gamma_h)\delta_h(r^* - \theta_h)Z_h - \frac{dm_h}{dr} g_h(1-\pi)K \\
&= 0 \\
&(1-\gamma_l)\delta_l g_l Z_l \frac{m_l}{m_r} + (1-\gamma_h)\delta_h g_h Z_h \frac{m_h}{m_r} + |\epsilon_l| g_l (1-\gamma_l)\delta_l(\theta_l - r^*)Z_l + |\epsilon_l| g_l(1-\pi)K \\
&+ |\epsilon_h| g_h (1-\gamma_h)\delta_h(r^* - \theta_h)Z_h - |\epsilon_h| g_h(1-\pi)K \\
&= dI + dS_l - dR_l \cdot r^* + dK_l + dR_h \cdot r^* - dS_h - dK_h \\
&= 0,
\end{aligned}
$$

where $\epsilon_l = \frac{dm_l}{dr}\frac{1}{m_r} < 0$ and $\epsilon_h = \frac{dm_h}{dr}\frac{1}{m_r} > 0$ by definition. The overall welfare effect can be decomposed into several terms after the default savings rate marginally increases from r^* to $r^* + dr$:

1. The aggregate weighted welfare gain to all passive savers on the intensive margin is $dI = (1-\gamma_l)\delta_l g_l Z_l \frac{m_l}{m_r} + (1-\gamma_h)\delta_h g_h Z_h \frac{m_h}{m_r}$.

2. The welfare gain to l-type workers for switching to their preferred rate θ_l under the new default $r^* + dr$ is $dS_l = |\epsilon_l| g_l (1-\gamma_l)\delta_l \theta_l Z_l$.

3. The welfare loss to l-type workers for opting out of the default rate is $dR_l = |\epsilon_l| g_l (1-\gamma_l)\delta_l Z_l$.

4. The welfare gain to l-type workers for making an active choice is $dK_l = |\epsilon_l| g_l (1-\pi)K$.

5. The welfare gain to h-type workers for starting to save at the default rate is $dR_h = |\epsilon_h| g_h (1-\gamma_h)\delta_h Z_h$.

6. The welfare loss to h-type workers for no longer saving at their preferred rate θ_h is $dS_h = |\epsilon_h| g_h (1 - \gamma_h) \delta_h \theta_h Z_h$.

7. The welfare loss to h-type workers for no longer making an active choice is $dK_h = |\epsilon_h| g_h (1 - \pi) K$.

Rearranging the last equation, we solve for the optimal default rate r^*:

$$r^* = \frac{dI + dS_l - dS_h + dK_l - dK_h}{dR_l - dR_h}.$$

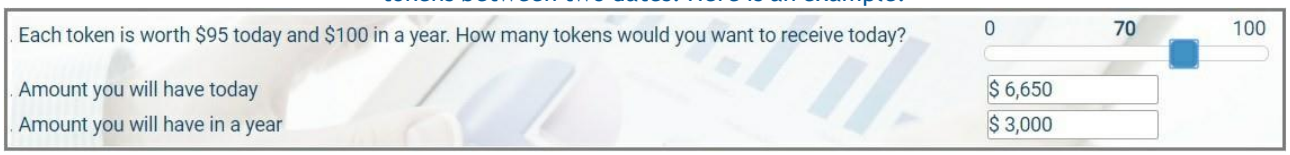
C Survey Questions to Elicit Time Preferences

I designed two waves of survey to collect workers' demographic information and time preferences. The surveys were sent out by the Oregon State Treasury. The first wave is a baseline survey that was sent out in June 2018 to collect demographic information. There are 440 workers (including enrollees and non-enrollees) who completed the baseline survey. These workers had access to OregonSaves in the first few months since its launch. The second wave is a follow-up survey that was sent out in June 2019 to collect time preferences. The 440 workers who completed the baseline survey were reached out to answer the follow-up survey. Among these workers, 143 of them completed the follow-up survey. The time preference questions were displayed in this section.

Survey design and the baseline results are explained in Section 4.2, Table 5, and Table 6. Appendix Section E presents supplementary results. Appendix Section D provides evidence that survey respondents truthfully answered survey questions and they are representative of all workers had access to OregonSaves.

OregonSaves Follow-Up Survey

Instructions: The following questions are all hypothetical, and your answers will not affect the amount of the gift card you will receive by completing the survey. In each of the following questions, please tell us how you think about tradeoffs between today and the future, by moving the slider. We ask you in each case to click the slider dividing 100 tokens between two dates. Here is an example:



This example shows how someone could divide 100 tokens between 70 today and 30 for a year from today. Each token today is worth $95, while each token for a year from today is worth $100. So this person would choose to receive 70*$95=$6,650 today and 30*$100=$3,000 a year from today.

Please use the slider to select the number of tokens you would like to receive today.

1. Each token is worth $100 today and $100 in a year. How many tokens would you want to receive today? 0 0 100
1. Amount you will have today
1. Amount you will have in a year

2. Each token is worth $99 today and $100 in a year. How many tokens would you want to receive today? 0 0 100
2. Amount you will have today
2. Amount you will have in a year

3. Each token is worth $98 today and $100 in a year. How many tokens would you want to receive today? 0 0 100
3. Amount you will have today
3. Amount you will have in a year

4. Each token is worth $95 today and $100 in a year. How many tokens would you want to receive today? 0 0 100
4. Amount you will have today
4. Amount you will have in a year

Survey navigation:
Next will advance you to the following question. After the last question, be sure to select Submit to complete the survey.

OregonSaves is overseen by the Oregon Retirement Savings Board. Ascensus College Savings Recordkeeping Services, LLC ("ACRS") is the program administrator. ACRS and its affiliates are responsible for day-to-day program operations. Participants saving through OregonSaves beneficially own and have control over their Roth IRAs, as provided in the program offering set out at saver.oregonsaves.com.

OregonSaves' Portfolios offer investment options selected by the Oregon Retirement Savings Board. For more information on OregonSaves' Portfolios go to saver.oregonsaves.com. Account balances in OregonSaves will vary with market conditions and are not guaranteed or insured by the Oregon Retirement Savings Board, the State of Oregon, the Federal Deposit Insurance Corporation (FDIC) or any other organization.

OregonSaves is a completely voluntary retirement program. Saving through a Roth IRA will not be appropriate for all individuals. Employer facilitation of OregonSaves should not be considered an endorsement or recommendation by your employer of OregonSaves, Roth IRAs, or these investments. Roth IRAs are not exclusive to OregonSaves and can be obtained outside of the program and contributed to outside of payroll deduction. Contributing to an OregonSaves Roth IRA through payroll deduction offers some tax benefits and consequences. You should consult your tax or financial advisor if you have questions related to taxes or investments.

The OregonSaves mark and OregonSaves logo are trademarks of the Oregon Retirement Savings Board and may not be used without permission.

Completed:

OregonSaves Follow-Up Survey

Please use the slider to select the number of tokens you would like to receive today.

1. Each token is worth $100 today and $100 in two years. How many tokens would you want to receive today? 0 0 100

1. Amount you will have today

1. Amount you will have in two years

2. Each token is worth $99 today and $100 in two years. How many tokens would you want to receive today? 0 0 100

2. Amount you will have today

2. Amount you will have in two years

3. Each token is worth $98 today and $100 in two years. How many tokens would you want to receive today? 0 0 100

3. Amount you will have today

3. Amount you will have in two years

4. Each token is worth $95 today and $100 in two years. How many tokens would you want to receive today? 0 0 100

4. Amount you will have today

4. Amount you will have in two years

Survey navigation:
Next will advance you to the following question. After the last question, be sure to select Submit to complete the survey.

OregonSaves is overseen by the Oregon Retirement Savings Board. Ascensus College Savings Recordkeeping Services, LLC ("ACRS") is the program administrator. ACRS and its affiliates are responsible for day-to-day program operations. Participants saving through OregonSaves beneficially own and have control over their Roth IRAs, as provided in the program offering set out at saver.oregonsaves.com.

OregonSaves' Portfolios offer investment options selected by the Oregon Retirement Savings Board. For more information on OregonSaves' Portfolios go to saver.oregonsaves.com. Account balances in OregonSaves will vary with market conditions and are not guaranteed or insured by the Oregon Retirement Savings Board, the State of Oregon, the Federal Deposit Insurance Corporation (FDIC) or any other organization.

OregonSaves is a completely voluntary retirement program. Saving through a Roth IRA will not be appropriate for all individuals. Employer facilitation of OregonSaves should not be considered an endorsement or recommendation by your employer of OregonSaves, Roth IRAs, or these investments. Roth IRAs are not exclusive to OregonSaves and can be obtained outside of the program and contributed to outside of payroll deduction. Contributing to an OregonSaves Roth IRA through payroll deduction offers some tax benefits and consequences. You should consult your tax or financial advisor if you have questions related to taxes or investments.

The OregonSaves mark and OregonSaves logo are trademarks of the Oregon Retirement Savings Board and may not be used without permission.

Completed:

OregonSaves Follow-Up Survey

Please use the slider to select the number of tokens you would like to receive in a year.

1. Each token is worth $100 in a year and $100 in two years. How many tokens would you want to receive today? 0 0 100

1. Amount you will have in a year

1. Amount you will have in two years

2. Each token is worth $99 in a year and $100 in two years. How many tokens would you want to receive today? 0 0 100

2. Amount you will have in a year

2. Amount you will have in two years

3. Each token is worth $98 in a year and $100 in two years. How many tokens would you want to receive today? 0 0 100

3. Amount you will have in a year

3. Amount you will have in two years

4. Each token is worth $95 in a year and $100 in two years. How many tokens would you want to receive today? 0 0 100

4. Amount you will have in a year

4. Amount you will have in two years

Survey navigation:
Next will advance you to the following question. After the last question, be sure to select Submit to complete the survey.

OregonSaves is overseen by the Oregon Retirement Savings Board. Ascensus College Savings Recordkeeping Services, LLC ("ACRS") is the program administrator. ACRS and its affiliates are responsible for day-to-day program operations. Participants saving through OregonSaves beneficially own and have control over their Roth IRAs, as provided in the program offering set out at saver.oregonsaves.com.

OregonSaves' Portfolios offer investment options selected by the Oregon Retirement Savings Board. For more information on OregonSaves' Portfolios go to saver.oregonsaves.com. Account balances in OregonSaves will vary with market conditions and are not guaranteed or insured by the Oregon Retirement Savings Board, the State of Oregon, the Federal Deposit Insurance Corporation (FDIC) or any other organization.

OregonSaves is a completely voluntary retirement program. Saving through a Roth IRA will not be appropriate for all individuals. Employer facilitation of OregonSaves should not be considered an endorsement or recommendation by your employer of OregonSaves, Roth IRAs, or these investments. Roth IRAs are not exclusive to OregonSaves and can be obtained outside of the program and contributed to outside of payroll deduction. Contributing to an OregonSaves Roth IRA through payroll deduction offers some tax benefits and consequences. You should consult your tax or financial advisor if you have questions related to taxes or investments.

The OregonSaves mark and OregonSaves logo are trademarks of the Oregon Retirement Savings Board and may not be used without permission.

Completed:

OregonSaves Follow-Up Survey

Please use the slider to select the number of tokens you would like to receive in a year.

1. Each token is worth $100 in a year and $100 in three years. How many tokens would you want to receive today? 0 (0 – 100)

1. Amount you will have in a year

1. Amount you will have in three years

2. Each token is worth $99 in a year and $100 in three years. How many tokens would you want to receive today? 0 (0 – 100)

2. Amount you will have in a year

2. Amount you will have in three years

3. Each token is worth $98 in a year and $100 in three years. How many tokens would you want to receive today? 0 (0 – 100)

3. Amount you will have in a year

3. Amount you will have in three years

4. Each token is worth $95 in a year and $100 in three years. How many tokens would you want to receive today? 0 (0 – 100)

4. Amount you will have in a year

4. Amount you will have in three years

Survey navigation:
Next will advance you to the following question. After the last question, be sure to select Submit to complete the survey.

OregonSaves is overseen by the Oregon Retirement Savings Board. Ascensus College Savings Recordkeeping Services, LLC ("ACRS") is the program administrator. ACRS and its affiliates are responsible for day-to-day program operations. Participants saving through OregonSaves beneficially own and have control over their Roth IRAs, as provided in the program offering set out at saver.oregonsaves.com.

OregonSaves' Portfolios offer investment options selected by the Oregon Retirement Savings Board. For more information on OregonSaves' Portfolios go to saver.oregonsaves.com. Account balances in OregonSaves will vary with market conditions and are not guaranteed or insured by the Oregon Retirement Savings Board, the State of Oregon, the Federal Deposit Insurance Corporation (FDIC) or any other organization.

OregonSaves is a completely voluntary retirement program. Saving through a Roth IRA will not be appropriate for all individuals. Employer facilitation of OregonSaves should not be considered an endorsement or recommendation by your employer of OregonSaves, Roth IRAs, or these investments. Roth IRAs are not exclusive to OregonSaves and can be obtained outside of the program and contributed to outside of payroll deduction. Contributing to an OregonSaves Roth IRA through payroll deduction offers some tax benefits and consequences. You should consult your tax or financial advisor if you have questions related to taxes or investments.

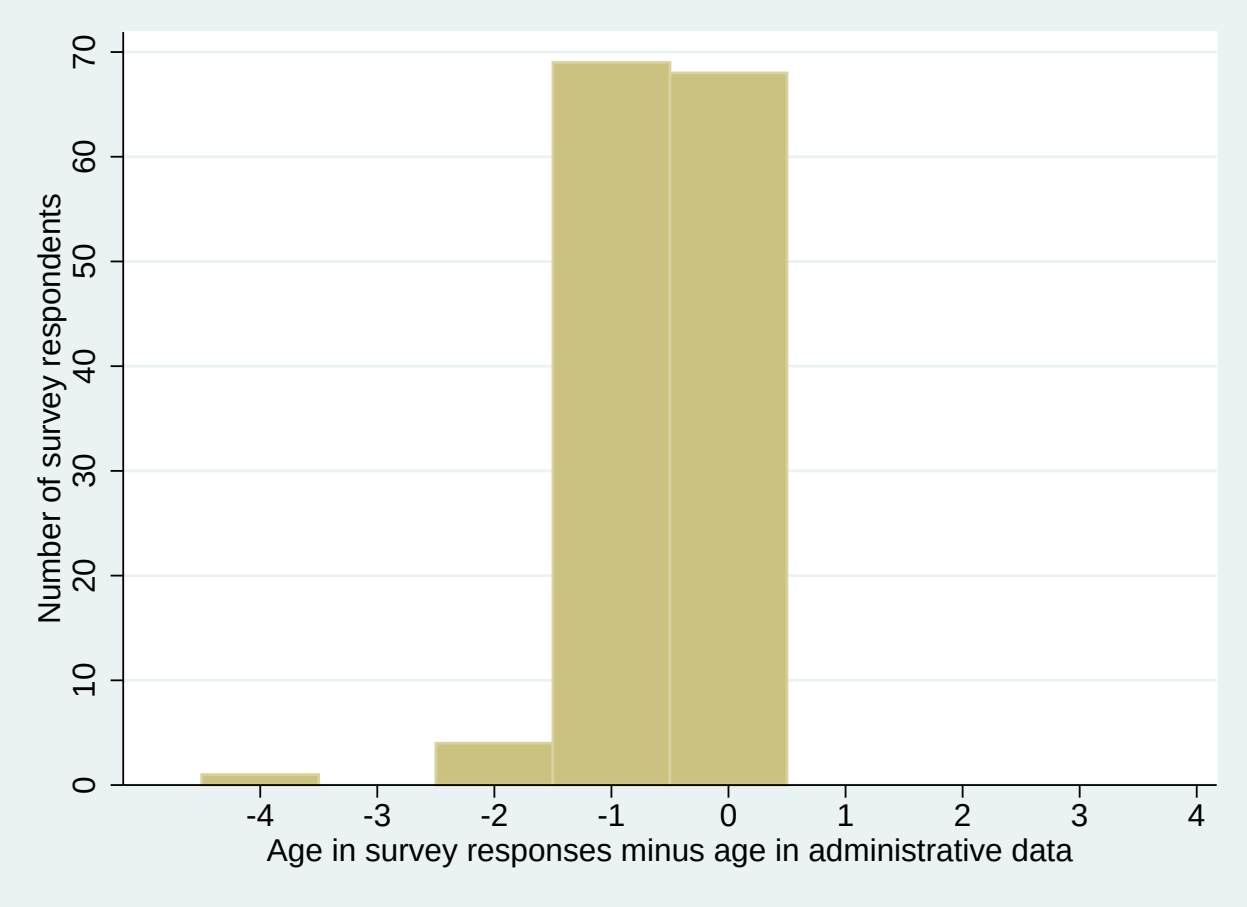
The OregonSaves mark and OregonSaves logo are trademarks of the Oregon Retirement Savings Board and may not be used without permission.

Completed:

D Supplementary Figures of Survey Respondents

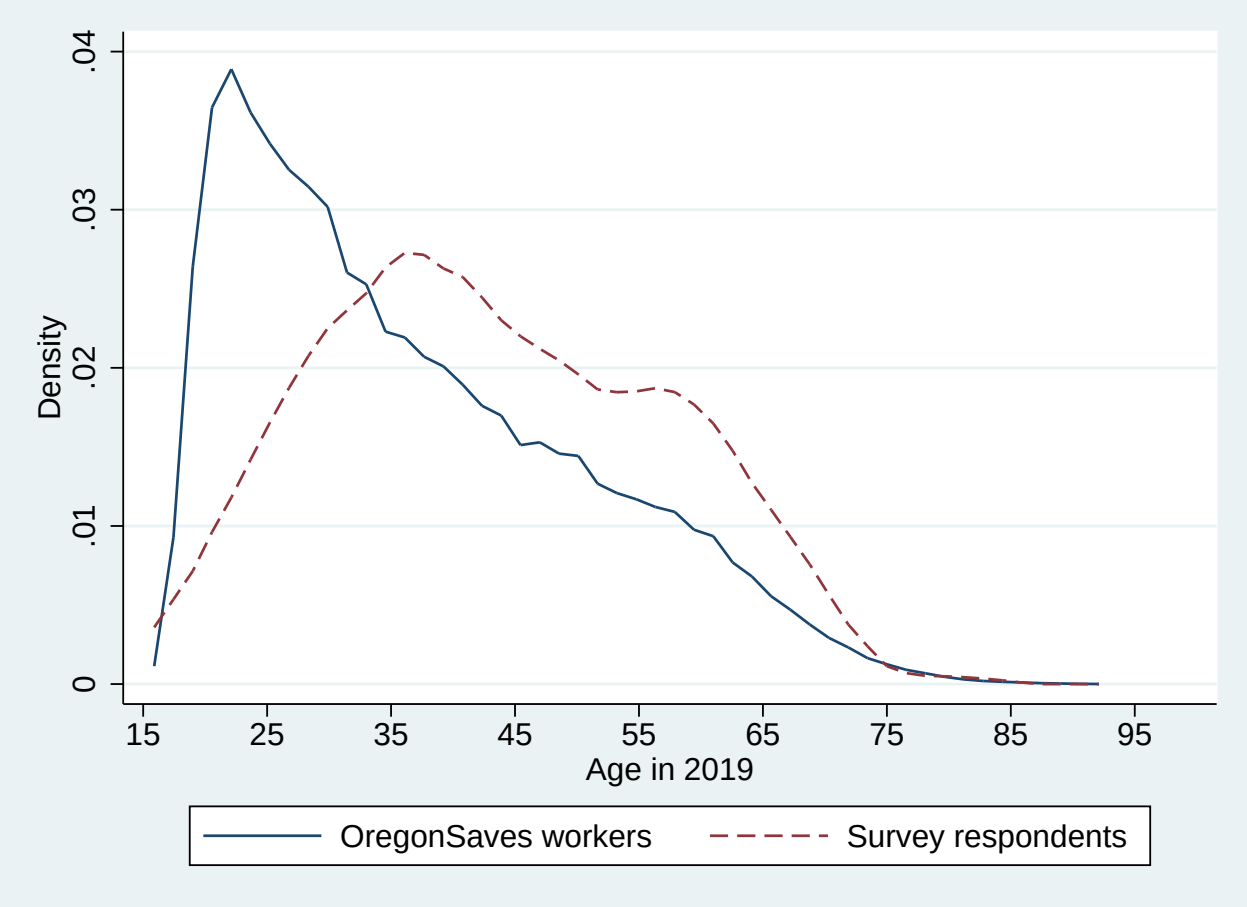
Figure 8 provides evidence that survey respondents truthfully answered survey questions. Figures 9 and 10 show that survey respondents are representative of all workers eligible for OregonSaves.

Figure 8: Difference in Self-Reported Age from Survey and from Administrative Data, June 2019



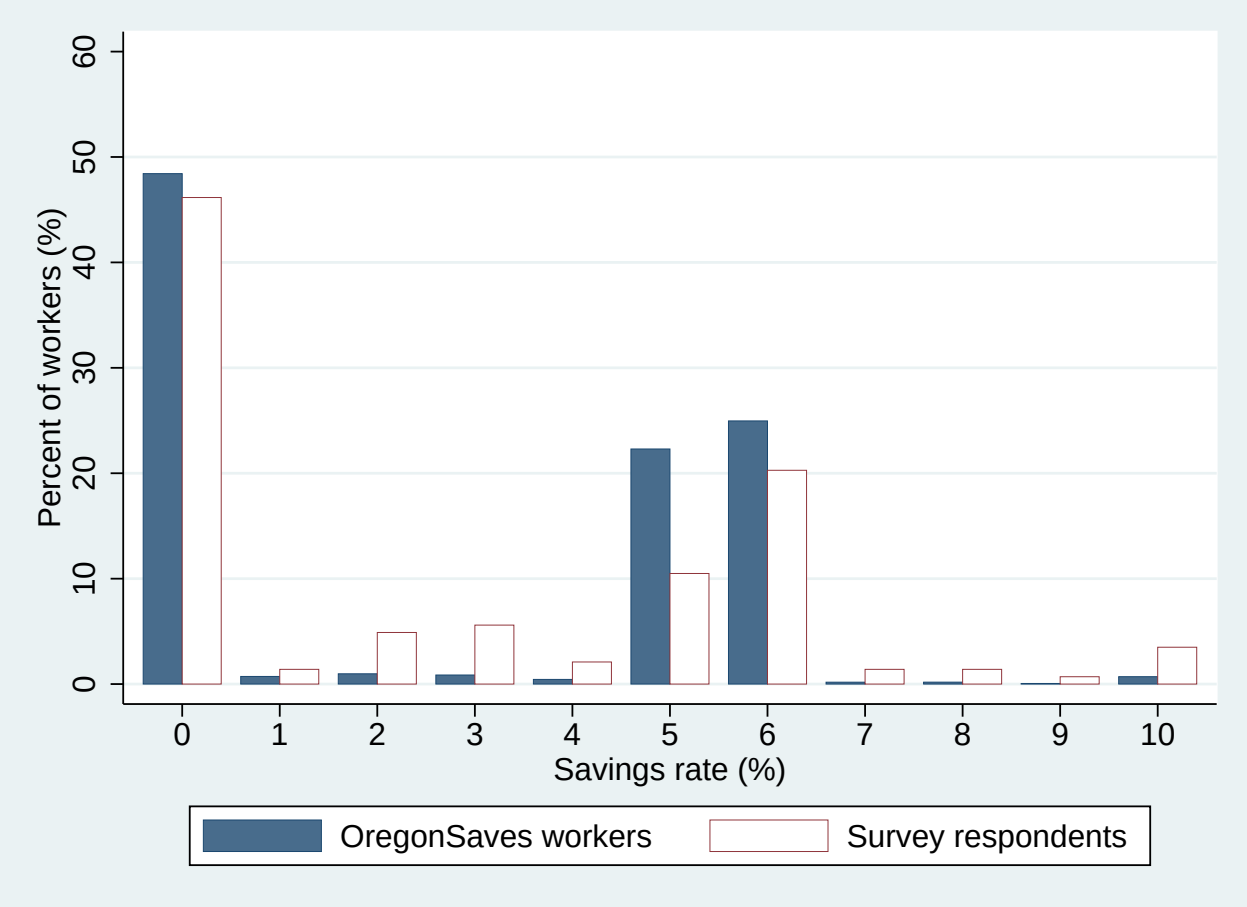
Notes: This figure presents the difference in self-reported age from survey responses and from administrative data for 142 survey respondents who answered the time preference questions in June 2019. By comparing their age in survey responses with the age in the administrative data, this figure suggests that most survey respondents truthfully answered survey questions. One survey respondent was excluded as the age from survey response significantly deviates from the age in administrative records (survey age = 55, administrative data age = 29). All time preference questions are presented in Appendix C. Estimates of time preferences are presented in Table 6. Self-reported age in survey data was collected from survey responses to the following question:"How old were you on your last birthday?" Age from the administrative data is also self-reported. When workers open an OregonSaves account for the first time, they are asked to report their birth year so that part of their contributions would be invested into an age-appropriate fund. Age in June 2019 from the administrative data was calculated from the birth year. Without knowing the exact birth date, the estimated age from the administrative data could be one-year different from the age in survey. There are 137 survey respondents whose age in survey responses is within one-year difference from the estimated age using the administrative record. Five people reported their age in the survey that is more than one year younger than the estimated age from the administrative record.

Figure 9: Comparison of the Distribution of Age between the OregonSaves Population and Survey Respondents, June 2019



Notes: The blue solid curve shows the distribution of age for all workers ever had access to OregonSaves (including enrollees and non-enrollees) by the end of June 2019 (N = 186,888). The red dashed curve is the distribution of age for all survey respondents (including enrollees and non-enrollees) (N = 143). These two curves are close to each other although survey respondents are slightly older than the entire OregonSaves population. The average age of all OregonSaves workers is 37 in 2019. The average age of survey respondents is 43 in 2019. The survey to collect demographic information was fielded in June 2018. These survey respondents are the first wave of workers who had access to OregonSaves before June 2018. Most workers were automatically enrolled in OregonSaves in 2019 and early 2020.

Figure 10: Comparison of the Distribution of the Savings Rate between the OregonSaves Population and Survey Respondents, June 2019



Notes: The blue solid bar displays the discrete distribution of the savings rate in June 2019 for workers had access to OregonSaves before the end of June 2018 (N = 27,417). The blank bar presents the distribution of the savings rate in June 2019 for survey respondents. All survey respondents had access to OregonSaves before the end of June 2018. This figure shows the comparison of these two groups because they had access to OregonSaves during the same period of time and they both experienced automatic escalation on January 1, 2019. Survey respondents made similar savings choices as all workers who had access to OregonSaves during the same period. When survey respondents were asked time preference questions in June 2019, their answers should be representative of the larger OregonSaves population.

E Supplementary Results of Time Preferences

This section shows estimates of workers' time preferences and risk preferences in addition to the baseline estimates presented in Table 6.

Table 9: Additional Parameter Estimates of Time Preferences with the CRRA Utility Function

Alternative background consumption W	$W = 0.01$			$W = 2,572$		
	(1)	(2)	(3)	(4)	(5)	(6)
	Andreoni and Sprenger (2012)	OregonSaves data		Andreoni and Sprenger (2012)	OregonSaves data	
Annual discount factor δ	0.324	0.991	0.991	0.275	0.987	0.987
	(0.173)	(0.004)	(0.004)	(0.162)	(0.005)	(0.005)
Present bias parameter γ	1.023	1.003	1.002	1.026	0.994	0.993
	(0.010)	(0.005)	(0.005)	(0.010)	(0.006)	(0.007)
CRRA curvature α	0.977	1.077	1.077	0.873	0.198	0.195
	(0.004)	(0.012)	(0.012)	(0.018)	(0.145)	(0.146)
ln(annual income)			0.053			0.0002
			(0.046)			(0.005)
Age			-0.002			0.001
			(0.006)			(0.0006)
Number of dependents			-0.107			0.015
			(0.112)			(0.010)
Observations	4,365	1,765	1,765	4,365	1,765	1,765
N. unique subjects	97	143	143	97	143	143
Log likelihood	–	-2,870.98	-2,858.72	–	616.07	628.89
R^2	–	0.01	0.02	–	-0.05	-0.07

Note: This table provides supplementary estimates of the regression specification in the form of Eq.(10) assuming constant relative risk aversion utility (CRRA) in addition to the baseline results in Table 6. All columns present estimation results from two-limit Tobit maximum likelihood regressions. The dependent variable indicates the percentage change of consumption at a sooner time relative to a later time scaled by the background consumption: $\ln\left(\frac{C-W}{R-W}\right)$. Columns 1-3 show estimates using the choice of the background consumption in Andreoni and Sprenger (2012): $W = 0.01$. Column 1 presents estimates in Column 4 of Table 2 in Andreoni and Sprenger (2012). These are the estimates using the same CRRA regression specification as Columns 2-3 and the baseline results. Columns 4-6 use the average annual background consumption in Andreoni and Sprenger (2012): $W = 2,571$. Column 4 presents estimates in Column 6 of Table 2 in Andreoni and Sprenger (2012). Columns 3 and 6 control for covariates as in Table 2 of Meier and Sprenger (2015). Individual-level data of covariates are collected from survey responses. Standard deviations clustered on the individual level in parentheses.

Table 10: Additional Parameter Estimates of Time Preferences with the CARA Utility Function

	CARA Eq.(11)			CARA Eq.(12)		
	(1)	(2)	(3)	(4)	(5)	(6)
	Andreoni and Sprenger (2012)	OregonSaves data		Andreoni and Sprenger(2012)	OregonSaves data	
Annual discount factor δ	0.254	0.987	0.987	0.335	0.988	0.988
	(0.159)	(0.005)	(0.005)	(0.136)	(0.005)	(0.005)
Present bias parameter γ	1.028	0.993	0.992	1.017	0.994	0.993
	(0.010)	(0.007)	(0.007)	(0.008)	(0.006)	(0.006)
CARA curvature ρ	0.008	0.00001	0.00002	0.007	0.00001	0.00001
	(0.001)	(0.000)	(0.000)	(0.001)	(0.000)	(0.000)
ln(annual income)			27.362			13.529
			(261.975)			(129.541)
Age			57.100			28.364
			(38.015)			(18.826)
Number of dependents			798.738			394.440
			(561.633)			(278.043)
Observations	4,365	1,765	1,765	4,365	1,765	1,765
N. unique subjects	97	143	143	97	143	143
Log likelihood	–	-16,152.86	-16,139.76	–	-15,075.17	-15,062.09
R^2	–	0.002	0.003	–	0.002	0.003

Note: This table provides supplementary estimates of the regression specification assuming constant absolute risk aversion utility (CARA) in addition to baseline results in Table 6. All columns present estimation results from two-limit Tobit maximum likelihood regressions. Columns 1-3 present results of the CARA tangency condition in the form of Eq.(11) where the dependent variable indicates the difference in consumption level at a sooner time relative to a later time: $C-R$. Column 1 presents estimates in Column 7 of Table 2 in Andreoni and Sprenger (2012). Columns 4-6 show results of the CARA solution function in the form of Eq.(12) where the dependent variable is the consumption level at the sooner time: C. Column 4 presents estimates in Column 8 of Table 2 in Andreoni and Sprenger (2012). Columns 3 and 6 control for covariates as in Table 2 of Meier and Sprenger (2015). Individual-level data of covariates are collected from survey responses. Standard deviations clustered on the individual level in parentheses.

Table 11: Additional Parameter Estimates of Time Preferences with Linear Utility

	Meier and Sprenger (2015)		OregonSaves data		
	(1)	(2)	(3)	(4)	(5)
Annual discount factor δ	0.760	0.690	0.892	0.892	0.894
	(0.0002)	(0.047)	(0.048)	(0.048)	(0.048)
Present bias parameter γ	0.712	0.973	1.040	1.026	1.029
	(0.024)	(0.012)	(0.065)	(0.065)	(0.065)
ln(annual income)				0.052	
				(0.046)	
Annual income/10K		0.000			0.004
		(0.003)			(0.001)
Age		-0.001		-0.002	-0.003
		(0.000)		(0.006)	(0.006)
Number of dependents		0.006		-0.108	-0.111
		(0.005)		(0.113)	(0.113)
Observations	31,812	31,812	1,765	1,765	1,765
N. unique subjects	1,446	1,446	143	143	143
Log likelihood	-20,490.403	-20,374.52	-2,903.365	-2,891.584	-2,892.424
R^2	–	–	0.0007	0.005	0.004

Note: This table provides supplementary estimates assuming linear utility. All columns present estimation results from two-limit Tobit maximum likelihood regressions. Columns 1-2 present my calculation using the estimates presented in Columns 1-2 of Table 2 in Meier and Sprenger (2015). The dependent variable is the log of the ratio of consumption at a sooner time relative to a later time: $\ln(C/R)$. Columns 3-5 show results using the same linear utility specification using the OregonSaves survey data. The estimates of δ in two studies are close. The estimates of γ is higher using the OregonSaves data than in Meier and Sprenger (2015). Individual-level data of covariates are collected from survey responses. Standard deviations clustered on the individual level in parentheses.