THE IMPACT OF ECONOMIC POLICIES ON HOUSEHOLD FINANCIAL AND LABOR SUPPLY BEHAVIOR

Mingli Zhong

A DISSERTATION

in

Applied Economics

For the Graduate Group in Managerial Science and Applied Economics

Presented to the Faculties of the University of Pennsylvania

in

Partial Fulfillment of the Requirements for the

Degree of Doctor of Philosophy

2020

Supervisor of Dissertation

Olivia S. Mitchell, International Foundation of Employee Benefit Plans Professor, Business Economics and Public Policy

Graduate Group Chairperson

Nancy R. Zhang, Ge Li and Ning Zhao Professor, Professor of Statistics

Dissertation Committee
Benjamin Lockwood, Assistant Professor of Business Economics and Public Policy
Hanming Fang, Class of 1965 Term Professor of Economics

THE IMPACT OF ECONOMIC POLICIES ON HOUSEHOLD FINANCIAL AND LABOR SUPPLY BEHAVIOR

Dedicated to

my parents, Shuzhen Zhu and Yongsheng Zhong, who always believe in me,

my husband, Ting Qian, who stopped my daughter from editing my dissertation,

and my daughter, Madeline Qian, who only tried to edit my dissertation once.

ACKNOWLEDGEMENT

I would like to thank my advisors Olivia S. Mitchell, and Benjamin Lockwood, and Hanming Fang for their expertise, assistance, and patience throughout the process of writing this dissertation. I thank Olivia for her continuous support from the beginning of my graduate study. There is no way I could even get my degree without her guidance and encouragement. I want to thank Ben for spending lots of time patiently answering (and re-answering) many of my questions and writing emails with answers much longer than my questions. I also want to thank Hanming for telling me, at the very beginning of my research life, that writing a paper is about telling a story. That has inspired me a lot. I am very thankful to my department chair Joseph Harrington and graduate coordinators Fernando Ferreira and Gilles Duranton for their support and encouragement during my entire graduate study.

I am extremely grateful to Alex Rees-Jones, Mike Abito, Katja Seim, Lin Shen, Junyuan Zou, Betty Wang, Amanda Chuan, Matt Davis, Andrew Johnston, and Lindsay Relihan for their extensive discussions that help me complete the dissertation. I thank Sendhil Mullainathan, James Poterba, Jonathan Skinner, James Choi, Corinne Low, Judd Kessler, Juuso Toikka, and many seminar and conference participants for their feedback that greatly helped me improve my research. I thank my coauthors John Chalmers, Olivia S. Mitchell, and Jonathan Reuter for their invaluable help and guidance. I am grateful to Jennie Huang, Ben Keys, Ben Hyman, Clayton Featherstone, Corinne Low, Judd Kessler, Eduardo Azevedo, Juuso Toikka, Shing-Yi Wang, Kent Smetters, Maisy Wong, Todd Sinai, and Jessie Handbury for their help to prepare me for the job market season.

The OregonSaves project (Chapters 1 and 2) 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). 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; the Quartet program at the University of Pennsylvania; and the TIAA Institute. 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 findings and conclusions are solely those of my own and do not represent the views of SSA, any agency of the federal government, the MRRC, OregonSaves, or any other institutions with which I am affiliated.

For the unemployment insurance project (Chapter 3), I am thankful to Kory Kroft and Matthew Notowidigdo for sharing their data and code with me. I acknowledge research funding from the Pension Research Council/Boettner Center for Pensions and Retirement Research at the Wharton School of the University of Pennsylvania. Opinions and views are solely my responsibility.

Last but not least, I thank many friends and family for their continuing support. I am extremely grateful to my parents and parents-in-law for their sacrifices, unlimited patience and love. Thanks to my daughter for inspiring me to be a stronger and better person. Thanks to my dog for being a low-maintenance dog. Thanks to my husband for all the great things that are beyond words.

ABSTRACT

THE IMPACT OF ECONOMIC POLICIES ON HOUSEHOLD FINANCIAL AND LABOR SUPPLY BEHAVIOR

Mingli Zhong

Olivia S. Mitchell

Economic policies significantly influence household financial and labor supply decisions. In return, household responses to the policy change inform policy design. In this dissertation, I document how individuals respond to economic policies by adjusting their critical economic decisions: how much to save and how much to work. I also combine the welfare analysis of optimal economic policies that maximize individual welfare with empirical evidence on household and individual responses to economic policies.

In the first chapter, I document the differential impacts of unemployment insurance on unemployment duration for older and younger workers. I find that older workers tend to receive longer periods of unemployment insurance than younger workers during economic downturns. This is partly because older workers have a harder time finding a job during economic downturns than younger workers. This age-specific difference is not salient in boom periods. I conclude that, since the purpose of unemployment insurance is to provide financial support for unemployed workers until they find a job, the optimal design of unemployment insurance needs to consider the differential job prospects across age groups.

The policy implication is that it could be socially welfare improving to extend unemployment insurance for older workers so that they have a longer time period of time to find a new job. This could prevent older workers from leaving the workforce too early while they are still capable of working. Additionally, older workers might start claiming Social Security or withdrawing from their retirement accounts earlier than planned. Claiming Social Security at the earliest possible time could lead to a lower level of Social Security benefits for the rest of their life. Either withdrawing retirement savings too early or claiming a lower level of Social Security benefits potentially increases retirees' chances of late-life poverty and their reliance on means-tested social transfers.

In the second chapter based on a working paper co-authored with John Chalmers, Olivia S. Mitchell, and Jonathan Reuter, we investigate the impact of a savings mandate imposed on private sector employers on expanding access to automatic enrollment retirement plans. Starting in 2017, the state of Oregon required that all private sector employers provide either an employer-sponsored retirement plan or enroll their employees into a state-sponsored retirement plan, called OregonSaves. We find about half of the workers who were automatically enrolled in OregonSaves chose to participate in the plan. The majority of participants stayed at the 5 percent default savings rate.

In the third chapter, I investigate the impact of the default savings rate in automatic enrollment retirement plans on individual welfare. I propose a unified framework to analyze the welfare effects of the default savings rate and derive a formula for the optimal default savings rate that depends on observable statistics. Using individual-level administrative and survey data and an exogenous increase in the default savings rate from 5% to 6% in the OregonSaves program, I estimate key statistics in the optimal default formula and find that the optimal default savings rate to be 8%.

Contents

List of Tables

List of Figures

CHAPTER 1 : AGE-SPECIFIC EFFECTS OF UNEMPLOYMENT INSURANCE ON UNEMPLOYMENT DURATION OVER THE BUSINESS CYCLE

1.1. Abstract

This paper investigates whether the effect of unemployment insurance benefits on unemployment duration varies by age over the business cycle. When tested individually, the unemployment duration of younger workers are significantly raised by the same level of increase in UI benefits more in a boom than in a recession, while those of older workers are equally affected over the business cycle. This difference between age groups is *not* significant when tested as an interaction effect in a more stringent regression model. Similarly, we also find that the age effect reported in previous study is non-significant when subjected to the same procedure of regression analysis. The current findings suggest that incorporating age into the design of UI benefits should require further study and more credible evidence.

1.2. Introduction

A great deal of research has documented the social welfare gains of extending the duration of unemployment insurance (UI) during recessions (e.g. Krueger and Mueller (2010); Piketty and Saez (2013)). An important finding is that unemployment benefits raise unemployment durations to a lesser degree in a recession than in a boom. For example, Rothstein (2011) finds that the impact of UI on raising unemployment durations in the U.S. was smaller during the Great Recession than during booms. Using a more quantitative approach, Kroft and Notowidigdo (2016) show that a one standard deviation increase in the unemployment rate reduces the impact of unemployment benefits on duration by 50 percent.

This paper investigates whether the impact of UI benefits on unemployment duration varies by age over the business cycle. The question is motivated by the finding that UI benefits, on average, raise longer unemployment durations for older workers than for younger workers

(Michelacci and Ruffo, 2015). One interpretation of this finding is that younger workers want jobs not only to increase their current income but also to acquire labor market skills and so improve career prospects and lifetime income (Michelacci and Ruffo, 2015). As a result, the additional income provided by UI does not reduce younger workers' incentive to return to work as much as older workers'. It remains to be explored, however, whether the weaker effect of UI benefits on the unemployment durations of younger workers changes with the overall economy.

Using data from the Survey of Income and Program Participation (SIPP), this paper shows that the impact of UI benefits on unemployment durations is more responsive to economic recessions for younger workers than for older workers. *Responsiveness* is measured by the elasticity of unemployment durations with respect to benefits, following methods used in previous research (Moffitt, 1985; Meyer, 1990; Chetty, 2008; Kroft and Notowidigdo, 2016; Michelacci and Ruffo, 2015). *Duration elasticity* is defined as the percentage increase in unemployment durations when benefit levels increase by 1%. In particular, we show that duration elasticity for younger workers is significantly larger in a recession than in a boom. By contrast, duration elasticity for older workers remains stable over the business cycle.

We begin by providing graphical evidence comparing the job finding rates of younger and older workers at different phases of the business cycle. Following the work of Kroft and Notowidigdo (2016), the overall economic environment is measured by the state-level unemployment rate of the first month of an individual's unemployment spell. We plot Kaplan-Meier survival curves for younger (20-40 years old) and older (41-60 years old) workers respectively (cf., Michelacci and Ruffo (2015). We show that job finding rates are higher for younger workers than for older workers under weak economic conditions and are similar for the two age groups under strong economic conditions. For example, after being unemployed for 16 weeks in a recession, 60 percent of younger workers are still unemployed, compared to 68 percent of older workers who are still unemployed. After being unemployed for 37 weeks in a recession, 28 percent of younger workers are unemployed, compared to 41 percent of older workers. By contrast, job finding rates are not significantly different between the two age groups during a boom. Second, we estimate a set of Cox hazard models to evaluate how duration elasticity varies with unemployment rates for each age group. The SIPP dataset includes actual UI benefits to an individual for only a small subset of all subjects. To alleviate the problem of data sparsity, we use two *proxies* for actual UI benefit levels so that each individual's UI benefits can be estimated: the average benefits of all claimants for each state/year pair, and the age-specific state/year average benefits for each age group. The latter can be a more accurate proxy for actual benefits if there exist sizable differences in actual benefits between the two age groups. In three Cox regression models, we find that the effect of unemployment rate on duration elasticity is larger in magnitude for younger workers than for older workers. In the baseline specification, a 10 percent increase in state-level unemployment rate significantly raises the hazard rate with respect to UI benefits by 16.5 percent for younger workers, and by 4 percent for older workers. The estimates imply that a one standard deviation increase in the unemployment rate (1.7 percentage points increase to 8.4%) significantly reduces the magnitude of duration elasticity for younger workers by 71 percent. By contrast, a one standard deviation increase in the unemployment rate has no significant effect for older workers.

In the following regression analysis, we seek the underlying mechanism to explain the results. We show that increasing unemployment rates reduces the attractiveness of UI benefits to younger workers. As a result, younger workers become more motivated in their job search during economic recessions compared to boom times. By contrast, the attractiveness of UI benefits for older workers does not vary with unemployment rates. Consequently, younger workers increase their job finding rates compared to older workers in economic recessions.

Understanding the variation in behavior between how older and younger workers respond to unemployment benefits over the business cycle has potentially important implications for the aging workplace and for the improvement of public policies concerning workers of different ages. Age-specific social insurance systems can be designed with considerations of different labor supply responses by older and younger workers for social welfare improvement.

The chapter proceeds as follows. Section 3.3 reviews the relevant literature. Section 3.4 describes the data used for analyses. Section 3.5 presents the model for evaluating how duration elasticity varies with unemployment rate by age. Section 3.6 presents graphical and statistical results. Section 3.7 concludes.

1.3. Literature Review

Previous research on unemployment insurance has proposed a range of formulas for the average level of the optimal unemployment insurance benefits. Baily (1978) and Chetty (2006) derive consumption-based formulas for the average level of optimal UI benefits in terms of three parameters: elasticity of unemployment duration with respect to benefits, the drop in consumption as a function of UI benefits, and the coefficient of relative risk aversion. The UI benefit level is measured by the UI replacement rate, which refers to the ratio of UI benefits to the worker's previous wage. The optimality condition equates the benefit of consumption smoothing to the cost of behavioral distortions. The first parameter, elasticity of unemployment duration with respect to UI benefits, captures the moral hazard cost due to behavioral distortions. The moral hazard cost is defined as a distortionary cost caused by UI recipients to stay away from work longer than those who do not receive unemployment benefits, since UI recipients face lower cost of being unemployed. The second parameter, drop in consumption, measures the gains due to consumption smoothing. The third parameter, risk aversion, reflects the worker's value of having a smoother consumption path.

Several papers have sought to estimate these three parameters, to allow empirical implementation of consumption-based UI formulas. Moffitt (1985); Meyer (1990); Krueger (2002) estimate the elasticity parameter, which quantifies the effect of UI benefits on durations. The consensus estimate is 0.5, meaning that a 10% increase in UI replacement rates leads to 5% longer unemployment duration, where replacement rate is defined as the ratio of UI benefits to pre-unemployment wage. Gruber (1997) evaluates the consumption smoothing benefits of UI. He concludes that a 10% increase in the UI replacement rate causes a 2.8% lower consumption decline. Using Meyer's and his own estimates, Gruber calibrates the optimal replacement rate, which is much lower than the observed benefit level. If the relative risk aversion were 2, for instance, the optimal replacement rate would be 5%, meaning that the optimal UI benefits should be 5% of the previous wage. In fact, however, the observed average UI rate is 50% of the previous wage. One reason why the implied optimal benefit level is far different from the observed could be that it is difficult to measure risk preferences in the context of unemployment insurance. Chetty and Szeidl (2007) suggest that there is also an amplifying impact of consumption commitments on risk aversion. Consumption commitments refer to consumption goods that cannot be adjusted frequently due to fixed adjustment costs, such as housing. This differs from the standard expected utility approach which assumes people can cut back on consumption goods freely. In fact, Chetty and Szeidl's models of consumption commitments suggest that relative risk aversion might exceed 4 for unemployment shocks. Consequently, the corresponding optimal replacement rate would be much higher, around 45%.

Other studies have devised alternative ways to calculate optimal UI benefits. Shimer and Werning (2007) use the responsiveness of reservation wages to unemployment benefits to test for the optimality of UI. After-tax reservation wages suggest that the take-home pay required to make a worker indifferent between working and remaining unemployed. The optimality condition depends on the gain from benefits and the cost of tax through reservation wage. On the one hand, higher benefits reduce the cost of remaining unemployed and consequently raise the pretax reservation wage. The magnitude of increase in reservation wage is measured by the elasticity of reservation wages to benefits. On the other hand, the increase in benefit must be funded by an increase in the employment tax which therefore reduces the after-tax reservation wage. The primary advantage of the Shimer/Werning approach over the consumption-based formula is that they do not need to independently estimate risk aversion and consumption drop. These two parameters are difficult to identify since risk aversion is context-dependent and measuring consumption declines requires a long and ideally administrative consumption panel dataset. One problem with their reservation wage approach is that it is challenging to find a reliable way to measure the reservation wage.

The third approach to calculating the average level of optimal UI benefits is a revealed preference formula by Chetty (2008). Using Baily (1978)'s formula, Chetty decomposes the effect of UI benefits on duration into two parts: a welfare-enhancing liquidity effect of benefits on duration, and a welfare-reducing moral hazard effect due to distorted incentives. The ratio of the liquidity effect to the moral hazard effect is a sufficient statistic for testing the optimality of UI benefits.

All three approaches to calculating the average optimal UI benefit recognize the disincentive effects of unemployment insurance as evidence of moral hazard. The elasticity of unemployment duration with respect to benefits is commonly used to measure how UI reduces labor supply. Additional evidence of moral hazard is the spike in the exit rate of unemployment at benefit exhaustion. Moffitt (1985); Katz and Meyer (1990a,b) interpret the sharp surge in the exit rate as evidence that UI recipients wait until their benefits run out to return to work. Instead of measuring the hazard rate of exiting the UI system, Card et al. (2007) measure the hazard rate of reemployment. They do not find a sharp spike in job finding at the point of benefit exhaustion, which implies that many individuals leave the unemployment system when their benefits expire without returning to work. Therefore, the spike in the unemployment exit rate overstates the degree of moral hazard induced by UI.

Most existing literature focuses on the average UI benefit rate, though two recent studies have investigated the optimal redistribution pattern of UI benefits over either the life cycle or the business cycle. Michelacci and Ruffo (2015) propose an optimal life cycle unemployment system that increases benefits for younger workers and decreases benefits for older workers. Their argument is that younger workers have little ability to smooth consumption during unemployment and show little evidence of moral hazard. Kroft and Notowidigdo (2016) propose an optimal unemployment system varying over the business cycle. Their finding shows that the moral hazard cost is low during times of high unemployment, and high during times of low unemployment in the United States. Schmieder et al. (2012) find similar results in Germany. Consequently, optimal UI benefit levels should be positively related to unemployment rate to minimize welfare loss due to moral hazard.

Some UI research presents empirical evidence to corroborate the hypothesis that UI benefits do not raise unemployment durations during times of high unemployment. Rothstein (2011) finds little evidence of moral hazard during the Great Recession, when extended UI payment spells were introduced. UI benefit extensions raised the unemployment rate in early 2011 by only about 0.1-0.5%. The effect of benefit extensions is more significant for staying in labor market and looking for jobs, than for reducing job search incentive. Farber and Valletta (2015) report similar results during the much milder downturn in the early 2000s and the 2009-2012 periods. They find a small reduction in the unemployment exit rate resulted from extended UI benefits in both periods. In addition, the reduction is primarily due to a decrease in exits from the labor force rather than a decrease in job finding rate. Both papers suggest that UI does not significantly raise unemployment durations or reduce job search incentive.

Other UI research further investigates whether UI benefits affect labor force attachment or labor search behavior by studying the event of the expiration of UI benefit extension. Farber et al. (2015) show consistent results during the phase-out period of benefit extensions in 2013-2014: this reduced the unemployment rate mainly by moving people out of the labor market instead of increasing the job finding rate. This result is also consistent with what Card et al. (2007) show, namely that the expiration of unemployment benefits has little effect on reemployment. Hagedorn et al. (2015) investigate the impact of extension expiration on employment from a macroeconomic perspective. They find that 1% drop in benefit duration causes an increase of employment by 0.0161 log points, implying that some 1.8 million additional people took jobs in 2014 due to the UI benefit cut. In addition, the expiration of benefit extensions led to almost 1 million non-participants entering the labor market. A plausible explanation is that non-participants were encouraged by the greater probability of finding jobs, since the expiration of UI extensions sends a positive signal on job creation. The macro evidence contradicts the micro evidence in Farber et al. (2015) that the expiration did not improve job search or job finding rate. The discrepancy between micro and macro evidence of the impact of UI extension expiration on employment needs to be further investigated.

Besides analyzing the transition of labor force status on the extensive margin, several papers document job search behavior of the unemployed on the intensive margin. Krueger and Mueller (2010) find that unemployed workers allocate 41 minutes per day to job search (on weekdays). Job search is inversely related to the generosity of unemployment benefits, with an elasticity between -1.6 and -2.2. Job search intensity for UI benefits recipients increases prior to benefit exhaustion, while remaining constant for non-recipients during unemployment spells. Aguiar et al. (2013) document a larger impact on job search time: they report that unemployed workers allocate 2-6 percent of their lost work hours to job search. Given that the average lost work time is about 33 hours per week, the extra time devoted to job search is 40-118 minutes per week. Besides job search, leisure absorbs 50 percent of foregone market work hours, mainly sleeping and television watching. Home production activities absorb 30 percent, including cooking, cleaning, laundry, shopping, home maintenance and repair, and child care. Health care and education investment absorb 12 percent of foregone market work hours. Kutyavina (2014) finds that, during the Great Recession, older unemployed workers spent less time on job search than younger workers.

In summary, previous studies document that elasticity of unemployment duration with respect to UI benefits is smaller in recessions than in booms. The contribution of this paper is to further explore whether the variation in duration elasticities over the business cycle is primarily driven by younger or older workers.

1.4. Data

In keeping with prior analysis in this field, we use data from the Survey of Income and Program Participation (SIPP) spanning 1985-2000. Each SIPP panel surveys households at 4-month intervals for 2-4 years. Compared to other widely-used data sets such as the Current Population Survey (CPS) and the Panel Study on Income Dynamics (PSID), the main benefits of the SIPP are the availability of weekly data on employment status to compute individual unemployment spells and UI benefit receipt, and its large sample size.

We follow the sample selection criteria adopted by Chetty (2008). Specifically, we focus on males 20-60 years old who (1) report searching for a job, (2) are not on temporary layoff, (3) have at least three months of work history in the survey, and (4) took up UI benefits within one month after job loss. We also censor unemployment spells at 50 weeks to reduce the influence of outliers and to focus on the search behavior for both younger and older workers in the year after job loss. These restrictions leave 4,380 unemployment spells in the core analysis sample, the same sample size as in Michelacci and Ruffo (2015)'s study on optimal age-dependent unemployment insurance. Our sample size (4,380 spells) is slightly smaller than Chetty (2008)'s (4,529 spells) since his age range is from 18-65 years old. The reason for choosing relatively strict age range is to exclude the effect of Social Security on labor supply decisions for older workers. Information on UI benefit levels was obtained from the Employment and Training Administration's Significant Provisions of State Unemployment Insurance Laws (various years). All dollar value in the data are adjusted to real 1990 dollars.

Panel A of Table 1 presents descriptive statistics for the SIPP sample. We divide the SIPP sample into two subsamples according to whether individuals began their unemployment spells in states with above-median unemployment rates or below-median unemployment rates. The median unemployment rate is defined across all unemployment spells in the sample. In keeping with prior studies, we define the unemployment rate assigned to each unemployment spell is the monthly state-level unemployment rate at the starting month of each spell. Panel B of 1 1 shows summary statistics for two age groups with above-median unemployment rates, where two age groups are 20-40 and 41-60 years old following Michelacci and Ruffo (2015). Panel C of 1 shows summary statistics for the two age groups with below-median unemployment rates. Individuals with above-median unemployment rate (Panel B) have longer unemployment durations and slightly higher pre-unemployment annual wage on average than those with below-median unemployment rate (Panel C). As expected, the average unemployment rate in Panel B (8%) is much higher than that in Panel C (5%). The weekly unemployment benefits and replacement rate are close between the two subsamples. Other demographic characteristics such as age, years of education, and marital status are also similar between the two subsamples.

[Table 1 here]

1.5. Empirical Strategy

The objective of our empirical analysis is to estimate how unemployment duration elasticity with respect to UI benefits varies by age, over the business cycle. Our empirical strategy closely follows Kroft and Notowidigdo (2016), extended to investigate heterogeneity between age groups. We divide the SIPP sample into two age groups, 20-40 years old (the younger workers) and 41-60 years old (the older workers). We use state-level monthly unemployment rates to determine the state of the economy. To address the potential concern that the unemployment rate would be endogenous to UI benefits and durations, we follow Kroft and Notowidigdo (2016) using the unemployment rate in the month at the start of an unemployment spell, instead of the actual unemployment rate at each point during an unemployment spell.

We estimate unemployment duration elasticities for each age group using cross-state and time variation in unemployment benefit levels. The unemployment duration data are collected in the SIPP, which reports the weekly employment status of every individual in the sample. We compute each individual's duration of unemployment by summing the number of consecutive weeks without a job, starting from the date of job separation and ending when the individual finds a job that lasts for at least one month. It is not possible to estimate each individual's UI benefit amount precisely, so we use two proxies for the actual UI benefit level: the average benefit of all claimants for each state/year pair obtained from the Department of Labor, and the age-specific state-year average benefits for each age group. The latter can be a more accurate proxy for the actual benefits if there are sizable differences in actual benefits between two age groups.

We estimate a set of Cox hazard models to evaluate how duration elasticity varies by age, with the unemployment rate. Let $h_{i,s,t}$ denote the unemployment exit hazard rate for individual i in state s at time t of an unemployment spell, $\alpha_{s,t}$ the baseline hazard rate in state s at time t, $b_{i,s,t}$ the state-level unemployment benefits for individual i at the start of the spell, $u_{s,t}$ the state-level monthly unemployment rate for individual i at the start of the spell, and $X_{i,s,t}$ a set of control variables. We estimate the following hazard baseline model:

$$\log\, h_{i,s,t} = \alpha_{s,t} + \beta_1\, \log(b_{s,t}) + \beta_2\, \log(u_{s,t}) + \beta_3\Big(\log(b_{s,t}) *\, \log(u_{s,t})\Big) + \beta_4 \mathrm{X}_{i,s,t}. \quad (1.1)$$

All results report standard errors clustered by state. Like Kroft and Notowidigdo (2016), all independent variables are de-meaned so that the coefficient β_1 corresponds to the elasticity of unemployment durations with respect to the UI benefits at the average state unemployment rate. The coefficient β_2 represents the elasticity of unemployment durations with respect to the unemployment rate for an average individual receiving the average UI benefits. The key coefficient of interest, β_3, is the additional change in duration elasticity for a one log point change in the state unemployment rate, holding other independent variables constant.

The baseline model (1) includes the same controls as Kroft and Notowidigdo (2016): state, year, industry and occupation fixed effects; a 10-piece log-linear spline for the claimant's pre-unemployment wage; age, education; dummies for marital status and being on the seam between interviews to adjust for the seam effect; and the interaction between log of

UI benefits and year fixed effect to control for any time-varying effect of the benefits on durations. All control variables are also de-meaned.

Besides the baseline model, we estimate the following extended regression model including a three-way interaction term between unemployment rate, UI benefits, and age dummy:

$$\log h_{i,s,t} = \alpha_{s,t} + \beta_1 \log(b_{s,t}) + \beta_2 \log(u_{s,t}) + \beta_3 \mathbf{1}\{41\text{-}60\} + \beta_4 \Big(\log(b_{s,t}) * \log(u_{s,t})\Big)$$
$$+\beta_5 \Big(\mathbf{1}\{41\text{-}60\} * \log(b_{s,t})\Big) + \beta_6 \Big(\mathbf{1}\{41\text{-}60\} * \log(u_{s,t})\Big) + \beta_7 \Big(\mathbf{1}\{41\text{-}60\} * \log(b_{s,t}) * \log(u_{s,t})\Big) + \beta_8 \mathrm{X}_{i,s,t}, \tag{1.2}$$

where the indicator variable, $\mathbf{1}\{41\text{-}60\}$, equals one for older workers 41-60 years old. The key coefficient of interest, β_7, represents the difference in the effect of unemployment rate on duration elasticity between younger and older workers. The reason to propose the extended model is that it tests a more stringent hypothesis and reaches a more direct conclusion on differential age effect on duration elasticity compared to the baseline model: the baseline model tests whether unemployment rate significantly affects duration elasticity for younger and older workers respectively. The extended model tests whether the effect of unemployment rate on duration elasticity is significantly different between younger and older workers.

In addition, we also estimate the following extended regression model including a two-way interaction between unemployment rate and UI benefits:

$$\log h_{i,s,t} = \alpha_{s,t} + \beta_1 \log(b_{s,t}) + \beta_2 \mathbf{1}\{41\text{-}60\} + \beta_3 \Big(\mathbf{1}\{41\text{-}60\} * \log(b_{s,t})\Big) + \beta_4 \mathrm{X}_{i,s,t}. \tag{1.3}$$

Equation (3) is an extended regression equation of Michelacci and Ruffo (2015)'s baseline model. The conclusion from their baseline model is that duration elasticity with respect to UI benefits on average is significantly different from zero for younger workers, not for older workers. The purpose of estimating equation (3) is to further evaluate whether duration elasticity on average is significantly different between older and younger workers. The key coefficient of interest, β_3, represents the difference in the average duration elasticity between older and younger workers. Combining the results from equation (2) and (3) can identify the robustness of age-specific responses to UI benefits.

1.6. Results

1.6.1. Graphical Evidence and Nonparametric Tests

To compare job finding rates between the two defined age groups (young versus old) under high and low unemployment rates. Kaplan-Meier survival curves are plotted to illustrate the relation between the fraction of workers who remain unemployed with respect to the total number of unemployed workers at week zero (the y-axis in Figures 1a and 1b) and the number of weeks of unemployment spell (the x-axis in Figures 1a and 1b)[1].

[Figure 1 here]

By definition, the fraction of unemployed workers at week zero is one, both the younger and the older groups have the same starting point in Figures 1a and 1b. Figure 1a shows that when the unemployment rate is above the median, suggesting a recession is in effect, the fraction of unemployed workers in the younger group decreased faster than that in the older group (nonparametric Wilcoxon test $p < 0.05$)[2]. For example, in a recession, 16 weeks after job loss 60 percent of younger workers are unemployed, while 68 percent of older workers are still unemployed; In a recession, 37 weeks after job loss, 28 percent of younger workers are unemployed, versus 41 percent of older workers. This results suggest that younger workers can find jobs much more quickly than older workers in a recession given the same level of

[1]These survival curves plotted using the SIPP data are adjusted for the seam effect in panel surveys following Chetty (2008). The seam effect arises because individuals are interviewed at 4-month intervals in the SIPP and tend to repeat answers about weekly job status in the past 4 months. Consequently, we observe a disproportionately large number of transitions in labor force status on the "seam" between interviews, which leads to artificial spikes in the hazard rate at 4 and 8 months. To smooth out these spikes, we add a time-varying indicator for being on a seam between interviews as a control variable in a Cox model. The resulting seam-adjusted survival curves represents the probability of remaining unemployed after t weeks of a spell for an individual who never crosses an interview seam.

[2]Since the job finding rates are not normally distributed over time, a nonparametric test is appropriate when the parametric assumption of normally distribution is unlikely to satisfy.

UI benefits. By contrast, when the unemployment rate is below the median, the two age groups behave similarly in their tendencies to find new jobs once unemployed (see Figure 1b; Wilcoxon test $p > 0.5$).

The results presented in Figure 1a and 1b suggest that the job finding rates for younger workers are higher than those for older workers under weak economic conditions, while they are similar under strong economic conditions. It is not yet clear whether the difference is caused by younger or older workers responding to UI benefits differently over the business cycle. The next section investigates the question.

1.6.2. Hazard Model Estimation

To evaluate the robustness of the graphical results presented above, we estimate the hazard model (1.1) to identify how duration elasticity with respect to UI benefits varies with unemployment rate by age. We report three coefficients in all specifications: β_1 is the elasticity of unemployment durations with respect to the UI benefits at the average state unemployment rate; β_2 is the elasticity of unemployment durations with respect to the unemployment rate for an average individual receiving the average UI benefits; and β_3, the key coefficient of interest, is the additional change in duration elasticity for a one log point change in the state unemployment rate, holding other independent variables constant.

We first estimate the hazard model (1.1) on the full sample from 20-60 years old to replicate the results from Kroft and Notowidigdo (2016). Column 1 of Table 2 reports Kroft and Notowidigo's baseline model estimates. Column 2 reports my own estimates of equation (1.1). My findings indicate that a 10 percent increase in UI benefit rate reduces hazard rate by 15.2 percent, at average unemployment rate. In addition, a 10 percent increase in state-level unemployment rate significantly boosts the hazard rate with respect to UI benefits by 13.8 percent. Results for the key coefficient of interest, β_3, are similar with the same significance level ($p<0.01$) between the two specifications. My own estimates of β_1 and β_2 are close to but slightly larger than Kroft and Notowidigdo's, due to the fact that they had a slightly different sample size from mine (their model dropped some observations for an unknown reason). The bottom two rows of Table 2 report duration elasticity when state unemployment rate is one standard deviation (1.7%) above and below the mean unemployment rate (6.7%) for the pooled sample. A one standard deviation increase in the unemployment rate (8.4%) reduces the magnitude of duration elasticity by 48 percent (from 1.518 to 0.796). A one standard deviation decrease in the unemployment rate (5%) raises the magnitude of duration elasticity by 48 percent (from 1.518 to 2.239). The magnitudes are similar to the results in Kroft and Notowidigdo (2016). Their estimates indicate that one standard deviation change in the unemployment rate is correlated with the change in the magnitude of duration elasticity by roughly 50 percent.

[Table 2 here]

The main results are reported in Tables 3 and 4. Columns 1-2 report results for the pooled sample, Columns 3-4 for the 20-40 age group, and Columns 5-6 for the 41-60 age group. Columns 1, 3, and 5 report the baseline results of the regression equation excluding unemployment rate and the interaction term between unemployment rate and UI benefits as independent variables. The only key variable in Column 1, 3, and 5 is UI benefits. Columns 2, 4, and 6 report results of equation 1.1. Table 3 uses average UI benefits as a measure of individual UI benefit levels. Table 4 uses age-specific average UI benefits.

[Table 3 here]

[Table 4 here]

The main results from Tables 3 and 4 can be summarized as follows: although duration elasticity is larger for older than for younger workers on average, duration elasticity for younger workers is more responsive to the overall economic conditions. In other words, duration elasticity for younger workers in a recession is significantly larger than that in a boom. By contrast, duration elasticity for older workers remains stable over time. In both panels, the estimate of β_1, the effect of UI benefits on hazard rate, is smaller in magnitude for younger than for older workers. This result is consistent even when the interaction term is not included. By contrast, the estimates of the key coefficient β_3, the effect of unemployment rate on duration elasticity, is larger in magnitude for younger workers than for older workers. In Table 3, for example, the estimate of β_3 in Column 4 indicates that a 10 percent increase in state-level unemployment rate raises the hazard rate with respect to UI benefits for younger workers by 16.5 percent. The elasticities derived from the estimates in Column 4 suggest that a one standard deviation increase in the unemployment rate (8.4%) reduces the magnitude of duration elasticity for younger workers by 71 percent (from 1.225 to 0.361). A one standard deviation decrease in the unemployment rate (5%) raises the magnitude of duration elasticity for younger workers by 71 percent (from 1.225 to 2.088). The elasticities in Column 6 suggest that a one standard deviation increase in the unemployment rate (8.4%) reduces the magnitude of duration elasticity for older workers by 7 percent (from 2.923 to 2.716). Overall, the main results suggest that the negative relationship between unemployment rate and duration elasticity documented in Kroft and Notowidigdo (2016) is primarily driven by younger workers. Older workers have a persistently high unemployment duration elasticity that is insensitive to aggregate economic conditions. By contrast, younger workers have a lower duration elasticity during a weak labor market.

The results of the extended regression model in Table 5 show that the differential impacts of unemployment rate on duration elasticity between older and younger workers are not robust. Column 1 in Table 5 displays the results of the regression equation (1.2) with a three-way interaction between unemployment rate, UI benefits, and age dummy, in which the estimate of the coefficient of the three-way interaction is not significant. The insignificance of the three-way interaction coefficient suggests that although the effect of unemployment rate on duration elasticity is significantly different from zero for younger workers, the effect is not significantly different between older and younger workers. In addition, Column 2 in

Table 5 displays the results of estimating equation (1.3), which is an extended regression equation of Michelacci and Ruffo (2015)'s baseline model. The estimate of the coefficient of the two-way interaction between age dummy and UI benefits suggests that their baseline results also fail the robustness test: although the average duration elasticity is significantly different from zero for younger workers, the average duration elasticity is not significantly different between older and younger workers. Overall, the differential age effect on duration elasticity is not statistically significant.

[Table 5 here]

1.7. Conclusion

This paper has investigated how unemployment duration elasticity with respect to UI benefits varies with unemployment rates by age. When tested in each age group individually, a one standard deviation increase in unemployment rate reduces the magnitude of duration elasticity for younger workers by 71 percent. By contrast, a one standard deviation increase in unemployment rate reduces the magnitude of duration elasticity for older workers by 7 percent.

The difference in the responsiveness of duration elasticity to unemployment rates between younger and older workers is not statistically significant when tested as an interaction effect in a more stringent regression model. In addition, the difference in the average duration elasticity reported in Michelacci and Ruffo (2015) is not significant either when subjected to the same procedure of regression analysis. Since the current results do not support the hypothesis that additional UI benefits will raise unemployment durations to a less degree for younger workers than for older workers over the business cycle, we cannot conclude that redistributing UI benefits from older workers to younger workers, as which is proposed by Michelacci and Ruffo (2015), can improve social welfare. Incorporating age into the design of UI benefits should require further study and more credible evidence.

CHAPTER 2 : AUTO-ENROLLMENT RETIREMENT PLANS FOR THE PEOPLE: CHOICES AND OUTCOMES IN OREGONSAVES [1]

2.1. Abstract

Insuring retirement security is an important challenge for our aging society, and many policymakers are seeking ways to help individuals save more for retirement. The state of Oregon recently launched an auto-enrollment retirement savings program for private sector workers who lack access to workplace retirement plans; many of these workers are lower-paid employees working at smaller firms. Our paper investigates early results from the OregonSaves program using data through June 2019. We find that OregonSaves is serving firms across many industries, including food services, health care, retail trade, and agriculture. In June 2019, approximately 24,000 contributing participants deposited an average of \$110 per month, or about 5% of their pay, which is the default savings rate. To date, over 40,000 individuals have accumulated combined assets over \$22.7 million. We also find that OregonSaves has provided access to workplace retirement accounts for employees of small to mid-sized firms (average firm size 36 employees), with participating employees' earning an average of \$2,182 per month.

2.2. Introduction

Only about half of the U.S. private-sector workforce is currently covered by an employer-sponsored retirement plan. This fact has sparked debates about a national "retirement crisis,"[2] and it has also prompted over half of all U.S. states to at least consider mandating that private-sector firms offer their employees retirement saving accounts. Oregon has led the way with its OregonSaves program, launched in 2017, with the goal of increasing workers' personal savings and strengthening retirement security beyond Social Security and means-

[1] This chapter is based on a manuscript authored by John Chalmers, Olivia S. Mitchell, Jonathan Reuter, and Mingli Zhong.

[2] See for instance, Miller et al. (2015), and a rebuttal by Biggs (2015, 2019a,b); Biggs and Schieber (2015); also Bee and Mitchell (2017).

tested social transfers. OregonSaves works under state law by requiring private-sector firms that lack existing employer-sponsored retirement plans to register and participate in the program. By requiring employers to participate in a pre-designed program that removes employers' fiduciary responsibility, the program reduced two barriers to employers offering a plan: set-up and monitoring costs.

OregonSaves is structured as a Roth Individual Retirement Account (IRA), with automatic enrollment, a default (after-tax) contribution rate of 5%, and employee-only contributions. Once an employer registers and provides OregonSaves with employees' data, employees enter into a 30-day enrollment period during which time their identity is verified and employees may choose to opt out. A Roth account 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 contribution rate, though employees can choose to save at a different contribution rate (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. One appealing feature of the plan is that participants may access a substantial portion of their money without risk of penalty. Although similar to privately-managed employer-sponsored retirement plans such as 401(k) and 403(b) plans, OregonSaves differs by permitting workers' retirement savings accounts to follow them as they move from one job to another. This portability feature is potentially important to employees lacking retirement plans through their employers, since this labor market is characterized by small firms offering lower pay, typically with high worker turnover.

A key rationale for state-based auto-enrollment retirement plans is the fact that the vast majority of workers lacking access to employer-sponsored retirement plans have no dedicated retirement saving vehicles Chen and Munnell (2017). In other words, while workers could have responded to the lack of employer-sponsored retirement plans by opening and funding their own Traditional or Roth IRAs, the vast majority have not done so.

We propose three, non-mutually exclusive explanations for this inaction. First, lower income workers at predominantly smaller firms may not be able to afford (or perceive that they cannot afford) to save for retirement. This explanation is consistent with the 2013 Survey of Consumer Finance, which found that only 4% of workers with bottom-quintile income had a defined contribution retirement plan, versus 68% for workers with top-quintile income Morrissey (2016). Furthermore, many households report that they have difficulty meeting even basic expenses; for example, the Board of Governors of the Federal Reserve System (2019) found that "17 percent of adults are not able to pay all of their current month's bills in full." Such statistics suggest that the marginal utility of income is high for many low-income workers. Moreover, the fact that these workers are so close to the margin highlights the importance of allowing savers to withdraw funds from a portable retirement savings plan.

A second rationale for employee inaction considers employees' expectations about how they might benefit from retirement saving. In particular, the Social Security replacement rate is relatively high for low-income workers, so workers at firms lacking employer-sponsored retirement plans may rationally perceive that they have little need for additional retirement savings. These first two explanations would predict low participation rates in OregonSaves, low contribution rates among those who do participate, or both.

A third explanation for the lack of retirement saving by workers without employer-sponsored retirement plans are search costs (Bronchetti et al., 2013). Specifically, it may be that workers face high search costs when they consider opening IRA accounts in the absence of an employer sponsored plan, discouraging them enrolling on their own. Research has shown that earnings, retirement planning, and financial literacy are positively correlated (e.g. Lusardi and Mitchell (2007); Clark et al. (2017)); as a result, it is likely that workers at firms without an employer-sponsored retirement plan are less financially literate than other employees. In turn, less financially literate individuals may lack the confidence and knowledge required to research and select IRAs, and to successfully manage their own retirement portfolios. In such circumstances, a state-based program that reduces search costs, both in terms of the enrollment process and in offering a simple set of default investment options, may lead to higher participation than in the absence of a plan like OregonSaves.

Since little is known about why few lower-income workers fail to set up IRAs, we evaluate how workers previously lacking access to workplace retirement plans are responding to the OregonSaves plan. Our analysis of administrative data from OregonSaves allows us to provide a preliminary look at the characteristics of eligible employees and their employers in a segment of the U.S. labor market that has not yet been studied with account-level data. These data allow us to provide preliminary data on how participation decisions, contribution rates, and account balances vary with employee and employer characteristics and, to study the reasons that employees give for opting out.

Importantly, these data allow us to explore the relative importance of the three explanations listed above. If the lack of retirement saving is mainly due to peoples' inability (or perceived lack of need) to save for retirement, then we would anticipate finding low participation in the OregonSaves program, particularly among workers with low and volatile earnings profiles. In contrast, if having no retirement savings is primarily due to search costs associated with lower levels of financial literacy, we will anticipate patterns qualitatively similar to Madrian and Shea (2001) who examined savings behavior of employees in a "large, publicly-traded Fortune 500 company in the health care and insurance industry" when the firm introduced automatic enrollment in its 401(k) plan. They found that plan participation rates averaged about 85% for new hires who joined the firm under automatic enrollment. Furthermore, the largest participation increases due to auto-enrollment were for younger and less highly-compensated employees. More generally, studies of participant behavior in employer-provided 401(k) plans find that the younger, lower-paid, and less educated workers are more likely to adopt default savings rates and invest through default investment options, especially target date funds (e.g., Madrian and Shea (2001); Mitchell and Utkus (2012);

Chalmers and Reuter (2012)). Accordingly, to the extent that low levels of retirement saving are the result of high search costs, we would anticipate finding relatively high participation rates under the OregonSaves program. Moreover, to the extent that the 5% default savings rate is perceived by plan participations to be the 'recommended' savings rate, we expect little variation in observed savings rates, especially among younger workers with lower wages.

Our preliminary findings rely on data through June 29, 2019 and suggest the following preliminary inferences. First, OregonSaves is generating appreciable retirement savings for a substantial number of employees. Approximately 40,000 employees have contributed over $22.7 million dollars to the program through June 2019. Second, most contributing employees (71.5%) are saving at the 5% default rate. Third, participation in the plan spans the state of Oregon and beyond, but the largest concentration of assets is located in the larger urban areas. One challenge we face in characterizing the OregonSaves outcomes pertains to how we define 'participation' in this program. Specifically, it is difficult to define who is eligible to contribute at any given time, given that the employers are quite heterogeneous and the employees experience frequent turnover. Moreover, participation can be measured in terms of anyone who ever participated in the program, or in terms of current contributors. In any event, given participants' relatively low earnings levels, it is not surprising that participation rates are lower than observed in Fortune 500 firms. Indeed, a common explanation employees offer for opting out is "I cannot afford to save," indicated by around 30% of those opting out during their 30-day enrollment period.

2.3. Data and Descriptive Statistics

Using individual-level administrative data from the first two years of the OregonSaves program, we present empirical evidence on (a) the impact of OregonSaves on expanding access to workplace retirement savings programs, (b) the characteristics of employers providing access to OregonSaves for their employees, (c) the characteristics of workers covered by OregonSaves, and (d) the impact of OregonSaves on retirement savings.

OregonSaves Expansion and Characteristics of Registered Employers

As of June 29, 2019, 4,970 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 171,243 individuals had information provided to OregonSaves by an employer, and median firm size was 16. Employees' average age in 2019 was 37. As of June 29, 2019, OregonSaves had accumulated $22.7 million in total assets.

Food services and retail trade are two of the largest industries represented in OregonSaves in terms of the number of registered employers. Food services and health care are two of the largest industries in terms of the number of employees ever had access to OregonSaves. It is our understanding that the large number of workers in health care can best described as home-health care workers. Finance, insurance, and management firms are some of the smallest industries in OregonSaves, in terms of both the number of registered employers and employees.

Characteristics of Workers Having Access to OregonSaves and Their Participation and Contribution Decisions

Panel A of Table 6 itemizes the status of the 171,243 employees who had a chance to enroll in OregonSaves. During the enrollment window, 41,757 (24.4%) employees opted out, while another 21,600 (12.6%) employees opted out after the enrollment window. There were 29,332 (17.1%) employees awaiting the background check, which, in many cases, extends their pending status. There were 12,630 (7.4%) employees who enrolled and passed their background checks, but their employers had not yet submitted payroll. Finally, 65,924 (38.5%) 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 29,332 pending cases and the 12,630 employees still to contribute are also potential participants for whom I cannot yet observe their choices. Of the original 171,243 names submitted, approximately 37% elected to opt out; this does not, however, imply that the complement of this group represents participants.

In principle, the program participation rate refers to the employees who are making or have made contributions to OregonSaves, as a percent of employees eligible to participate, working, and who have an employer cooperating with OregonSaves. Yet when measuring the participation rate, there are two challenges to defining the denominator. Given the data to which I have access, I cannot distinguish between someone who is working and not contributing, from someone who is not working. It is also difficult to identify employees not participating because of actions taken by their employers, rather than actions they themselves took. As result, I must define a group of potential participants which is eligible and active using a set of imperfect but necessary assumptions.

Panel B of Table 6 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 76,438 people. In this group, 23,503 individuals received a monthly contribution to their accounts in June 2019, with a mean contribution amount of $110. For employees with a positive contribution amount and a positive savings rate, I estimate their monthly incomes (=contribution/contribution rate) to be $2,182. 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 7 presents data for the 40,652 OregonSaves participants having a positive

OregonSaves account balance. Given total assets of $22.7 million, the average balance per account stood at $558 as of June 2019. Panel B of Table 7 shows that 28,083 of the 40,652 with a positive balance are classified as eligible active workers. When averaged over accounts with a positive balance, the average account balance for EAW is $653.20.

To illustrate some of the challenges in defining participation rates, I refer to Table 7 where 40,652 individuals have a positive account balance. Some of these, however, are not defined as active. One might argue that the participation rate could include all people who have participated as a fraction of current EAW or 53% (40,652/76,438), which is the ratio of anyone with a positive balance relative to the EAW group. If I focus on EAW workers who are eligible for contributions and actively working, the participation rate includes EAWs who ever had a positive balance relative to all EAWs, or 41.3% (31,573/76,438). In June 2019, there were 23,503 people contributing to the program. If one were interested in the number of contributing employees in June 2019 as a fraction of EAW, this would produce a contribution participation rate of 30.7% (23,503/76,438). Benchmarking the numbers of participants is difficult relative to prior studies, because our data include multiple employers, multiple jobs for some employees, and months in which no contributions are paid, along with our limited ability to discern workers' employment status from our data especially when people opted out or set their savings rates to zero.

Table 8 presents summary statistics for the eligible active workers (EAW) by industry, at the end of June 2019. Consistent with our expectations, the largest industries represented in OregonSaves are food services, which employed 26,787 employees or about one-third of the EAW employees, across 499 employers. In food services, the average participation rate as a fraction of EAW is 46.4%. The next largest industry is healthcare, which employs about 8,860 workers, across 159 employers; the participation rate in healthcare is 43.6%. It is our understanding that a large number of workers in healthcare can best described as home-health care workers. Business support and retail trade are the next two largest industries in terms of employees; we understand that business support includes a large number of temporary employers, which explains the large number of employees per employer (69.5) and the average number of jobs being among the highest. Agriculture is noteworthy as it has the lowest participation rate (23.6%) and includes workers with the highest average number of jobs (1.2). This category includes many of the employers who hire temporary farm workers. Overall, Table 2 shows us that the firms being served are small, averaging 34.7 employees per firm, having an average participation rate of 41.3% of EAWs.

Table 9 presents program participation rates by industry and firm size. It is important to recall that firms with fewer than 10 people have not yet been required to enroll in the program. As a result, the 72.3% participation rate for firms with 1-4 people almost certainly reflects firms' choosing to enroll early, perhaps due to employee enthusiasm about the program. Across enrollment waves 1-3, the highest participation rates are evident in the largest firms, while the patterns across industry are similar to those observed in Table 8.

Figure 2 presents the geographic distribution of assets accumulated under the OregonSaves program to date, by zip code. Regions on the map that are shaded have at least $10,000 of assets under management in OregonSaves, and the darkest red areas have up to about $3.2 million assets under management in that zip code. Not surprisingly, the darkest regions are located in and around the largest cities in Oregon, including Portland, Salem, Eugene, Bend, Roseburg, and Medford/Ashland. Nonetheless it is also worth noting that participation is dispersed throughout the state.

Table 10 presents the distribution of savings rates for eligible active workers (EAW). About 22.1% of the EAW elected the default rate of 5% in June 2019. About 4,000 or 5.3% of EAW had savings rates of 6%, a large fraction of which may be attributed to the auto-escalation feature of the plan. Details and evidence about auto-escalation are shown in Table 14 which will be discussed in the following section. Of the 76,438 EAW, 69% had savings rates of zero, a tally that includes people who opted out (closed their accounts), along with EAWs who later set their rates to zero without closing their accounts, to leave open the possibility of saving later. Of the remaining employees, few had savings rates other than 0, 5%, and 6%.

2.4. Predicting Participation

In Table 11 we present results from multivariate regressions predicting participation in the OregonSaves program. At present we are limited to three sets of characteristics to predict participation: participant age, their employers' industry and number of employees, and estimated average employee earnings for firms having at least one employee contributing to OregonSaves (the latter allows us to estimate an average wage for that firm). We estimate multinomial Logit marginal effects and OLS coefficients for a binary dependent variable equal to 1 if the employee participates, and 0 otherwise. This sample comprises EAWs employed at firms with more than 10 workers, as this group has been mandated to participate in OregonSaves; we also include firms for which we have industry classifications. In robustness analysis, we also estimate focus on the subset of employees having only one job (columns 2-4 and 6-8), a criterion that slightly reduces the sample size.

Column 1 of Table 11 indicates that, relative to the reference age variable (66+), younger workers are more likely to participate, and the point estimates are larger for younger employees. Additionally, people that already have an existing OregonSaves account are also more likely to participate. Column 4 presents the Logit results when industry controls are included and the sample is limited to employees with a single employer. The coefficient estimates for the age variables and the presence of a pre-existing OregonSaves account are similar in size and significance to those in Column 1. The industry coefficients may be interpreted relative to the reference category of wholesale trade. Thus, higher participation rates are seen in the Arts Entertainment, Business Support, Food Services, Healthcare, Management, Other Services, and Retail Trade sectors. Our OLS estimates in Columns 5-8 are similar both in significance and in magnitude. Columns 3,4,7 and 8 report coefficient estimates on firm size (=number of employees) and monthly employee income: here the results are noteworthy in that prior research has found that firm size is significantly related to participation rates, while in our dataset, we cannot reject the null hypothesis that the coefficients are zero.

Table 12 offers a summary of the reasons people offered for opting out of the OregonSaves program. Panel A tallies answers provided by users where they had to select one of a set of choices: the most common reason given was that people felt they could not afford to save as 29% of those who opted out offered that explanation. Another 20.6% of those opting out said that they already had their own retirement plans, and 25% gave "other" reasons. Additionally, 8% suggested they were not interested in contributing through their current employers. Panel B offers additional insights into the more than 5,000 responses given when an employee elected the "other" category indicated in Panel A. The three most prominent rationales for opting out include that fact that people were no longer employed, were not interested, and were already near or in retirement.[3]

2.5. Discussion

Our preliminary analysis has indicated that OregonSaves does provide access to workplace retirement accounts for employees of small to mid-sized firms, with the average participant monthly earnings estimated at about $2,100. Accordingly, this program is serving a demographic that has not traditionally been served by retirement saving accounts. Employees making contributions in June 2019 contributed an average of $110, representing approximately 5% of their pay. While the opt out rates in the program are on the order of 35% relative to the total number of employees who entered their 30-day enrollment windows, the participation rates relative to eligible and active workers are approximately 41% using our definitions. We have also outlined challenges in characterizing participation rates, namely defining who the relevant group of potential participants is, and who participates. We offer our definitions with the understanding that these ratios are subjective. The number of participants is much easier to characterize: as of June 2019, approximately 40,000 people

[3] Other themes included opposition to the government and to auto-enrollment plans, as well as anti-social comments such as "none of your dam business." Example comments are provided as-is with the exception that curse words have been slightly disguised.

have participated in OregonSaves and almost 24,000 made contributions that month. In the absence of OregonSaves, it is likely that these participants would not have otherwise chosen to begin saving in a retirement account.

By way of comparison, the United Kingdom's National Employment Savings Trust (NEST) program is also an auto-enrollment retirement plans targeted at firms lacking retirement plans. Yet NEST has a significantly longer history relative to OregonSaves: large U.K. firms were required to begin enrolling workers in NEST (or another plan) in October 2012, and there was a staged rollout for smaller firms. A recent analysis of NEST members' behavior by Vanguard (2019) reported that between March 2013 and January 2018, "a total of a total of 612,000 employers and over 6 million unique members joined NEST." That program's overall opt-out rate for ongoing enrollments was 6%; opt-out rates were 3% for those younger than age 25 to over 30% for those age 65+; and they stood at 13% for the smallest firms (1-4 employees) rising to 6% for the largest firms (5,000+) Vanguard (2018). Nevertheless, participation rates were much higher in the U.K. (around 94%), compared to the OregonSaves program (around 40%).

As the present report is preliminary, we leave several important questions for future research. First, we will investigate whether OregonSaves participants use the first $1,000, allocated to the Capital Preservation Fund, as a form of rainy-day account to cover unexpected expenses. To address this question, we will analyze the timing and magnitude of withdrawals from OregonSaves accounts, including the link between job turnover and partial withdrawals. Second, we seek to investigate whether participation in OregonSaves helps improve household balance sheets. Beshears et al. (2017) used administrative information linked to credit bureau data to study whether automatic enrollment of federal civilian employees into the Thrift Savings Plan was associated with greater debt. If we are able to link credit bureau data (while retaining anonymity of the individual records), we will study the effect of OregonSaves participation on the borrowing levels and credit scores of participants with different earnings and/or in different industries. We will start by comparing

OregonSaves participants (or eligible employees) to a matched sample of Oregonians not covered by the program. We will also exploit any variation in opt out rates between large and small firms within the same industry, under the assumption that this variation is more likely to reflect differences in the intensity of firm-level outreach and education, rather than within-industry variation in employee preferences.

Another important question is whether participation in OregonSaves boosts workers' total retirement savings. To answer this question, we will exploit longitudinal survey data on the level of retirement assets outside of OregonSaves. And, in the longer-term, we would like to evaluate whether and to what extent OregonSaves reduces workers' reliance on other social welfare programs, along the lines of Bernheim et al. (2015).

Finally, in the introduction we suggested three reasons that an auto enrollment plan may or may not work. Currently, we can offer observations but only interim conclusions. With respect to 1) Not enough money to save: we find substantial evidence that this is a common reason given by approximately 30% of the people who opt out of OregonSaves. With respect to 2) Replacement rates from Social Security will be closer to current income: we observe that income for the contributors in our sample are on the order of $2,000 per month. As a result, Social Security payments will very likely represent a larger proportion of $2,000, than they will compared to average workers' retirement incomes. With respect to 3) Search costs deter enrollments: we do find that participants land on the default deferment rate of 5% with the highest frequency when we observe a positive contribution. Overall, the rate of 0% is most common and applies to those with an explicit rate of 0% and those that opt out. This suggests that participants who desire a retirement plan, benefit from the lower search costs embedded in the OregonSaves default structures. Nevertheless, it is also encouraging to see that many potential participants are able and willing to make a choice not to participate, while others contribute above the default rate. These facts suggest that reducing search costs while allowing a sufficiently clear path to opting out has the potential to generate a larger social benefit. In other words, OregonSaves is reducing the costs of building retirement saving for those that wish to save, but it also allows those who do not wish to save the opportunity to avoid participation.

CHAPTER 3 : OPTIMAL DEFAULT RETIREMENT SAVING POLICIES: THEORY AND EVIDENCE FROM OREGONSAVES

3.1. Abstract

Automatically enrolling individuals into a retirement plan has been widely adopted around the world. In this paper, I develop a unified framework to characterize the counterbalancing welfare effects of the default savings rate in auto-enrollment retirement plans, and derive a formula for the optimal default savings rate that maximizes individual welfare. The optimal default savings rate is shaped by the social benefits of increased savings due to adherence to the default that keeps workers from actively undersaving. The optimal default savings rate is also counterbalanced by the social benefits of action when an undesirable default option compels workers to make an active decision. I derive a formula for the optimal default savings rate that depends on reduced-form statistics. To empirically estimate these statistics, I exploit an exogenous increase in the default savings rate from 5% to 6% in OregonSaves, the first state-sponsored auto-enrollment plan in U.S., to estimate individual adherence to the default savings rate. I also combine individual-level survey data with administrative data to estimate the degree of active undersaving. The formula suggests that the optimal default savings rate 8%.

3.2. Introduction

A large literature has documented that automatically enrolling individuals into a retirement plan significantly increases their savings. Moreover, a substantial fraction of participants passively save at the default savings rate determined by the plan designer. Automatic enrollment and the default savings rate are considered to be cost-effective policy tools to help individuals better allocate resources over the life cycle and potentially improve their lifetime welfare. This argument has led to a widespread adoption of automatic enrollment in retirement plans around the world. Although many studies have found a causal effect that defaulting individuals into a predetermined savings rate effectively increases savings, little is known about the welfare effects of the default savings rate.

In this paper, I theoretically analyze and empirically identify the welfare effects of the default savings rate. I then provide an explicit formula for the optimal default savings rate that maximizes individual welfare. 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 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 they decide to opt out of the default rate. A growing literature in behavioral economics and household finance, however, suggests that individuals might underestimate the amount of retirement savings they need due to financial illiteracy, misinformation, or myopia.

Motivated by these findings, this paper proposes an additional welfare effect of the default savings rate, namely, to correct 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 that maximizes their life cycle utility. In contrast, if the non-default preferred rate that individuals actively choose does not maximize their life cycle 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 level of the default savings rate is shaped by its welfare effects when correcting the two types of behavioral biases, I propose a unified framework that incorporates both types of behavioral biases. 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 elect a non-default preferred savings rate. Given individual binary choices of either staying or opting out of the default rate, the policymaker with more information or who is more forward-looking than individuals maximizes the normative utility. That describes the *reality* of how individual choices actually affect their lifetime welfare.

Based on the unified welfare framework, I derive a formula for the optimal default savings rate as a function of reduced-form statistics that can be empirically identified. The traditional approach to such welfare and optimal policy questions usually requires structural estimation of an underlying model's primitives, and then a numerical simulation of the effects of policy changes. I instead identify a set of sufficient statistics directly from exogenous policy variations and survey responses. These sufficient statistics measure how individuals respond to different levels of the default savings rate by opting out of the defaults, and how their responses affect their lifetime welfare.

The welfare analysis and the sufficient statistic formula proposed in this paper differ from the traditional approach because my approach does not rely on specific underlying behavioral models of why the default options affect individual behavior. I instead directly characterize how individual behavior impacts their welfare. This approach allows my welfare analysis to incorporate a range of underlying behavioral models that explain the default effect on individual behavior. Additionally, since the sufficient statistics in the formula for the optimal default savings rate measure individual responsiveness to the default rate and individual revealed preference, my approach provides a non-paternalistic method to incorporate behavioral biases into the optimal design of default saving policy.

I implement the formula for the optimal default savings rate empirically by estimating three sets of key statistics. The first are two semi-elasticities 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 U.S. launched in 2017, OregonSaves, I find that 12% of passive savers no longer stayed at the default rate as the default rate increased by one percentage point. Additionally, using data from Beshears et al. (2012), I find that 13% of active savers chose to stay at the default rate as 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 where 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 are revealed time preferences to infer 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 rates improves individual welfare by protecting them 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 of OregonSaves-eligible workers, in which I find that on average workers weakly prefer spending most of their income now over spreading out the income between now and the future.

The third statistic is the average perceived cost of opting out of the default savings rate. This statistic reflects the extent to which the default savings rate promotes individual welfare by protecting them from inaction. Choukhmane (2018) estimates that on average individuals believe that opting out of the default rate costs about \$250. This estimate of the perceived opt-out cost suggests that, instead of spending \$250 from public funds to compensate workers for voluntarily selecting their non-default preferred rates, a policymaker can avoid the public expenditure while achieving the same goal by setting the default rate optimally. Plugging the point estimates into the optimal formula, I find the optimal default savings rate to be 8%.

Contributions to the Literature. The welfare analysis 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 in previous 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 that the optimal default savings rate should be set at the employer matching cap in an employer-sponsored retirement plan. My analysis and the formula for the optimal default savings rate complement 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 the welfare analysis.

The present paper also contributes to the literature on a sufficient statistics approach for optimal public policy starting with Saez (2002). I extend the applicability of the sufficient statistics approach to "nudge" policies, and in particular the default retirement saving rate. Farhi and Gabaix (2020) discussed optimal nudges in taxation, while I focus on optimal nudges in the context of default saving policies.

The present paper also sheds light on two long-standing questions in household finance. The first question is why a substantial fraction of American households saves 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 because they are not *automatically* enrolled in a savings plan. 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, expected life expectancy, 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 net programs which increase the fiscal burden imposed on all taxpayers to finance these social 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 chapter is organized as follows. Section 1.3 describes the institutional background of the OregonSaves program and provides descriptive statistics for the first two years of the program as of June 2019. Section 1.4 discusses the welfare impact of the default rate in a sufficient statistic framework and presents an explicit formula of the optimal default rate. Section 1.5 describes the identification strategies and estimation results of the key statistics in the optimal default formula using OregonSaves administrative and survey data. In Section 1.6, I calculate the optimal default rate using the estimation results from Section 1.5. Section 1.7 concludes.

3.3. 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.

3.3.1. 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 13 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[1] 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

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

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, and the smallest firms (4 or fewer employees) will start enrolling May 12, 2020. In practice, however, some smaller firms did register earlier than required, and some unknown number of larger firms may not have registered to date. As of June 29, 2019, OregonSaves was still rolling out to smaller employers. 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.

3.3.2. Data and Descriptive Statistics

Using individual-level administrative data from the first two years of the OregonSaves program, I present early evidence on the effect of automatic escalation in savings rates on participation and contributions.[2]

Automatic Escalation

Table 14 reports the impact of automatic escalation on participating and contribution decisions. On January 1, 2019, workers who had open accounts for six months were eligible

[2] Additional descriptive results are documented in Chalmers et al. (2019), which is also Section 2 in this dissertation.

for auto-escalation. Additionally, workers who initially elected any non-zero savings rates (default or non-default) were eligible for auto-escalation. Similar to the initial default savings rate, eligibles could actively opt out of the auto-escalation arrangement; if they did not, savings rates automatically increased on January 1 by 1 percent, and would continue to do so until they reached 10%. Panel A shows subgroups of all individuals eligible for auto-escalation. Panel B shows how eligible active workers (EAW) who were eligible for auto-escalation responded. Panel B is more informative because EAWs are active workers with positive earnings who make non-zero contributions. Individuals counted in Panel A include inactive workers not making positive contributions even with a non-zero savings rate recorded. Results for EAW are similar to the full sample in Panel A except for two numbers. First, only 10.7% of EAW opted out of the program after auto-escalation took into effect. This suggests that most individuals opted out of the program after auto-escalation, because they no longer worked for the employer offering OregonSaves. Second, 66.4% of EAW eligible for auto-escalation adhered to the new rates at the end of June, 2019, higher than 51.1% for all individuals.

In summary, our early findings for the first two years of OregonSaves suggest several preliminary conclusions. First, 4,970 employers complied with the mandate to register their businesses in OregonSaves as of June 29, 2019, most of which were small businesses with fewer than 20 employees. The largest industries represented in OregonSaves were food services and health care (mostly home health care workers). Second, through these registered employers, OregonSaves provided 171,243 private sector workers access to workplace retirement saving plans. Moreover, among these covered workers, 76,438 were eligible to contribute and were actively working so that they could make positive contributions to their OregonSaves accounts if they elected a positive savings rate. Of these, 41.3% participated in the program and 30.7% made a positive contribution in June 2019. The leading rationales for opting out were being unable to afford to save or having an existing retirement plan. About 72% of participating EAWs accepted the 5% default rate. The average monthly positive contributions were $110. Finally, when the savings rates automatically increased by 1 percent on January 1, 2019, about 70% of EAWs eligible for automatic escalation accepted the rate increase.

3.4. 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.

3.4.1. 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 λ. 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 λ 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, $\theta := (Z, \lambda, \delta) \in \Theta$, 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 $R(\theta) = rZ(\theta)$. Each type of workers chooses a pension saving amount $P(\theta)$ from two discrete options: the preferred saving amount $S(\theta)$ or the default saving amount $R(\theta)$. The preferred saving amount equals the preferred savings rate θ times their income: $S(\theta) = \theta Z(\theta)$. 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:

$$U(C(\theta), P(\theta); \theta, r, K) = u(C(\theta)) + \lambda(\theta)\delta(\theta)v(P(\theta)) - K\mathbf{1}\{P(\theta) = S(\theta)\}, \tag{3.1}$$

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,[3] 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:

$$\begin{aligned} N(C(\theta), P(\theta); \theta, r, K, \pi) &= u(C(\theta)) + \delta(\theta)v(P(\theta)) - \pi K\mathbf{1}\{P(\theta) = S(\theta)\} \\ &= U + (1 - \lambda(\theta))\delta(\theta)v(P(\theta)) + (1 - \pi)K\mathbf{1}\{P(\theta) = S(\theta)\}, \end{aligned} \tag{3.2}$$

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 homogeneous across the population.

Equation (3.2) 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, $(1-\lambda(\theta))\delta(\theta)v(P(\theta))$, 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

[3] 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.

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 $\alpha(\theta)$:

$$W(r) = \max_{r} \int_{\Theta} \alpha(\theta) N(C(\theta), P(\theta); \theta, r, K, \pi) dm(\theta) \tag{3.3}$$

subject to individual optimization

$$\{C(\theta), P(\theta)\} = \arg \max_{\{C,P\}} U(C, P; \theta, r, K) \tag{3.4}$$

where

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

3.4.2. Optimal default savings rate

Let r^* denote the optimal default savings rate. Next I consider the welfare effect of a marginal increase in the optimal default rate from r^* to $r^* + dr$. Based on the individual optimization problem characterized in Equation (3.4), workers of the same type select the same contribution amount $P(\theta) \in \{S(\theta), R(\theta)\}$. For a continuum of types $\theta \in [0, 1]$, workers whose preferred saving rates are between θ_l and θ_h will adhere to the default saving rate where $\theta_l < d < \theta_h$. The density of workers saving at the default $m_r = m_l + m_h = \int_{\theta=\theta_l}^{d} dm(\theta) + \int_{\theta=d}^{\theta_h} dm(\theta)$. I define workers who remain at the default as passive savers, where m_l is the fraction of passive savers (in the population) whose preferred rates are below the default, and m_h is the fraction of passive savers whose preferred rates are above the default. I refer to m_l as l-type passive savers, and m_h as h-type passive savers. Figure 3 displays how each type of passive savers responds to a marginal perturbation of the optimal default rate.

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 savings from a type-θ worker as much as \$$g(\theta)$ from public funds. The welfare weights can be formally expressed as:

$$g(\theta) := \frac{\alpha(\theta) v'_{P(\theta)}}{\lambda}. \tag{3.5}$$

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

1. Individuals make their saving decisions once at the beginning of their working lives.[4]

[4] 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).

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: $v(P(\theta)) = P(\theta)$.

Next I characterize the optimal default rate r^* based on the policymaker's problem described in Equations (3.3) - (3.4). A marginal increase in r^* induces three welfare effects on passive savers whose preferred rates are between s_l and s_h. First, passive savers marginally increase their savings by $\dfrac{dN(R(\theta))}{dr}$. Second, a fraction of h-type passive savers whose preferred rates are above r^* start saving at the increased new default $r^* + dr$, because it is now closer to their preference. This welfare effect on the extensive margin is proportional to $\dfrac{dm_h}{dr}$, and the savings amount per worker decreases from $S_h (= \theta_h Z_h)$ to $R_h (= rZ_h)$. Third, a fraction of l-type passive savers whose preferred rates are below r^* stop saving at the increased new default $r^* + dr$ because it is farther from their preference. This welfare effect on the extensive margin is proportional to $\dfrac{dm_l}{dr}$, and the savings amount per worker decreases from $R_l (= rZ_l)$ to $S_l (= \theta_l Z_l)$. The first-order condition for the social welfare function W equals zero at the optimum:[5]

$$\begin{aligned}\frac{dW(r^*)}{dr} &= \frac{d}{dr}\int_{\theta=\theta_l}^{\theta_h} \alpha(\theta) N(P(\theta)) dm(\theta) \\ &\approx \int_{\theta=\theta_l}^{\theta_h} \alpha(\theta) \frac{dN(R(\theta))}{dr} dm(\theta) + \frac{dm_h}{dr}\alpha_h(N(R_h) - N(S_h)) - \frac{dm_l}{dr}\alpha_l(N(S_l) - N(R_l)) \qquad (3.6) \\ &= 0.\end{aligned}$$

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

[5] I use $N(P(\theta))$ to represent $N(C(\theta), P(\theta); \theta, r, K, \pi)$ in Equation (3.3) subject to the budget constraint $C(\theta) + P(\theta) = Z(\theta)$. Retirement savings $P(\theta)$ is chosen between the default saving amount $R(\theta)(= rZ)$ and the preferred saving amount $S(\theta)(= \theta Z)$. The differentiation under the integral sign employs the Leibniz integral rule where the end points of the interval of the integral θ_l and θ_h are functions of the derivative argument r.

equation at the optimum:

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

Proof. See Appendix B. The overall welfare effect can be decomposed into several terms after the optimal initial default rate marginally increases from r^* to $r^* + dr$:

1. The aggregate weighted social welfare gain to all passive savers on the intensive margin is $\frac{dI}{dr} = \frac{m_l}{m_d} \cdot g_l \cdot (1-\lambda_l)\delta_l Z_l + \frac{m_h}{m_d} \cdot g_h \cdot (1-\lambda_h)\delta_h Z_h$. For example, as the initial default rate increases by dr, l-type passive savers on the intensive margin increases their savings by $dr \cdot Z_l$. Although they might feel indifferent to the marginal policy change, there is an increase in the welfare internality of savings, which is the realized welfare gain to passive savers that they do not internalize. Based on Equation 3.1, the marginal increase in the welfare internality of savings for a l-type worker is $(1-\lambda_l)\delta_l drZ$, and the marginal increase is weighted by g_l to evaluate its impact on social welfare. The social value of the marginal increase in the welfare internality of savings is then weighted by the fraction of l-type passive savers $\frac{m_l}{m_d} \cdot g_l \cdot (1-\lambda_l)\delta_l drZ_l$. Analogously, the social welfare gain to h-type passive savers is $\frac{m_h}{m_d} \cdot g_h \cdot (1-\lambda_h)\delta_h drZ_h$. The aggregate weighted social welfare gain on the intensive margin equals $dI = \frac{m_l}{m_d} \cdot g_l \cdot (1-\lambda_l)\delta_l drZ_l + \frac{m_h}{m_d} \cdot g_h \cdot (1-\lambda_h)\delta_h drZ_h$.

2. The welfare gain to l-type workers for switching to their preferred rate s_l under the new default $r^* + dr$ is $\frac{dS_l}{dr} = |\epsilon_l| g_l (1-\lambda_l)\delta_l s_l Z_l$. As the new default rate is farther from their preferred rate, the fraction of the l-type workers on the margin of opting out of the default is $\frac{dm_l}{m_r} = dr|\epsilon_l|$. Each l-type worker opting out of the default enjoys the welfare internality of saving at their preferred rate $(1-\lambda_l)\delta_l s_l Z_l$ weighted by g_l. The total social welfare gain is $dS_l = dr|\epsilon_l| \cdot g_l \cdot (1-\lambda_l)\delta_l s_l Z_l$.

3. The social welfare loss to h-type workers for no longer saving at their preferred rate s_h is $\frac{dS_h}{dr} = |\epsilon_h| g_h (1-\lambda_h)\delta_h s_h Z_h$. As the new default moves closer to h-type workers'

preference, the fraction of h-type workers on the margin of starting to save at the default ($dr|\epsilon_h|$) no longer enjoy the welfare internality of saving at their preference, $(1-\lambda_h)\delta_h s_h Z_h$, weighted by g_h. Similar to the size of dS_l with an opposite direction, the welfare loss to h-type workers for no longer saving at their preference equals $dS_h = dr|\epsilon_h| \cdot g_h \cdot (1-\lambda_h)\delta_h s_h Z_h$.

4. The social welfare loss to l-type workers for no is $\frac{dR_l}{dr} \cdot r^* = |\epsilon_l| g_l (1-\lambda_l)\delta_l r^* Z_l$. As l-type workers on the margin ($dr|\epsilon_l|$) stop saving at the default, the social welfare loss equals the welfare internality of saving at the default $(1-\lambda_l)\delta_l r^* Z_l$ weighted by its social marginal Iight g_l. The total social welfare loss to l-type workers on the margin for no longer saving at the default equals $dR_l = dr|\epsilon_l| \cdot g_l \cdot (1-\lambda_l)\delta_l Z_l$.

5. The social welfare gain to h-type workers for starting to save at the default rate is $\frac{dR_h}{dr} \cdot r^* = |\epsilon_h| g_h (1-\lambda_h)\delta_h r^* Z_h$. As h-type workers on the margin ($dr|\epsilon_h|$) start saving at the default, the social welfare gain equals the welfare internality of saving at the default $(1-\lambda_h)\delta_h r^* Z_h$ weighted by g_h. The total social welfare gain to l-type workers on the margin for starting to save at the default equals $dR_h = dr|\epsilon_h| \cdot g_h \cdot (1-\lambda_h)\delta_h Z_h$.

6. The social welfare gain to l-type workers for making an active choice is $\frac{dK_l}{dr} = |\epsilon_l| g_l (1-\pi)K$. For each l-type worker on the margin of electing their preferred rate, they enjoy the positive welfare internality of action measured by $(1-\pi)K$. The welfare internality of action has social consequences, because the marginal personal welfare gain can improve social welfare by g_l. The social welfare gain to all l-type workers on the margin ($dr|\epsilon_l|$) for taking action equals $dK_l = dr|\epsilon_l| \cdot g_l \cdot (1-\pi)K$.

7. The social welfare loss to h-type workers for no longer making an active choice is $\frac{dK_h}{dr} = |\epsilon_h| g_h (1-\pi)K$. For each h-type worker on the margin of accepting the default, they become inactive and lose the welfare internality of action, $(1-\pi)K$ weighted by g_h. The social welfare loss to all h-type workers on the margin of no longer taking action equals $dK_h = dr|\epsilon_h| \cdot g_h \cdot (1-\pi)K$.

3.5. Estimating Key Parameters for the Optimal Default Savings Rate

In this section, I outline an empirical strategy to identify key statistics to calculate the optimal default savings rate in Proposition 1 using OregonSaves data described in Section 3.3. Table 19 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 fraction of l-type passive savers (with preferred rates below the default, denoted by dm_l) with respect to the default rate
- ϵ_h: the (observed) semi-elasticity of the percentage change in the fraction of h-type passive savers (with preferred rates above the default, denoted by dm_h) with respect to the default rate;
- δ_l, δ_h: the normative time preference for l- and h-type passive savers; and
- λ_l, λ_h: the behavioral time preference for l- and h-type passive savers.

3.5.1. Semi-elasticities ϵ_l and ϵ_h

The (observed) semi-elasticity ϵ_l measures the percentage change in the fraction 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 one percentage point (dr), equal to $\frac{dm_l}{m_r}\frac{1}{dr}$. Similarly, ϵ_h is 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 one percentage point (dr), equal to $\frac{dm_h}{m_r}\frac{1}{dr}$. Suppose the default rate increased from r to r', then ϵ_l and ϵ_h can be formally expressed as:

$$\begin{aligned}\epsilon_l(r) &= \frac{dm_l}{m_r}\frac{1}{dr} \\ &= \frac{m_{l'} - m_l}{m_r}\frac{1}{r' - r},\end{aligned} \tag{3.7}$$

and

$$\begin{aligned}\epsilon_h(r) &= \frac{dm_h}{m_r}\frac{1}{dr} \\ &= \frac{m_{h'} - m_h}{m_r}\frac{1}{r' - r},\end{aligned} \tag{3.8}$$

where m_l' is the fraction of l-type passive savers under the new default rate r', m_l is the fraction of l-type passive savers under the original default rate r, $m_r = m_l + m_h$ is the total fraction of passive savers, m_h' is the fraction of h-type passive savers under the new default rate r', and m_h is the fraction of h-type passive savers under the original default rate r. when the default rate is d' and m_d is the fraction of passive savers when the default rate is d.

Based on Equations (3.7) and (3.8), two strategies are available to estimate ϵ_l and ϵ_h with distinct advantages. A key assumption these two strategies rely on is that the semi-elasticities are constant across default rates: $\epsilon_l(r) = \epsilon_l$ and $\epsilon_h(r) = \epsilon_h$. [6]

Identification from Automatic Escalation in OregonSaves

I exploit the exogenous variation in the default rate resulting from automatic escalation to identify the (observed) semi-elasticity for l-type passive savers ϵ_l. I use workers' responses to auto-escalation to proxy how they would respond differently to two initial default rates. Section 3.3 and Table 14 describe the institutional details and summary statistics of automatic escalation.

[6] This empirical assumption can be relaxed when I observe long-term data from OregonSaves.

I start by identifying $m_{l'} - m_l$ in Equation (3.7), which is the change in the fraction of l-type passive savers when the default rate automatically increased from 5% to 6%. Although I do not directly observe the fraction of l-type passive savers with an underlying preferred rate below the default rate, I can infer the change in the fraction of l-type passive savers from how many of them become active savers after auto-escalation. The increase in the fraction of l-type active savers is the same size as the decrease in the fraction of l-type passive savers, based on the theoretical assumption 3 in Section 3.4.2 that the underlying preferred rate is invariant.

Table 15 presents the distribution of savings rates for eligible active workers (EAW) eligible for auto-escalation at the end of November 2018 and at the end of June 2019. I exclude EAW eligible for auto-escalation who opted out of the auto-escalation arrangement before it took effect on January 1, 2019. Panel B of Table 14 shows that, among 5,694 eligible EAW, 1,186 (= 410 + 776) opted out of auto-escalation before it took effect. This leaves the sample for estimating the elasticity ϵ of 4,508 (= 5,694 - 1,186). The reason I exclude these is that I need a precise estimate of individual responses *after* the exogenous rate increase. The 1,186 eligible EAW who opt out of auto-escalation in advance were done so for various other reasons. November 2018 is the last month before individuals received notifications about auto-escalation that would take into effect on January 1, 2019. June 2019 is six months after auto-escalation occurred, so that eligible workers could have had enough time to adjust their savings rates in response to the rate increase.

Table 15 shows that 6.6% are l-type active savers saving between 1% - 4% under the 5% original default rate before auto-escalation, and 17.2% are l-type active savers between 1% - 5% under the 6% new default rate after auto-escalation. As a result, l-type active savers increase by 10.6% after auto-escalation. This suggests that l-type passive savers decreases by 10.6% after auto-escalation: $m_{l'} - m_l = -10.6\%$. I also observe that there are 91.9% of passive savers under the 5% default rate: $m_r = 91.9\%$ in Equation (3.7). It is worth noting that in the November distribution, no eligible EAW opted out of the program because workers had to participate in OregonSaves to be eligible for auto-escalation. Additionally, their accounts had to be open for at least 6 months to be eligible (before June 30, 2018). As the OregonSaves program is still rolling out and most workers were registered after June 30, 2018, only a small fraction of EAW were eligible for auto-escalation. I will be able to observe more workers eligible for auto-escalation in the future. I plug in the numbers into Equation (3.7) and get:

$$\begin{aligned} \epsilon_l &= \frac{m_{l'} - m_l}{m_d} \cdot \frac{1}{r' - r} \\ &= \frac{-10.6\%}{91.9\%} \cdot \frac{1}{6-5} \\ &= -0.12. \end{aligned} \tag{3.9}$$

The estimate of ϵ_l suggests that 12% of l-type passive savers (whose preferred rates are below the default) stopped saving at the default rate when it increased by 1 percentage point. Although I can use auto-escalation to identify ϵ_l, I cannot identify ϵ_h, which quantifies the fraction of h-type active savers becoming passive savers as the initial default rate increases by 1 percentage point (h-type are savers with a preferred rate higher than the default). Since h-type active savers opted out of the original 5% default rate before auto-escalation, they were unaffected by the increase in the default rate. I do not know how they would respond to a default rate other than 5%.

Identification Using Data from Related Literature

I use data from Beshears et al. (2012) to estimate ϵ_h. That analysis studied differential responses to the default rate by income in three employer-sponsored retirement saving plans. They found that the low-paid were more likely to save at the default than the high-paid. Using their data, I investigate two groups of employees in the same firm who were assigned two different default rates. Firm C in their paper had a 3% default savings rate for 2,785 full-time employees hired at the firm between January 1, 2003 and February 29, 2004. The same firm C had a 5% default savings rate for 3,765 full-time employees hired between June 1, 2005 and July 31, 2006. Employers provided matching contributions in both time periods. The maximum employer match was 7%, meaning that employers matched employees' contributions up to 7% of their earnings if employees contributed 7% or more.

The key underlying assumption required to exploit the variation in default rates to estimate ϵ_h is that the characteristics of the two cohorts facing different default rates must be similar. This assumption largely holds based on the summary statistics provided by Beshears et al. (2012): the mean age for both groups was 33-34 years and the mean annual income was \$42,000 - \$44,000. Employees in Firm C on average earned more than eligible workers in OregonSaves whose average annual income is \$26,212.8 (in 2019 dollars) as shown in Table 6. Appendix C provides the distributions of employee savings rates at Firm C when the default rate was 3% and 5%. Based on Equation (3.8), I compute the change in the fraction of h-type passive savers $(m_{h'} - m_h)$ with respect to the fraction of all passive savers (m_r). Similar to the calculation of ϵ_l in Section 3.5.1, the increase in the fraction of h-type passive savers is the same size as the decrease in the fraction of h-type active savers. Data from Beshears et al. (2012) show a decrease by 11% of h-type active savers when the default rate increased from 3% to 5%. That is equivalent to a 11% increase in h-type passive savers: $m_{h'} - m_h = 11\%$. I also observe 32% total passive savers under the 3% default rate: $m_r = 32\%$. Plugging these numbers into Equation (3.8), I get:

$$\begin{aligned}\epsilon_h &= \frac{m_{h'} - m_h}{m_r} \cdot \frac{1}{r' - r} \\ &= \frac{11\%}{32\%} \cdot \frac{1}{5-3} \\ &= 0.17. \end{aligned} \tag{3.10}$$

The value of ϵ_h suggests that 17% of active savers would start saving at the default rate if the initial default rate increased by 1 percentage point. I can also use data from Beshears et al. (2012) to obtain an estimate for ϵ_l:

$$\begin{aligned}\epsilon_l' &= \frac{m_{l'} - m_l}{m_d} \cdot \frac{1}{r' - r} \\ &= \frac{-8\%}{32\%} \cdot \frac{1}{5-3} \\ &= -0.13. \end{aligned} \tag{3.11}$$

I find $\epsilon_l'(= -0.13)$ close to $\epsilon_l(= -0.12)$ estimated from the OregonSaves data in Section 3.5.1. One caveat of using any data from employer-sponsored retirement plans is that the estimates could be confounded by the employer matching cap. Specifically in firm C studied by Beshears et al. (2012), that firm offered matching up to 7%. As the default rate moveed closer to 7% (from 3% to 5%), employees were more likely to actively switch to 7% to take full advantage of the matching benefit than they would do without matching. Consequently, when the default rate is 5%, I expect more active savers with matching than without matching. Equivalently, I expect fewer passive savers with matching than without matching, which makes the observed $m_{h'}$ biased upwards and ultimately biases ϵ_h downwards in Equation (3.10).

3.5.2. Normative and Behavioral Time Preferences δ and λ

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 one method to experimentally elicit present-biased discount rates.

Estimation Strategy Using Survey Data

Besides the individual-level administrative records of OregonSaves savings data, I surveyed a subgroup of OregonSaves eligible workers in June 2019, including those who opted out and who were participating. I sent the survey to 441 workers and 143 responded (32.4% response rate). Survey respondents had two weeks to answer the survey through an email link and all respondents received a \$40 Starbucks gift card for completing the survey.

Our identification strategy, called the Convex Time Budget (CTB) approach, follows Andreoni and Sprenger (2012) to simultaneously estimate the time preferences $\lambda - \delta$ and the curvature of the utility function. Survey participants answered 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. They choose C tokens to receive at a sooner time and R tokens to receive at a later time continuously along a convex budget set:

$$(1+r)C + R = 100. \tag{3.12}$$

I used variations in starting times t to identify respondents' behavioral discount rates λ. I used variations in delay length k and interest rates $(1+r)$ to identify the normative discount rates δ and utility function curvature. 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 16 shows the time periods, token budgets, token unit values, and annual interest rates for all 16 combinations. Appendix D 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, participants had a budget of 100 tokens. 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 × 60), and the total dollar amount she would have in a year, \$4,000 (= \$100 × 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 (3.12).

Given consumption at a sooner time C and consumption at a later time R, I express decision utility U as a multi-period time separable CRRA (constant relative risk aversion) utility function subject to budget constraint (3.12):

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

The parameter α is the CRRA curvature parameter, λ is the behavioral time preference, δ is the normative time preference, and k is the delay length between the two time points. The variable W is background consumption which is the negative of the minimum consumption level in a typical year. 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:

$$\ln\Big(\frac{C-W}{R-W}\Big) = \Big(\frac{\ln\lambda}{\alpha-1}\Big)\mathbf{1}\{t=0\} + \Big(\frac{\ln\delta}{\alpha-1}\Big)k + \Big(\frac{1}{\alpha-1}\Big)\ln(1+r). \tag{3.13}$$

W is the negative of minimum annual consumption level asked in the survey. C and R are survey responses to the intertemporal allocation questions described in Appendix D; $\mathbf{1}\{t=0\}$ 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 16; and $\ln(1+r)$ is the natural log of the annual interest rate in Table 16. I use a two-limit Tobit maximum likelihood regression to estimate parameters λ, δ, and α.

I also estimate these parameters using an alternative utility function, constant absolute risk aversion (CARA). The decision utility U in this formulation subject to budget constraint (3.12) is expressed as:

$$U(C,R) = -\exp(-\rho C) - \lambda\delta^{k}\exp(-\rho R),$$

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

$$C - R = \left(\frac{\ln \lambda}{-\rho}\right)\mathbf{1}\{t = 0\} + \left(\frac{\ln \delta}{-\rho}\right)k + \left(\frac{1}{-\rho}\right)\ln(1 + r). \tag{3.14}$$

Results

Table 17 shows estimates of λ and δ from two-limit Tobit maximum likelihood regressions. There were 143 survey respondents who answered the time preference survey questions, and they made 1,765 intertemporal choices in total. Column 1 shows estimates of the CRRA regression (Equation (3.13)). The annual background consumption $w = -1,040$, equal to the negative of the minimum consumption level among all survey respondents. The average normative discount factor δ is 0.995 (standard deviation 0.006), and the average behavioral discount factor is λ is 0.987 (s.d. 0.005). Column 2 shows estimates of the CARA regression (Equation (3.14)). The average δ is 0.987 (s.d. 0.005) and the average λ is 0.993 (s.d. 0.007). For a baseline calculation of the optimal default rate, I assume that the normative time preference is the same for l- and h-type passive savers: $\delta_l = \delta_h = 0.995$. The behavioral time preference for h-type passive savers is assumed to be the average level under CRRA utility: $\lambda_h = 0.987$. The behavioral time preference for l-type passive savers is assumed to be one standard deviation below the average: $\lambda_l = 0.982$.

3.6. Computing the Optimal Default Savings Rate

The optimal default rate is computed by plugging the values listed in Table 19 into Proposition 1. An additional 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 weighting of different types of passive savers is unobserved. The optimal default rate r^* using baseline estimates can be computed as follows:

$$
\begin{aligned}
r^*\% &= \frac{dI + dS_l - dS_h + dK_l - dK_h}{dR_l - dR_h} \\
&= \frac{3120.1 + 2792.7 - 645.8 + 39.6 - 38.3}{698.2 - 71.8}\% \\
&= 8.4\%.
\end{aligned}
$$

The optimal default is higher than the current 5% default rate in OregonSaves mainly for two reasons. First, the fraction of passive savers accepting the optimal default rate could be overestimated. I use individual responses to auto-escalation to proxy how two identical groups of workers would respond to two initial default rates. Since the initial default rate is more salient than auto-escalation, passive savers are more likely to opt out of a high initial default rate compared to a low initial default than opting out of auto-escalation. Second, our estimates suggest that passive savers greatly benefit from saving at the default. The actual benefit of default savings could be lower than calculated because the current welfare framework does not take into account Social Security benefits. Additional retirement income from Social Security could diminish the marginal benefit of default savings. The combination of these two reasons implies that the actual social welfare gain from saving at the default could be lower than estimated. The 8.4% baseline calculation should therefore be considered as an upper bound of the optimal default rate.

3.6.1. Sensitivity Analysis

Table 20 reports estimates of the optimal default rate under alternative assumptions. I first consider alternative assumptions about individual time preferences. The second row of Table 20 reports that, if individuals with a low preferred savings rate are highly present-biased, the optimal default rate (8.2%) would be lower than the baseline value (8.4%). This is mostly because the policymaker wants to lower the default rate to attract present-biased individuals who would actively undersave from opting out the default. The third row shows the alternative optimal default rate if individuals with a low preferred savings rate do not highly value personal savings due to some normative reasons. These normative reasons could be a second retirement income source from social insurance programs or family members. The third row reports that the optimal default rate (8.5%) is slightly higher than the baseline value (8.4%) when normative reasons cause individuals to elect a low preferred savings rate. When the default rate is relatively high, opting out of the default rate and actively saving less is not necessarily welfare reducing for individuals with a second retirement income source. Moreover, a higher default rate could be welfare improving for passive savers who stay at the default rate.

The fourth row considers an alternative assumption about the (observed) semi-elasticity of the fraction of passive savers with respect to the default rate. The default rate remains the same (8.4%) when passive savers are less responsive to the default rate by opting out of the default rate. The fifth row reports the optimal default rate when the average cost of opting out of the default rate perceived by individuals is large. Recent studies such as Bernheim et al. (2015) find that the perceived opt-out cost could be in thousands of dollars. The fifth row suggests that the optimal default rate is stable around 8.4% despite a wide range for perceived opt-out costs.

3.7. 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 that maximizes their life cycle utility. In contrast, if the non-default preferred rate that individuals actively choose does not maximize their life cycle 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, about an additional 12% 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 8%, 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 income transfer programs and nudge policy.

Table 1: Summary Statistics by Age for SIPP Sample

	20-40 Years Old		41-60 Years Old		Full Sample	
	mean	sd	mean	sd	mean	sd
Panel A: Pooled						
Unemployment duration (weeks)	17.3	13.7	19.7	14.8	18.1	14.2
Weekly unemployment benefits ($)	153.2	47.1	187.8	47.0	165.2	49.9
Individual replacement rate	50%	0.07	50%	0.09	50%	0.08
Age	30.3	5.7	48.7	5.6	36.7	10.4
Years of education	12.1	2.9	12.1	3.3	12.1	3.0
Percent married	50%	0.5	80%	0.4	60%	0.5
Pre-unemployment annual wage ($)	18,797.4	11,874.9	24,892.5	15,480.5	20,915.4	13,552.3
State unemployment rate	7%	0.02	7%	0.02	7%	0.02
Number of unemployment spells	2,858		1,522		4,380	
	20-40 Years Old		41-60 Years Old		Full Sample	
	mean	sd	mean	sd	mean	sd
Panel B: State Unemployment Rate above Median						
Unemployment duration (weeks)	18.3	14.2	21.0	15.2	19.2	14.6
Weekly unemployment benefits ($)	153.6	49.1	188.2	49.1	165.3	51.8
Individual replacement rate	50%	0.08	50%	0.09	50%	0.09
Age	30.4	5.6	48.8	5.6	36.6	10.3
Years of education	12.1	2.9	12.1	3.6	12.1	3.2
Percent married	50%	0.5	70%	0.4	60%	0.5
Pre-unemployment annual wage ($)	19,003.9	12,480.5	25,556.7	15,694.2	21220.7	13996.5
State unemployment rate	8%	0.01	8%	0.01	8%	0.01
Number of unemployment spells	1,344		682		2,016	
	20-40 Years Old		41-60 Years Old		Full Sample	
	mean	sd	mean	sd	mean	sd
Panel C: State Unemployment Rate below Median						
Unemployment duration (weeks)	16.4	13.3	18.6	14.4	17.2	13.7
Weekly unemployment benefits ($)	152.9	45.3	187.5	45.3	165.1	48.2
Individual replacement rate	50%	0.07	50%	0.09	50%	0.08
Age	30.2	5.8	48.6	5.5	36.8	10.5
Years of education	12.0	2.8	12.2	3.0	12.1	2.9
Percent married	50%	0.5	80%	0.4	60%	0.5
Pre-unemployment annual wage ($)	18,616.6	11,319.4	24,353.2	15,292.9	20,655.0	13,159.1
State unemployment rate	5%	0.009	5%	0.010	5%	0.010
Number of unemployment spells	1,524		840		2,364	

Notes: The data are individual-level unemployment spells from the 1985-2000 SIPP data set. All dollar values are in real 1990 dollars. See the main text for more details.

Table 2: Replication of Kroft and Notowidigdo's Baseline Results

	Dependent variable: log(hazard rate) = −log(duration)	
	20-60 years old	
	KN's baseline results (1)	My replication (2)
Demeaned log(average UI benefit)	-0.632* (0.332)	-1.518** (0.652)
Demeaned log(state unemployment rate)	0.035 (0.124)	0.053 (0.122)
Demeaned log(average UI benefit) x Demeaned log(state unemployment rate)	1.346*** (0.457)	1.381*** (0.480)
Number of spells	4307	4380
Implied elasticities:		
High unemployment elasticity (u=8.4%)	-0.277 (0.364)	-0.796 (0.586)
Low unemployment elasticity (u=5%)	-0.987*** (0.343)	-2.239*** (0.795)

Notes: Data are individual-level unemployment spells from 1985-2000 SIPP. All columns report semiparametric Cox hazard model from estimating equation (1.1). Column 1 of Table 2 reports Kroft and Notowidigo's baseline model estimates (see Column 1 of table 2 in Kroft and Notowidigdo (2016)). Column 2 reports my own estimates of equation (1.1). The average UI benefit amount is the average weekly benefit paid to claimants in a given state-year. Both specifications include state, year, industry, and occupation fixed effects, a 10-piece log-linear spline for the claimant's pre-unemployment wage, age, education, dummies for marital status and being on the seam between interviews, and interaction between year fixed effects and log of the UI benefit amount. The final two rows report linear combination of parameter estimates to produce the duration elasticity when the state unemployment rate is one standard deviation above or below the median. Standard error are reported in parentheses and are clustered by states. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3: Average Effect of UI Benefits on Unemployment Duration by Age

	Dependent variable: log(hazard rate) = −log(duration)					
	All		20-40 years old		41-60 years old	
	Baseline (1)	Interaction (2)	Baseline (3)	Interaction (4)	Baseline (5)	Interaction (6)
Demeaned log(average UI benefit)	-0.603**	-1.518**	-0.375	-1.225**	-1.340***	-2.923**
	(0.277)	(0.652)	(0.334)	(0.614)	(0.495)	(1.386)
Demeaned log(state unemployment rate)		0.053		0.140		-0.085
		(0.122)		(0.164)		(0.315)
Demeaned log(average UI benefit) x Demeaned log(state unemployment rate)		1.381***		1.653**		0.397
		(0.480)		(0.764)		(1.561)
Number of spells	4,380	4,380	2,858	2,858	1,522	1,522
Implied elasticities:						
High unemployment elasticity (u=8.4%)		-0.796		-0.361		-2.716*
		(0.586)		(0.558)		(1.501)
Low unemployment elasticity (u=5%)		-2.239***		-2.088**		-3.130*
		(0.795)		(0.872)		(1.709)

Notes: Data are individual-level unemployment spells from 1985-2000 SIPP. All columns report semiparametric Cox hazard model from estimating equation (1.1). The average UI benefit amount is the average weekly benefit paid to claimants in a given state-year. Column 1, 3, and 5 report elasticities of hazard rate with respect to UI benefits for the full sample and for each age group. The estimates of the interaction term in column 2, 4, and 6 report the effect of unemployment rate on duration elasticity with respect to benefits for the full sample and for each age group. All columns include state, year, industry, and occupation fixed effects, a 10-piece log-linear spline for the claimant's pre-unemployment wage, age, education, dummies for marital status and being on the seam between interviews, and interaction between year fixed effects and log of the UI benefit amount. The final two rows report linear combination of parameter estimates to produce the duration elasticity when the state unemployment rate is one standard deviation above or below the median. Standard error are reported in parentheses and are clustered by states. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 4: Age-Specific Effect of UI Benefits on Unemployment Duration by Age

	Dependent variable: log(hazard rate) = −log(duration)					
	All		20-40 years old		41-60 years old	
	Baseline	Interaction	Baseline	Interaction	Baseline	Interaction
	(1)	(2)	(3)	(4)	(5)	(6)
Demeaned log(age-specific average UI benefit)	-0.336*	-0.735	-0.191	-0.513	-1.363***	-2.220**
	(0.204)	(0.509)	(0.248)	(0.488)	(0.456)	(1.046)
Demeaned log(state unemployment rate)		-0.064		0.161		-0.507
		(0.113)		(0.172)		(0.354)
Demeaned log(age-specific average UI benefit) x Demeaned log(state unemployment rate)		0.646*		1.729**		1.022
		(0.365)		(0.799)		(1.201)
Number of spells	4,380	4,380	2,858	2,858	1,522	1,522
Implied elasticities:						
High unemployment elasticity (u=8.4%)		-0.398		0.390		-1.686
		(0.467)		(0.452)		(1.159)
Low unemployment elasticity (u=5%)		-1.073*		-1.417*		-2.753**
		(0.611)		(0.788)		(1.278)

Notes: Data are individual-level unemployment spells from 1985-2000 SIPP. All columns report semiparametric Cox hazard model from estimating equation (1.1). The age specific UI benefit amount is the average weekly benefit of each age group in a given state-year. Column 1, 3, and 5 report elasticities of hazard rate with respect to UI benefits for the full sample and for each age group. The estimates of the interaction term in column 2, 4, and 6 report the effect of unemployment rate on duration elasticity with respect to benefits for the full sample and for each age group. All columns include state, year, industry, and occupation fixed effects, a 10-piece log-linear spline for the claimant's pre-unemployment wage, age, education, dummies for marital status and being on the seam between interviews, and interaction between year fixed effects and log of the UI benefit amount. The final two rows of each panel report linear combination of parameter estimates to produce the duration elasticity when the state unemployment rate is one standard deviation above or below the median. Standard error are reported in parentheses and are clustered by states. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 5: Extended Model

	Dependent variable: log(hazard rate) = $-$log(duration)	
	20-60 years old	
	Three-way interaction	Two-way interaction
Demeaned log(average UI benefit)	-1.563**	-1.193**
	(0.639)	(0.578)
Demeaned log(state unemployment rate)	0.060	
	(0.119)	
Demeaned log(average UI benefit) x Demeaned log(state unemployment rate)	1.421***	
	(0.538)	
1{41-60 years old}	-0.258***	-0.255***
	(0.034)	(0.033)
1{41-60 years old} x Demeaned log(average UI benefit)	0.113	0.135
	(0.191)	(0.188)
1{41-60 years old} x Demeaned log(state unemployment rate)	-0.048	
	(0.137)	
1{41-60 years old} x Demeaned log(average UI benefit) x Demeaned log(state unemployment rate)	-0.602	
	(0.957)	

Notes 1: Data are individual-level unemployment spells from 1985-2000 SIPP. Column 1 reports estimates from equation (1.2). The last coefficient of column 1 is the estimate of the coefficient of a three-way interaction between unemployment rate, UI benefits, and age dummy. The three-way interaction coefficient is not significant, which suggests that the effect of unemployment rates on duration elasticity is not significantly different between older and younger workers. Column 2 report estimates from equation (1.3). The last coefficient of column 2 is the estimate of a 2-way interaction between UI benefits and age dummy. The two-way interaction coefficient is not significant, which suggests that the average duration elasticity is not significantly different between older and younger workers. Both columns include state, year, industry, and occupation fixed effects, a 10-piece log-linear spline for the claimant's pre-unemployment wage, education, dummies for marital status and being on the seam between interviews, and interaction between year fixed effects and log of the UI benefit amount. Standard error are reported in parentheses and are clustered by states. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 6: Summary Statistics for Individuals Ever Had Access to OregonSaves, June 2019

	N	%	$
Panel A: All individuals			
Total unique individuals entered by employers	171,243	100.0	–
Immediate opted-out individuals	41,747	24.4	–
Delayed opted-out individuals	21,600	12.6	–
Pending individuals	29,332	17.1	–
Enrolled individuals w/o payroll info	12,630	7.4	–
Enrolled individuals with payroll info	65,924	38.5	–
Panel B: Eligible active workers (EAW)			
Total EAW	76,438	100.00	–
Immediate opted-out workers	27,743	36.3	–
Delayed opted-out workers	17,122	22.4	–
EAWs with no balance	6,793	8.9	–
Suspended contributors	1,277	1.7	–
Contributors	23,503	30.7	–
Average monthly contributions if > 0, June 2019	–	–	110
Average monthly income	–	–	2,184

Note: Data from anonymized administrative records as of June 29, 2019. 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 June 29, 2019. 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 June 2019. Contributors are EAWs with a positive balance and positive monthly contributions in June 2019.

Table 7: Summary Statistics for Individuals with a Positive Account Balance, June 2019

	N	$
Panel A: All individuals with a positive balance		
All individuals with a positive balance	40,652	–
Opted-out individuals with a positive balance	3,409	–
Participating individuals with a positive balance	37,243	–
Average balance if positive	–	558
Total assets	–	22.7 million
Panel B: Eligible active workers (EAW) with a positive balance		
EAWs with a positive balance	28,083	–
Opted-out EAWs with a positive balance	3,303	–
Participating EAWs with a positive balance	24,780	–
Average balance if positive	–	653

Note: Data from anonymized administrative records on June 29, 2019. Panel A reports statistics for all individuals ever had access to OregonSaves with a positive balance on June 29, 2019. Opted-out individuals with a positive balance are persons who opted out of the program before June 29, 2019 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 June 29, 2019, had ever contributed, and did not withdraw all contributions. Panel B presents statistics for eligible active workers (EAW) with a positive balance on June 29, 2019. EAWs are persons eligible for contributions (at least one employer) and inferred to be actively working on June 29, 2019. 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 8: Summary Statistics of Firms for Eligible Active Workers (EAW), June 2019

	N. employers	Avg firm size	N. EAWs	Avg participation rate (%)
Agriculture	88	37	3,856	23.6
Arts & entertainment	58	29	2,187	43.4
Business support	89	70	7,901	39.1
Construction	96	18	2,153	37.9
Education	48	24	1,499	41.2
Finance & insurance	2	2	4	50.0
Food services	499	38	26,787	46.4
Health care	159	40	8,860	43.6
Information	17	13	284	38.7
Management	5	45	153	58.2
Manufacturing	132	35	5,389	38.5
Other services	101	26	3,080	42.1
Professional & scientific	48	30	1,503	31.7
Real estate	38	21	867	34.3
Retail trade	172	29	6,773	37.3
Transportation & storage	30	26	1,209	34.7
Wholesale trade	32	40	1,330	28.0
Not specified	53	34	2,603	46.9
Total	1,667	35	76,438	41.3

Note: Data from anonymized administrative records as of June 29, 2019. Firm size is measured by the number of active employees. Average participation rate in each industry equals the number of eligible active participants (EAW with no balance + EAW suspended contributors + EAW contributors) / the number of eligible active workers in each industry. 4,970 employers have registered their business in OregonSaves. Median firm size is 16. Mean firm size is 31. 1,667 firms have at least 1 eligible active workers.

Table 9: Average Participation Rates by Firm Size and Industry (%), June 2019

Industry \ Firm size	100+	50-99	20-49	10-19	5-9	1-4	All
Agriculture	16.1	22.4	27.3	31.2	46.4	82.4	23.6
Arts/entertainment	46.5	45.8	43.2	32.5	50.6	70.8	43.4
Business support	42.1	26.8	32.8	39.7	44.1	70.3	39.1
Construction	45.2	20.9	34.7	47.9	45.1	71.8	37.9
Education	45.8	43.5	39.1	41.9	37.1	55.6	41.2
Finance/insurance						50.0	50.0
Food services	49.0	41.2	44.1	48.1	54.4	68.6	46.4
Health care	48.4	36.7	42.8	37.8	52.3	75.0	43.6
Information		39.7	34.3	43.5	20.0	53.3	38.7
Management	61.4		50.0		40.0	25.0	58.2
Manufacturing	43.4	33.2	35.6	35.2	38.5	74.4	38.5
Other services	48.9	38.4	40.2	37.5	43.7	85.3	42.1
Professional/scientific	25.5	31.2	27.8	38.7	46.2	74.1	31.7
Real estate	23.1	44.9	30.8	23.8	61.4	75.0	34.3
Retail trade	35.6	39.8	33.9	37.9	57.9	79.0	37.3
Transportation/storage	28.1	38.3	28.7	50.8	35.5	71.4	34.7
Wholesale trade	43.2	13.5	20.8	26.7	58.1	0.0	28.0
Not specified	48.9	38.5	46.7	45.7	71.0	75.0	46.9
Total	43.8	36.9	38.7	41.4	50.0	72.3	41.3

Note: Data from anonymized administrative records as of June 29, 2019. Average participation rate equals the number of eligible active participants (EAW with no balance + EAW suspended contributors + EAW contributors) / the number of eligible active workers in each industry and each firm size category.

Table 10: Distribution of OregonSaves Contribution Rates for Eligible Active Workers (EAW), June 2019

Contribution rate (%)	N. EAW	Percent of EAW (%)
0	52,852	69.1
1	512	0.7
2	546	0.7
3	824	1.1
4	181	0.2
5	16,875	22.1
6	4,038	5.3
7	80	0.1
8	90	0.1
9	14	0.0
10	332	0.4
>10	94	0.1
Total	76,438	100

Note: Data from anonymized administrative records on June 29, 2019. The contribution rate refers to the average contribution rate of all current employers where employees are eligible and active workers. These include employees who have opted out in the zero contribution rate bin if they are EAW. About 22.1% of EAW had contribution rates of 5%. About 5.3% had contribution rates of 6%, a large fraction of which can be attributed to the automatic escalation feature of the plan. On January 1, 2019, workers who had opened their accounts for six months were eligible for auto-escalation. The rates automatically increased by 1 percent unless workers actively chose to opt out of the auto-escalation arrangement.

Table 11: Marginal Effects of Participation by Eligible Active Workers (EAW), June 2019

Participation decision (=1 if participating, =0 not)				
	(1)	(2)	(3)	(4)
	all EAW	EAW working for 1 ER		
Age 18-25	0.22***	0.22***	0.22***	0.19***
	(0.02)	(0.02)	(0.02)	(0.02)
Age 26-45	0.20***	0.19***	0.20***	0.18***
	(0.02)	(0.02)	(0.02)	(0.02)
age 46-65	0.15***	0.15***	0.15***	0.15***
	(0.02)	(0.02)	(0.02)	(0.02)
Working >1 ER	0.18***			
	(0.01)			
Had OS account prior to current ER(s)	4.38***	4.46***	4.26***	4.44***
	(0.06)	(0.04)	(0.03)	(0.15)
Ln(firm-level monthly income)			-0.00	0.02
			(0.01)	(0.01)
Ln(firm size)			0.01	0.01
			(0.01)	(0.01)
Agriculture				-0.04
				(0.06)
Arts/Entertainment				0.16***
				(0.05)
Business/Support				0.10*
				(0.06)
Construction				0.10*
				(0.06)
Education				0.16***

				(0.05)
Food services				0.17***
				(0.05)
Health care				0.16***
				(0.05)
Information				0.13**
				(0.06)
Management				0.29***
				(0.05)
Manufacturing				0.10*
				(0.05)
Other services				0.15***
				(0.05)
Professional/scientific				0.02
				(0.06)
Real estate				0.07
				(0.06)
Retail trade				0.10**
				(0.05)
Transportation/storage				0.07
				(0.06)
N	71,335	68,258	66,808	66,808
Pseudo R-squared	0.042	0.038	0.037	0.046
Mean of Dep Var.	0.41	0.40	0.40	0.40
SD of Dep Var.	0.49	0.49	0.49	0.49

Note: Sample includes eligible active workers (EAW) at firms with $\geq$ 10 employees and nonmissing industry. Reference categories are: Age 66-100, first-time access to OregonSaves through current ER(s), wholesale trade (industry fixed effect), and working for 1 employer (only for column 1). Standard errors in parentheses, clustered by firm. Coefficient significant at * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 12: Reasons for Opting Out Provided by Eligible Active Workers (EAW), June 2019

Panel A: Main reasons	N	%
I can't afford to save at this time	13,142	29.3
I don't qualify for a Roth IRA due to my income	214	0.5
I don't trust the financial markets	1,230	2.7
I have my own retirement plan	9,236	20.6
I would prefer a Traditional IRA	488	1.1
I'm not interested in contributing through this employer	6,468	14.4
I'm not satisfied with the investment options	773	1.7
Other	11,269	25.1
Did not specify	2,045	4.6
Total	44,865	100.0

Panel B: Sample explanations of "Other" reasons employees opted out

Characterization of responses	% of other	Verbatim quote
Left employment	Large	Quit job
Not interested	Large	Nunca
Retiring soon or already	Large	85 years old
Anti government	Noticeable	Babylon is falling. One Love. One Heart. What was built on the sand will not stand.
Anti opt out plan	Noticeable	Because you have no g#$am right to automatically sign me up for this bulls***.
Anti-social	Noticeable	not your dam business
Confused by plan	Noticeable	Because I don't want the government's ROTH IRA. Its going to be terrible compared to what I could get for the price with another competitor.
Fees are too high	Noticaeble	1.01% return on my money. but a 1% yearly fee....no thanks

Note: Data from anonymized administrative records as of June 29, 2019. Opted-out workers (N=44,865) include eligible active workers who opted out during the first 30-day enrollment window through all employers (immediate opted-out workers, N=27,743) and those who opted out anytime through all employers (delayed opted-out workers, N=17,122).

Table 13: State Legislation Establishing a State-Sponsored Retirement Plan, October 2019

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
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 10 states that have passed the legislation for a voluntary or a mandatory program, about another 21 states have introduced legislation but not yet enacted. AARP summarized the status of these 21 states: `https://www.aarp.org/ppi/state-retirement-plans/savings-plans/`.

Table 14: Summary Statistics for Individuals Eligible for Automatic Escalation, June 2019

Panel A: All individuals eligible for auto escalation	N	%
Opted out of auto increase before notification on Dec 1, 2018	505	6.8
Opted out of auto increase after notification before taking into effect on Jan 1, 2019	781	10.6
Rate auto increased, actively opted out of the program end of June 2019 (6 months after auto increase)	2,217	30.0
Rate auto increased, then actively lowered rate end of June 2019	71	1.0
Rate auto increased, then actively raised rate by more than 1 percent end of June 2019	46	0.6
Rate auto increased by 1 percent end of June 2019	3,781	51.1
Total	7,401	100.0
Panel B: EAW eligible for auto escalation	N	%
Opted out of auto increase before notification on Dec 1, 2018	410	7.2
Opted out of auto increase after notification before taking into effect on Jan 1, 2019	776	13.6
Rate auto increased, actively opted out of the program end of June 2019 (6 months after auto increase)	610	10.7
Rate auto increased, then actively lowered rate end of June 2019	71	1.2
Rate auto increased, then actively raised rate by more than 1 percent end of June 2019	46	0.8
Rate auto increased by 1 percent end of June 2019	3,781	66.4
Total	5,694	100.0

Note: Data from anonymized administrative records as of June 29, 2019. On January 1, 2019, workers who had accounts open for six months were eligible for auto-escalation. Additionally, workers who initially elected any non-zero contribution rate (default or non-default) were eligible for auto-escalation. Contribution rates are automatically increased by 1 percent until they reached 10%, every year on January 1 for all eligible workers. Auto-escalation eligibles may actively opt out of the auto-escalation arrangement any time. Panel A shows subgroups of individuals eligible for auto-escalation. About 6.8% of them opted out of the auto-escalation option before the OregonSaves administrator sent a notification a month before auto escalation took into effect (Dec 1, 2018). About 10.6% opted out after they received the notification, and before it took into effect. About 30% opted out of the OregonSaves program at the end of six months after auto-escalation occurred (June 30, 2019). One percent lowered their rates while still contributing at the end of June 30, 2019; 0.6% raised their rates at the end of June 30, 2019; and 51.1% were unresponsive to auto-escalation. Panel B shows how eligible active workers eligible for auto-escalation responded. 66.4% accepted the auto-escalation arrangement.

Table 15: Distribution of Contribution Rates for Eligible Active Workers (EAW) Eligible for Automatic Escalation Before and After Automatic Escalation

	Before: Nov 2018		After: June 2019	
Contribution rate (%)	N	%	N	%
0	–	–	409	9.1
1	81	1.8	16	0.4
2	85	1.9	81	1.8
3	110	2.4	97	2.2
4	22	0.5	69	1.5
5	4,142	91.9	97	2.2
6	18	0.4	3,632	80.6
7	17	0.4	26	0.6
8	8	0.2	29	0.6
9	3	0.1	6	0.1
10	19	0.4	34	0.8
>10	3	0.1	12	0.3
Total	4,508	100.0	4,508	100.0

Note: Data from anonymized administrative records as of June 29, 2019. This table identifies the fraction of active and eligible workers (EAW) eligible for automatic escalation who did not opt out of the auto-escalation arrangement before it took into effect. There were 4,508 of EAWs included in this table, equal to the total EAWs eligible for auto-escalation (N=5,694 in Panel B of Table 14) minus EAWs eligibles who opted out of auto-escalation before it occurred (N = 1,186 = 410 + 776, first two rows in Panel B of Table 14, so that 4,508 = 5,694 - 1,186). November 2018 was the last month unaffected by auto-escalation. The OregonSaves administrator notified participants eligible for auto-escalation on December 1, 2018. Auto-escalation happened on January 1, 2019. The contribution rate refers to the average contribution rate of current employers where employees are eligible and active workers. Columns 2-3 present that, at the end of November 2018, 91.9% saved at the initial 5% default rate. Columns 3-4 show that, at the end of June 2019, 80.6% saved at the new 6% default rate in June 2019.

Table 16: Choice Sets to Identify Time Preferences from Survey Responses

Start date t (unit: year)	Delay length k (unit: 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

Note: 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 D. Parameters of interest are identified using regression models specified in Equations (3.13) and (3.14). Estimation results are presented in Table 17.

Table 17: Parameter Estimates of Time Preferences and Utility Function Curvature

	(1) Estimates from Eq.(3.13)	(2) Estimates from Eq.(3.14)
Normative time preference δ	0.995	0.987
	(0.006)	(0.005)
Behavioral time preference λ	0.987	0.993
	(0.005)	(0.007)
CRRA curvature: α	0.501	
	(0.089)	
CARA curvature: ρ		2.033
		(0.374)
Observations	1,765	1,765
N. unique subjects	143	143

Note: Data from anonymized survey responses collected in June 2019. An online experimental survey was sent to 441 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 D, and these respondents made 1,765 intertemporal decisions in total. Both columns present estimation results from two-limit Tobit maximum likelihood regressions. Column 1 shows estimates of the regression specification in the form of Equation (3.13) assuming constant relative risk aversion utility (CRRA). 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 normative discount factor δ under CRRA is 0.995, and the average behavioral discount factor β under CRRA is 0.987. Column 2 shows estimates of the regression specification in the form of Equation (3.14) assuming constant absolute risk aversion utility (CARA). The average δ under CARA is 0.987 and the average β under CARA is 0.993. Standard deviations are in parentheses.

Table 18: Social Marginal Welfare Weight g Calculations

	l-type savers	h-type savers
Average annual income Z_s	\$24,487	\$36,257
Percent of type h_s	71.8%	6%
Primitive Pareto weight $\alpha_s = \frac{1}{Z_s}$	0.000041	0.000028
Aggregate weighted Pareto weight $\bar{\alpha} = \sum_{s=\{l,h\}} \alpha_s h_s$	0.000031	0.000031
Social marginal welfare weight $g_s = \frac{\alpha_s}{\bar{\alpha}}$	1.32	0.90

Notes: This table reports estimates of the social marginal welfare weights for l-type passive savers (preferred rates below the default) and for h-type passive savers (preferred rates above the default). The welfare weight for a given type g_s 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. These calculations are based on two empirical assumptions. First, we use observed data when the default rate is 5% to estimate the welfare weights at the optimal default. Second, statistics on annual income and the percent of type for l-type passive savers are inferred from the average level of all savers who elected a rate below the default; statistics for h-type passive savers are inferred by the average level of all savers who elected a rate above the default. The income information for each type Z_s is imputed from the OregonSaves savings data in June 2019, where 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 taken into account due to the limitation of the imputation calculation. Imputed average annual income equals the average monthly income times 12. Following Saez (2002), the third row shows that the primitive Pareto weight α_s equals the inverse of income $\frac{1}{Z_s}$. 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.

Table 19: Baseline Optimal Default Contribution Rate Calculations

Statistics	Values
Panel A: Statistics for l-type passive savers	
Semi-elasticity ϵ_l	-0.12
Normative time preference δ_l	0.995
Behavioral time preference β_l	0.982
Annual income Z_l	\$24,487
Social marginal welfare weight g_l	1.32
Preferred rate of passive savers on the margin s_l	0.04
Panel B: Statistics for h-type passive savers	
Semi-elasticity ϵ_h	0.17
Normative time preference δ_h	0.995
Behavioral time preference β_h	0.987
Annual income Z_h	\$36,257
Social marginal welfare weight g_h	0.90
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^*	8.4%

Notes: Estimates of key statistics used to compute the optimal default contribution rate in Proposition 1: All statistics in Panel A and Panel B are estimated from the OregonSaves data (see text) except that ϵ_h uses data from Beshears et al. (2012). Estimates for δ_l, β_1, δ_h, and β_h are identified using survey data collected from OregonSaves-eligible workers in Table 17 (see Section 3.5.2). Estimation procedures for g_l and g_h are are provided in Table 18. 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 3.6.

Table 20: Optimal Default Contribution Rate Under Alternative Assumptions

	Optimal default contribution rate (%)
Baseline	8.4
High present bias ($\beta_l = 0.45$)	8.2
Low long-run discount factor ($\delta_l = 0.7$)	8.5
Low elasticity to the default rate ($\epsilon_h = 0.12$)	8.4
Large perceived opt-out cost ($K = \$1,000$)	8.4

Notes: This table reports the optimal default contribution rate r^*, as computed using the sufficient statistics formula in Proposition 1 under different assumptions. The first row is the baseline calculation, which uses the estimates of statistics displayed in Table 19. The second row reports the optimal default rate under the assumption that individuals electing a low preferred rate are very present biased ($\beta_l = 0.45$ lower than the baseline value). The third row reports the optimal default rate under the assumption that individuals electing a low preferred rate have a low long-run discount factor ($\delta_l = 0.7$). The fourth row assumes that individuals with a high preferred rate are less responsive to the default rate than those in the baseline calculation ($\epsilon_h = 0.12$). The last row assumes that the average perceived cost of opting out of the default rate is higher than the value in the baseline calculation ($K = \$1,000$).

Figure 1: Survival Curves under High/Low Unemployment Rate by Age

(a) High unemployment

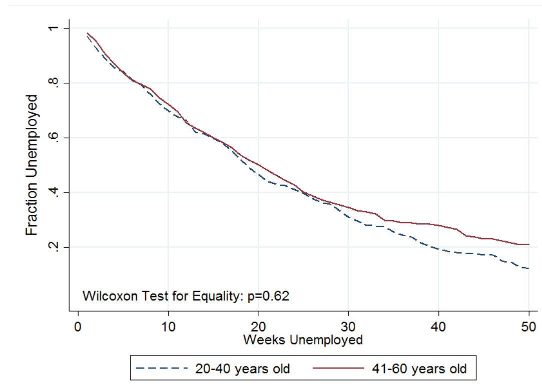

(b) Low unemployment

Note: Data are taken from individual-level unemployment spells from 1985-2000 SIPP. In Figure 1a, the sample includes spells in states with unemployment rates above the median across states. See Panel B of Table 1 for the descriptive statistics of the above-median sample. In Figure 1b, the sample includes spells in states with unemployment rates below the median across states. See Panel C of Table 1 for the descriptive statistics of the below-median sample. Each figure plots Kaplan-Meier survival curves for two age groups. The survival curves are adjusted for a seam effect by fitting a Cox model with a seam dummy and recovering baseline hazards as in Chetty (2008).

Figure 2: Total OregonSaves Participant Assets by Employee Zip Code, June 2019

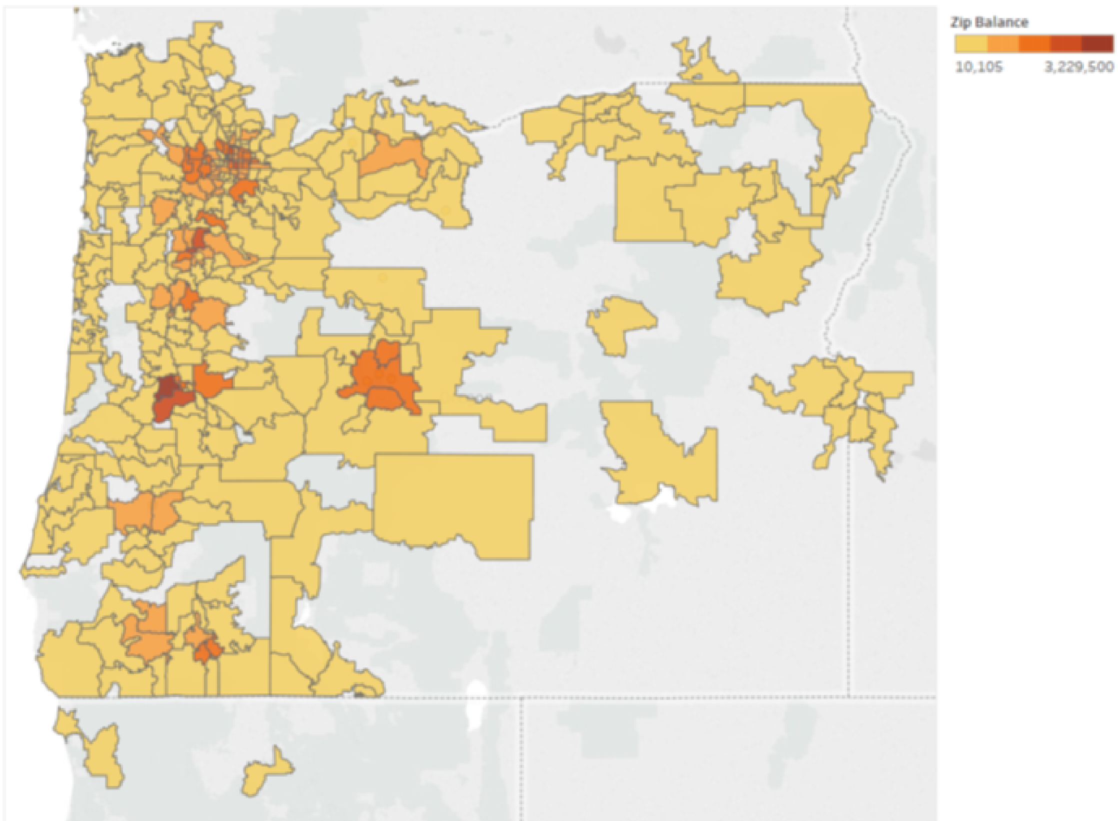

Figure 3: Impact of a Marginal Perturbation of the Default Savings Rate r

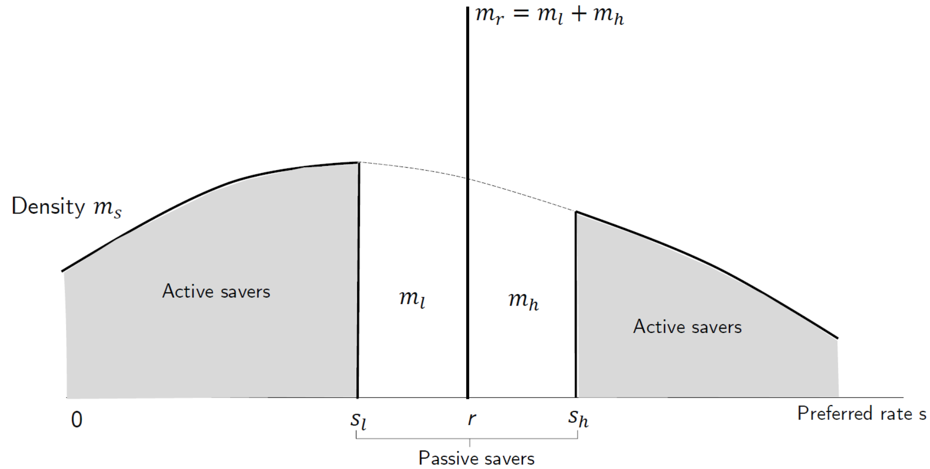

(a) In this figure, r is the default savings rate. Workers with a preferred savings rate between s_l and s_h save at the default rate r because the default is close to their preferred rates. These workers are defined as passive savers, with density $m_r = m_l + m_h$, where m_l are the fraction of passive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

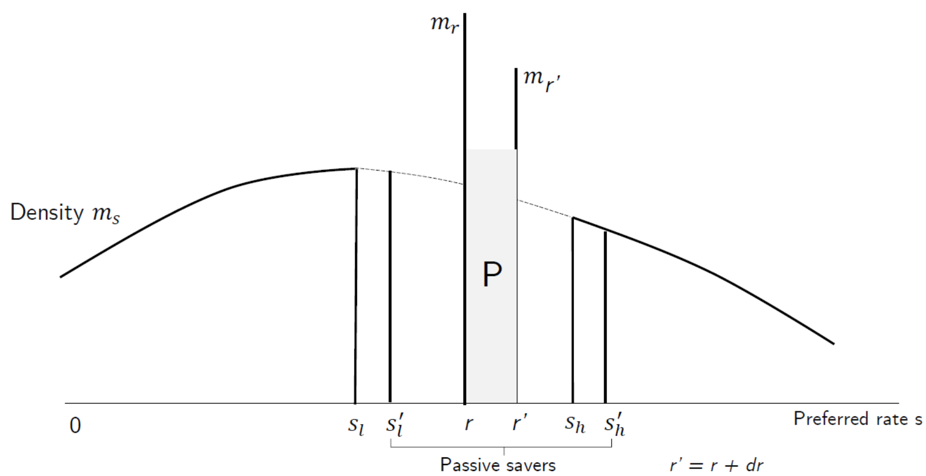

(b) Suppose the policymaker sets a default rate at r' instead of r, where $r^{'} = r + dr$. When the default rate is r', passive savers are workers with a preferred savings rate between $s_l^{'}$ and $s_h^{'}$. Their preferred rates are close to the new default rate r'. The total density of passive savers 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 3: Impact of a Marginal Perturbation of the Default Savings Rate r (Continued)

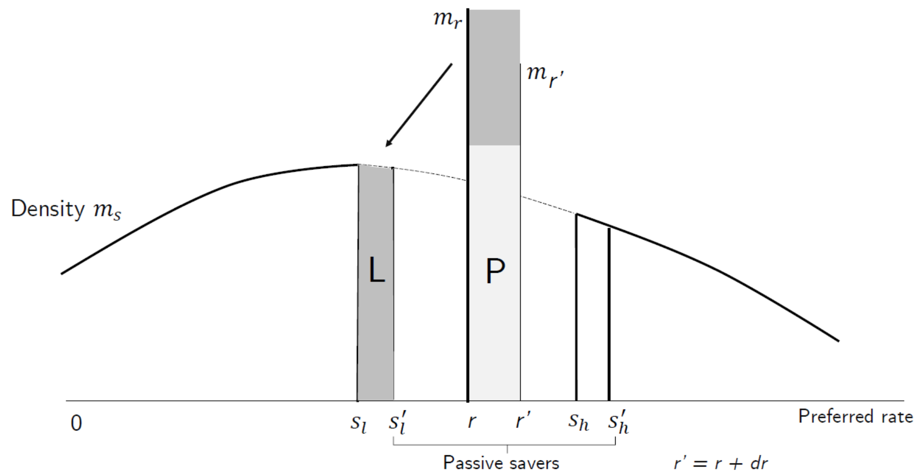

(c) The shaded area, denoted by L, indicates the fraction of workers who are passive savers under the low default rate r but active savers under the high default rate r'. These are workers with a preferred rate between s_l and $s_l^{'}$. They are passive savers under the low default r because it is close 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.

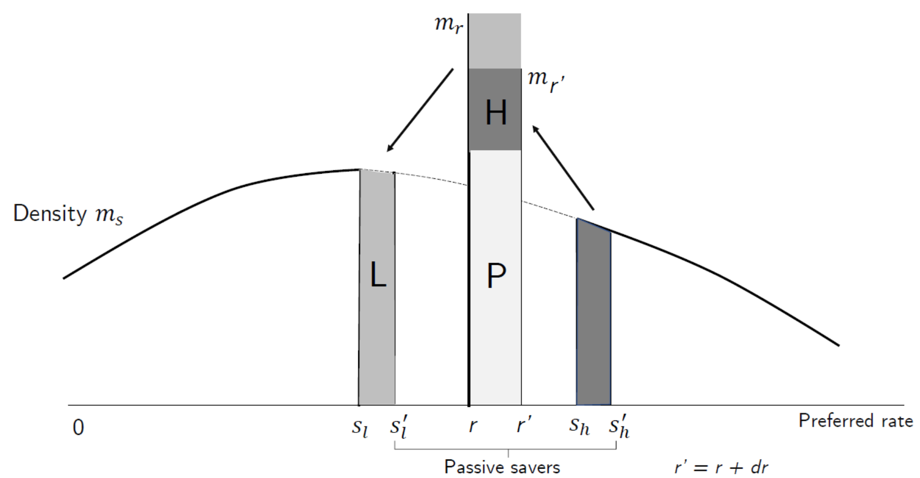

(d) The shaded area, denoted by H, shows the fraction of workers who are active savers under the low default rate r but passive savers under the high default rate r'. These are workers with a preferred rate between s_h and $s_h^{'}$. They are active savers under the low default r because the low default 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.

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:

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

where Z is the labor income, λ is the behavioral time preference, δ is the normative time preference, and K is the perceived cost of making an active decision. See Section 3.4 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 $R(\theta) = rZ(\theta)$ and her preferred saving amount $S(\theta)$. Her observed choice of saving amount in the presence of a default rate r maximizes Equation (3.1) in Section 3.4.

B. Proof of Proposition 1

The first-order condition for the social welfare function, Equation (3.3), equals zero at the optimal default rate r^*:

$$\begin{aligned}\frac{dW(r^*)}{dr} &= \frac{d}{dr}\int_{\theta=\theta_l}^{\theta_h} \alpha(\theta)N(P(\theta))dm(\theta) \\ &\approx \int_{\theta=\theta_l}^{\theta_h} \alpha(\theta)\frac{dN(R(\theta))}{dr}dm(\theta) + \frac{dm_h}{dr}\alpha_h(N(R_h) - N(\theta_h)) - \frac{dm_l}{dr}\alpha_l(N(\theta_l) - N(R_l)) \\ &= 0.\end{aligned} \tag{A.2}$$

The first term in Equation (A.2) can be decomposed into two terms:

$$
\begin{aligned}
&\int_{\theta=\theta_l}^{\theta_h} \alpha(\theta)\frac{dN(R(\theta))}{dr}dm(\theta) \\
\approx &\int_{\theta_l}^{r^*} \alpha_l \frac{dN(R_l)}{dr}dm_l + \int_{r^*}^{\theta_h} \alpha_h \frac{dN(R_h)}{dr}dm_h \\
&= \alpha_l \frac{dN(R_l)}{dr}m_l + \alpha_h \frac{dN(R_h)}{dr}m_h \\
&= \alpha_l \frac{dN}{dR_l}\frac{dR_l}{dr}m_l + \alpha_h \frac{dN}{dR_h}\frac{dR_h}{dr}m_h \\
&= \alpha_l \frac{dN}{dR_l}Z_l m_l + \alpha_h \frac{dN}{dR_h}Z_h m_h,
\end{aligned} \tag{A.3}
$$

where $R_l = r \cdot Z_l$ so that $\frac{dR_l}{dr} = Z_l$. Based on Equation (3.2) that $N = U + (1-\beta_l)\delta_l v(R_l)$, the partial derivative $\frac{dN}{dR_l}$ can be rewritten as:

$$
\begin{aligned}
\frac{dN}{dR_l} =& \frac{d}{dR_l}(U + (1-\beta_l)\delta_l v(R_l)) \\
=&(1-\beta_l)\delta_l v'_{R_l} \\
=&(1-\beta_l)\delta_l \frac{g_l \lambda}{\alpha_l},
\end{aligned} \tag{A.4}
$$

where $g_l := \frac{\alpha_l v'_{R_l}}{\lambda}$ by definition. Similarly, $\frac{dN}{dR_h} = (1-\beta_h)\delta_h \frac{g_h \lambda}{\alpha_h}$. Combining Equations (A.3) and (A.4), we rewrite the first term in Equation (A.2) as:

$$
\begin{aligned}
&\int_{\theta_l}^{\theta_h} \alpha(\theta)\frac{dN(D(\theta))}{dr}dm(\theta) \\
=&\alpha_l \frac{\partial N}{\partial R_l}Z_l m_l + \alpha_h \frac{\partial N}{\partial R_h}Z_h m_h \\
=&(1-\beta_l)\delta_l g_l \lambda Z_l m_l + (1-\beta_h)\delta_h g_h \lambda Z_h m_h.
\end{aligned} \tag{A.5}
$$

Based on Equation (3.2) that $N(P(\theta)) = U(P(\theta)) + (1-\beta(\theta))\delta(\theta)v(P(\theta)) + (1-\pi)K\mathbf{1}\{P(\theta) \neq R(\theta)\}$, where $s \in \{h, l\}$ and $P(\theta) \in \{R(\theta), S(\theta)\}$, the second term in Equation (A.2) can be rewritten as:

$$
\begin{aligned}
&\frac{dm_h}{dr}\alpha_h(N(R_h) - N(\theta_h)) \\
=&\frac{dm_h}{dr}\alpha_h\Big(U(R_h) + (1-\beta_h)\delta_h v(R_h) - U(\theta_h) - (1-\beta_h)\delta_h v(\theta_h) - (1-\pi)K\Big) \\
=&\frac{dm_h}{dr}\alpha_h\Big((1-\beta_h)\delta_h(v(R_h) - v(\theta_h)) - (1-\pi)K\Big). \qquad \text{(A.6)}
\end{aligned}
$$

Workers on the margin of switching to their preferred saving amount $\theta_h(=\theta_h Z_h)$ are indifferent from saving at the default or their preference in terms of the decision utility. Therefore, $U(R_h) = U(\theta_h)$. Based on Assumption (4) that $v(R_h) = R_h$ and the definition of g_h in Section 3.4.1 that $\alpha_h = \frac{g_h\lambda}{v'_{R_h}} = g_h\lambda$, Equation (A.6) can be expressed as:

$$
\begin{aligned}
&\frac{dm_h}{dr}\alpha_h(N(R_h) - N(\theta_h)) \\
=&\frac{dm_h}{dr}\alpha_h\Big((1-\beta_h)\delta_h(R_h - \theta_h) - (1-\pi)K\Big) \\
=&\frac{dm_h}{dr}g_h\lambda\Big((1-\beta_h)\delta_h(r^* - \theta_h)Z_h - (1-\pi)K\Big). \qquad \text{(A.7)}
\end{aligned}
$$

Similarly, the third term in Equation (A.2) can be expressed as:

$$
\begin{aligned}
&\frac{dm_l}{dr}\alpha_l(N(\theta_l) - N(R_l)) \\
=&\frac{dm_l}{dr}\alpha_l\Big(U(\theta_l) + (1-\beta_l)\delta_l\theta_l + (1-\pi_l)K_l - U(R_l) - (1-\beta_l)\delta_l R_l\Big) \\
=&\frac{dm_l}{dr}g_l\lambda\Big((1-\beta_l)\delta_l(\theta_l - r^*)Z_l + (1-\pi)K\Big). \qquad \text{(A.8)}
\end{aligned}
$$

Combining Equations (A.5), (A.7), and (A.8), we get

$$
\begin{aligned}
\frac{dW(r^*)}{dr} =&(1-\beta_l)\delta_l g_l\lambda Z_l m_l + (1-\beta_h)\delta_h g_h\lambda Z_h m_h \\
&+ \frac{dm_h}{dr}g_h\lambda\Big((1-\beta_h)\delta_h(r^* - \theta_h)Z_h - (1-\pi)K\Big) \\
&- \frac{dm_l}{dr}g_l\lambda\Big((1-\beta_l)\delta_l(\theta_l - r^*)Z_l + (1-\pi)K\Big) \qquad \text{(A.9)} \\
=&0.
\end{aligned}
$$

We rearrange Equation (A.9) and plug in semi-elasticities $\epsilon_l = \frac{dm_l}{dr}\frac{1}{m_d} < 0$ and $\epsilon_h = \frac{dm_h}{dr}\frac{1}{m_d} > 0$:

$$
\begin{aligned}
&\frac{dW(r^*)}{dr} \\
&= (1-\beta_l)\delta_l g_l Z_l m_l - \frac{dm_l}{dr} g_l(1-\beta_l)\delta_l(\theta_l - r^*)Z_l - \frac{dm_l}{dr} g_l(1-\pi)K \\
&+ (1-\beta_h)\delta_h g_h Z_h m_h + \frac{dm_h}{dr} g_h(1-\beta_h)\delta_h(r^* - \theta_h)Z_h - \frac{dm_h}{dr} g_h(1-\pi)K \\
&= (1-\beta_l)\delta_l g_l Z_l \frac{m_l}{m_d} + |\epsilon_l| g_l(1-\beta_l)\delta_l(\theta_l - r^*)Z_l + |\epsilon_l| g_l(1-\pi)K \\
&+ (1-\beta_h)\delta_h g_h Z_h \frac{m_h}{m_d} + |\epsilon_h| g_h(1-\beta_h)\delta_h(r^* - \theta_h)Z_h - |\epsilon_h| g_h(1-\pi)K \\
&= 0.
\end{aligned}
$$

The overall welfare effect can be decomposed into several terms after the optimal initial default rate marginally increases from r^* to $r^* + dr$:

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

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

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

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

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

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

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^*:

$$r^* = \frac{dI + d\theta_l - d\theta_h + dK_l - dK_h}{dR_l - dR_h}.$$

C. Distributions of Contribution Rates from Beshears et al. (2012) to Identify Semi-Elasticities in Section 3.5.1

Description from Beshears et al. (2012): Figure 3. The Distribution of Employee Contribution Rates at Firm C with a 3% Default. This figure gives the distribution of employee contribution rates at one year of tenure at Firm C when there was a 3% default contribution rate. The sample is the 2,785 full-time employees who were hired at the firm between January 1, 2003 and February 29, 2004, who remained at the firm for at least one year, and who were not Highly Compensated Employees. The default contribution rate was 3%, and the minimum contribution rate necessary to obtain the full employer match was 7%.

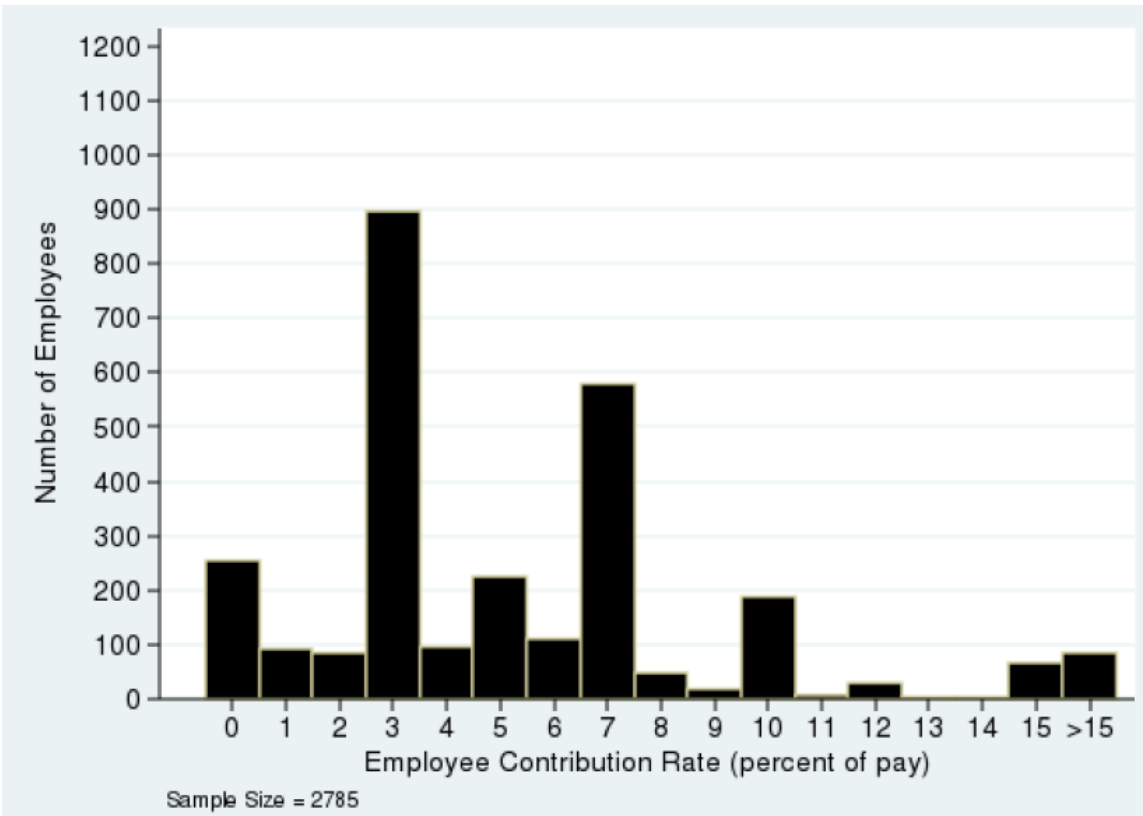

Figure 4. The Distribution of Employee Contribution Rates at Firm C with a 5% Default. This figure gives the distribution of employee contribution rates at one year of tenure at Firm C when there was a 5% default contribution rate. The sample is the 3,765 full-time employees who were hired at the firm between June 1, 2005 and July 31, 2006, who remained at the firm for at least one year, and who were not Highly Compensated Employees. The default contribution rate was 5%, and the minimum contribution rate necessary to obtain the full employer match was 7%.

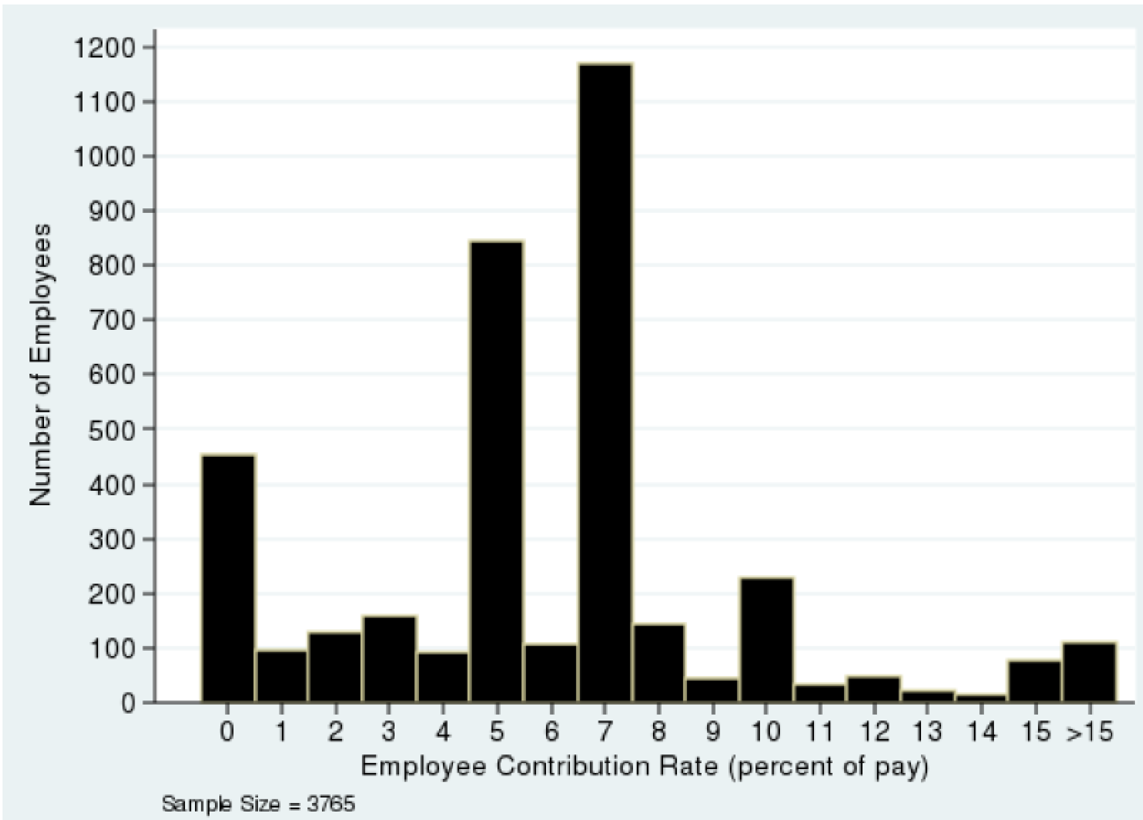

D. Survey Questions to Elicit Time Preferences

Survey design and results are explained in Section 3.5.2, Table 16, and Table 17.

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:

Each token is worth $95 today and $100 in a year. How many tokens would you want to receive today?	0	70	100
Amount you will have today	$ 6,650		
Amount you will have in a year	$ 3,000		

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.

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.

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 $99 today and $100 in a year. 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 $98 today and $100 in a year. 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 $97 today and $100 in a year. 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 today and $100 in a year. 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.

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.

Completed:

BIBLIOGRAPHY

Aguiar, Mark, Erik Hurst, and Loukas Karabarbounis (2013). Time Use during the Great Recession. *American Economic Review 103*(5), 1664–1696.

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

Baily, Martin Neil (1978). Some Aspects of Optimal Unemployment Insurance. *Journal of Public Economics 10*(3), 379–402.

Bee, Charles Adam and Joshua Mitchell (2017). Do Older Americans Have More Income Than We Think? *Proceedings. Annual Conference on Taxation and Minutes of the Annual Meeting of the National Tax Association 110*, 1–85.

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 (2012). Default Stickiness among Low-Income Individuals. *Working Paper*.

Beshears, John, James J. Choi, Brigitte C. Madrian, and William L. Skimmyhorn (2017). Borrowing to Save? The Impact of Automatic Enrollment on Debt. *Working Paper*.

Biggs, Andrew (2015). Good News: Retirement Income Still Being Undercounted. *Forbes*.

Biggs, Andrew (2019a). The "Retirement Coverage Gap" is Vastly Exaggerated. *Forbes*.

Biggs, Andrew G. (2019b). How Hard Should We Push the Poor to Save for Retirement? *Journal of Retirement 6*(4), 20–29.

Biggs, Andrew G. and Sylverster J. Schieber (2015). Why Americans Don't Face a Retirement Crisis. *AEI Economic Perspectives*.

Board of Governors of the Federal Reserve System (2019). Report on the Economic Well-Being of U.S. Households in 2018.

Bronchetti, Erin Todd, Thomas S. Dee, David B. Huffman, and Ellen Magenheim (2013). When a Nudge Isn't Enough: Defaults and Saving Among Low-Income Tax Filers. *National Tax Journal 66*(3), 609–634.

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

Card, David, Raj Chetty, and Andrea Weber (2007). The Spike at Benefit Exhaustion: Leaving the Unemployment System or Starting a New Job? *American Economic Review 97*(2), 113–118.

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.

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

Chalmers, John and Jonathan Reuter (2012). Is Conflicted Advice Better Than No Advice? *NBER Working Paper 18158*.

Chen, Anqi and Alicia H Munnell (2017). Who Contributes to Individual Retirement Accounts? *Center for Retirement Research Issue Brief* (17-8).

Chetty, Raj (2006). A General Formula for the Optimal Level of Social Insurance. *Journal of Public Economics 90*(10), 1879–1901.

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

Chetty, R and A Szeidl (2007). Consumption Commitments and Risk Preferences. *Quarterly Journal of Economics* (May).

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

Clark, Robert, Annamaria Lusardi, and Olivia S. Mitchell (2017). Employee Financial Literacy and Retirement Plan Behavior: a Case Study. *Economic Inquiry 55*(1), 248–259.

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

Farber, Henry S., Jesse Rothstein, and Robert G. Valletta (2015). The Effect of Extended Unemployment Insurance Benefits: Evidence from the 2012-2013 Phase-Out. *American Economic Review Papers and Proceedings 105*(5), 171–176.

Farber, Henry S. and Robert G. Valletta (2015). Do Extended Unemployment Benefits Lengthen Unemployment Spells? Evidence from Recent Cycles in the US Labor Market. *Journal of Human Resources 50*(4), 873–909.

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

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

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

Gruber, J (1997). The Consumption Smoothing Benefits of Unemployment Insurance. *American Economic Review 87*(1), 192–205.

Hagedorn, Marcus, Kurt Mitman, and Marcus Hagedorn (2015). The Impact of Unemployment Benefit Extensions on Employment: the 2014 Employment Miracle? *NBER Working Paper No. w20884*.

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

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.

Katz, Lawrence F. and Bruce D. Meyer (1990a). The Impact of the Potential Duration of Unemployment Benefits on the Duration of Unemployment. *Journal of Public Economics 41*(1), 45–72.

Katz, Lawrence F. and Bruce D. Meyer (1990b). Unemployment Insurance, Recall Expectations, and Unemployment Outcomes. *Quarterly Journal of Economics 105*(4), 973–1002.

Kroft, Kory and Matthew J. Notowidigdo (2016). Should Unemployment Insurance Vary with the Unemployment Rate? Theory and Evidence. *Review of Economic Studies 83*(3), 1092–1124.

Krueger, Alan B. and Andreas Mueller (2010). Job Search and Unemployment Insurance: New Evidence from Time Use Data. *Journal of Public Economics 94*(3-4), 298–307.

Krueger, Alan B., Meyer Bruce D. (2002). Labor Supply Effects of Social Insurance. In *Handbook of Public Economics*, Volume 4, pp. 2327–2392.

Kutyavina, Marina (2014). Variation in Job Search Intensity across Age Groups: Evidence from the Great Recession. *Working Paper* (2013), 1–33.

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

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

Lusardi, Annamaria and Olivia S. Mitchell (2007). Baby Boomer Retirement Security: The Roles of Planning, Financial Literacy, and Housing Wealth. *Journal of Monetary Economics 54*(1), 205–224.

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.

Meyer, Bruce D. (1990). Unemployment Insurance and Unemployment Spells. *Econometrica 58*(4), 757–782.

Michelacci, Claudio and Hernán Ruffo (2015). Optimal Life Cycle Unemployment Insurance. *American Economic Review 105*(2), 816–859.

Miller, K., David Madland, and Christian E. Wellar (2015). The Reality of the Retirement Crisis. *Center for American Progress Working Paper*, 1.

Mitchell, Olivia S and Stephen Utkus (2012). Target-Date Funds in 401(K) Retirement Plans. *NBER Working Paper 17911*.

Moffitt, R (1985). Unemployment Insurance and the Distribution of Unemployment Spells. *Journal of Econometrics 28*(1), 85–101.

Morrissey, Monique (2016). The State of American Retirement: How 401(k)s Have Failed Most American Workers. *Retirement Inequality Chartbook 3.*

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

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

Piketty, Thomas and Emmanuel Saez (2013). Optimal Labor Income Taxation. *Handbook of Public Economics*, 391–474.

Rothstein, Jesse (2011). Unemployment Insurance and Job Search in the Great Recession. *NBER Working Paper No. w17534*.

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

Schmieder, Johannes F, Till Von Wachter, and Stefan Bender (2012). The Long-Term Effects of UI Extensions on Employment. *American Economic Review 102*(3), 514–519.

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

Shimer, Robert and Iván Werning (2007). Reservation Wages and Unemployment Insurance. *Quarterly Journal of Economics 122*(3), 1145–1185.

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

Vanguard (2018). How the Uk Saves 2018.

Vanguard (2019). How America Saves 2019.