Default Options and Retirement Saving Dynamics

Taha Choukhmane[1]

December 2024

Abstract

Using data from over 100 U.S. retirement plans and a representative U.K. panel, I document the impact of auto-enrollment on retirement savings at different horizons. I replicate the impact of auto-enrollment on participation and contributions in the short–run (at 12 months), but I show that these gains are attenuated over the medium-run (at 36 months). At this longer horizon, the average savings increases are modest, though auto-enrollment significantly lowers inequality in savings. To assess auto-enrollment's lifetime impact, I estimate a lifecycle consumption-savings model that can fit the observed patterns with a switching cost of approximately $250, smaller than previous estimates.

[1]Choukhmane: MIT Sloan School of Management (tahac@mit.edu). I am deeply indebted to my dissertation committee members—James Choi, Nick Barberis, Costas Meghir, and Cormac O'Dea—for their generous help at all stages of this project and to the editor, Stefano DellaVigna, and three anonymous reviewers for their guidance and suggestions that have significantly improved the paper. For helpful discussions and suggestions, I thank Jason Abaluck, Joseph Altonji, Antoine Arnoud, Doug Bernheim, Nicolas Coeurdacier, Eric French, Xavier Gabaix, John Eric Humphries, David Laibson, Patrick Moran, Yusuke Narita, Matthew Notowidigdo, Michaela Pagel, Michael Peters, Daniel Reck, Jesse Shapiro, Jonathan Skinner, Jaya Wen, Nicolas Werquin, and Fabrizio Zilibotti. I thank an anonymous U.S. financial institution for making available the data used in this paper and for helpful discussions and technical support. I acknowledge financial support from the Whitebox Advisors Fellowship, the Boston College Center for Retirement Research, and the Alfred P. Sloan Foundation through Grant G-2017-9813 to the National Bureau of Economic Research. This work contains Crown copyright statistical data from the U.K. Office of National Statistics (ONS). The use of ONS statistical data does not imply endorsement of the ONS in relation to the interpretation or analysis of the statistical data.

1 Introduction

Automatic enrollment has transformed the way millions of people around the world save for retirement. Starting with Madrian and Shea (2001), several studies have found that moving from an opt-in regime (where new hires must actively elect to participate) to an automatic enrollment regime (where new hires are enrolled by default with the option to opt out) can double participation in a retirement savings plan.[2] Motivated by this evidence, auto-enrollment has become perhaps the most common policy application of behavioral economics and is already the norm in retirement savings plans in the U.S., U.K., New Zealand, and Turkey. Beginning in 2025, almost all newly established 401(k) and 403(b) plans in the U.S. will be required to auto-enroll their employees.

Does auto-enrollment increase lifetime wealth accumulation? Despite its widespread adoption, the answer to this question remains unknown. The adoption of auto-enrollment policies is too recent for researchers to directly observe the effect of enrolling a 25-year-old on her retirement wealth 40 years later. In this paper, I show that considering auto-enrollment in a fully specified dynamic setting is crucial for understanding its long-term impact. I first replicate the finding that auto-enrollment increases participation and contributions in the short run (at 12 months). Then, I examine the impact on medium-term savings and use a lifecycle consumption-savings model to estimate long-term effects.

I begin by documenting seven stylized facts that illustrate the dynamic and heterogeneous effect of auto-enrollment policies and provide motivation for the model. First, using a new proprietary dataset of administrative records from 46 U.S. 401(k) plans, I replicate in an event-study setting the existing evidence that auto-enrollment increases participation and average contributions in the short run, at 12 months of tenure (Fact I). Then, I show that these participation and savings benefits are attenuated over the medium run, by 36 months of tenure (Fact II). This empirical pattern holds not only for the most common auto-enrollment plan design (with a 50% match rate and a 3% default option) but also for plans with a higher default option (6%), a higher match rate, and both with or without an auto-escalating default contribution rate. The impact of higher default contribution rates is estimated to be lower largely because of selection. The plans selecting a high default option tend to have higher levels of participation and contributions prior to the adoption of auto-enrollment (Fact III). The average effects documented across the first three facts mask substantial heterogeneity:

[2]Large participation and contribution effects of auto-enrollment have been documented in the U.S. (Choi et al., 2004 and 2006; Beshears et al., 2009; Clark et al., 2015), in the U.K. (Cribb and Emmerson, 2016), and in an experimental setting in Afghanistan (Blumenstock et al., 2018).

while cumulative contributions converge at the median, the saving gains persist at the bottom of the savings distribution (Fact IV).

Next, I explore *how* employees undo some of the positive effects of auto-enrollment on savings over time. I show that the non-auto-enrolled catch up to the auto-enrolled by contributing at higher rates (Fact V). Furthermore, I find that increasing the default in 107 U.S. 401(k) plans increased the average contribution rate, albeit not one-for-one. Employees partially offset this positive effect by reducing their contributions or opting out of participation entirely (Fact VI). Finally, using representative data from the U.K., I estimate the effect of auto-enrollment after employees leave their employer and move to a new job. I find no evidence that auto-enrollment creates long-lasting saving habits; if anything, I find that previously auto-enrolled workers are less likely to participate when required to make an active opt-in decision in their next opt-in job. However, when their new job has auto-enrollment, I find no negative effect of previous auto-enrollment on contributions in the new job (Fact VII).

Motivated by this empirical evidence, I build and structurally estimate a lifecycle consumption-savings model with switching costs to extrapolate the effect of auto-enrollment across many job transitions and over a lifetime. In the model, individuals can save in a realistic retirement account and a liquid asset, borrow through unsecured credit, and face realistic labor market risk and a progressive tax and benefit system. I identify the model's preference parameters using quasi-experimental variation in the default contribution rate in the U.S. sample. Specifically, I compare the contribution decisions of employees hired the year before and the year after the adoption of auto-enrollment for new hires.

I estimate a switching cost of approximately $250 every time an employee changes her contribution rate. Previous work, which has not studied auto-enrollment in a fully specified dynamic model, has inferred switching costs on the order of thousands of dollars to rationalize the fact that, under an opt-in regime, many individuals choose to forego the tax and matching benefits of retirement saving (DellaVigna 2009 and 2018; Bernheim et al., 2015).[3] The required switching cost is smaller in my fully specified dynamic setting because workers can compensate for low initial saving rates by contributing more later. In turn, a smaller switching cost implies that the long-term impact of auto-enrollment on asset accumulation may be smaller than previously thought. The model predicts

[3]DellaVigna (2009, 2018), based on a back-of-the-envelope calculation of annual matching benefits, suggests the minimum switching cost to be above $1,000 with time-consistent preferences. Bernheim et al. (2015) estimate an average switching cost of approximately $2,000 (and above $3,000 for the 60% of workers with positive switching costs).

that an auto-enrollment policy adopted by every employer over a lifetime would increase the present value of consumption in retirement by 0.3% to 0.6% under a 3% default (or 0.7% to 1.4% under a 6% default), depending on the assumption about the incidence of increased employer matching contributions.[4] The model predicts that low-income individuals are more passive because they have less stable employment and because the switching cost represents a larger share of their earnings.[5] Consequently, a typical default saving rate of 3% or 6% of income adopted by all employers raises retirement wealth at the bottom of the lifetime earnings distribution.

In addition to showing that the switching costs and lifetime wealth gains from auto-enrollment are smaller than previously thought, the model serves three other important roles. First, I use the model to correct for bias in the data arising from plans with ex ante higher contributions selecting higher defaults. Absent this source of selection bias, the model predicts that a 6% default would create larger savings gains over three years than observed in the data. Second, I use the model to analyze the interaction between auto-enrollment and employer matching contributions. The model predicts that the retirement consumption gains from a universal auto-enrollment policy are cut in half if employers reduce the matching rate to satisfy their budget constraint. Finally, the model is useful for extrapolating to other policies, populations and institutional environments. I find that the model—with preference parameters estimated in the U.S. sample—replicates the effect of the U.K. national auto-enrollment policy rollout. The U.S. and U.K. environments vary along multiple dimensions, and the model is calibrated to capture the differences in their retirement savings plan features, income profiles, and tax and benefit systems.

This paper relates to a number of literatures. While several studies have documented the short-run effects of auto-enrollment on saving (Madrian and Shea, 2001; Choi et al. 2004 and 2006; Beshears et al., 2009; Clark et al., 2015), there is less evidence on the policy's medium-run effect, and there are no studies on its long-term impact. For the medium run, my empirical results confirm the findings of Choi et al. (2004), who document convergence at the median in two 401(k) plans over 2 to 4 years.[6] I expand on previous work and provide new evidence from a larger sample of plans,

[4]Because auto-enrollment increases employee contributions in the retirement plan, it may also increase the amount that employers have to spend on matching contributions. I simulate auto-enrollment policies under three alternative means of balancing employers' budget constraint: reducing employer profits, reducing the retirement plan match rate, and reducing wages.

[5]In the baseline model, the switching cost is fixed and identical across individuals. I explore in extensions a specification where the switching cost is proportional to income and one with two types of individuals with different earnings and preferences.

[6]More recently, Beshears et al. (2022) and Falk and Karamcheva (2023) have found that saving gains are attenuated within three to four years after the adoption of auto-enrollment in the Federal Thrift Savings plan. Derby et al. (2023) also find similar patterns of attenuation when they compare private-sector workers hired four years prior to auto-enrollment adoption to those hired one year after adoption.

including plans with different default options, with and without auto-escalation, and with different matching formulas, and measure the effect of auto-enrollment after employees move to their next job. Going beyond these medium-term effects, the estimated model allows me to evaluate the effect of auto-enrollment over a lifetime. I show that the model can replicate the heterogeneity in effects of auto-enrollment in the medium term and predicts that these heterogeneous treatment effects persist in the long run. My finding that the treatment effect of auto-enrollment attenuates over time extends to the effect of a default asset allocation (as shown by Choukhmane and de Silva (2023)).

The paper also relates to the literature studying inertia in retirement savings plans outside the context of auto-enrollment policies. Carroll et al. (2009) study a policy where employees are required to make an active contribution choice. Consistent with my findings for auto-enrollment, they find that "active decision enrollment only affects asset accumulation in the bottom half of the accumulation distribution". Thaler and Benartzi (2004) document that employees in one firm who chose to enroll in the Save More Tomorrow auto-escalation program after a consultation with a financial advisor increased their contribution from 3.5% to 13.6% of salary over 40 months.[7] However, in the absence of a control group, it can be challenging to disentangle selection effects, the treatment effect from the effect of the financial advising, and the treatment effect from auto-escalation.[8] Finally, Chetty et al. (2014) document that approximately 85% of savers in Denmark are passive and do not adjust their savings in response to increased employer contributions or a small change in mandatory government savings. There are two key differences between their setting and mine: (i) auto-enrolled employees may be more likely to notice and react to higher *employee* contributions because these contributions reduce their take-home pay (unlike increased employer contributions), and (ii) the increase in savings is an order of magnitude larger than the discontinuity in mandated savings studied in Chetty et al. (2014).

The paper also contributes to the literature modeling default effects in saving. Starting with Carroll et al. (2009), which characterizes the optimal enrollment policy in a model with switching costs and present-biased preferences, the literature has found support for and analyzed models with cognitive and adjustment costs (Bernheim et al. (2015), Blumenstock et al. (2018), Goldin and Reck (2022), Bernheim and Gastell (2024)). Relative to the existing literature, I show that a modest

[7]This study also has evidence on two other program implementations, but for the two other firms, the authors have data only for the first six months.

[8]The auto-escalation implementations that I study in this paper differ in design from the original proposal of Thaler and Benartzi (2004) in key aspects: auto-escalation occurs annually instead of being tied to pay raises, and enrollment into auto-escalation is done automatically for new hires (rather than following a session with a financial advisor).

switching cost can account quantitatively for the dynamics and heterogeneity in responses to auto-enrollment, and I highlight the importance of studying auto-enrollment in a fully dynamic setting.

The rest of the paper proceeds as follows. In Section 3, I document seven empirical facts about auto-enrollment that provide motivation for the model introduced in Section 4. I estimate the model's parameters in Section 5, and Section 6 presents the model-predicted effect of an auto-enrollment policy adopted by all employers over a lifetime. Section 7 documents the out-of-sample performance of the model, and Section 8 concludes.

2 Data and Identification

2.1 U.S. 401(k) administrative data

I use administrative 401(k) data from a large U.S. financial institution. The data set contains 401(k) administrative records from nearly 600 firms with individual-level panel data for the period between December 2006 and December 2017. For each employee, demographic characteristics, contribution rates (as a share of salary), and 401(k) balances are recorded annually in December. Employee and employer contribution dollar amounts and asset allocations are recorded at the end of every month. Salary information is available only for employees with a positive contribution rate.

Table 1: Summary statistics: U.S. 401(k) event-study sample

	3% Default		6% Default	
	No. of plans	No. obs. in event window	No. of plans	No. obs. in event window
Full sample:	29	8,321	17	5,925
— Plans w/ 50% match up to 6%	19	6,473	10	3,419
— Plans w/ 100% match up to 6%	3	625	4	706
— Plans w/ auto-escalation	17	5,696	1	204
Demographics:				
— Average age	38.4		38.3	
— Median earnings in 2006 $	$40,987		$59,332	

Notes: Sample of employees hired in the 12 months before and 12 months following the month of auto-enrollment adoption. Plan and sample selection criteria are described in Appendix B.1. I assume that all nonparticipating employees (for whom earnings information is not available) have below-median earnings.

Event-study sample. My main identification strategy relies on event studies by month of hire relative to the month of auto-enrollment adoption.[9] I compare employees hired in the 12 months

[9] A similar identification strategy has been used by Madrian and Shea (2001) and Choi et al. (2004) with data from

before and in the 12 months following the adoption of auto-enrollment in 46 401(k) plans. Both cohorts of workers face the same financial incentives to contribute to the plan, but those hired before the change have a different default option from those hired after. Twenty-nine plans adopted auto-enrollment with a 3% default, while 17 adopted auto-enrollment with a 6% default (Table 1). In Appendix B.1, I give more details about the sample construction and show that the employees hired each month before and after the policy change are similar in terms of number of observations, age, and salary (Appendix Figure A.1).

2.2 U.K. national auto-enrollment policy and representative data

The U.K. Pension Act of 2008 requires employers to automatically enroll their employees (with the option to opt out) in a workplace pension scheme. The law sets the initial minimum default contribution rate to 2% of salary with at least 1% of salary in the form of an employer contribution. The policy was rolled out in phases by employer size between 2012 and 2017, and unlike in the U.S. setting, both new hires and nonparticipating seasoned employees were automatically enrolled (see Appendix H for details on the policy rollout). I exploit the timing of the phased rollout as a source of quasi-experimental variation.

To study this policy, I use individual-level panel data from the U.K. Annual Survey of Hours and Earnings (ASHE) from 2006 to 2017. ASHE is a 1% panel of all U.K. workers, sampled based on the last digits of their National Insurance number. Demographic characteristics, labor earnings, and retirement contributions are collected from administrative tax records and from an annual survey of employers. The data are collected every year in April. There are two advantages to using ASHE. First, it is a nationally representative sample with approximately 200,000 private-sector employees per year. Second, it follows the same individuals over time as they change jobs. However, unlike the U.S. 401(k) administrative data, it does not contain data on workers' retirement account balances or the detailed features of each firm's retirement savings plans.

3 Seven Facts on Auto-Enrollment's Effect on Retirement Savings

3.1 Fact I: Auto-enrollment raises retirement savings in the short run

In line with the existing empirical literature (Madrian and Shea, 2001; Choi et al. 2004 and 2006; Beshears et al., 2009; Clark et al., 2015), I show that auto-enrollment significantly increases both

four 401(k) plans.

participation and average cumulative contributions during the initial year of employment. An employee is coded as participating in the plan if she makes a positive employee contribution during the observation month, while cumulative contributions are computed as the sum of all employee (or combined employee and employer) contributions made to the 401(k) plan since the date of hire divided by the employee's annual salary.

In Figure 1, I compare employees hired every month in a 24-month window centered on the adoption of auto-enrollment at a default contribution rate of either 3% (top panel) or 6% (bottom panel). Each dot corresponds to the average participation rate (left panel), average cumulative employee contributions (middle panel) and average cumulative contributions from both employees and employers (right panel) for workers hired the same number of months before or after the adoption of auto-enrollment. At 12 months of tenure, both participation and cumulative contributions jump for employees hired after the adoption of the policy. As shown in Panel A of Table 2, auto-enrollment at 3% increases the participation rate by 43 p.p. (or an 83% increase from a baseline of 52% in the opt-in group), cumulative employee contributions by 0.73 p.p. (or 27%), and total contributions including the employer contributions by 1.7 p.p (40%). The effects are similar for plans adopting a 6% default (columns 4, 5 and 6 of Table 2). This evidence of large short-run participation and contribution increases under auto-enrollment has motivated the rapid adoption of auto-enrollment in retirement savings plans, which is required of almost all new U.S. 401(k) plans starting in 2025.

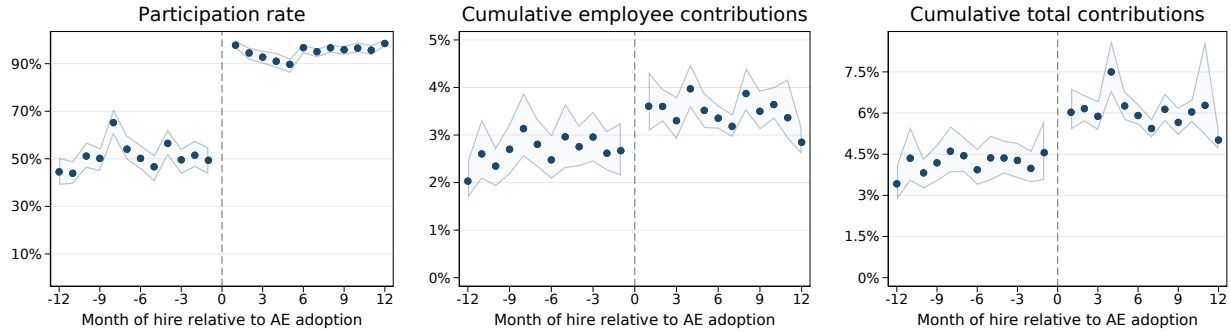
3.2 Fact II: The average positive effects are attenuated in the medium run

Next, I document in Figure 2 that the positive average effect of auto-enrollment on retirement savings is attenuated over the medium run. Auto-enrolled employees observed in their 36th month of tenure still have higher participation rates (left panel), but their cumulative employee contributions are not statistically different from those of the employees hired under the opt-in regime (middle panel). Importantly, total contributions (which include employer contributions) remain higher for the auto-enrolled group at 36 months of tenure (right panel). This is because the auto-enrolled begin contributing earlier and are less likely to contribute beyond the employer's matching limit.

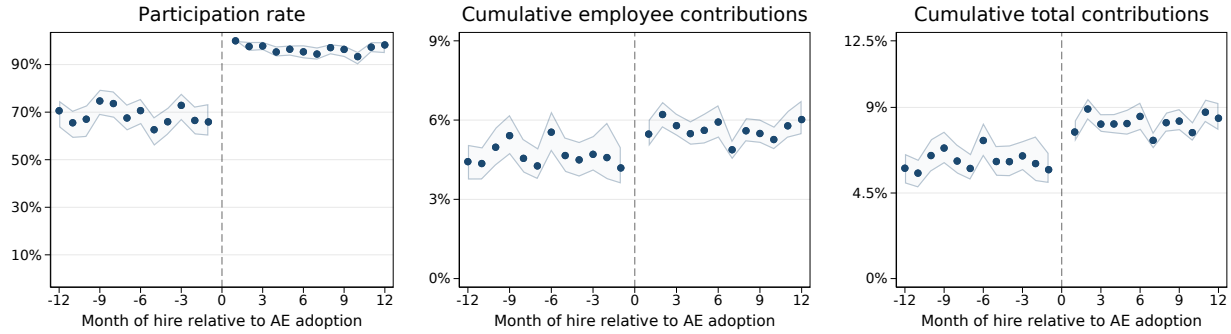
Panel B of Table 2 shows that the positive impact of auto-enrolment on cumulative employee contributions is not statistically significant at 36 months. However, under a 3% default, the magnitude of the increase at 36 months (0.58 p.p. of annual salary, s.e. 0.79) is similar to that measured at 12 months, albeit with a larger standard error (0.73 p.p., s.e. 0.31). This pattern suggests that, on average, auto-enrollment ceases to create additional savings beyond the first year of tenure, though

Figure 1: Effect of auto-enrollment in the short run (12th month of tenure)

(a) Plans with a 3% default contribution rate



(b) Plans with a 6% default contribution rate



Notes: The horizontal axis shows the normalized month of hire relative to the time of auto-enrollment adoption. The left panel shows the average participation rate of employees observed in their 12th month of tenure. The middle (right) panel shows average cumulative employee (combined employee and employer) 401(k) contributions observed in the 12th month of tenure divided by individuals' annual salary. Shaded areas correspond to the bootstrapped 95% confidence intervals stratified at the plan level for each series.

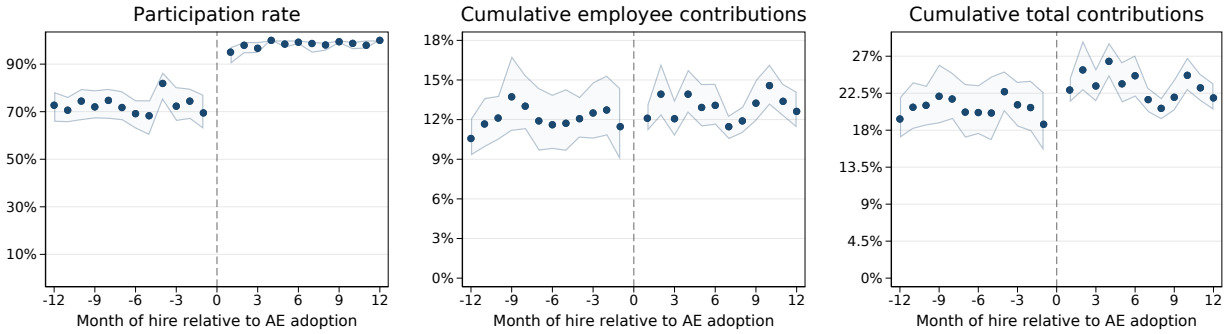
the earlier gains are not reversed later on. Since cumulative contributions grow over time, the relative difference in savings between those hired before and after auto-enrollment becomes smaller: the cumulative employee contributions of the auto-enrolled go from being 27% higher at 12 months to being 5% higher at 36 months under a 3% default. The evidence is broadly similar for plans adopting auto-enrollment at 6% (see columns 4, 5, and 6 of Table 2). The magnitude of the increase is similar at 12 and 36 months, but in relative terms, the cumulative employee contributions of the auto-enrolled go from being 25% higher at 12 months to being 7% higher at 36 months.

Evidence that the savings increases from auto-enrollment become attenuated in the medium run has been documented before, starting with Choi et al. (2004) in the context of two 401(k) plans with an auto-enrollment default of 2% and 3%.[10] I document this medium-run attenuation of the

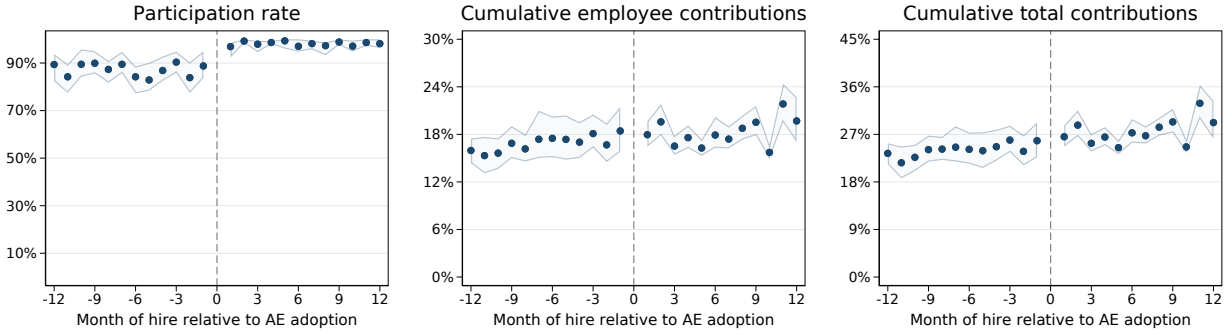
[10]The cited authors document that the median 401(k) balance-to-pay ratio of non-auto-enrolled workers caught up

Figure 2: Effect of auto-enrollment in the medium run (36th month of tenure)

(a) Plans with a 3% default contribution rate



(b) Plans with a 6% default contribution rate



Notes: The horizontal axis shows the normalized month of hire relative to the time of auto-enrollment adoption. The left panel shows the average participation rate of employees observed in their 36th month of tenure. The middle (right) panel shows average cumulative employee (combined employee and employer) 401(k) contributions observed in the 36th month of tenure divided by individuals' annual salary. Shaded areas correspond to the bootstrapped 95% confidence intervals stratified at the plan level for each series.

positive effect of auto-enrollment for a larger set of 401(k) plans and across plans with different designs, including plans with a higher default (panel (b) of Figure 2), plans with either a 50% or a 100% match rate (Appendix Figure A.2), and plans with an auto-escalating default (Appendix Figure A.3). A similar pattern is also observed in the absence of the baseline sample compliance restrictions (Appendix Figure A.4). The catching-up behavior observed in the opt-in group is similar when I look at employees hired more than 12 months prior to the policy, suggesting that it is not driven by peer effects.[11]

with that of the auto-enrolled after four years in company A (with a 2% default) and over 2 years in company B (with a 3% default), the same company studied previously by Madrian and Shea, 2001.

[11]Because employees hired in the 12 months prior to the policy interact with an increasing number of auto-enrolled new colleagues, their catching-up behavior could be partly driven by peer effects. However, the evolution of participation and contributions over three years is very similar when I look at employees hired between 12 and 24 months, between 24 and 36 months, and between 36 and 60 months prior to the policy. These prior cohorts of hires have less

Table 2: Effect on participation and contributions in the U.S. 401(k) event-study sample

	Plans with a 3% default			Plans with a 6% default		
	(1) Particip. rate	(2) Employee contrib.	(3) Total contrib.	(4) Particip. rate	(5) Employee contrib.	(6) Total contrib.
Panel A: Auto-enrollment impact at 12 months						
Hired post-AE	0.427	0.00732	0.0170	0.283	0.0115	0.0186
	(0.0723)	(0.00313)	(0.00390)	(0.0185)	(0.00299)	(0.00498)
Constant	0.516	0.0271	0.0425	0.685	0.0457	0.0625
	(0.0381)	(0.00165)	(0.00205)	(0.00949)	(0.00153)	(0.00256)
R-squared	0.291	0.048	0.078	0.160	0.081	0.153
N	6,549	6,549	6,549	5,372	5,372	5,372
Panel B: Auto-enrollment impact at 36 months						
Hired post-AE	0.245	0.00581	0.0148	0.106	0.0119	0.0239
	(0.0615)	(0.00794)	(0.0101)	(0.0231)	(0.00934)	(0.0142)
Constant	0.732	0.122	0.213	0.874	0.169	0.246
	(0.0327)	(0.00422)	(0.00539)	(0.0125)	(0.00507)	(0.00769)
R-squared	0.201	0.060	0.198	0.068	0.042	0.069
N	3,428	3,428	3,428	3,223	3,223	3,223
Plan FE	✓	✓	✓	✓	✓	✓

Notes: Standard errors, clustered by retirement plan, in parantheses. Table entries are coefficients from ordinary least squares (OLS) regressions with retirement plan fixed effects. "Hired post-AE" corresponds to the difference between workers hired in the 12 months after and 12 months before the adoption of auto-enrollment and observed at 12 months of tenure (Panel A) and 36 months of tenure (Panel B). The outcome variables are a dummy for participating in the plan (columns 1 and 4), cumulative employee 401(k) contributions divided by individuals' annual salary. (columns 2 and 5), and cumulative combined employee and employer 401(k) contributions divided by individuals' annual salary.

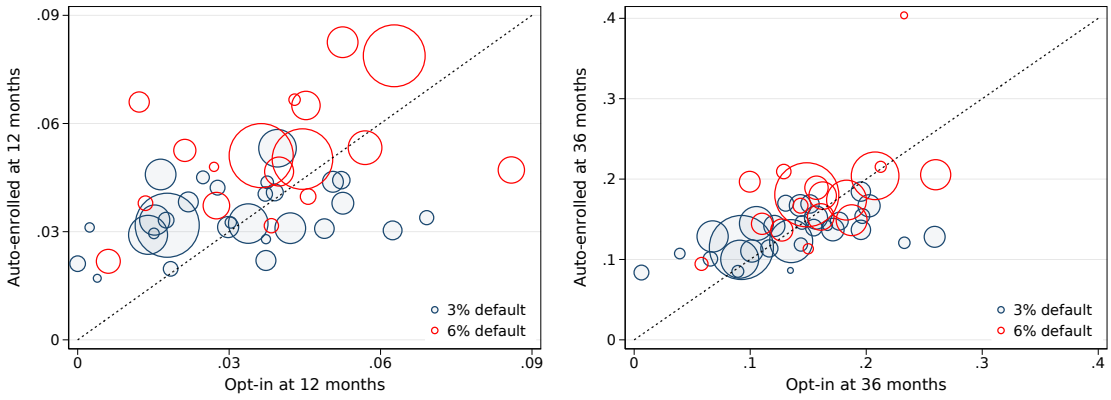
3.3 Fact III: The effect of a higher default is estimated to be lower because plans with higher contributions prior to auto-enrollment select higher defaults

Figure 3 plots the change in contributions pre- and post-adoption of auto-enrollment for each retirement plan in the event-study sample. Each circle is one retirement plan, with the size of each circle proportional to the number of observations in the plan. On the x-axis are the average cumulative employee contributions for those hired prior to auto-enrollment , and on the y-axis are the average cumulative employee contributions for those hired after the adoption of auto-enrollment. A plan on the 45-degree line is one in which the average contributions of those hired before and after auto-enrollment are similar. At 12 months of tenure, 70% of plans (accounting for 78% of employees) are above the 45-degree line, which indicates a positive impact of auto-enrollment on

overlap (or no overlap at all) with the auto-enrollment cohort and, therefore, could not have been subject to peer effects from this group in their first three years of tenure. These figures are available on request.

average cumulative employee contributions at the plan level. After 36 months of tenure, only 52% of plans (accounting for 58% of employees) are above the 45-degree line and, thus, exhibit a positive impact of auto-enrollment on average cumulative employee contributions. This evidence suggests that the patterns documented in Facts I and II are broad based across the different plans in the sample: auto-enrollment has a positive effect on employee contributions in the first year of tenure , but this positive effect is attenuated after three years.

Figure 3: Heterogeneity in the effect of auto-enrollment across plans



Notes: Each circle corresponds to the average cumulative employee contributions for the auto-enrollment cohort (y-axis) and the opt-in cohort (x-axis) in each 401(k) plan.The size of the circle is proportional to the number of observations in each plan. Blue (red) circles correspond to plans with a 3% (6%) default contribution rate.

In Figure 3, the distribution of plans with a 6% default (red circles) is shifted to the right of the distribution of plans with a 3% default (blue circles). This indicates that plans selecting a higher default (6%) have higher contribution levels under the opt-in regime. On average, after 36 months of tenure, opt-in employees in plans with a 6% default accumulate 4.8 p.p. of salary (or 40%) more in cumulative employee contributions than opt-in employees in plans adopting a 3% default. This suggests that the impact of higher default contribution rates is estimated to be lower largely because the plans selecting a high default option tend to have higher levels of participation and contributions prior to the adoption of auto-enrollment.

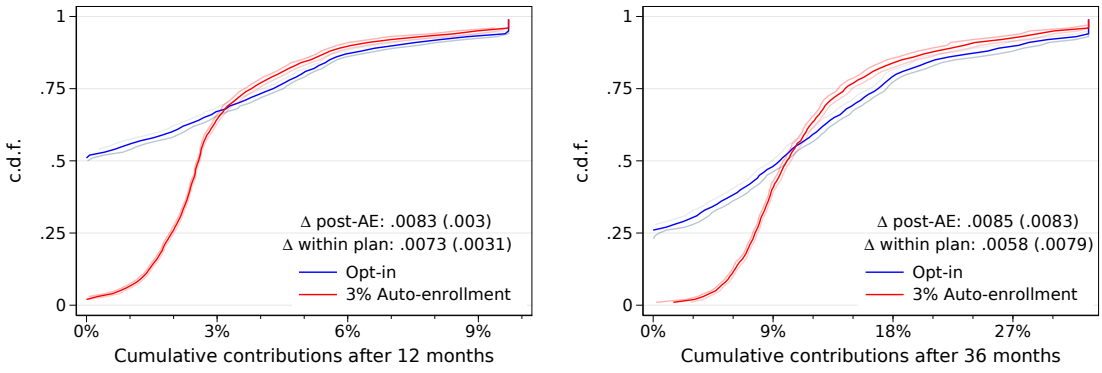
3.4 Fact IV: The average effect masks substantial heterogeneity, and the saving gains persist at the bottom of the savings distribution

The average effect of auto-enrollment on savings, documented in the previous three facts, masks substantial heterogeneity. In Figure 4, I plot the cumulative distribution functions (CDFs) of employee

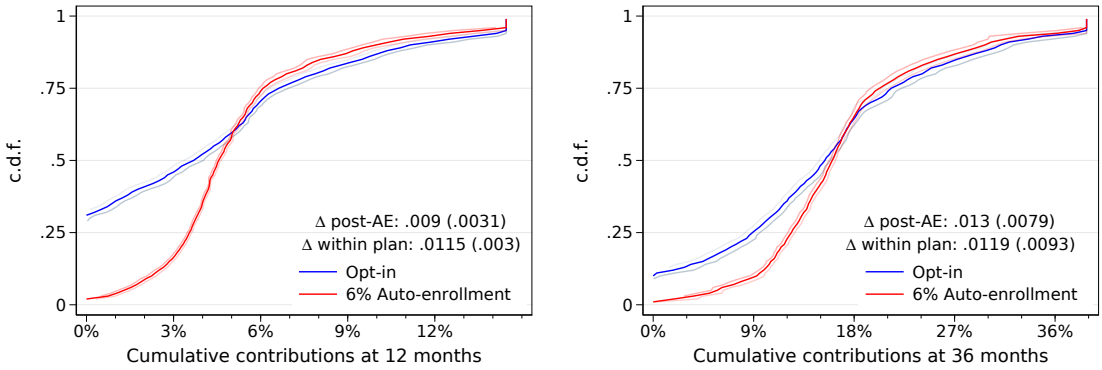
contributions at 12 and 36 months of tenure under the two enrollment regimes.[12] The horizontal gap between the CDF for the auto-enrolled (red series) and the opt-in group (blue series) measures the effect of auto-enrollment on cumulative contributions at each percentile of the savings distribution (y-axis). That is, auto-enrollment increases cumulative contribution at a given percentile of the savings distribution if the CDF of the auto-enrolled is to the right of the CDF of the opt-in group.

Figure 4: Heterogeneity in the effect of auto-enrollment on cumulative contributions

(a) Plans with a 3% default contribution rate



(b) Plans with a 6% default contribution rate



Notes: Each line corresponds to the cumulative distribution function (CDF) of cumulative employee 401(k) contributions divided by individuals' annual salary observed in the 12th month (left panels) and 36th month (right panels) of tenure. Values are winsorized above the top 5% of cumulative contributions. Shaded areas correspond to the bootstrapped 95% confidence intervals stratified at the plan level for each series.

[12]Note that the bunching around the cumulative contribution levels implied by saving at the default option is diffuse for two reasons. The first reason is that there is variation with respect to when an employee is automatically enrolled. Most new hires are typically automatically enrolled sometime during their first three months of tenure. This implies that, by the 12th month of tenure, some employees have been automatically enrolled for 12 months while others have been automatically enrolled for only 9 months. The second reason is that, as income changes over time, the denominator of the cumulative contributions–to–earnings ratio changes. Hence, even if an employee contributes 3% every month, the ratio will vary as her income varies.

Under both default options (3% in panel (a) and 6% in panel (b)), auto-enrollment has a small negative effect on cumulative contributions at the top of the savings distribution. Some employees who would have contributed above the default under an opt-in regime accumulate less savings when auto-enrolled. Around the median, auto-enrollment increases savings in the short run, especially for the 3% default option (i.e., the median non-auto-enrolled worker makes no contribution in her first 12 months of tenure, while the median auto-enrolled employee accumulates close to 3%), but the savings gains at the median are attenuated within 36 months. Finally, at the bottom of the savings distribution, auto-enrollment creates persistent savings gains after 3 years. These opposing trends—increasing savings gains at the bottom and decreasing gains around the median—can explain why the savings gains estimated in Table 2 are similar at 12 and 36 months. This suggests that auto-enrollment reduces inequality in retirement savings (in terms of both employee and employer contributions). Specifically, it increases the savings of workers at the bottom of the savings distribution who otherwise would not have saved and would have forgone the employer match.

3.5 Fact V: The non-auto-enrolled partially catch up to the auto-enrolled by contributing at higher rates

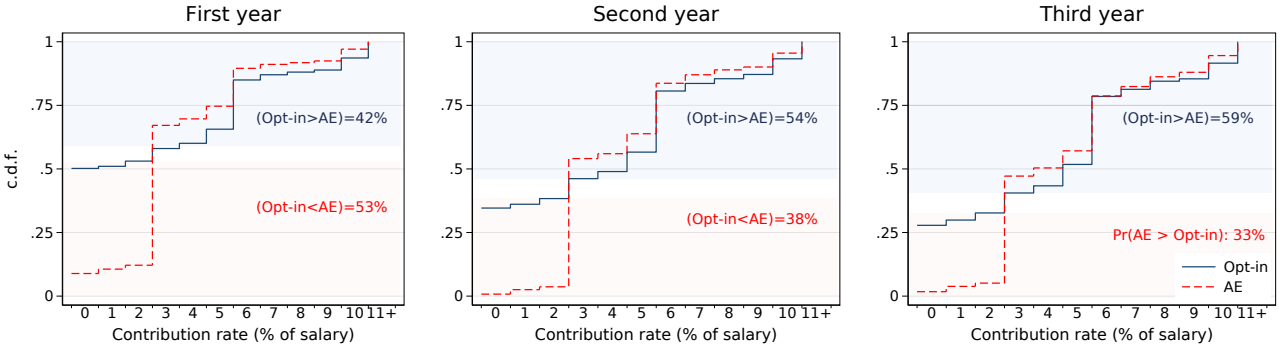
Having documented the effect of default options on cumulative contributions, I explore *how* employees undo some of the positive effects of auto-enrollment. I restrict the sample to employees with at least 36 months of tenure such that I can follow the contribution rates of a balanced sample of employees over three years. First, I show that the non-auto-enrolled catch up over time by contributing at higher rates. Panel (a) of Figure 5 illustrates how an employee contributing at the median would accumulate the same cumulative employee contributions after three years before and after auto-enrollment across three plan designs. Panel (b) shows the CDFs of contribution rates in the first, second, and third years of tenure in plans with a 3% default without auto-escalation. In the first year, the top 42 percentiles of contributions are higher under opt-in than under auto-enrollment. By the third year, the opt-in cohorts contribute at higher rates, and the top 59 percentiles of contributions are higher than for the auto-enrollment cohorts. I show similar evidence for plans with auto-escalation and with a 6% default in Appendix B.3. The evidence on contribution rates aligns with the previous evidence on cumulative contributions: savings gains are attenuated at the median over 36 months and persist at the bottom. In addition, Figure 5 highlights that this empirical pattern also holds for employees who stayed with the same employer for 36 months and across plans with and without auto-escalation.

Figure 5: Evolution of contribution rates over three years of tenure

(a) Saving path for an employee contributing at the median over 3 years

	Opt-in				Auto-enrolled			
	year 1	year 2	year 3	**Total**	year 1	year 2	year 3	**Total**
3% default without auto-escalation	0%	5%	5%	**10%**	3%	3%	4%	**10%**
3% default with auto-escalation	0%	6%	6%	**12%**	3%	4%	5%	**12%**
6% default	6%	6%	6%	**18%**	6%	6%	6%	**18%**

(b) Distribution of contribution rates: plans with a 3% default without auto-escalation



Notes: Panel (a) illustrates the accumulated employee contributions of an employee choosing the median contribution rate of her enrollment cohort and plan type every year. Panel (b) shows the cumulative distribution functions (CDFs) of contribution rates as a % of salary recorded in December of the employee's first, second, and third years of tenure in plans with a 3% default without auto-escalation.

3.6 Fact VI: Higher defaults raise average contributions but lead more employees to contribute at the lowest rates or stop participating

Next, I show that increasing the default contribution rate increases the average contribution rate in the first 12 months but that employees partially offset this positive effect by reducing their contributions or opting out of participation entirely. I use a sample of 764,006 employees, all of whom were auto-enrolled, where some faced a higher default contribution rate than others.[13] I compare those hired before to those hired after the default contribution rate in 107 plans increased from one positive number to a higher one (for instance, from 3% to 6% of salary).[14] To control for time trends, I include a sample of 245 never-treated plans, with auto-enrollment but with no change to their default option. In Figure 6, I show regression coefficients comparing those hired before to those hired after the default increase. I find that increasing the auto-enrollment default reduces

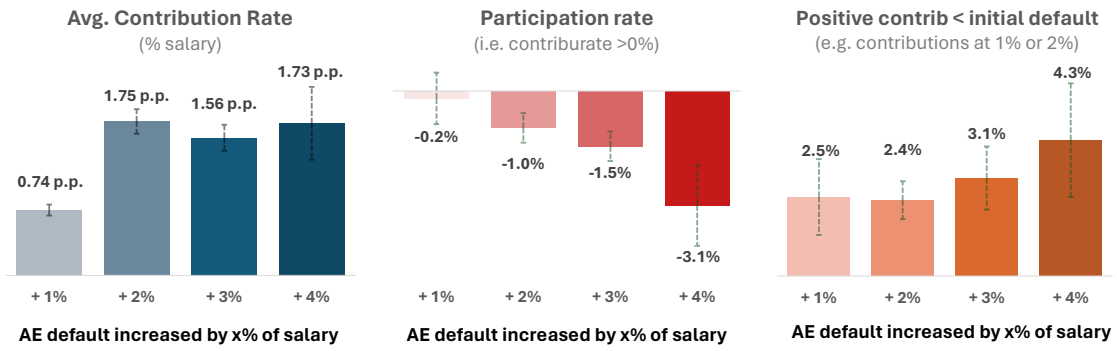
[13] I restrict the sample to employees in their first year of tenure to avoid confounding the direct effect of an employee's facing a higher default by the indirect wealth effect. Workers enrolled at different default rates accumulate different amounts of retirement wealth and, over time, have different incentives to save.

[14] The sample includes 88 plans in which the default rate increased once and 19 plans in which the default rate increased twice.

participation and increases contributions at the lowest rates—those lower than the initial default.

I estimate the effect of a higher default contribution rate in a linear regression with plan fixed effects, year fixed effects, and controls for log age and log tenure. The coefficients of interest are plotted in Figure 6, and full regression results are reported in Appendix Table A.2. The first outcome variable is the average employee contribution rate (in the left panel of Figure 6): increasing the default increases the average contribution rate in the first 12 months, albeit not one-for-one. For example, raising the default contribution rate by 3 percentage points is estimated to increase employee contributions by 1.56 p.p. of salary. The next panels highlight two mechanisms that partially explain why this increase is less than one-for-one. In the middle panel, the outcome variable is a dummy equal to 1 if the worker contributes a positive amount to the plan. For each percentage-point increase in the auto-enrollment default contribution rate, participation drops by, on average, 0.6 p.p. (Appendix Table A.2). In addition, increasing the default contribution rate not only reduces participation but also increases contributions at the lowest rates. The outcome variable in the right panel is a dummy variable for participating in the plan at a contribution rate strictly below the initial default contribution rate. I find that for each percentage-point increase in the default contribution rate, positive contributions below the initial default increase by 1.2 p.p. For example, when the default increases from 3% to 6% of salary, workers are more likely to contribute either 1% or 2% of their salary to the plan. I obtain similar patterns using the dynamic event-study approach of Sun and Abraham (2021), which is robust to treatment effect heterogeneity (Appendix Figure A.6 and Appendix Table B.4).

Figure 6: Effect of a higher auto-enrollment default on participation and contributions



Notes: Each column reports coefficients from a linear regression for employees observed in their first 12 months of tenure and after the end of the auto-enrollment grace period. The independent variables are dummies equal to one for employees hired after the default contribution rate was increased by 1, 2, 3, or 4 percentage points in 101 US 401(k) plans. The sample also includes employees in 245 never-treated plans with no change to the auto-enrollment default option. Controls include year fixed effects, plan fixed effects, log tenure (in days), and log age. Standard errors are clustered at the plan level.

3.7 Fact VII: Previously auto-enrolled workers save less in their subsequent opt-in job, with no negative impact when their next job has auto-enrollment

Finally, I exploit the phased rollout of the U.K. national auto-enrollment policy to study the participation and contributions of employees moving from a firm subject to the auto-enrollment requirement (larger firms) to a firm not yet affected by the policy (smaller firms). I compare these employees' saving decisions in their first year of tenure to the saving decisions of workers who moved to the same firm in the same year from an employer not yet affected by the policy (i.e., from a smaller to another smaller firm).

I use data from 12 successive waves of ASHE from April 2006 to April 2017. The policy rollout started in 2012 with only the largest employers affected, and in every ASHE wave between 2012 and 2017, I observe more firms subject to the auto-enrollment requirement (see Appendix H for details of the rollout). By the last wave in April 2017, all employers with more than 30 employees were subject to the auto-enrollment mandate. I include data from the period prior to the beginning of the policy rollout to control for time-invariant differences between workers moving between firms of different sizes. The sample contains 37,120 job switchers in their first year of tenure at their new employer. I estimate the following specification for a worker i who transitioned in year t from employer e_{n-1} to employer e_n:

$$y_{i,t} = \beta_0 \, AEtoNoAE_{i,t} + \beta_1 \, AEtoAE_{i,t} + \sum_{j=1}^{7} \phi_j \, Size_j^{e_{n-1}} + \sum_{j=1}^{7} \sum_{g=1}^{7} \gamma_{j,g} \, Size_j^{e_{n-1}} \times Size_g^{e_n} + \varphi \, Z_{i,t} + \alpha_{e_n,t} + \varepsilon_{i,t} \tag{1}$$

I consider two outcome variables $y_{i,t}$: the employee's participation status in the retirement savings plan and her contribution rate as a share of salary. The variable $AEtoNoAE_{i,t}$ is a dummy equal to 1 for worker i in year t if both (i) her current employer e_n is not subject to the auto-enrollment requirement in year t and (ii) her previous employer e_{n-1} was subject to the auto-enrollment requirement the last time employee i was observed at firm e_{n-1}. In the presence of firm-by-year fixed effects $\alpha_{e_n,t}$, the coefficient β_0 captures the difference in participation and contribution rates among new hires in the same opt-in firm who were or were not subject to the auto-enrollment requirement in their previous job. Similarly, the variable $AEtoAE_{i,t}$ is a dummy equal to 1 if the worker was subject to automatic enrollment in both her current and her previous job, and the coefficient β_1 captures the difference in participation and contribution rates among new hires in the same auto-enrollment firm who were or were not subject to the auto-enrollment requirement in their previous job. Since the policy was rolled out on the basis of employer size, one concern could be that new hires coming to firm e_n from larger previous employers (which were subject to auto-enrollment

earlier) are systematically different from those joining from smaller employers. I use data from the periods before and after the policy rollout to identify the effect of auto-enrollment separately from the effect of an employee's moving from a larger to a smaller employer. To do so, I control for a set of interaction dummies $Size_{e_{n-1}} \times Size_{e_n}$ for each pair of current and previous employer size categories (i.e., with 7 firm-size categories, I have 49 pairs of current and previous employer sizes).[15] I use data from prior to the beginning of the auto-enrollment rollout to identify these interaction dummies separately from the effect of an employee's moving from a non-auto-enrollment to an auto-enrollment employer. Finally, $Z_{i,t}$ is a set of individual controls (a third-order polynomial in age, gender, log salary and log tenure in the current job e_n and log salary and log tenure when the employee is last observed in the previous job e_{n-1}).

My baseline estimate for β_0 is reported in the first column of Table 3. Workers who were previously auto-enrolled participate 12.8 percentage points less and contribute 0.55% less of their salary than workers who were not previously auto-enrolled and were hired in the same year by the same next non-auto-enrollment employer. In contrast, when workers move to a new employer subject to the automatic enrollment mandate, there is no statistically significant difference between those who were previously auto-enrolled and those who were not. To verify that these estimates capture the effect of the policy rollout and not differences across employers, I conduct a number of placebo exercises. I run the same specification under the assumption that the policy rollout started before the actual policy implementation date in 2012. The results of these placebo tests, reported in columns 2 to 8 of Appendix Table A.3, are all insignificant at the 5% level for both the participation and contribution outcomes, and a single coefficient is significant at the 10% level.

There are several possible explanations for these findings: Individuals induced to accumulate more savings in their previous job could have less incentive to contribute later on because of a standard wealth effect. Alternatively, previously auto-enrolled individuals may choose not to opt into their new employer's retirement plan because they are better informed about the future rollout of the policy and expect that their new employer will soon enroll them automatically.[16] I discuss and assess these alternative explanations in Appendix H.

[15] The seven employer size categories are firms with more than 30,000 workers, 6,000–29,999 workers, 350–5,999 workers, 160–349 workers, 58–159 workers, 50–57 workers, and 30–49 workers. Firms with fewer than 30 employees are the omitted category.

[16] Note that if all new hires have rational expectations about the future policy rollout, we should not expect those who were previously auto-enrolled and those who were not to have different expectations about the future rollout of the policy.

Table 3: Effect of auto-enrollment after a job transition

Beginning of policy rollout	(1) Participation rate	(2) Contribution rate (% of pay)
AE to non-AE	-0.128	-0.546
(i.e., difference relative to non-AE to non-AE)	(0.055)	(0.226)
AE to AE	0.010	-0.039
(i.e., difference relative to non-AE to AE)	(0.019)	(0.068)
Employee and job transition controls	✓	✓
Employer-by-year FEs	✓	✓
Observations	37120	37120

Notes: Each cell reports coefficients from a regression whose dependent variable is a dummy for workers' participation status in the retirement savings plan (column 1) or the contribution rate in percentage points of pensionable pay (column 2). The coefficients correspond to the difference between workers whose previous employer was and those whose previous employer was not subject to the auto-enrollment mandate when their subsequent employer is not (AE job to non-AE job) and is (AE job to AE job) subject to the auto-enrollment requirement. I assign each employer's enrollment regime (i.e,. AE job or non-AE job) based on the rollout schedule of the U.K. national auto-enrollment policy (in Appendix H). Controls for employee characteristics include log salary and log tenure in the current job, log salary and log tenure when the employee is last observed in the previous job, a dummy for gender, and a third-order polynomial in age. The sample is restricted to private-sector job switchers between the ages of 22 and 60 observed in their first year of tenure. Standard errors clustered by current employer are in parentheses. Data source: U.K. Annual Survey of Hours and Earnings waves 2006 to 2017.

My findings go against the hypothesis that auto-enrollment creates lasting saving habits, which should cause treated workers to contribute more in their next job's savings plan. Instead, the evidence suggests that under partial adoption of auto-enrollment by employers, some of the savings gains from auto-enrollment could be undone after a job transition to an opt-in employer. However, as more employers adopt auto-enrollment, this negative impact could become less relevant, given my estimate of no negative effect for employees transitioning to a new job with auto-enrollment. I show in Section 7.2 that the magnitude of these effects and the difference accross the two types of transitions (AE to non-AE and AE to AE) are consistent with the model-predicted responses.[17]

3.8 Discussion and implications for modeling the long-term effect of auto-enrollment

How should the default effect be modeled? As shown by both Bernheim et al. (2015) and Goldin and Reck (2022), the fact that increasing the default option leads more individuals to opt out (Fact VI) supports a model in which deviating from the default option involves a switching cost. For instance, consider an employee whose ideal contribution rate is 1%. When the default contribution rate is 3%, she may stay at the default to avoid incurring the switching cost because the default

[17]Interestingly, the model prediction is closer to the empirical estimate for job switchers whose employer will not become subject to the auto-enrollment mandate in the following 12 months (and who are therefore less subject to the confounding effect of expectations about imminent adoption of auto-enrollment).

option is close enough to her preferred contribution rate. However, if the contribution rate is raised to 6%, the same employee may choose to incur the cost and switch to 1% because the higher default is further away from her preferred contribution rate. Therefore, this employee reduces her contribution rate from 3% to 1% when the default is increased from 3% to 6%. In contrast, a model with loss aversion or psychological anchoring yields the opposite prediction: increasing the default should lead to a higher saving reference point (or saving psychological anchor) and therefore lead employees to contribute and participate more (not less). Formal results are discussed in Appendix C.

Why do previous studies estimate large switching costs? The top-left panel of Figure 4 shows that, at the median, non-auto-enrolled workers have made no contributions by the end of the first year of tenure while the median cumulative employee contributions of auto-enrolled workers are close to 3% of salary (i.e., the auto-enrollment default in this sample). This implies that, by the end of their first year of tenure, auto-enrolled workers have earned substantially more employer matching and tax benefits than non-auto-enrolled workers. Because the difference in benefits earned by auto-enrolled and non-auto-enrolled workers is large, previous studies have suggested that implausibly large switching costs are required to rationalize inertia around the default contribution rate. DellaVigna (2009, 2018) sets the minimum switching cost above $1,000 with time-consistent preferences, based on a back-of-the-envelope calculation of annual matching benefits for the median worker. Bernheim et al. (2015) estimate a median switching cost of $2,200 using data on workers' first few quarters of tenure.

Why might the switching cost be smaller in a fully dynamic setting? By the end of the third year, the median cumulative contributions of non-auto-enrolled workers catch up to those of auto-enrolled workers (top-right panel of Figure 4). The fact that many non-auto-enrolled workers catch up with the auto-enrolled by contributing more later implies that non-auto-enrolled workers are not fully forgoing the match and tax benefits of retirement savings by not contributing early in their tenure. Hence, applying the previous back-of-the-envelope calculation with data from the end of the third year of tenure would imply that the switching cost is negligible at the median. Because the estimated size of the switching cost is sensitive to the observation period in models that do not allow for the dynamic adjustment of contributions, it is necessary to use a fully dynamic model to obtain an estimate of the switching cost that is consistent across observation periods. I show in Section 5.6 that, in a dynamic lifecycle model, a switching cost of approximately $250 generates the patterns of behavior documented in Figure 4. In turn, a relatively modest switching cost implies that the long-term average effect of auto-enrollment on both wealth accumulation and welfare may

be small for most people.

Why does modeling job-transitions matter? On average, U.S. workers change jobs more than ten times over their lifetime (BLS 2019), and the evidence presented in Fact VII shows that auto-enrollment in one's current job can change contributions and participation in the retirement plan of the next job. Whether the savings gains documented for employees at the bottom of the savings distribution persist all the way to retirement depends on how savings change across many job transitions. In the next section, I build and estimate a structural lifecycle model with switching costs to extrapolate the effect of auto-enrollment across many job transitions.

4 A Lifecycle Model with Default Effects

Motivated by the empirical evidence presented in the previous section, I develop a lifecycle model with default effects. Before detailing the model, I summarize here its key features. Every period, individuals choose a level of consumption, retirement wealth, liquid wealth, and unsecured debt. Individuals incur a utility cost every time they choose a retirement contribution rate different from the default. Their contribution choice becomes, in turn, the next period's default contribution rate.

In what follows, I introduce the baseline model. Table A.1 in Appendix A gives a summary of all the parameters introduced in this section. In Appendix F, I explore three extensions to the baseline model with, respectively, present-biased preferences, a switching cost proportional to employees' earnings, and no offset of unemployment benefits by retirement withdrawals.

4.1 Demographics and preferences

Age is indexed by a, and working life starts at age 22. Each period t is one quarter. Mortality is stochastic; individuals face a death probability m_a every period and die with certainty at age A. The period utility function is of the isoelastic form with a constant elasticity of intertemporal substitution σ such that $u_a(\cdot) = n_a \frac{\left(\frac{\cdot}{n_a}\right)^{1-\frac{1}{\sigma}}}{1-\frac{1}{\sigma}}$. The term n_a is an equivalence scale that captures the evolution of household composition over the lifecycle. In the baseline, individuals maximize their discounted intertemporal utility with an exponential discount factor δ. I introduce and estimate in Appendix F.1 an extension of the model with quasi-hyperbolic time preferences.[18]

[18] While present bias does not, by itself, generate inertia around the default option, it can exacerbate the effect of a small switching cost (O'Donoghue and Rabin, 1998; Carroll et al., 2009; DellaVigna 2009 and 2018).

4.2 Labor market

Individuals in the model face uninsurable earnings and employment risk. Every period, they can be in one of four employment states emp_t: continuously employed (E), newly employed following a job-to-job transition (JJ), unemployed (U), or retired (Ret). Individuals retire deterministically at age A^{ret} and receive a Social Security–style public pension in retirement, described in Section 4.6.

Continuous employment. When employed, individuals receive uncertain labor earnings. The log of labor earnings is equal to the sum of a deterministic component (a cubic in age) and a stochastic component θ_t, such that $\ln w_t = \delta_0 + \delta_1 a_t + \delta_2 a_t^2 + \delta_3 a_t^3 + \theta_t$.

The stochastic component θ_t follows an AR(1) process with normally distributed innovations ξ such that $\theta_t = \rho\theta_{t-1} + \xi_t \quad \text{if } emp_t = E$. The variance of these innovations in the first period, $\xi_1 \sim \mathcal{N}\left(0, \sigma^2_{\xi_1}\right)$, differs from that in subsequent periods, $\xi_t \sim \mathcal{N}\left(0, \sigma^2_{\xi}\right)\ \forall\, a > 1$, to account for initial heterogeneity in income at age 22.[19]

Job-to-job transitions. Job-to-job transitions are modeled separately from continuous employment because pension arrangements change when employees change jobs. An employee switches employers with a probability $\pi^{JJ}(a, \theta, ten)$ that depends on her age, her labor productivity, and her tenure in her current job. A job-to-job transition is associated with a new earnings draw, such that $\theta_t = \rho\theta_{t-1} + \zeta_t \quad \text{if } emp_t = JJ$. The labor earnings innovation $\zeta_t \sim \mathcal{N}\left(\mu^{JJ}, \sigma^2_{\xi}\right)$ has mean $\mu^{JJ} > 0$ to capture the earnings premium associated with job-to-job transitions.

Unemployment. The probability of moving into unemployment $\pi^{EU}(a, \theta, ten)$ depends on both age, labor productivity, and tenure. Unemployed individuals receive unemployment insurance benefits $ui(\theta_t)$, described in Section 4.6. The stochastic component of earnings θ_t stays constant during an unemployment spell. An unemployed individual transitions back to employment with a probability $\pi^{UE}(a, \theta)$ that depends on age and labor productivity. The earnings innovation in the new job $\varsigma_t \sim \mathcal{N}\left(\mu^{UE}, \sigma^2_{\xi}\right)$ has mean $\mu^{UE} < 0$ to capture the earnings loss associated with job displacement such that $\theta_t = \rho\theta_{t-1} + \varsigma_t \quad \text{if } emp_t = U$.

4.3 Decisions

In the model, people save for both precautionary motives (to insure against bad labor earnings realizations and unemployment risk) and lifecycle motives (to complement Social Security in retirement).

[19] To simulate the income process, I need to take a stand on the initial distribution of earnings in period 0 (i.e., at age 22) since it cannot be defined in terms of period t-1 realizations. I choose the initial condition of the income process to match the empirical distribution of incomes at age 22.

Every period, they make two decisions: how much to hold in liquid wealth (l_t) and how much to contribute to or withdraw from a tax-deferred defined contribution (DC) retirement account (dc_t).

Liquid taxable wealth. People start their lives with zero liquid wealth. Every period thereafter, they choose a level of liquid asset holdings l_{t+1}. Contributions to the liquid asset are made out of after-tax income. Positive liquid asset holdings grow at a constant rate $1+r$, and these returns are subject to a capital tax τ^k. Borrowing up to a limit $\underline{l}_a$ is allowed, reflecting the availability of unsecured credit. Borrowing is subject to a higher interest rate $1+r^{CC}$. The after-tax interest rate on liquid assets is thus equal to $R^{liq}(l) = 1 + r\left(1-\tau^k\right)$ if $l \geq 0$, and $R^{liq}(l) = 1 + r^{cc}$ if $l < 0$.

DC contributions. Employees can contribute a fixed percentage of their labor earnings $s_t \geq 0$ to a DC account. Contributions are subject to an age-dependent maximum contribution limit set by the government, such that $w_t s_t \leq limit_a$. DC contributions are tax deferred and, thus, reduce taxable income in the contribution period. DC wealth dc_t earns the same riskless return r as the liquid asset, but these returns are not subject to capital taxation ($R^{dc} = 1+r$).

DC withdrawals. People can withdraw a fraction $draw_t \in [0,1]$ of their DC wealth when unemployed or retired.[20] Withdrawals are subject to both income taxation and—depending on the individual's age—an additional tax penalty pen_a for early withdrawals.

4.4 Retirement savings plans

Employers (indexed by e) make contributions to the DC account, denoted ec^e. The employer contribution formula consists of a match rate ($match^e$) and a cap on matching (cap^e). Contributions above this threshold do not earn an employer match contribution. In some plans, an employee who separates from her employer before the end of a vesting period may lose part (or all) of the employer matching contributions. To capture this effect without introducing an additional state variable, I assume that the employer match is adjusted by a factor $\Upsilon_e(a,\theta,ten) \leq 1$ that depends both on the vesting schedule and on the probability that an employee—based on her age, labor productivity, and tenure—separates from her employer before the end of the vesting period (see Appendix D.2.2 for more details). An employee contributing s_t receives an employer DC contribution equal to: $ec^e(s_t,a_t,ten_t) = \left(\Upsilon_e(a_t,\theta_t,ten_t) \times match^e \times \min\{s_t, cap^e\}\right) w_t$.

[20] Many 401(k) plans offer hardship withdrawals and loan options for employed participants, but the tax code discourages these early distribution options. Withdrawals after a documented hardship are often subject to a 10% tax penalty for workers younger than $59^{1/2}$ and, during the sample period, triggered a minimum six-month suspension from contributing to the plan. 401(k) loans are exempt from the tax penalty only if repaid in full before seperation.

4.5 Default mechanism

Default contribution rate. In the first period in a job, the default contribution rate $\bar{d}^e$ is drawn from the exogenous distribution of employer types $E(.)$. It is equal to zero under the opt-in enrollment regime and strictly positive under auto-enrollment. Later in a worker's tenure, the default equals her contribution rate in the previous period; it is therefore endogenous to the individual's saving decisions ($d_t = s_{t-1} \quad \text{if } ten_t > 1$).

Switching cost. Individuals in the model act as if they are required to pay a utility cost every time they change their contribution rate. To change their contribution rate from the default option, people have to incur a utility cost k. The period utility of an employee contributing s_t to her DC plan is given by $u_a\left(c_t - \mathbb{1}_{(s_t \neq d_t)}k\right)$. In the baseline model, the switching cost is the same for all individuals and in all periods. This switching cost captures both real transaction costs and several behavioral explanations for the default effect, such as the cost of seeking advice from a financial adviser or the cognitive costs of paying attention or finding an optimal choice.[21]

4.6 Government

Taxes. Agents face a nonlinear income tax schedule $tax^i(.)$. DC accounts are treated favorably by the tax system: contributions are not subject to income taxation, but withdrawals (in either unemployment or retirement) increase taxable income by the withdrawal amount.[22] Taxable income is equal to labor earnings plus DC withdrawals minus DC contributions. Returns on capital are taxed at a constant rate τ^k, but DC wealth is exempt from capital taxation.

Unemployment benefits. People receive a benefit $ui(\theta_t)$ when unemployed. This benefit depends on labor productivity in the last period of previous employment. In the baseline model, the employer contribution portion of withdrawals from the retirement account is treated as compensation and may offset unemployment benefits (details in Appendix D.2.2).

Public pension benefits. After retirement, individuals receive Social Security–style public

[21] Goldin and Reck (2022) show a correspondence between these switching costs and models of inattention, status quo bias and defaults as implicit advice. Consistent with the importance of cognitive and complexity costs, both Choi et al. (2009) and Blumenstock et al. (2018) find that more employees make an active decision when the enrollment process is simplified or when they receive help computing the implications of alternative savings scenarios. However, this specification does not encapsulate other forms of (nonrational) inattention that have been linked to default effects in other settings (Brot-Goldberg et al. (2023)).

[22] This is the traditional tax-deferred DC model. I abstract from the more recent, and less common, Roth 401(k) tax model, where DC contributions are taxed up-front.

pension benefits. Public pension benefits $ss(.)$ are computed according to a nonlinear schedule with an income floor. The amount received depends on the individual's average lifetime earnings ae_{T^R}, which evolve according to $ae_{t+1} = \frac{w_{t+1}+a_t ae_t}{1+a_t}$ if $a_t < A^{ret}$ and $ae_{t+1} = ae_{A^{ret}}$ if $a_t \geq A^{ret}$.

4.7 Dynamic optimization problem

I describe the recursive optimization problem and the intertemporal budget constraint in periods of retirement, employment, and unemployment in Appendix D. This dynamic optimization problem is characterized by a vector X_t of 9 state variables: age (a), employment status (emp), tenure (ten_t), employer DC plan type (e), labor productivity (θ), average lifetime earnings (ae), liquid assets (l), DC wealth (dc) and the default contribution rate (d). There is uncertainty over survival to the next period (with mortality probability (m_a)), the stochastic component of earnings (θ), employment status (emp), and the type of the future employer after a job change (e). The joint distribution of earnings, employment and employer types is denoted $F(\theta_t, emp_t, e_t)$. This dynamic optimization problem does not admit a closed-form analytical solution; the numerical solution procedure is discussed in Appendix I.

5 Estimation and Model Fit

5.1 Method of simulated moments

I estimate the model parameters in two stages.[23] In the first stage, I set a number of parameters outside of the model, either estimated directly from the data or taken from previous literature. In the second stage, I use the method of simulated moments to estimate the model's three preference parameters: the intertemporal discount factor (δ), the elasticity of intertemporal substitution (σ), and the switching cost (k).

5.2 Identification of preference parameters

The three model preference parameters are jointly estimated. In what follows, I discuss which type of variation in the data allows the identification of each parameter.

Time preferences. The model has two parameters governing saving preferences: the elasticity of intertemporal substitution (σ) and the intertemporal discount factor (δ). Since both parameters

[23] This two-step procedure is typical in the estimation of lifecycle models (see, among others, Gourinchas and Parker (2002), Low et al. (2010), Laibson et al. (2024), and O'Dea (2018)).

affect saving decisions, it is often difficult to obtain separate identification from lifecycle saving profiles alone. To overcome this challenge, I exploit a discontinuity in the savings interest rate at the employer matching contribution threshold. In the estimation sample, contributions up to 6% of income are matched by the employer, while contributions above the threshold are not. Bunching of employee contributions at 6% can therefore be used to identify the elasticity of intertemporal substitution separately from the discount factor, which governs the overall level of saving. This approach is similar to that of Bernheim et al. (2015) in a similar setting and that of Best et al. (2020) in a different setting (these authors use jumps in the mortgage interest rate at thresholds of the loan-to-value ratio to identify the elasticity of intertemporal substitution).

Switching cost. I compare the contribution behavior of employees hired before versus after the adoption of auto-enrollment to identify the size of the switching cost. Since both groups of employees face similar economic incentives, absent the switching cost, they should make similar contribution choices. Bunching of contributions at the default contribution rate in the auto-enrollment group (and its evolution with tenure) identifies the size of the switching cost. This identification strategy is similar to that of Bernheim et al. (2015): we should expect no excess mass at the default if the switching cost is zero and 100% bunching at the default if the switching cost is infinitely large.

5.3 First-stage parameters

I use data from the Survey of Income Programs and Participation (SIPP) to estimate the parameters of the labor earnings process and labor market transition probabilities at the quarterly level. The remaining first-stage parameters (including demographic variables, the parameters of the tax and benefits system, the vesting schedule, and the rate of return on assets) are estimated directly from the data or calibrated with reference to other papers. The estimation of the earnings process, treatment of the data, and calibration of other first-stage parameters are discussed in Appendix D.

5.4 Second-stage parameter estimation

The three second-stage parameters (δ, σ, and k) are estimated with the method of simulated moments. The parameters are chosen to minimize the distance between the model-simulated and empirical contribution patterns for employees hired before and after the introduction of auto-enrollment.

Estimation sample. I estimate the model's second stage parameters in the event-study sample described previously in Section 2.1, and I further restrict the sample to 401(k) plans (i) that have no auto-escalation feature, (ii) that offer a single-tier matching formula with either a 50% or 100%

employer match on employee contributions up to 6% of earnings, and (iii) for which the exact date of auto-enrollment adoption is available. Of the 46 plans in the event-study sample, 16 satisfy these inclusion criteria. I restrict the sample to employees hired by the sponsors of these 16 plans in the 12 months before or 12 months after the introduction of automatic enrollment (2,283 employees before and 2,174 after the policy change) and focus on auto-enrolled workers observed after the end of the auto-enrollment grace period.[24] I estimate preferences separately for three types of plans: (i) plans with a 3% default and a 50% match rate up to 6% of salary, (ii) plans with a 6% default and a 50% match rate up to 6% of salary and (iii) plans with a 6% default and a 100% match rate up to 6% of salary. I use the sample with the 3% default as the baseline because this is the most common default option and because the income distribution of employees in this sample better matches that in nationally representative data (see Appendix Figure A.10). In contrast, employees in plans with a 6% default and a 50% match rate, and to a greater extent those with a 6% default and a 100% match rate, have significantly higher earnings, and their earnings—which are 63% and 114% higher, respectively, at the median than their counterpart in the 3% sample—are not well captured by the model's income process estimated in nationally representative data. To address this issue, I reweight observations in the simulated data to match the distribution of earnings in the two samples with a 6% default. That is, in computing the simulated moments, I give more weight to observations with a higher simulated income (see Appendix E.1 for details). In a separate extension in Appendix E.1, I explore an alternative approach with two types of agents (standard and high-earners) having different preferences.

Simulation. To perform the estimation, I simulate 5,000 individuals under two scenarios. In the opt-in scenario, individuals face a 0% default contribution at the start of every job. In the auto-enrollment counterfactual, they experience the same sequence of employment and productivity draws. However, each simulated individual—at one randomly drawn job transition—*unexpectedly* faces a positive default contribution rate when starting a new job. I compare the simulated contribution behavior under these two scenarios to the empirical patterns of employees hired before and after the introduction of auto-enrollment.

Estimation moments. I use two sets of moment conditions (56 in total). First, I use the distribution of contribution rates among employees in their first year of tenure: the fraction of workers contributing 0%, 3%, 6%, and 10% or above of salary under both the opt-in and the auto-

[24]Employees can withdraw, without a tax penalty, any amount automatically contributed to a 401(k) plan if they opt out before the end of a grace period, typically corresponding to one to three months from the date of hire. Workers forfeit any matching contribution that would have been made for the automatic enrollment contributions.

enrollment regimes (8 moment conditions).[25] Second, I use the evolution of bunching at the opt-in default (0%) and bunching at the auto-enrollment default (either 3% or 6% of salary) over the first 4 years of tenure. That is the participation rate under opt-in, excess participation under auto-enrollment, the share of contributions at the default contribution rate under opt-in, and the excess share of contributions at the default under auto-enrollment every quarter between the 5^{th} and 16^{th} quarters of tenure (48 moment conditions).

Weighting matrix. Given that all moments are in the same unit (i.e., percentages of participants), I use the identity matrix as my baseline weighting matrix in the estimation, and I show that the results are similar when I use the inverse of the diagonal terms of the variance–covariance matrix of the second-stage moment conditions obtained by bootstrap. I do not use the optimal weighting matrix because it has been shown to have poor small-sample properties (Altonji and Segal, 1996).

5.5 Results and model fit

Estimation results. I report second-stage parameter estimates in Table 4. The time preference estimates in my baseline specification (column 1) are broadly in line with values estimated in other lifecycle settings. The quarterly discount factor $\delta = 0.989$ (or $\delta^4 = 0.957$ annually) is similar to the estimates from Gourinchas and Parker (2002), who find annual values of between 0.930 and 0.962. The estimated elasticity of intertemporal substitution $\sigma = 0.520$ is also standard in the literature.[26] The switching cost is estimated at \$235 every time a worker changes her contribution rate, with a standard error of \$ 31. The estimate of the switching cost across the different samples and model specifications are of a similar order of magnitude: they range from \$152 to \$295. Consistent with the evidence in Fact III that plans selecting a higher default have ex ante higher contribution levels, I estimate a larger intertemporal discount factor in the sample with a 6% default of $\delta = 0.998$ (or $\delta^4 = 0.992$ annually) than in the sample with a 3% default (columns 6 to 9). In contrast, my estimates of the elasticity of intertemporal substitution (EIS), which range from 0.52 to 0.66, are quite similar across the different samples and specifications.

[25] I do not observe workers who never joined the 401(k) plan and left their employer before the month of December of their first year of tenure. This is not an issue in the empirical section of the paper because I measure balances after at least 12 months of tenure, at which point every nonparticipating employee has been observed at least once. To account for this data availability constraint in the model, I compute simulated moments at a randomly drawn quarter from the simulated workers' first year of tenure. Depending on these draws, I use the simulated behavior in the first, second, third, or fourth quarter of tenure. These correspond, respectively, to workers in the data hired between October and December, January and March, April and June, or July and September and observed in December of their first year of tenure. I observe both participants and nonparticipants after the first year of tenure and can target quarterly data in the estimation.

[26] Havránek (2015) reports a mean estimate of the elasticity of intertemporal substitution in micro studies centered on $0.3 - 0.4$, after correcting for publication bias, in a meta-analysis of 169 published studies.

Table 4: Preference parameter estimates

	Auto-enrollment at 3%				Auto-enrollment at 6%			
	50% match rate				50% match rate		100% match rate	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Baseline	Alternative weighting	Immediate vesting	1st year of tenure only	Baseline calibration	Immediate vesting	Baseline calibration	Immediate vesting
Switching cost (k)	\$ 235	\$ 262	\$ 221	\$ 156	\$ 289	\$ 295	\$ 176	\$ 152
	(31)	(27)	(28)	(38)	(31)	(19)	(29)	(25)
Quarterly disc. factor (δ)	0.989	0.989	0.988	0.983	0.998	0.998	0.998	0.998
	(0.002)	(0.002)	(0.002)	(0.005)	(0.002)	(0.002)	(0.001)	(0.001)
EIS (σ)	0.520	0.682	0.660	0.531	0.558	0.575	0.536	0.583
	(0.085)	(0.076)	(0.069)	(0.070)	(0.043)	(0.062)	(0.053)	(0.061)

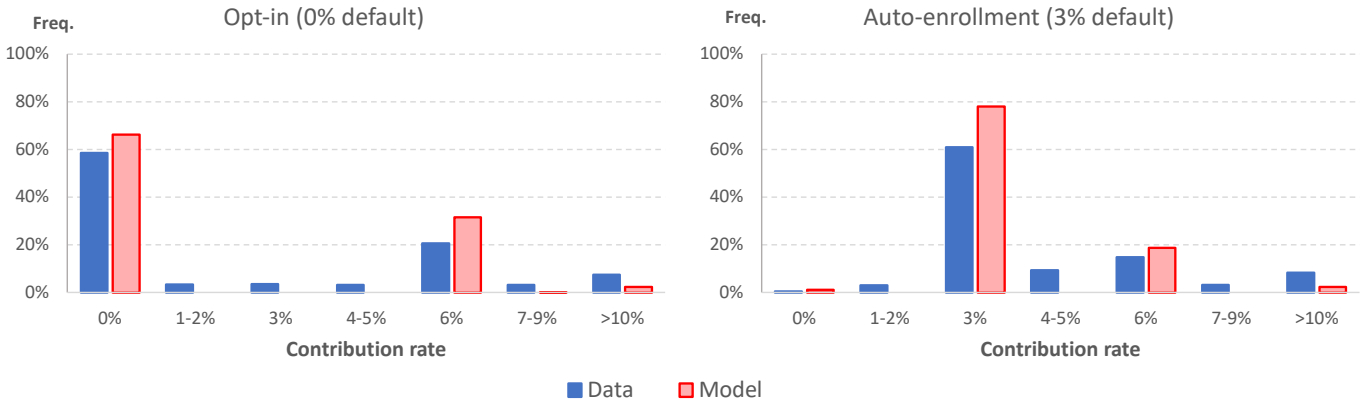
Notes: Columns 1 to 5 are estimated with data from plans with a 3% default and a 50% match rate up to 6% of salary. Columns 6 and 7 (8 and 9) are estimated with data from plans with a 6% default and a 50% (100%) match rate up to 6% of salary. In columns 6 to 9, the simulated moments are reweighted to match the empirical distribution of earnings in the corresponding sample (see Appendix E.1). I use the identity weighting matrix, except in column 2, where I use the inverse of the diagonal terms of the variance-covariance matrix of second-stage moments. Standard errors adjusted for measurement error in the estimation of the income process are in parentheses.

Robustness and sensitivity. In column 2 of Table 4, I show that the esimates are similar when I use a different weighting matrix (the inverse of the diagonal terms of the variance–covariance matrix of empirical moments) or under the assumption that vesting is immediate (columns 3, 6 and 8). In column 4, I show that my estimating a smaller switching cost than the estimates in previous studies is not due to my using longer panel data: the estimated switching cost is, if anything, smaller when I use only the moments from the first year of tenure. In Appendix E.4, I estimate a larger switching cost when I fix δ or σ to be lower, and a smaller cost when I fix σ to be higher. In Appendix J, I discuss the sensitivity of the estimates using the measure of Andrews et al. (2017).

Model fit in the sample with a 3% default. I report in Figures 7 and 8 the simulated behavior for the baseline estimates (column 1 of Table 4) alongside the empirical moments used in estimation. Figure 7 shows that, under auto-enrollment, the model generates realistic bunching of contributions at the default contribution rate (3% of salary) and at the match threshold (6% of salary). However, the model overpredicts contributions at the match threshold among workers without auto-enrollment. Figure 8 shows that the model replicates the evolution of contributions and participation in the first four years of tenure. Participation increases with tenure for non-auto-enrolled workers (Figure 8, left panel) and stays constant at a high rate for auto-enrolled workers (Figure 8, middle panel). Over time, auto-enrolled workers move away from the default contribution rate, and by the end of the fourth year of tenure, only approximately 30% of them contribute exactly 3% (Figure 8, right panel). While the model is estimated with 401(k) data only, I show in Appendix G.1 that my estimates are also consistent with the evolution of total wealth accumulation over the

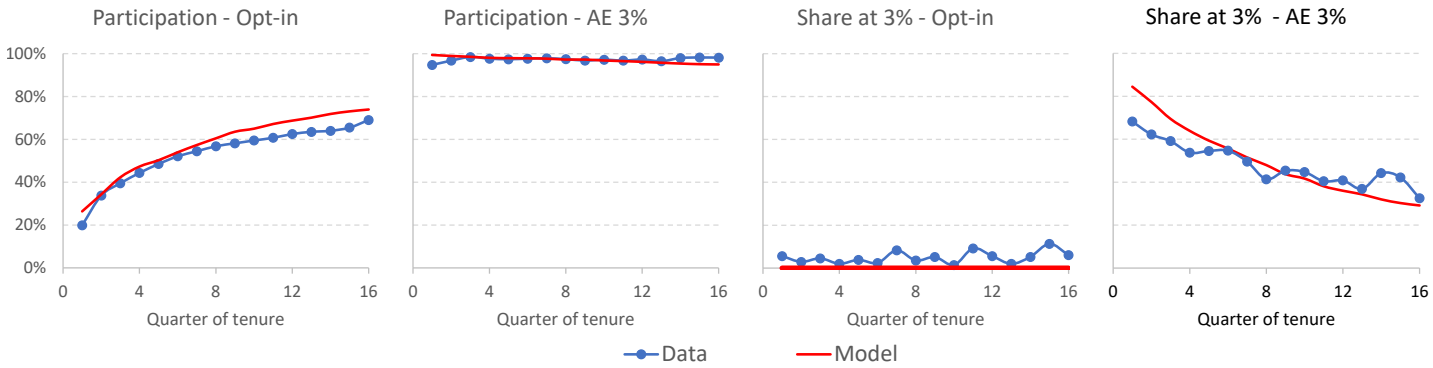
lifecycle in the Survey of Consumer Finances.

Figure 7: Empirical and simulated distributions of contribution rates in the first 12 months of tenure



Notes: The left (right) panel shows contribution rates of individuals hired in the 12 months before (after) the introduction of auto-enrollment for new hires. The model series corresponds to simulations for identical individuals under a 0% and a 3% default.

Figure 8: Empirical and simulated contribution behavior over 4 years of tenure



Notes: The left (right) two panels show the 401(k) participation rate (the share who contribute exactly 3% of salary) for workers hired in the 12 months before and 12 months after the month of auto-enrollment adoption. The model series are simulations for identical individuals under a 0% and a 3% default.

Model fit in the sample with a 6% default. Workers in plans with a 6% default contribution rate are estimated to have higher discount factors than those in plans adopting a 3% default. The switching cost estimated under a 6% default is lower in plans with a 100% match rate and larger in plans with a 50% match rate, and the model fits the data well, as shown in Appendix Figures A.11 to A.14. As an extension, I show that a version of the model with heterogeneous agents generates similar results. I assume two types of agents: (i) standard agents with the baseline preference parameter estimates and the baseline earnings distribution and (ii) high-earners, who are allowed to

have different preferences (see Appendix E).

Standard errors and overidentification test. Standard errors reported in Table 4 are adjusted for measurement error in the estimation of the income process in the first stage. The baseline model with 53 degrees of freedom is rejected by the χ^2 overidentification test with a goodness-of-fit value of 323 (or 297 when adjusting for measurement error in estimating the income process) .[27]

5.6 Heterogeneity in short- and medium-run effects

The model is quite parsimonious, with only 3 preference parameters estimated on population averages, but I show that it generates the heterogeneity in auto-enrollment effects observed in the data. Figure 9 shows that, in line with Fact IV in Section 3, auto-enrollment's impact on retirement savings is largest at the bottom of the savings distribution, and the median cumulative contributions of non-auto-enrolled workers catch up to those of auto-enrolled workers after three years.

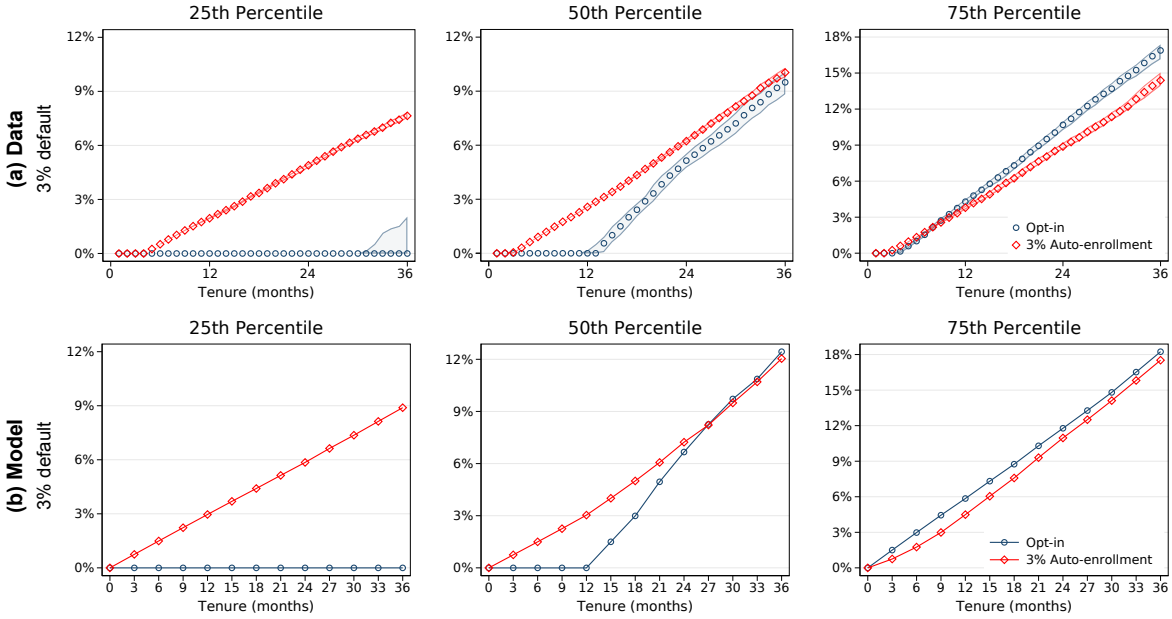
What causes the effects of auto-enrollment to differ across people? I run a regression to estimate the differences in the propensity of auto-enrolled workers to stay at the default option across different age and income groups. The outcome variable is the probability that an auto-enrolled employee in her first year of tenure contributes exactly 3% (the auto-enrollment default contribution rate) conditional on participating in the plan. I condition on participation because salary data are available only for employees with a positive contribution rate (i.e., 95% of auto-enrolled employees). Because age and income are correlated, I control for income using within-age-group quintiles of earnings in the following linear probability specification for a worker i in her first year of tenure:

$$Pr\left(s_i = 3 \,|\, s_i > 0\right) = \alpha_0 + \sum_{j=1}^{5} \beta_j age_i^j + \sum_{k=1}^{5} \gamma_k inc_{i,t}^k + \varepsilon_i \tag{2}$$

I run the same specification in the empirical and simulated samples. Estimates for β_j (default behavior by age group) and γ_k (default behavior by age-group earnings quintile) are reported in Figure 10. In line with previous evidence (Madrian and Shea, 2001), young and low-income workers are most likely to contribute at the auto-enrollment default contribution rate. The model generates qualitatively similar heterogeneity but tends to overestimate the extent to which inertia declines with income.

[27] I construct the same goodness-of-fit measure as Laibson et al., 2024. Failure of the overidentification test is common in the estimation of lifecycle models (Gourinchas and Parker, 2002; Laibson et al., 2024).

Figure 9: Empirical and simulated ratios of cumulative employee 401(k) contributions to salary

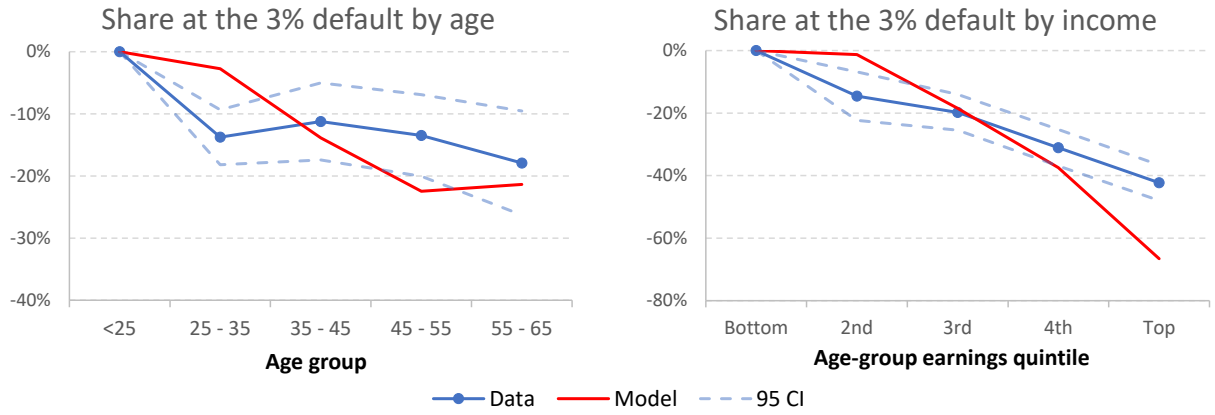


Notes: The empirical series correspond to employees hired within 12 months of the month of adoption of auto-enrollment with a 3% default contribution rate. The model series correspond to simulated data for identical individuals under a 0% and a 3% enrollment default using the baseline preference estimates from column 1 of Table 4.

Heterogeneity by age. Conditional on participating, workers in their 20s are 18 p.p. more likely to stay at the 3% default contribution rate than workers in their late 50s and early 60s (Figure 10, left panel). In the model, this heterogeneity reflects the fact that the option value of waiting to contribute decreases with age. A young employee choosing not to save has plenty of time to accrue the tax and matching benefits by contributing more later in life. In contrast, there is no option value to waiting in the last period prior to retirement. A worker about to retire permanently forgoes the tax and matching benefits of retirement saving if she does not contribute.

Heterogeneity by income. The lowest-income workers are 42 p.p. more likely to remain at the 3% default contribution rate than the highest-paid workers (Figure 10, right panel). In the model, this heterogeneity is explained by the following forces: (i) the switching cost represents a larger share of earnings for low-income individuals, (ii) low-income individuals face higher separation probabilities and, as a consequence, expect to stay in the firm for a shorter period, and (iii) high-income individuals have more to gain by contributing to the retirement savings plan because they face a higher marginal tax rate and expect a lower Social Security replacement rate. In Appendix F.2, I show that a specification of the model where the switching cost is proportional to individuals' earnings generates less heterogeneity in default effects by income.

Figure 10: Heterogeneity in default effects by age and income



Notes: Each point corresponds to the coefficient from a linear regression estimated according to equation (2). The dependent variable is a dummy equal to 1 if, conditional on participating, an auto-enrolled employee contributes exactly at the 3% default contribution rate in her first year of tenure. The "Data" series corresponds to the sample of plans with auto-enrollment at 3%. The "Model" series corresponds to simulated data using the baseline preference estimates from column 1 of Table 4. The area between the dashed lines corresponds to the 95% confidence intervals based on standard errors clustered at the plan level.

5.7 The predicted impact of a 6% default is larger after correcting for selection

The evidence presented in Fact III suggests that the impact of a higher default is estimated to be lower because plans with higher contributions prior to auto-enrollment select higher defaults. Indeed, the plans with a 6% default have higher earnings and more patient workers, as indicated by the larger estimated discount factor in this sample. I use the model to assess the effect of a higher default absent these two sources of selection. I set the preference parameters at their baseline estimated values from column 1 of Table 4 and simulate the effect of a 6% auto-enrollment policy on (i) employees with earnings drawn from the baseline income process estimated in SIPP data and (ii) employees reweighted to match the distribution of earnings observed in the plans with a 6% default and a 50% match rate. The model predicts an auto-enrollment policy with a 6% default implemented in plans with baseline preferences and earnings would raise average cumulative employee contributions by 5.0% of salary over three years, and the savings gains are predicted to persist at the median (right panel of Appendix Figure A.15). However, when the simulated observations are reweighted to match the observed distribution of earnings in plans adopting a 6% default (middle panel of Appendix Figure A.15), the model predicts that average cumulative contributions increase by only 1.2% of salary after three years, which is similar to the empirical estimate of 1.19% (column 5 in Table 2 Panel B).

6 Long-run Increases in Wealth

6.1 Policy experiment: Universal auto-enrollment policy

I use the estimated model to predict the effect on lifetime savings of an auto-enrollment policy with a default contribution rate of 3% or 6% of salary adopted by all employers. These are the two most common default contribution rates in U.S. 401(k) plans (Arnoud et al., 2021).

Model simulation. I simulate the effect of an auto-enrollment policy with a default d^{SP} adopted by all employers over a lifetime. I assume that, initially, all employers share the same matching formula (a 50% match for contributions up to 6% of income with immediate vesting). I define lifetime earnings as the sum of all labor income and unemployment benefits earned between ages 22 and 65. I measure the effect of auto-enrollment as the difference between each simulated individual's consumption and savings under the universal auto-enrollment policy and the same simulated individual's consumption and savings under an opt-in regime.

Limitations. While the model attempts to capture the main aspects of the retirement saving environment, it abstracts from a number of features. In Appendix Section F.4, I discuss how (i) the estimated switching costs could be smaller if agents can endogenously adjust their labor supply, (ii) the savings gains from auto-enrollment may be larger when coupled with a high-equity default asset allocation, and (iii) the general equilibrium impact on interest rates is likely to be small.

6.2 Incidence and means of balancing the budget

More workers participate and contribute under auto-enrollment, and therefore, the policy has the potential to increase employers' matching costs. I impose that counterfactual policies should satisfy an aggregate employer budget constraint such that the sum of aggregate employer profits $Pf\left(d^{SP}\right)$, aggregate wages $W\left(d^{SP}\right)$, and aggregate employer matching expenditures $Mtc\left(d^{SP}\right)$ is constant across auto-enrollment default options[28]: $Pf\left(d^{SP}\right) + W\left(d^{SP}\right) + Mtc\left(d^{SP}\right) = \overline{Y}$. I assume that employers respond to the policy in one of three ways.

Full incidence on employers. In the first scenario, I assume that the policy's incidence falls entirely on employers, who reduce their profits to cover higher matching costs. This reduction in profits is assumed to have no effect on the interest rate (which is calibrated to match the long-term yield on government bonds). This assumption is implicit in previous theoretical characterizations of

[28]Note that since employers are ex ante identical in this setup, the aggregate employer budget constraint is equal, in expectation, to each individual employer's budget constraint.

default effects that do not feature employers' budget constraints (Bernheim et al., 2015; Goldin and Reck, 2022). A drawback of this specification is that the direct effect of auto-enrollment on wealth accumulation cannot be disentangled from the effect of employers' spending more on matching, which increases workers' total compensation.

Full incidence on employees. In the second scenario, I assume that the incidence of the policy falls fully on workers. Butrica and Karamcheva (2012) document that, in the cross-section, total compensation and total retirement plan costs do not differ between firms with and without auto-enrollment but that auto-enrollment correlates with a less generous matching formula. Motivated by this evidence, I assume that employers adjust the match rate $match^e$ such that the total matching expenditure satisfies the aggregate budget constraint. The match rate adjusts to keep employers' total matching expenditure constant under alternative default contribution rates.[29] Finally, in Appendix F.5, I consider the possibility that employers adjust wages to balance their budget.

6.3 Results

The effect of auto-enrollment on lifetime savings is small at the top. In the top half of the distribution of lifetime earnings, financial wealth at retirement, defined as the sum of liquid wealth and DC retirement assets at age 65, is changed by less than 2% under a national auto-enrollment policy at 3% relative to that under an opt-in regime (top-left panel of Figure 11). A national auto-enrollment policy at 6% creates more savings gains in the top half of the distribution, but these gains are significantly reduced under the assumption that employers reduce the match rate to balance their budget (Figure 11, bottom right panel). An additional benefit of auto-enrollment in the model that is not captured by the change in balances at retirement is that agents are, on average, less likely to bear the switching cost under the 6% default than under the opt-in regime.

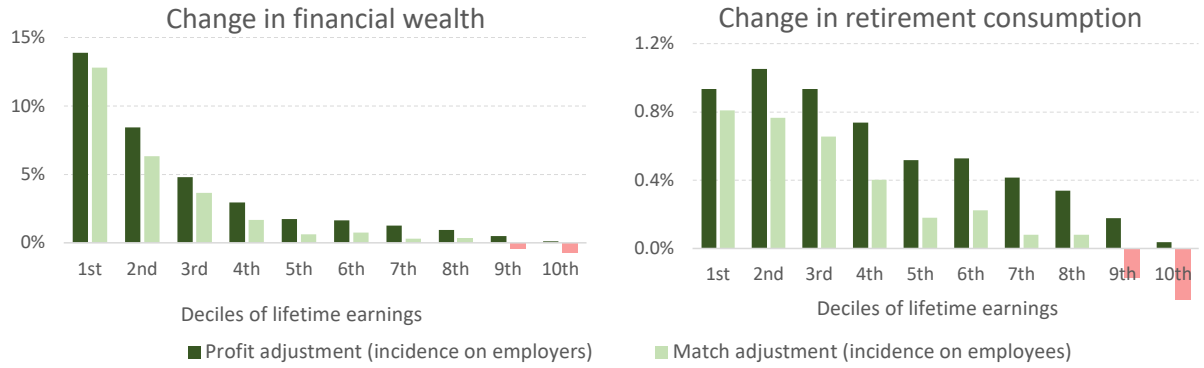
Lifetime savings increase at the bottom. Auto-enrollment increases financial wealth at retirement for those in the bottom decile of the lifetime earnings by up to 14% under a 3% default and up to 29% under a 6% default adopted by all employers during the lifetime (right panels of Figure 11). However, financial wealth represents a smaller share of total wealth for lower-income individuals, who rely more on Social Security income. To measure the effect of auto-enrollment on total financial resources available in retirement, I show in the right two panels of Figure 11 the

[29] In principle, employers could reduce their total expenditures on matching and increase profits after adopting auto-enrollment. Because auto-enrollment raises contributions at the bottom, plans are more likely to satisfy nondiscrimination testing—an IRS requirement that the benefits of a qualified retirement savings plan do not accrue disproportionately to highly compensated employees—without having to offer generous matching incentives. Bubb and Warren (2020) study a related setting in which employers select a default option to maximize profits.

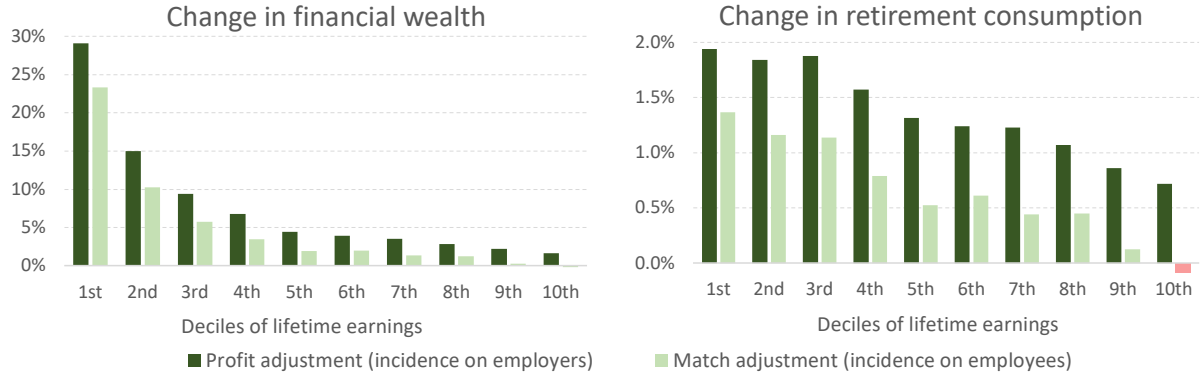
change in the present discounted value of consumption in retirement.[30] A 3% (6%) auto-enrollment policy raises the present discounted value of retirement consumption at the bottom of the lifetime earnings distribution by up to 1% (up to 1.9%).

Figure 11: Effect of auto-enrollment on financial wealth at age 65 and on retirement consumption

(a) Universal auto-enrollment policy at 3% over a lifetime



(b) Universal auto-enrollment policy at 6% over a lifetime



Notes: In the left (right) two panels, each bar corresponds to the model-predicted percentage change in financial wealth at age 65 (present discounted value of consumption between ages 65 and 90) under an auto-enrollment policy adopted by all employers relative to that under an opt-in regime. I use the risk-free interest rate (3%) for discounting retirement consumption. The policies are simulated under the baseline preference estimates from column 1 of Table 4. Dark-shaded bars correspond to the results under the assumption that the incidence of increased matching costs falls on employers, while the light-colored bars correspond to the assumption that employers reduce the 401(k) matching rate to satisfy their budget constraint.

Average consumption gains in retirement are halved when the policy's incidence falls fully on workers. The average increase in the present value of retirement consumption from a universal auto-enrollment is cut from 0.6% to 0.3% under a 3% default (or from 1.4% to 0.7% under a 6% default) if employers reduce the matching rate to satisfy their budget constraint.

[30] I compute the present discounted value of consumption in retirement using the baseline calibration of the discount rate ($r = 3\%$). This measure also addresses the fact that liquid wealth and retirement wealth are taxed differently and converts changes in both measures of wealth into comparable units of consumption.

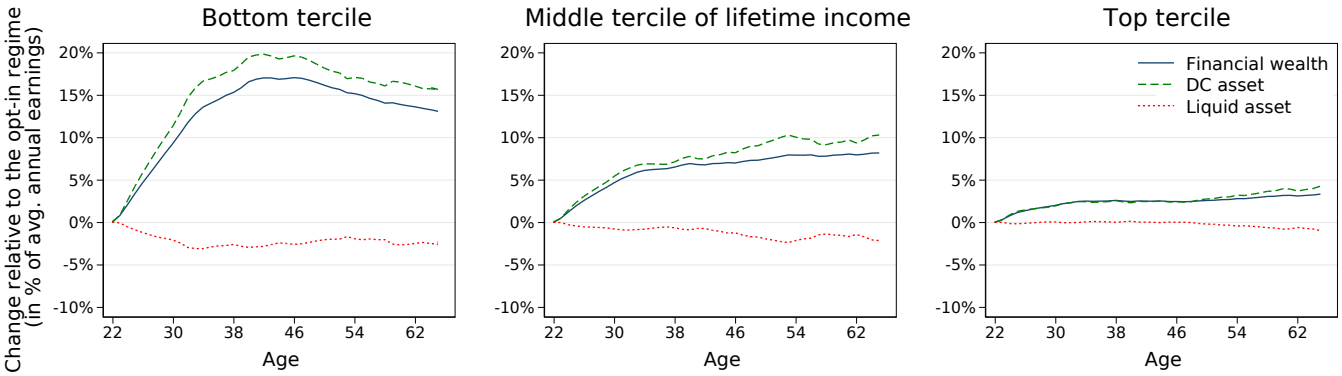
Lifecycle profile of auto-enrollment savings gains. Figure 12 shows that the savings increases created by auto-enrollment are concentrated in the first 20 years of working life and are largest around age 45: individuals in the bottom tercile of lifetime earnings hold an additional 20% of DC retirement wealth under auto-enrollment (Figure 12). In the second half of working life, however, the wealth effect dominates; auto-enrolled workers contribute less to the retirement savings account because they have accumulated more wealth early in life. At the bottom of the lifetime earnings distribution, nearly one-fourth of the increased savings created by auto-enrollment is undone by lower contributions between the ages of 45 and 65. This result suggests that, even with long-term data following individuals over 20 years, empirical estimates of the auto-enrollment effect on wealth may overstate the policy's impact on lifetime savings. A caveat to note here is that there is no direct empirical evidence on the long-term wealth effect from auto-enrollment, and the attenuation of savings increases over the course of a few years need not imply that long-term wealth effects are large.[31] In Appendix F.6, I show similar evidence for the 6% default policy and under different assumptions about the policy incidence.

Modest crowding-out of liquid savings. As shown in Figure 12, there is modest crowding-out of liquid assets by the auto-enrollment policy—most of the increases in retirement savings created by the policy result in increased total wealth accumulation. In the model, the liquid asset (i.e., liquid wealth and high-interest-rate unsecured borrowing) serves the precautionary saving motive, while the DC account serves the lifecycle saving motive. Because the two assets serve different purposes, there is limited substitution between the two. As in Kaplan and Violante (2014), the marginal propensity to consume out of liquid assets is high in the model, so auto-enrolled workers reduce their consumption in response to a decrease in disposable income created by the policy. This effect is strongest for low-income individuals, who have fewer liquid assets and are more liquidity constrained. For individuals in the bottom tercile of lifetime earnings, I find that every dollar of additional DC retirement wealth at age 65 created by a 3% auto-enrollment policy reduces liquid savings by only 16 cents. In contrast, while still modest, there is more crowding-out of liquid savings for individuals in the top tercile of lifetime earnings, as these individuals face less severe liquidity constraints. The model predicts a drop of 22 cents in liquid savings for every dollar of additional DC wealth at 65.[32]

[31] For example, the attenuation over the course of a few years might be driven by liquidity needs—an employee who is automatically enrolled has less liquid wealth in the short run and might seek to build liquid wealth by reducing their contribution rate (or refraining from increasing it) over the next few years.

[32] This finding of modest crowding-out aligns with the empirical evidence from Beshears et al. (2022) and Chetty et al. (2014). Beshears et al. (2022) find that auto-enrollment caused no significant increase in unsecured debt after

Figure 12: Effect of auto-enrollment on lifecycle DC savings, liquid assets, and net financial wealth



Notes: Each series corresponds to the model-predicted difference between assets under an auto-enrollment policy at 3% and an opt-in regime, adopted by all employers. The policies are simulated under the baseline preference estimates from column 1 of Table 4 and the earnings process estimated in SIPP data. Changes are expressed as a fraction of each simulated individual's average annual earnings between the ages of 22 and 65. The incidence of the policy is assumed to fall on employers, who reduce their profits to cover higher matching costs under auto-enrollment.

Welfare and optimal policy implications. I study the welfare impact of these counterfactual exercises in a companion paper (Choukhmane (2024)). I characterize the optimal default policy under alternative assumptions about social preferences and about the incidence of higher employer matching costs. While the assumptions in my dynamic model differ from those in previous studies, the resulting optimal policies—namely, either an opt-in regime or auto-enrollment with a default equal to the employer match threshold—are quite similar to those identified by Bernheim et al. (2015), Bernheim and Gastell (2024) and Goldin and Reck (2022). Specifically, if the policymaker shares the same time preferences as individuals and has no redistributive motive beyond declining marginal utility of consumption, I find that an opt-in regime is preferable to auto-enrollment because existing tax and matching incentives for retirement saving already cause individuals to consume more in retirement than implied by their time preferences alone. Hence, auto-enrollment, by inducing people to save even more for retirement, reduces welfare even when the incidence of higher employer matching costs falls fully on employers. On the other hand, if the policymaker is paternalistic (i.e., the social discount rate is lower than the private discount rate) or inequality averse (i.e., the planner puts more weight on the welfare of low-income individuals), a default contribution rate near the employer matching threshold can maximize social welfare. This result holds even when employers reduce wages or the match rate to balance their budget in response to the policy.

four years. Similarly, Chetty et al. (2014), studying a different policy in Denmark, estimate an 80% pass-through of retirement savings contributions to total savings. Choukhmane and Palmer (2024), using merged U.K. pension and bank account data, show that a similar model can replicate the heterogeneity in the consumption response to an increase in retirement contributions in the U.K.

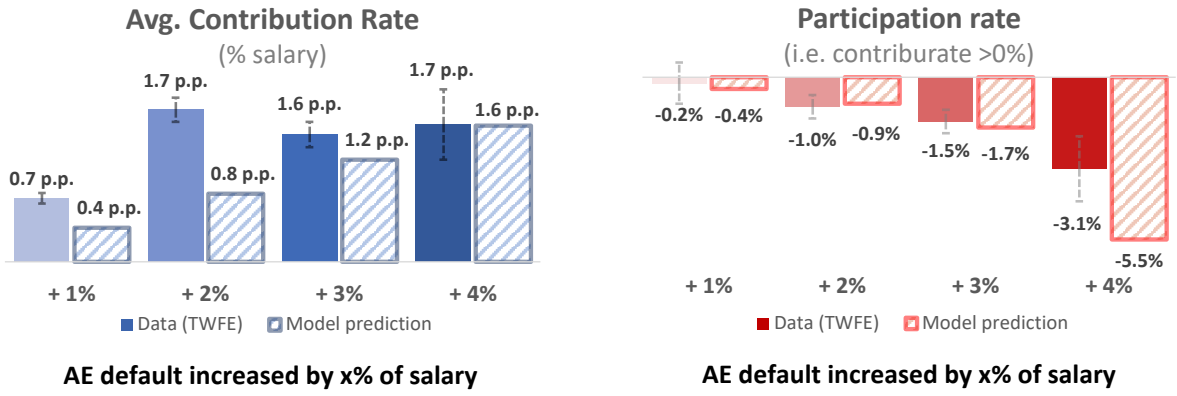
7 Stable Auto-enrollment Effects across Different Settings

The main advantage of structural estimation is the ability to extrapolate effects to a different policy, population, or institutional environment. In this section, I show that my estimates predict behavior reasonably well in two different contexts. The out-of-sample performance of the model gives credibility to its long-term predictions.

7.1 Out-of-sample validation I: Increasing the default contribution rate

The model—estimated with the introduction of auto-enrollment at 3% or 6% of salary—can predict behavior out of sample among the 101 U.S. firms used to document Fact VI of Section 3. These 101 U.S. firms increased their default contribution rate from a positive default to an even higher default. The model and data series are not directly comparable because the plans in the empirical sample vary in terms of their employer matching formula, while the model simulations assume a 50% match rate up to 6%. Nevertheless, the model replicates the main pattern observed in the data.[33] Specifically, the model predicts a modest rise in average contribution rates and a consistent drop in participation when default contribution rates increase (13). In the model, too few individuals contribute at positive amounts below the initial default, and the model cannot replicate the increase in contributions at the lowest rates observed in the data (right panel of Figure 6).

Figure 13: Empirical and simulated response to a higher auto-enrollment default



Notes: The "Data" series corresponds to regression coefficients from the dynamic event-study approach of Sun and Abraham (2021) in the sample described in Section 3.6. Standard errors are clustered at the plan level. The model series correspond to comparisons of simulated data for identical individuals in their first year of tenure under a 1%, 2%, 3%, or 4% auto-enrollment default using the baseline preference estimates from column 1 of Table 4 and baseline income process.

[33] Formally, the null hypothesis that the four coefficients on participation rates are equal to the model-implied effects is not rejected at conventional significance levels (p-value of the F-test: 0.24). However, the null hypothesis regarding the impact of increasing the default on contribution rates is rejected (p-value of the F-test: 0.00).

7.2 Out-of-sample validation II: U.K. national auto-enrollment policy

I exploit the rollout of the U.K. national auto-enrollment policy to validate my estimates in a different population and institutional environment (details of the policy are described in Appendix H).

U.K. calibration. I set the time preference parameters equal to their baseline values estimated in the U.S. sample ($\delta = 0.989$ and $\sigma = 0.52$), and I set the switching cost $k =$ £171 to maintain the same ratio between the switching cost and median earnings as in the baseline U.S. estimate. I recalibrate the model to match the parameters of the U.K. tax, unemployment insurance, and pension systems as well as the heterogeneity in employers' contribution formulas across firms and the parameters of the age-income profile, both estimated in the ASHE data. Details on the data and calibration of the U.K. counterfactual are discussed in Appendix H.

Empirical evidence. I compare workers' participation and contribution behavior before and after the policy rollout. The policy year corresponds to the first survey year in which the policy is binding for a given firm.[34] I compare the contributions after the policy implementation to the contributions in the same firms two years prior to the policy.[35]

Model exercise. In the model, the policy is implemented at a randomly chosen quarter of employment t^{exp}. Prior to t^{exp}, all employers offer an opt-in regime. After t^{exp}, all employees receive a 1% employer contribution, and those with a zero default contribution rate (both new hires and seasoned employees who did not contribute in period $t^{exp} - 1$) see their default contribution rate increase to 1% of income. They can choose to contribute this amount or incur the switching cost and change their contribution rate. The policy is unexpected, and after t^{exp}, all individuals expect the new auto-enrollment regime to persist in the future.

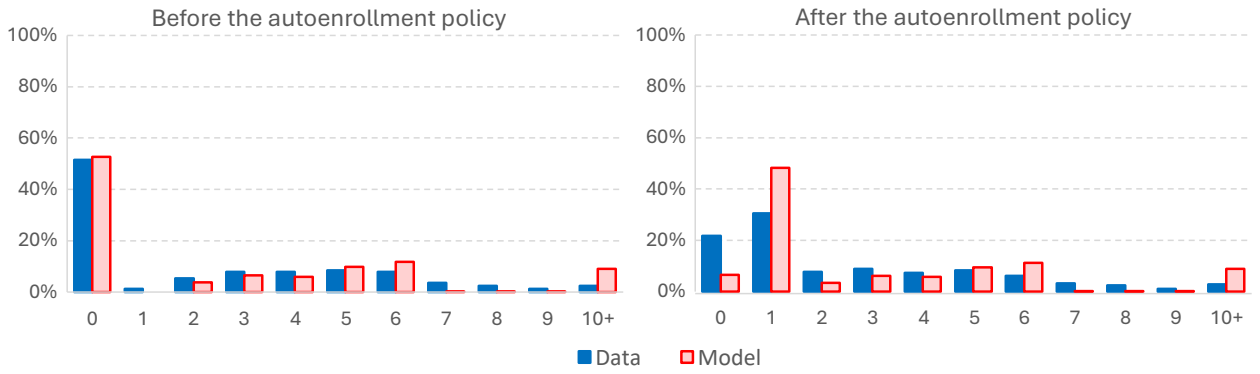
Results. The model with parameters estimated in the U.S. predicts reasonably well workers' contribution behavior in the context of the U.K. national auto-enrollment policy. As shown in Figure 14, the model matches the change in the distribution of contribution rates but overestimates the bunching at the 1% auto-enrollment default.[36]

[34] An alternative approach to this event study is to exploit the staggered implementation of the policy by employer size. Cribb and Emmerson (2016) document similar patterns in a difference-in-difference specification taking advantage of this policy variation. I favor a single-difference approach because it is easier to compare to model simulations.

[35] The results are similar when I use data from three, four or five years prior to the policy. I exclude the year prior to the policy because some firms could have adopted auto-enrollment a few months ahead of the legal deadline.

[36] The Kolmogorov–Smirnov statistic, which represents the maximum absolute difference between the empirical and model-implied CDFs, is 6.6% for the opt-in distribution and 15.1.% for the auto-enrollment distribution. The null hypothesis that the empirical and model-implied distributions are identical is rejected.

Figure 14: Distributions of employees' contribution rates before and after the U.K. national policy



Notes: The empirical series correspond to workers observed two years before the auto-enrollment mandate became binding at a given firm (left panel) and in the first survey year after the policy rollout (right panel). Data source: U.K. ASHE waves 2010 to 2016.

Across-job evidence. In Appendix Section H.3, I compare the model-predicted effects after a job transition to the empirical evidence in Fact VII of Section 3. The model predicts that auto-enrollment causes workers to participate 8.4% less and contribute 0.1 p.p. of salary less in their next employer's opt-in retirement savings plan. However, in line with the empirical evidence, there is almost no negative impact on contributions when the new job has auto-enrollment: contributions decline by only 0.01 p.p. of salary and participation by 2.8%.

8 Conclusion

A growing literature has documented, in a variety of settings, that individuals stick with default options and make infrequent switching decisions.[37] In the context of contributions to retirement savings plans, this paper has shown the importance of studying decision-making inertia in a fully dynamic setting. When individuals can compensate for their inaction with other actions in the future, small switching costs generate large inertia, and the observed short-run effects may overstate the long-term consequences of inaction. These insights are likely to extend to different degrees to other settings where large switching costs have been estimated. However, even though the average effect of inertia on savings is largely attenuated over time, changing saving default options can have desirable distributional properties by helping those at the bottom of the savings distributions increase their retirement savings and receive more employer contributions.

[37]For instance, in the retirement savings context, significant inertia and large switching costs have been estimated for switching between retirement funds (Illanes, 2017; Dahlquist et al., 2018) or between a defined benefit and a defined contribution plan (Dobrescu et al., 2018).

References

Altonji, J. G. and L. M. Segal (1996). Small-sample bias in gmm estimation of covariance structures. *Journal of Business & Economic Statistics 14*(3), 353–366.

Andrews, I., M. Gentzkow, and J. M. Shapiro (2017). Measuring the sensitivity of parameter estimates to estimation moments. *The Quarterly Journal of Economics 132*(4), 1553–1592.

Arnoud, A., T. Choukhmane, J. Colmenares, C. O'Dea, and A. Parvathaneni (2021). The evolution of u.s. firms retirement plan offerings: Evidence from a new panel data set. *Working Paper*.

Bernheim, B. D., A. Fradkin, and I. Popov (2015). The welfare economics of default options in 401 (k) plans. *American Economic Review 105*(9), 2798–2837.

Bernheim, D. and J. M. Gastell (2024). Optimal default options: The case for (weighted) opt-out minimization. *Journal of Public Economics 237*(105203).

Beshears, J., J. J. Choi, D. Laibson, and B. C. Madrian (2009). The importance of default options for retirement saving outcomes. pp. 167–195. Social security policy in a changing environment.

Beshears, J., J. J. Choi, D. Laibson, B. C. Madrian, and W. L. Skimmyhorn (2022). Borrowing to save? the impact of automatic enrollment on debt. *Journal of Finance 77*(1), 403–447.

Best, M., J. Cloyne, and E. Ilzetzki (2020). Estimating the elasticity of intertemporal substitution using mortgage notches. *Review of Economic Studies 87*(2), 656–690.

BLS (2019). Number of jobs, labor market experience, and earnings growth:results from a national longitudinal survey. Technical report, Bureau of Labor Statistics.

Blumenstock, J., M. Callen, and T. Ghani (2018). Why do defaults affect behavior? experimental evidence from afghanistan. *American Economic Review 108*(10), 2868–2901.

Brot-Goldberg, Z., T. Layton, B. Vabson, and A. Y. Wang (2023). The behavioral foundations of default effects: theory and evidence from medicare part d. *American Economic Review 113*(10), 2718–2758.

Bubb, R. and P. L. Warren (2020). An equilibrium theory of retirement plan design. *American Economic Journal: Economic Policy 12*(2), 22–45.

Butrica, B. and N. Karamcheva (2012). Automatic enrollment, employee compensation, and retirement security. *Working Paper*.

Carroll, G. D., J. J. Choi, D. Laibson, B. C. Madrian, and A. Metrick (2009). Optimal defaults and active decisions. *The quarterly journal of economics 124*(4), 1639–1674.

Chetty, R., J. N. Friedman, S. Leth-Petersen, T. H. Nielsen, and T. Olsen (2014). Active vs. passive decisions and crowd-out in retirement savings accounts: Evidence from denmark. *The Quarterly Journal of Economics 129*(3), 1141–1219.

Choi, J., D. Laibson, B. C. Madrian, and A. Metrick (2006). Saving for retirement on the path of least resistance. In *Behavioral Public Finance: Toward a New Agenda*, pp. 304–351.

Choi, J. J., D. Laibson, and B. C. Madrian (2009). Reducing the complexity costs of 401 (k) participation through quick enrollment. In *Developments in the Economics of Aging*, pp. 57–82.

Choi, J. J., D. Laibson, B. C. Madrian, and A. Metrick (2004). For better or for worse: Default effects and 401 (k) savings behavior. In *Perspectives on the Economics of Aging*, pp. 81–126.

Choukhmane, T. (2024). Optimal saving nudges. *Working Paper*.

Choukhmane, T. and T. de Silva (2023). What drives investors portfolio choices? separating risk preferences from frictions. Technical report, Working paper.

Choukhmane, T. and C. Palmer (2024). How do consumers finance increased retirement savings? *Working Paper*.

Clark, J., S. Utkus, and J. Young (2015). Automatic enrollment: The power of the default. *Vanguard*.

Cribb, J. and C. Emmerson (2016). What happens when employers are obliged to nudge? automatic enrolment and pension saving in the uk. *IFS Working Paper*.

Dahlquist, M., O. Setty, and R. Vestman (2018). On the asset allocation of a default pension fund. *The Journal of Finance 73*(4), 1893–1936.

DellaVigna, S. (2009). Psychology and economics: Evidence from the field. *Journal of Economic literature 47*(2), 315–372.

DellaVigna, S. (2018). Structural behavioral economics. In *Handbook of Behavioral Economics: Foundations and Applications*, Volume 1, pp. 613–723.

Derby, E., K. Mackie, and J. Mortenson (2023). Worker and spousal responses to automatic enrollment. *Journal of Public Economics 223*(104910).

Dobrescu, L. I., X. Fan, H. Bateman, B. R. Newell, A. Ortmann, and S. Thorp (2018). Retirement savings: A tale of decisions and defaults. *The Economic Journal 128*(610), 1047–1094.

Falk, J. and N. S. Karamcheva (2023). The impact of an employer match and automatic enrollment on the savings behavior of public-sector workers. *Journal of Pension Economics & Finance*.

Goldin, J. and D. Reck (2022). Optimal defaults with normative ambiguity. *Review of Economics and Statistics 104*(1), 17–33.

Gourinchas, P.-O. and J. Parker (2002). Consumption over the life cycle. *Econometrica 70*(1), 47.

Havránek, T. (2015). Measuring intertemporal substitution: The importance of method choices and selective reporting. *Journal of the European Economic Association 13*(6), 1180–1204.

Illanes, G. (2017). Switching costs in pension plan choice. *Working Paper*.

Kaplan, G. and G. L. Violante (2014). A model of the consumption response to fiscal stimulus payments. *Econometrica 82*(4), 1199–1239.

Laibson, D., P. Maxted, S. Chanwook Lee, A. Repetto, and J. Tobacman (2024). Estimating discount functions with consumption choices over the lifecycle. *The Review of Financial Studies*.

Low, H., C. Meghir, and L. Pistaferri (2010). Wage risk and employment risk over the life cycle. *American Economic Review 100*(4), 1432–1467.

Madrian, B. C. and D. F. Shea (2001). The power of suggestion: Inertia in 401 (k) participation and savings behavior. *The Quarterly journal of economics 116*(4), 1149–1187.

O'Dea, C. (2018). Insurance, efficiency and the design of public pensions. *Working Paper*.

O'Donoghue, T. and M. Rabin (1998). Procrastination in preparing for retirement. In *Behavioral Dimensions of Retirement Economics*.

Sun, L. and S. Abraham (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics 225*(2), 175–199.

Thaler, R. H. and S. Benartzi (2004). Save more tomorrow: Using behavioral economics to increase employee saving. *Journal of political Economy 112*(1), 164–187.

APPENDIX FOR ONLINE PUBLICATION

A Parameters Definition

Table A.1 summarizes all the model parameters introduced in the paper (excluding the appendices).

Table A.1: Summary of model parameters

Preference Parameters		**Assets**	
δ	Discount factor	R^{dc}	Rate of return on DC wealth
σ	Elasticity of intertemporal substitution	$R^{liq}(\,)$	Rate of return on liquid assets
k	Switching cost	$1+r$	Interest rate on positive asset holdings
State variables		$1+r^{cc}$	Interest rate on unsecured debt
X_t	Vector of all state variables	$\underline{l}_a$	Borrowing limit
a	Age	**Defined Contribution (DC) Account**	
emp_t	Employment status	$E(\,)$	Distribution of DC account types
ten_t	Tenure	$ec^e(\,)$	Employer DC contribution
e	Employer DC plan type	$match^e$	Employer matching rate
θ	Labor productivity	cap^e	Threshold on employer matching
ae	Average lifetime earnings	$\Upsilon_e(\,)$	Vesting risk-adjustment
l	Liquid assets	**Labor Market**	
dc	DC wealth stock	$\pi^{JJ}(\,)$	Job-to-job transition probability
d	Default contribution rate	$\pi^{EU}(\,)$	Unemployment transition probability
$F(\,)$	Joint distribution of θ, emp, e	$\pi^{UE}(\,)$	Out-of-unemployment transition probability
Choices		$\{\delta_i\}_{i=0}^3$	Deterministic component of earnings
c	Consumption	w	Labor earnings
s	DC contribution rate	ρ	Autocorrelation in earnings shocks
$draw$	DC withdrawal rate	ξ	Earnings innovation if continuously employed
Utility		$\sigma^2_{\xi_1}$	Variance of the first earnings innovation
$u_a(\,)$	Utility function	σ^2_{ξ}	Variance of subsequent innovations
$V(\,)$	Value function	ζ	Earnings innovation after job-to-job transition
$V^S(\,)$	Value function cdt. on DC contrib. s	μ^{JJ}	Avg. wage gain after a job-to-job transition
Demographics		ς_t	Earnings innovation out of unemployment
A	Maximum age	μ^{UE}	Avg. wage loss out of unemployment
A^R	Retirement age	ι	Measurement error in earnings
m_a	Mortality risk	σ^2_ι	Variance of measurement error
n_a	Equivalence scale	η	Earnings innovation plus measurement error
Aggregate budget constraint		**Tax and benefit system**	
$\overline{Y}$	Employers' aggregate budget constraint	$tax^i(\,)$	Tax on income
$Pf(\,)$	Aggregate employer profits	τ^k	Tax on capital returns
$W(\,)$	Aggregate wages	$limit_a$	Tax limit on DC contributions
$Mtc(\,)$	Agg. employer matching contributions	pen_a	Tax penalty for early DC withdrawals
$\bar{y}(\,)$	Avg. individual income under opt-in	$ui(\,)$	Unemployment insurance benefit
		$ss(\,)$	Public pension income

B Additional Evidence for the Empirical Section

B.1 Additional details on the event-study sample contruction

Plan selection. The event-study sample includes 401(k) plans that satisfy the following criteria:

- The plan must have adopted auto-enrollment for new hires with a default contribution rate of either 3% or 6% of salary.
- The plan must have a minimum of five years of data, starting from the year prior to the adoption of the auto-enrollment policy (t-1) and a minimum of three years of data after the policy adoption (t+3). This ensures that the opt-in and auto-enrollment cohorts are observed for at least 36 months of tenure.
- The plan must have new hires in at least six of the twelve months prior to and six of the twelve months after the adoption of the auto-enrollment policy. This excludes smaller plans with few employees hired during the event-study window and plans with hiring concentrated in a small number of months.
- The plan must offer an employer match for employee contributions up to 6% of salary throughout the event-study period. This is the most common matching cap both in this dataset and for nearly half of U.S. 401(k) plans (Arnoud et al., 2021).

Sample restrictions. I further restrict the sample to include continuously-employed workers who satisfy the following three criteria. The goal of these restrictions is to ensure compliance with the auto-enrollment treatment.

(i) I exclude employees hired in the month of auto-enrollment adoption (month 0).

(ii) I exclude employees hired under auto-enrollment but auto-enrolled more than 3 months after their date of hire.

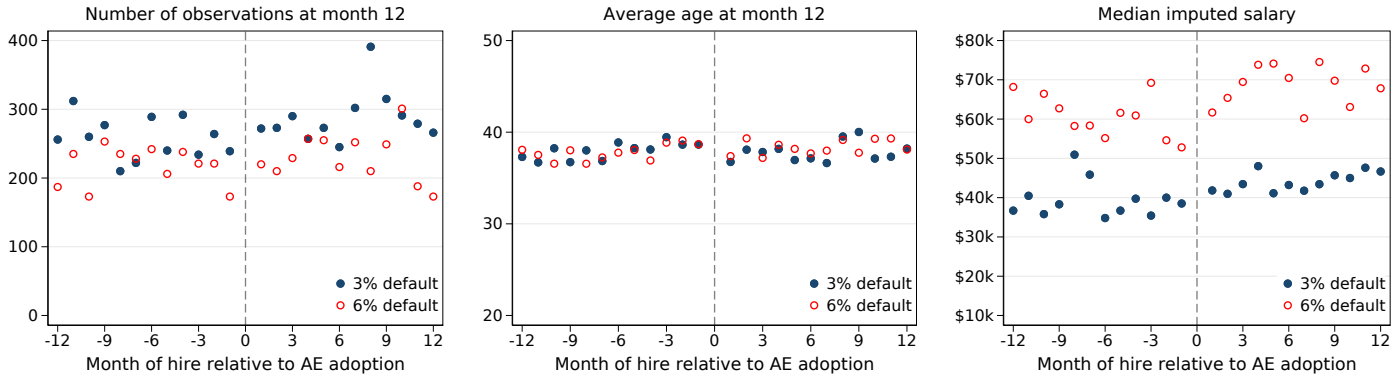
(iii) I exclude employees who are auto-enrolled despite being hired prior to the adoption of the auto-enrollment policy.

The resulting sample is described in Table 1. I show in Appendix B.2.3 that the empirical results are similar without these three sample restrictions, but the effects are more difficult to interpret due to imperfect compliance with the policy.

Balance check. I define the month of hire relative to the adoption window as each succesive 31 calendar days relative to the month of auto-enrollment adoption. Following the approach initated

by Madrian and Shea (2001), I compare employees hired in the 12 months prior to the month of auto-enrollment adoption (opt-in cohorts) to those hired in the 12 months after the policy adoption (auto-enrollment cohorts). Figure A.1 shows that both cohorts are balanced in terms of number of observations, age and salary.

Figure A.1: Observable characteristics: Auto-enrollment vs. opt-in cohort (12th month of tenure)



Notes: The horizontal axis shows the normalized month of hire relative to the time of auto-enrollment adoption. In the left panel, each dot corresponds to the number of observations for employees in their 12th month of tenure by date of hire relative to the month of the policy adoption. The middle panel corresponds to the average age. The right panel corresponds to estimated annual median salary. Salary information is available only for employees with a positive contribution rate, and I assume that all nonparticipating employees have below-median salaries.

Main variables. Below, I discuss the definition of the main variables:

- Participation is defined as an employee's making a positive contribution to the 401(k) plan in a given month. Participation rates in the first few months of tenure are biased upward because I do not observe workers who never joined the 401(k) plan and left their employer before the month of December of their first year of tenure. After 12 months of tenure, every nonparticipating employee has been observed at least once in December.

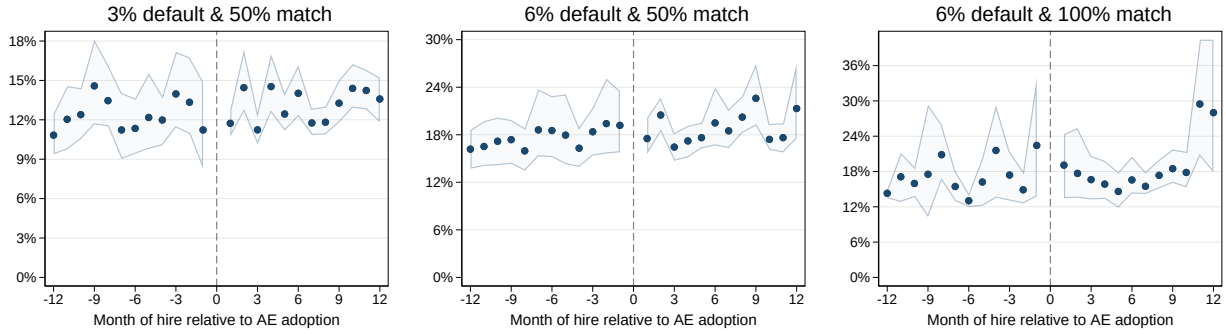
- Cumulative contributions are computed as the sum of all employee (or combined employee and employer) contributions made to the 401(k) plan since the date of hire divided by the employee's current annual salary. Values are winsorized when the ratio of cumulative contributions to cumulative earnings is above 1 (i.e., implying a savings rate above 100%).

- Earnings information is available only for employees with a positive contribution rate. When computing median earnings, I assume that all nonparticipating employees have below-median earnings.

B.2 Additional evidence on Fact II

B.2.1 Plans with different matching formulas

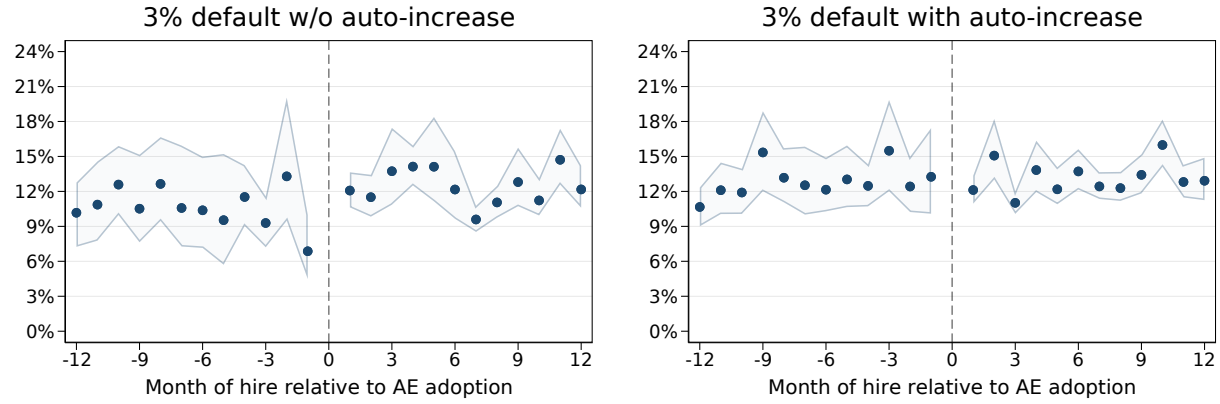
Figure A.2: Effect of auto-enrollment on cumulative employee contributions after 36 months of tenure by employer matching formula



Notes: The horizontal axis shows the normalized month of hire relative to the month of auto-enrollment adoption. Each panel shows average cumulative employee 401(k) contributions observed in the 36th month of tenure divided by individuals' annual salary for plans with different employer matching formulas. The left (middle) panel corresponds to plans with a 50% match rate up to 6% of salary and a 3% (6%) auto-enrollment default contribution rate. The right panel corresponds to plans with a 100% match rate up to 6% of salary and a 6% auto-enrollment default contribution rate. Shaded areas correspond to the bootstrapped 95% confidence intervals stratified at the plan level for each series.

B.2.2 Plans with and without an auto-escalating default option

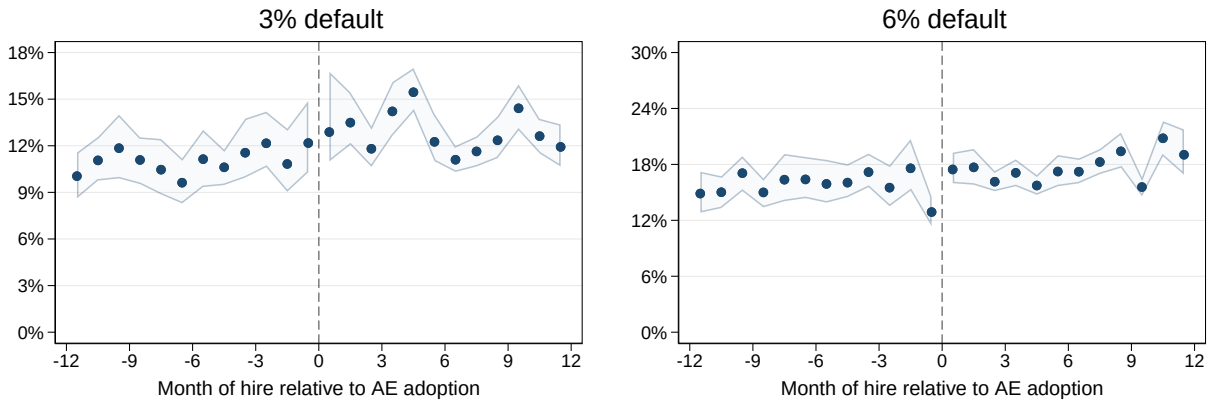
Figure A.3: Effect of auto-enrollment on cumulative employee contributions after 36 months of tenure with and without auto-escalation



Notes: The horizontal axis shows the normalized month of hire relative to the month of auto-enrollment adoption. The left (right) panel shows average cumulative employee 401(k) contributions observed in the 36th month of tenure divided by individuals' annual salary in plans adopting auto-enrollment at 3% without (with) an auto-esclating default contribution rate. Shaded areas correspond to the bootstrapped 95% confidence intervals stratified at the plan level for each series.

B.2.3 Results without baseline sample restrictions

Figure A.4: Effect of auto-enrollment on cumulative employee contributions after 36 months of tenure without the baseline sample restrictions

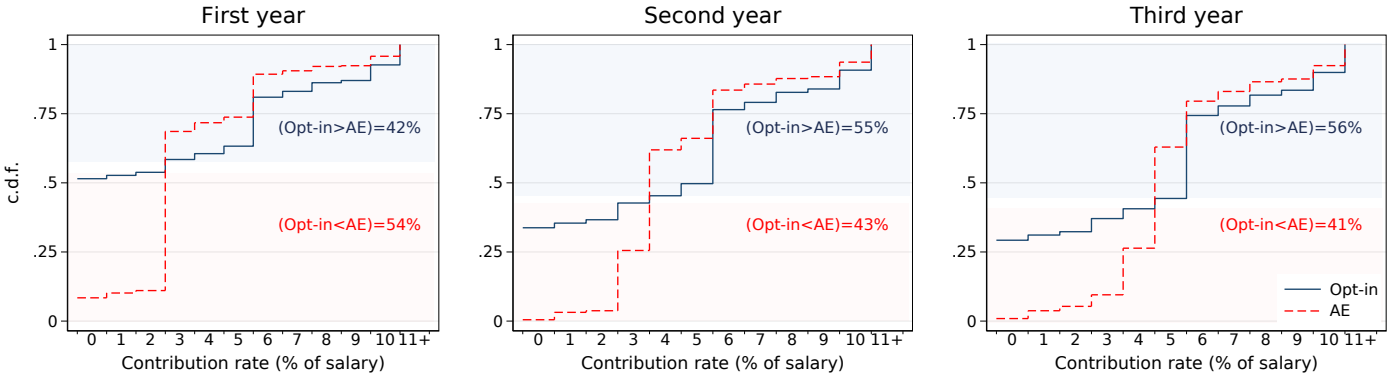


Notes: The horizontal axis shows the normalized month of hire relative to the calendar date of auto-enrollment adoption. The left (right) panel shows average cumulative employee 401(k) contributions observed in the 36th month of tenure divided by individuals' annual salary in plans adopting auto-enrollment with a 3% (6%) default contribution rate. Event-study without the three sample restrictions described in Appendix B.1. Shaded areas correspond to the bootstrapped 95% confidence intervals stratified at the plan level for each series.

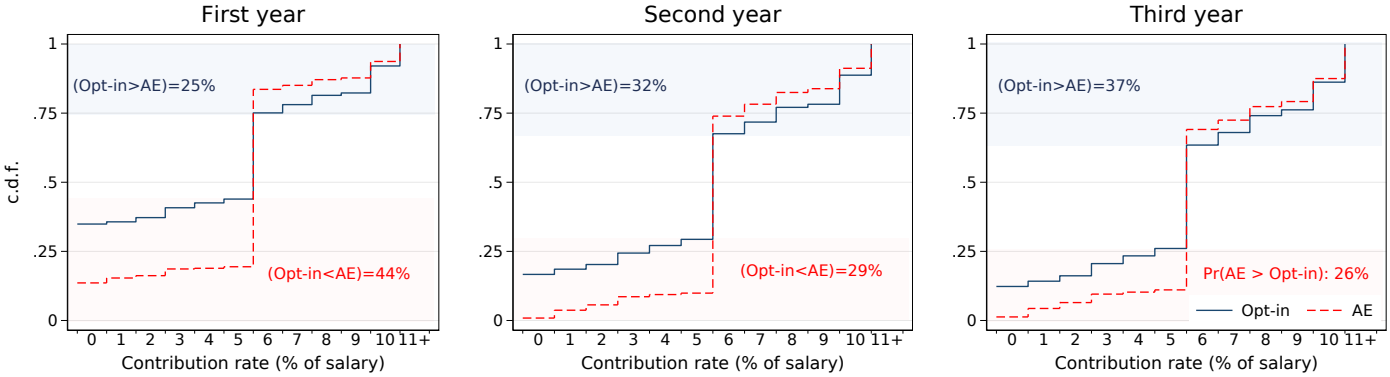
B.3 Additional evidence on Fact V

Figure A.5: Distribution of contribution rates over tenure

(a) Plans with a 3% default with auto-escalation



(b) Plans with a 6% default



Notes: Each line corresponds to the cumulative distribution function (CDF) of contribution rates as a % of salary recorded in December of the employee's first year (left panels), second year (middle panels), and third year (right panels) of tenure. The solid line corresponds to employees hired in the 12 months prior to the adoption of auto-enrollment (the opt-in cohort), and the dashed lines correspond to the distribution for employees hired in the 12 months following auto-enrollment adoption (the AE cohort).

B.4 Additional evidence for Fact VI

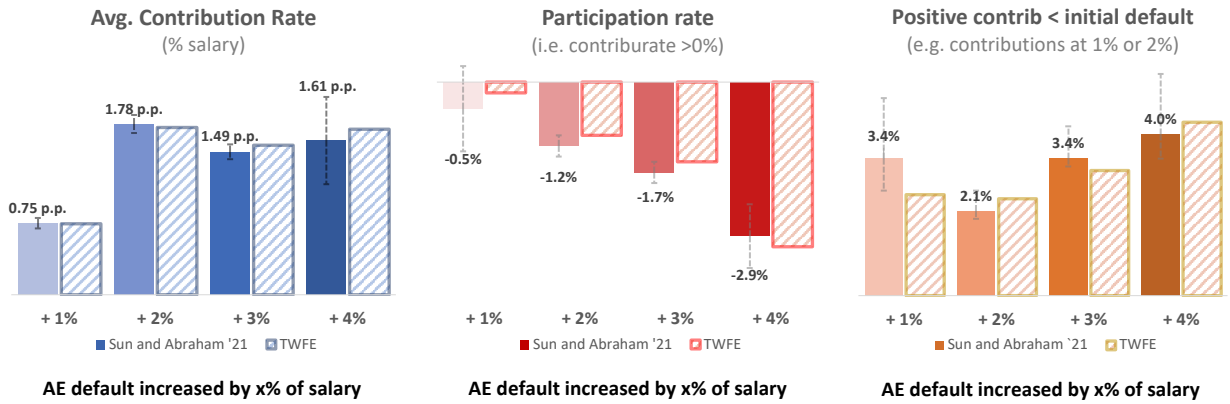
Figure A.6 shows that the evidence presented in Fact VI is robust to my using the dynamic event-study approach of Sun and Abraham (2021), which is robust to treatment effect heterogeneity. Table A.2 shows the regression coefficient estimates under both approaches.

Table A.2: Effect of a higher auto-enrollment default on participation and contributions

	Average contribution rate (share of salary)			**Participation (i.e. contribution > 0)**			**Contribution below initial default (inc. 0)**		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	TWFE	TWFE	SunAbraham	TWFE	TWFE	SunAbraham	TWFE	TWFE	SunAbraham
Increase in the default (p.p.)	0.613			-0.006			0.012		
	(0.071)			(0.001)			0.003		
Increase in default = 1%		0.742	0.747		-0.002	-0.005		0.025	0.034
		(0.063)	(0.054)		(0.007)	(0.008)		(0.012)	(0.015)
Increase in default = 2%		1.748	1.784		-0.01	-0.012		0.024	0.021
		(0.14)	(0.095)		(0.004)	(0.002)		(0.006)	(0.005)
Increase in default = 3%		1.563	1.494		-0.015	-0.017		0.031	0.034
		(0.148)	(0.078)		(0.004)	(0.002)		(0.01)	(0.008)
Increase in default = 4%		1.728	1.611		-0.031	-0.029		0.043	0.040
		(0.413)	(0.456)		(0.011)	(0.006)		(0.018)	(0.015)
Ln(age)	1.155	1.159	1.158	-0.005	-0.005	-0.005	-0.004	-0.004	-0.004
	(0.173)	(0.173)	(0.172)	(0.002)	(0.002)	(0.01)	(0.003)	(0.003)	(0.002)
Ln(tenure)	0.328	0.331	0.332	0.014	0.014	0.014	0.01	0.01	0.01
	(0.050)	(0.051)	(0.051)	(0.01)	(0.01)	(0.002)	(0.002)	(0.002)	(0.003)
Year FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Plan FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Observations	764,006	764,006	764,006	764,006	764,006	764,006	764,006	764,006	764,006
R-squared	0.149	0.149	0.149	0.014	0.014	0.014	0.074	0.074	0.075

Notes: Each column reports coefficients from a regression where the dependent variable is a dummy for the outcome in the column heading. The coefficients correspond to the difference between employees hired before and after each firm increased its auto-enrollment default contribution rate. Each data point corresponds to an employee observed in her first 12 months of tenure, and after the end of the auto-enrollment grace period. Columns 1, 2, 4, 5, 7 and 8 are estimated using a two-fixed effect regression, while columns 3, 6 and 9 are the result of the dynamic event-study estimator of Sun and Abraham (2021). Standard errors clustered at the firm level are in parentheses.

Figure A.6: Effect of a higher auto-enrollment default on participation and contributions



Notes: Each column reports coefficients from a two-way fixed effect regression (solid shading) or the dynamic event-study approach of Sun and Abraham (2021) (diagonal stripes) for employees observed in their first 12 months of tenure and after the end of the auto-enrollment grace period. The independent variables are dummies equal to one for employees hired after the default contribution rate was increased by 1, 2, 3, or 4 percentage points in 101 US 401(k) plans. The sample also includes employees in 245 never-treated 401(k) plans with no change to their auto-enrollment default option. Controls include year fixed effects, plan fixed effects, log tenure (in days), and log age. Standard errors are clustered at the plan level.

B.5 Additional evidence on Fact VII

In this appendix, I provide additional evidence on the effect of auto-enrollment after a job transition, as described in Section 3.7.

First, I run the same specification as in Table 3 under the assumption that the policy rollout started before the actual policy implementation date in 2012. The results of these placebo tests, reported in columns 2 to 8 of Appendix Table A.3, are all insignificant at the 5% level for both the participation and contribution outcomes, and a single coefficient is significant at the 10% level. These results help validate the empirical strategy and suggest that the estimates capture the effect of the auto-enrollment rollout rather than the effect of an employee's having moved from a larger to a smaller employer.

Table A.3: Auto-enrollment effect after a job transition

	Actual policy	**Placebo tests**						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Beginning of policy rollout	2012	2005	2006	2007	2008	2009	2010	2011
				Panel A – Participation rate				
AE to non-AE employer	-0.128	0.077	0.034	0.014	0.016	0.053	0.016	-0.061
	(0.055)	(0.058)	(0.037)	(0.051)	(0.049)	(0.058)	(0.051)	(0.086)
AE to AE employer	0.010	0.001	0.007	0.014	0.008	0.016	0.028	0.021
	(0.019)	(0.017)	(0.014)	(0.014)	(0.014)	(0.014)	(0.014)	(0.017)
				Panel B – Contribution rate (in percentage of pay)				
AE to non-AE employer	-0.546	-0.001	-0.018	0.261	-0.139	0.003	-0.269	-0.209
	(0.226)	(0.209)	(0.187)	(0.541)	(0.221)	(0.244)	(0.217)	(0.336)
AE to AE employer	-0.039	-0.013	-0.026	-0.04	-0.001	0.052	-0.026	-0.035
	(0.068)	(0.093)	(0.073)	(0.064)	(0.062)	(0.068)	(0.099)	(0.069)
Employee characteristics	✓	✓	✓	✓	✓	✓	✓	✓
$\text{Size}_{e'} \times \text{Size}_e$	✓	✓	✓	✓	✓	✓	✓	✓
Employer×Year FEs	✓	✓	✓	✓	✓	✓	✓	✓
Observations	37120	37120	37120	37120	37120	37120	37120	37120

Notes: Each cell reports coefficients from a regression whose dependent variable is a dummy for workers' participation status in the retirement savings plan (Panel A) or contribution rate in percentage points of pensionable pay (Panel B). The sample contains job switchers in their first year of tenure. The coefficients correspond to the effect of the automatic enrollment mandate in the previous employer on contributions when the current employer is not subject to the auto-enrollment requirement (AE to non-AE employer) and when the current employer is subject to automatic enrollment (AE to AE employer). In column 1, I assign an enrollment regime to each employer (i.e., opt AE or non-AE) based on the U.K. national auto-enrollment policy rollout schedule (in Appendix H). In columns 2 to 8, I assign enrollment regimes based on a (placebo) policy rollout schedule that starts earlier than the actual policy. I assume that the policy rollout started in the year reported in the column heading. Controls for employees' characteristics include controls for log salary and log tenure in the current job, log salary and log tenure when the employee is last observed in the previous job, a dummy for gender, and a third-order polynomial in age. The sample is restricted to private-sector workers between the ages of 22 and 60. Standard errors clustered by current employer are in parentheses. Data source: U.K. ASHE waves 2006 to 2017.

In the model, I interpret the difference in participation and contribution rates between those who were previously auto-enrolled and those who were not as reflecting a wealth effect (more saving in the previous job reduces the need to save in the current job). However, previously auto-enrolled wokers may be better informed about the future roll-out of the policy and may expect to eventually be automatically enrolled by their new employer. They may therefore choose not to opt into their new employer's retirement savings plan because they expect to be automatically enrolled soon anyway. This mechanism should affect mainly those whose new employer will soon become subject to the auto-enrollment mandate. To test this hypothesis in Table A.4, columns 3 and 6, I interact the

variable for a transition to an opt-in employer with a dummy equal to 1 if the new employer is expected to become subject to automatic enrollment within the next 12 months. While the results are noisy and not statistically significant, previously auto-enrolled workers who do not expect to be auto-enrolled in their new job within a year reduce their participation by 7.5 percentage points and their contributions by 0.12% of income, in line with the model predictions in 7.2 (a model-implied drop in participation of 9.6% and in contributions of 0.11% of salary). The drop in participation and contributions for those whose employer was scheduled to adopt auto-enrollment within 12 months is significantly larger (participation drops by an additional 11.6 percentage points and contributions by an additional 0.93% of income), which suggests that differential expectations about the policy rollout may play a role.

Table A.4: Auto-enrollment effect after a job transition

	Participation rate			**Contribution rate**		
	(1)	(2)	(3)	(4)	(5)	(6)
AE to non-AE employer	-0.065	-0.128	-0.075	-0.267	-0.538	-0.122
	(0.031)	(0.055)	(0.052)	(0.108)	(0.227)	(0.266)
AE to non-AE employer ×			-0.116			-0.927
AE to be adopted within a year			(0.092)			(0.517)
$\text{Size}_{e'} \times \text{Size}_e$		✓	✓		✓	✓
Employee characteristics	✓	✓	✓	✓	✓	✓
Employer×Year FE	✓	✓	✓	✓	✓	✓
Observations	37120	37120	37120	37120	37120	37120

Notes: Each cell reports coefficients from a regression whose dependent variable is a dummy for workers' participation status in the retirement savings plan (columns 1, 2 and 3) or contribution rate in percentage points of pensionable pay (columns 4, 5, and 6). The sample contains job switchers in their first year of tenure. I assign an enrollment regime to each employer (i.e., AE or non-AE) based on the U.K. national auto-enrollment policy rollout schedule (in Appendix H). The coefficients for "AE to non-AE employer" correspond to the effect of the automatic enrollment mandate in the previous employer on contributions when the current employer is not subject to the auto-enrollment requirement. In columns 3 and 6, the variable "AE to be adopted within a year" is a dummy equal to 1 if the current employer is expected to be subject to automatic enrollment in the 12 months following the date of observation. Controls for employees' characteristics include controls for log salary and log tenure in the current job, log salary and log tenure when the employee is last observed in the previous job, a dummy for gender, and a third-order polynomial in age. The sample is restricted to private-sector workers between the ages of 22 and 60. Standard errors clustered by current employer are in parentheses. Data source: U.K. ASHE waves 2006 to 2017.

C A theoretical framework for comparing alternative theories

I begin with a simple theoretical framework to compare the predictions of different behavioral theories that can potentially explain inertia at the savings default. This is a simplified version of the full

model developed in Section 4.

A worker contributes a fraction s of her wage w to a retirement savings account. She faces a default contribution rate d and selects an optimal contribution choice s^* from a discrete set S such that:

$$s^* = \underset{s \in S}{\operatorname{argmax}} \ \{U(s, d) + \delta V_{t+1}(s)\}$$

where δ is the intertemporal discount factor, U is flow utility, and V_{t+1} is the continuation value, which is strictly increasing in contributions to the retirement savings account. I use this framework to compare the predictions of three possible explanations of the default effect: switching costs, loss aversion and psychological anchoring.

C.1 A model of switching costs fits the empirical evidence

A first explanation for the default effect is that employees have to incur real or behavioral costs to depart from the savings default. As discussed later in Section 4, these switching costs can capture a broad class of neoclassical and behavioral explanations of the default effect such as the cost of completing a form or seeking out advice from a professional adviser, as well as the cognitive cost of paying attention or finding the optimal saving choice.

I assume that workers incur a cost k to choose a contribution rate other than the default:

$$U(s, d) = \begin{cases} u((1-s)w - k) & \text{if } s \neq d \\ u((1-s)w) & \text{if } s = d \end{cases}$$

where the utility function u is increasing and strictly concave in take-home pay $(1-s)w$.

Proposition 1 shows that, in line with the empirical evidence, increasing the default contribution rate in a model with a fixed switching cost causes more people to drop out of the retirement savings plan entirely or to contribute at the lowest rates.[1]

Proposition 1. *If V_{t+1} is concave, increasing the default contribution rate from $\underline{d}$ to $\overline{d}$ weakly increases contributions strictly below $\underline{d}$ in a model with a fixed switching cost.*

To understand the intuition behind Proposition 1, consider an employee whose preferred contribution rate is 1% of salary. When the default contribution rate is 3%, this employee may stay at

[1] This pattern of behavior, dubbed the "drop-out" effect" by Caplin and Martin (2017), has been documented in other settings. For instance, in the case of charitable giving (Altmann et al., 2019) and tips for taxi rides (Haggag and Paci, 2014), increasing the default led more people to not contribute at all.

the default to avoid incurring the switching cost because the default option is close enough to her preferred contribution rate. However, if the contribution rate is raised to 6%, the same employee may choose to incur the cost and switch to 1% because the higher default is further away from her preferred contribution rate. Therefore, this employee reduces her contribution rate from 3% to 1% when the default is increased from 3% to 6%. [2] Later, in Section 7.1, I show that the switching cost model not only makes the right qualitative prediction but also predicts the magnitude of employees' response to increasing the default.

C.2 Loss aversion yields the opposite prediction

Loss aversion is a common explanation for individuals' propensity to stay at the status quo (Samuelson and Zeckhauser, 1988). Employees may perceive the default as a reference point around which gains and losses are evaluated. For instance, they may perceive an actively chosen additional contribution to the retirement savings plan as a loss in take-home pay, whereas resources that are automatically contributed to the savings account may not feel like losses. To capture this idea, I adopt a functional form similar to that in DellaVigna et al. (2017) and define flow utility as:

$$U(s,d) = \begin{cases} u((1-s)w) + \eta(v(s) - v(d)) & \text{if } s < d \\ u((1-s)w) + \eta\lambda(v(s) - v(d)) & \text{if } s \geq d \end{cases}$$

where $v(.)$ is decreasing in the contribution rate and reflects gain–loss utility.[3] An employee who contributes below the default contribution rate has more resources than her reference level and experiences gain utility $v(s) - v(d) > 0$ with weight $\eta \geq 0$. An employee who contributes above the default has fewer resources than her reference level and experiences loss utility $v(s) - v(d) < 0$ with a weight $\eta\lambda$. The parameter $\lambda \geq 1$ captures loss aversion.

Proposition 2 shows that, under loss-averse preferences, increasing the default contribution rate reduces the fraction of employees contributing at the lowest rates.

Proposition 2. *Under loss-averse preferences, increasing the default contribution rate from* $\underline{d}$ *to* $\overline{d}$

[2] The proof relies on the assumption that V_{t+1} is concave in the retirement contribution rate. This assumption may not always hold because of the discrete nature of the contribution choice decision. In practice, if there is enough uncertainty about the future, the continuation value V_{t+1} is generally concave.

[3] A natural candidate functional form is $v = u((1-s)w)$, where take-home pay serves as a consumption reference point. This specification satisfies the condition that v is decreasing in s.

weakly decreases contributions strictly below $\underline{d}$*:*

$$Pr\left(s^* < \underline{d} \,|\, d = \underline{d}\right) \geq Pr\left(s^* < \underline{d} \,|\, d = \overline{d}\right)$$

To understand the intuition behind Proposition 2, consider a plan that has increased its default contribution rate from 3% to 6% of salary. Contributions at 4%, 5%, or 6% of salary no longer feel like losses. Thus, these contributions become relatively more desirable than contributions at 0, 1%, and 2% of income (which are in the gain domain under both default contribution rates).

C.3 Psychological anchoring makes yields the opposite prediction

People may perceive the default contribution rate as a psychological anchor. Following Bernheim et al. (2015), I assume that the anchoring parameter χ shifts workers' preferences toward the value that would rationalize the default as an optimal choice. I assume that workers act as if they were more (less) patient when contributing below (above) the default. The value of contributing s when the default is d is equal to:

$$\begin{cases} u\left(\left(1-s\right)w\right) + \left(\delta + \chi\right) V_{t+1}\left(s\right) & \text{if } s < d \\ u\left(\left(1-s\right)w\right) + \delta V_{t+1}\left(s\right) & \text{if } s = d \\ u\left(\left(1-s\right)w\right) + \left(\delta - \chi\right) V_{t+1}\left(s\right) & \text{if } s > d \end{cases}$$

This specification also captures the idea of cognitive dissonance: participants shift their time preference to rationalize why the default may be optimal for them. Similarly to the case with loss aversion, Proposition 3 shows that a model of psychological anchoring predicts that increasing the default contribution rate reduces contributions at the lowest rates.

Proposition 3. *When the default serves as a psychological anchor, increasing the default contribution rate from* $\underline{d}$ *to* $\overline{d}$ *weakly decreases contributions strictly below* $\underline{d}$*:*

$$Pr\left(s^* < \underline{d} \,|\, d = \underline{d}\right) \geq Pr\left(s^* < \underline{d} \,|\, d = \overline{d}\right)$$

Proposition 3 demonstrates that anchoring contributions at a higher level should always lead employees to contribute and participate more (not less).

D Baseline Model and Calibration – Additional Details

D.1 Individual Maximization Problem

Individuals maximize their discounted intertemporal utility with an exponential discount factor δ. In what follows, I introduce the recursive optimization problem in periods of retirement, employment, and unemployment. This dynamic optimization problem is characterized by 9 state variables: age (a), employment status (emp), tenure (ten_t), employer DC plan type (e), labor productivity (θ), average lifetime earnings (ae), liquid assets (l), DC wealth (dc) and the default contribution rate (d). The vector of state variables is denoted $X_t = [a_t, emp_t, ten_t, e_t, \theta_t, ae_t, l_t, dc_t, d]$. There is uncertainty over survival to the next period (with mortality probability (m_a)), the stochastic component of earnings (θ), employment status (emp), and the type of the future employer after a job change (e). The joint distribution of earnings, employment and employer types is denoted $F(\theta_t, emp_t, e_t)$.

Individual problem in retirement (after age 65). The only source of uncertainty in retirement is mortality risk. Retired individuals decide every period how much to hold in the liquid asset and how much to withdraw from the DC account.

$$V_t(X_t) = \max_{draw_t \in [0,1],\, l_{t+1} \geq \underline{l}_a} u_a(c_t) + \delta(1 - m_a) V_{t+1}(X_{t+1})$$

$$s.t.\ l_{t+1} = R^{liq}(l_{t+1}) \left[ss(ae_{A^{ret}}) + draw_t dc_t + l_t - c_t - tax_t^i \right]$$

$$dc_{t+1} = R^{dc}((1 - draw_t) dc_t)$$

Individual problem when employed. People are employed following a job-to-job transition, an unemployment-to-employment transition or while in continuous employment. They face earnings and employment risk and decide how much to save in liquid wealth or borrow through unsecured credit and how much to contribute to the DC account. The decision is in two steps.

In the first step, for every possible DC contribution choice s in the discrete set S, the worker solves for the optimal amount of liquid wealth or unsecured debt to hold subject to the intertemporal budget constraint 3. The value of choosing a contribution rate s is thus given by:

$$V_t^s(X_t) = \max_{l_{t+1} \geq \underline{l}_a} u_a\left(c_t - \mathbb{1}_{(s_t \neq d_t)} k\right) + \delta(1 - m_a) \int V_{t+1}(X_{t+1}) dF(\theta_t, emp_t, e_t)$$

$$s.t.\ l_{t+1} = R^{liq}(l_{t+1}) \left[(1 - s_t) w_t + l_t - c_t - tax_t^i \right] \tag{3}$$

In the second step, the worker selects a DC contribution rate s subject to a tax limit on DC contributions:

$$V_t(X_t) = \max_{s_t \in S} \{V_t^s(X_{t+1})\}$$

$$s.t.\ s_t w_t \leq limit_a$$

$$dc_{t+1} = R^{dc}(dc_t + s_t w_t + ec_t^e)$$

Individual problem when unemployed. After an unemployment shock, people receive constant unemployment benefits and face uncertainty about finding a job and the wage and type of DC plan in their future job. They decide how much liquid wealth or unsecured debt to hold and how much to withdraw from the DC account.

$$V_t(X_t) = \max_{draw_t \in [0,1],\, l_{t+1} \geq \underline{l}_a} u_a(c_t) + \delta(1 - m_a) \int V_{t+1}(X_{t+1})\, dF(\theta_t, emp_t, e_t)$$

$$s.t.\ l_{t+1} = R^{liq}(l_{t+1}) \left[ui_t + draw_t(1 - pen_a)\, dc_t + l_t - c_t - tax_t^i\right]$$

$$dc_{t+1} = R^{dc}(1 - draw_t)\, dc_t$$

D.2 Labor market parameters

Data. I use the Survey of Income and Program Participation (SIPP) to estimate of the wage earnings process and labor market transition probabilities. For each individual in the panel, we observe data on income, employment status and her employer's identity for each week. I use the 1996 panel of the SIPP, which contains data from December 1995 to February 2000, and aggregate the data to quarterly frequency. I focus on an individual's primary job (defined as the job where she works the most hours). I restrict the sample to individuals aged 22 to 65 years old and exclude full-time students and business owners. Following Borella et al. (2018), I use data on both male and female respondents. I use the publicly available replication files provided by Menzio et al. (2016) to build the SIPP panel and use a broadly similar definition of labor market transition variables. I assign employment status based on individuals' responses in the first week of each quarter. An individual is classified as employed if she reports having a job. I record a job-to-job transition if the identity of an individual's employer is different in two successive quarters. I record a job separation if an individual is employed at the beginning of a quarter and not employed at the beginning of the subsequent quarter. Job separations include early retirement (before age 65) decisions.

Earnings process. I estimate the labor earnings process for workers staying in the same job using a standard two-step minimum distance approach similar to the approaches of Guvenen (2009) and Low et al. (2010). The empirical income process is given in equation (4). It is the

empirical counterpart of the model earning process with one additional term: serially uncorrelated measurement error $\iota_{i,t} \sim N\left(0, \sigma_\iota^2\right)$.

$$\ln w_{i,t} = \delta_0 + \delta_1 a_{i,t} + \delta_2 a_{i,t}^2 + \delta_3 a_{i,t}^3 + \underbrace{\eta_{i,t}}_{\theta_{i,t}+\iota_{i,t}} \tag{4}$$

The estimation has two steps. In the first step, I estimate the parameters of the deterministic component of earnings $\left(\{\delta_j\}_{j=0}^3\right)$—a cubic in age. In the second step, I use the residual from regression (4) to estimate the four parameters governing the stochastic component of earnings: the coefficient of autocorrelation in earnings shocks (ρ), the variances of the first earnings innovation ($\sigma_{\xi_1}^2$), the variance of subsequent innovations (σ_ξ^2), and the variance of measurement error (σ_ι^2). I estimate these four parameters by minimizing the distance between the empirical variance–covariance matrix of earnings residuals and its theoretical counterpart implied by the statistical model. The parameters of the income process that I estimate (given in Table A.5) are in line with other results from the literature. In particular, I estimate very persistent income innovations, almost following a unit root, with a coefficient of autocorrelation $\rho = 0.974$.

Table A.5: Earnings process estimates

Age-income profile				**Stochastic component of earnings**			
δ_0	δ_1	δ_2	δ_3	ρ	$\sigma_{\xi_1}^2$	σ_ξ^2	σ_ι^2
1.632	0.094	−0.00107	$3.9e^{-7}$	0.974	0.184	0.0125	0.10

Notes: Quarterly earnings process estimated with a two-step minimum distance estimator on a panel of workers continuously employed in the same job. Data source: U.S. Survey of Income and Program Participation.

Earnings after a transition. I estimate the median change in log salary following a job-to-job transition (μ^{JJ}) to be equal to 0.048. I estimate that job transitions following a period of unemployment are associated with a loss in earnings. I estimate the median change in log salary relative to the last salary prior to unemployment (μ^{UE}) to be equal to −0.078.

Numeraire. The average net compensation per worker in the U.S. was around $\$37,078$ in 2006 (from the Social Security Administration national average wage index). I calibrate average quarterly earnings to be equal to $\overline{w} = \$9,250$.

Labor transition probabilities. I use SIPP micro-data to estimate quarterly job-to-job (π^{JJ}) and job-to-nonemployment (π^{EU}) transition probabilities by age, income, and tenure and job finding rates (π^{UE}) by age and income. The initial unemployment rate is set equal to 22%, which is the share not employed at age 22 in SIPP. The probability that a worker switches to another job is

determined by equation (5), and the probability that she moves to nonemployment is determined by equation (6):

$$\ln\left(\frac{\pi^{JJ}\left(a_i,\theta_i,ten_i\right)}{1-\pi^{JJ}\left(a_i,\theta_i,ten_i\right)}\right)=\sum_{k=1}^{3}\alpha_k^{JJ}a_i^k+\varrho^{JJ}\ln\left(\frac{w_i}{\overline{w}_a}\right)+\sum_{j=1}^{9}\iota_k^{JJ}\mathbb{1}_{(ten_i=j)} \tag{5}$$

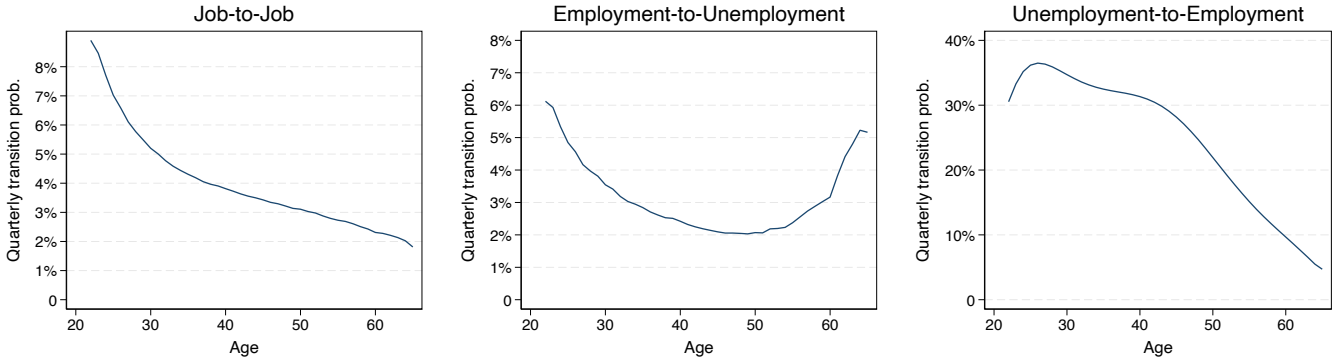
$$\ln\left(\frac{\pi^{EU}\left(a_i,\theta_i,ten_i\right)}{1-\pi^{EU}\left(a_i,\theta_i,ten_i\right)}\right)=\sum_{k=1}^{3}\alpha_k^{EU}a^k+\varrho^{EU}\ln\left(\frac{w_i}{\overline{w}_a}\right)+\sum_{j=1}^{9}\iota_k^{UE}\mathbb{1}_{(ten=j)} \tag{6}$$

The probability that an unemployed individual finds a job, given in equation (7), is defined as a fifth-order polynomial in age.

$$\pi^{UE}\left(a\right)=\sum_{k=1}^{5}\alpha_k^{EU}a^k \tag{7}$$

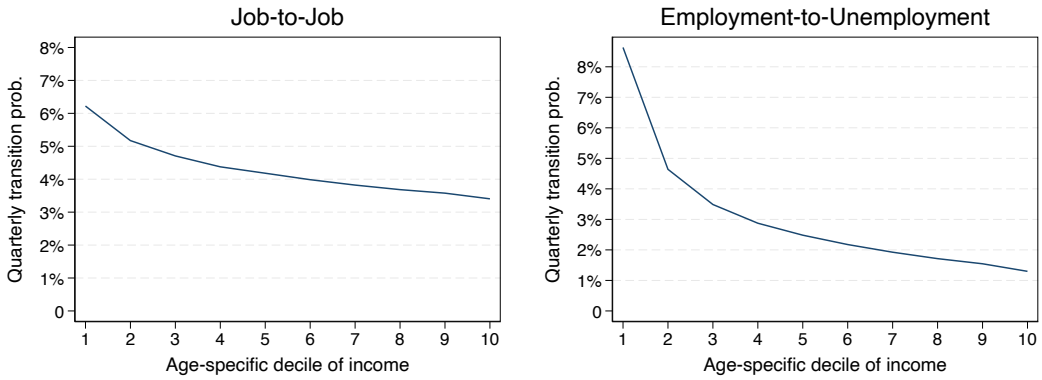
I estimate equations (5) and (6) with a logistic regression and equation (7) with a linear probability regression. Figures A.7, A.9, and A.9 show, respectively, the age, income and tenure profiles of labor market transitions in the SIPP that are implied by the estimated models.

Figure A.7: Age component of quarterly labor market transitions



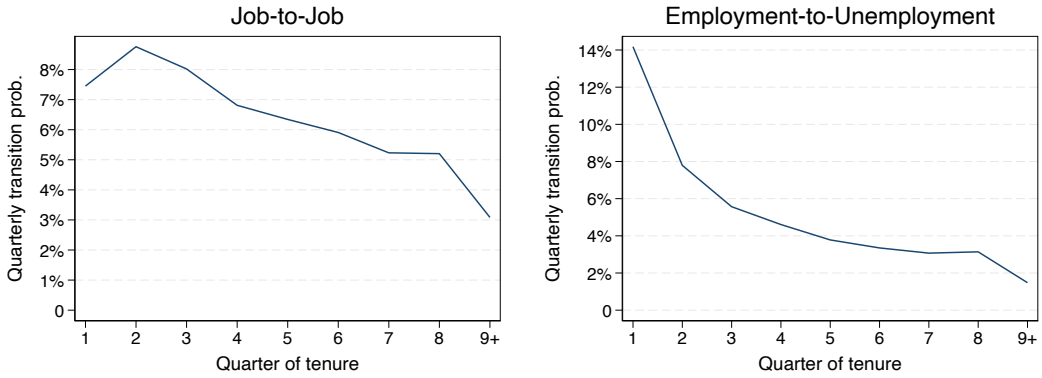
Notes: The graphed series plot the predicted age profile of labor market transitions in the SIPP sample estimated based on regression equations (5), (6), and (7).

Figure A.8: Income component of quarterly labor market transitions



Notes: The graphed series plot the predicted income profile of labor market transitions in the SIPP sample estimated based on regression equations (5) and (6).

Figure A.9: Tenure component of quarterly labor market transitions



Notes: The graphed series plot the predicted tenure profile of labor market transitions in the SIPP sample estimated based on regression equations (5) and (6).

D.2.1 Demographics

Survival probabilities. Survival probabilities for each age are calibrated to the U.S. Social Security 2014 period life table.[4]

Equivalence scale. Changes in household composition over the lifecycle are captured by an equivalence scale in the utility function. I use the equivalence scale by age estimated by Lusardi et al. (2017). Using Panel Study of Income Dynamics (PSID) data from 1984 to 2005, Lusardi et al. (2017) estimate $z(j_t, k_t) = (j_t + 0.7k_t)^{0.75}$, where j_t and k_t are, respectively, the average number of adults

[4]In the data, death probabilities are given at yearly frequency, while individuals make decisions every quarter in the model. I assume that the quarterly death probability is equal to 0.25 times the annual probability.

and children (under 18 years old) in a household with a head of age t. They normalize this measure by $z\,(2,1)$—the composition of a household with 2 adults and 1 child—to get the equivalence scale at age t equal to $n_t = \frac{z(j_t,k_t)}{z(2,1)}$. To estimate n_t, I use publicly available replication files from Lusardi et al. (2017) and aggregate the data across education groups.

D.2.2 Tax and benefit system

Income taxation. Taxable income is defined as the sum of labor earnings, Social Security and unemployment benefits, and DC withdrawals, less contributions to the DC account:

$$y_t^{tax} = \begin{cases} (1-s_t)\,w_t & if\ emp_t \in \{E, JJ\} \\ ui_t + draw_t dc_t & if\ emp_t = U \\ ss\,(ae_{A^{ret}}) + draw_t dc_t & if\ emp_t = Ret \end{cases}$$

Individuals' income tax liability is calculated according to the federal income tax schedule of 2006 (the first year of data and the base year for the calibration) for an individual filling as single and claiming the standard deduction. The tax formula has 5 annual income brackets: $\{\tilde{\kappa}_i^\tau\}_{i=1}^5 = \{\$\,5,150\,;\,\$\,7,550\,;\,\$\,30,650\,;\,\$\,74,200\,;\,\$\,154,800\}$.[5] Quarterly tax brackets are defined as: $\kappa_i^\tau = \frac{1}{4}\tilde{\kappa}_i^\tau$. The quarterly income tax liability is equal to:

$$tax_t^i = \begin{cases} 0 & if\ y^{tax} \leq \kappa_1^\tau \\ 0.10\,(y^{tax} - \kappa_1^\tau) & if\ \kappa_2^\tau \geq y^{tax} > \kappa_1^\tau \\ 0.10\,(\kappa_2^\tau - \kappa_1^\tau) + 0.15\,(y^{tax} - \kappa_2^\tau) & if\ \kappa_3^\tau \geq y^{tax} > \kappa_2^\tau \\ 0.10\,(\kappa_2^\tau - \kappa_1^\tau) + 0.15\,(\kappa_3^\tau - \kappa_2^\tau) + 0.25\,(y^{tax} - \kappa_3^\tau) & if\ \kappa_4^\tau \geq y^{tax} > \kappa_3^\tau \\ 0.10\,(\kappa_2^\tau - \kappa_1^\tau) + 0.15\,(\kappa_3^\tau - \kappa_2^\tau) + 0.25\,(\kappa_4^\tau - \kappa_3^\tau) + 0.28\,(y^{tax} - \kappa_4^\tau) & if\ \kappa_5^\tau \geq y^{tax} > \kappa_4^\tau \\ 0.10\,(\kappa_2^\tau - \kappa_1^\tau) + 0.15\,(\kappa_3^\tau - \kappa_2^\tau) + 0.25\,(\kappa_4^\tau - \kappa_3^\tau) + 0.28\,(\kappa_5^\tau - \kappa_4^\tau) + 0.33\,(y^{tax} - \kappa_5^\tau) & if\ y^{tax} > \kappa_5^\tau \end{cases}$$

Public pension. The amount of public pension benefit (ss) is computed according the 2006 Social Security formula at full retirement age, with an income floor guaranteed by the Supplemental Security Income program (with a monthly benefit $si = \$603$). Quarterly public pension benefits are equal to:

$$ss\,(ae_{A^{ret}}) = 3 \times \max\{si\,;\,\tilde{ss}\,(ae_{A^{ret}})\}$$

[5] Note that the first bracket corresponds to the standard deduction amount in 2006.

where $\tilde{ss}$, the monthly Social Security benefit, is increasing in average lifetime earnings $ae_{A^{ret}}$ up to a maximum monthly benefit:

$$\tilde{ss} = \begin{cases} 0.90 \times \frac{1}{3} ae_{A^{ret}} & if\ \frac{1}{3} ae_{A^{ret}} \leq \$656 \\ 0.90 \times \$656 + 0.32 \times (\frac{1}{3} ae_{A^{ret}} - \$656) & if\ \$3,955 > \frac{1}{3} ae_{A^{ret}} > \$655 \\ \min \left\{ 0.90 \times \$656 + 0.32 \times \$3,299 + \left(0.15 \times \frac{1}{3} ae_{A^{ret}} - \$3,299\right) ; \$2,053 \right\} & if\ \frac{1}{3} ae_{A^{ret}} > \$3,955 \end{cases}$$

Unemployment benefits. Unemployment insurance provides a constant replacement rate ω of labor earnings implied by the labor productivity level in the last period of employment. Labor productivity θ_t stays constant during an unemployment spell. I set $\omega = 0.40$, which is the average replacement rate across all U.S. states (U.S. Department of Labor, 2018). In the U.S., unemployment benefits may be reduced by a claimant's retirement income. The Unemployment Compensation Amendments of 1976 require that all retirement income be offset against unemployment compensation. There are differences in how states implement the offset, and generally, for withdrawals from a 401(k), only the amount contributed by the employer offsets unemployment benefits (see Franco (2004) for differences across U.S. states). For simplicity, I assume that the employer contribution portion of an early withdrawal is always equal to the employer match rate. This simplifying assumption is valid assuming that participants contribute below the matching threshold and that contributions are fully vested. The adjusted unemployment benefits for an individual unemployed since period $t - x$ are given by:

$$ui_t = \max \{0\,;\, \omega w_t (\theta_{t-x}) - draw_t dc_t \times match_e\}$$

D.2.3 Assets

Asset returns and limits. I set the annual (riskless) rate of return at 3%, which is equal to the long-term nominal yield on government bonds over the sample period between 2006 to 2017.[6] The annual interest rate on unsecured credit is set at $R^{cc} = 11.52\%$, the value estimated by Laibson et al. (2024) for the interest rate on credit card debt adjusted for both bankruptcy and inflation. The borrowing constraint prior to retirement is fixed and set equal to 74% of average quarterly earnings ($\underline{l}_{a<A^{ret}} = 0.74\overline{w}$). This is the average credit card limit estimated by Kaplan and Violante (2014) in

[6] R is assumed to be exogenous, and I abstract from the equilibrium determination of the interest rate for two reasons. First, the effect of auto-enrollment policies on aggregate savings is small, and second, the effect of the policy is concentrated on low-income individuals (who hold a relatively small share of aggregate assets).

the Survey of Consumer Finances. Unsecured borrowing is not allowed in retirement: $\underline{l}_{a \geq A^{ret}} = 0$.[7]

Asset taxation. In line with IRS rules for 2006, the maximum contribution limit for tax-deferred retirement contributions ($limit_a$) is set equal to $\$3,750$ per quarter ($\$15,000$ annually) for individuals younger than 50 years old and $\$5,000$ after that in 2006 dollars. The tax penalty for early DC withdrawals (pen_t) is equal to 10% before age $59^{1/2}$ and to zero afterward.[8] Capital income is taxed at a rate τ^k of 15%, which implies that the annual return (net of tax) on liquid wealth is 2.55%.

Vesting schedule. An employee who separates from her employer before the end of the vesting period may lose part (or all) of the employer matching contribution. A vesting schedule $vst_e(ten)$ determines the percentage of employer contributions that an employee keeps if she separates at a given tenure level. For instance, under a cliff vesting schedule, vst_e is equal to zero before the end of the vesting period and to one afterward. Modeling the vesting schedule explicitly would introduce an additional continuous state variable to the dynamic problem: the amount of nonvested DC wealth. Instead, I adjust employer contributions by a factor $\Upsilon_e(t, \theta, ten)$ proportional to the risk of losing unvested employer contributions. The adjustment factor $\Upsilon_e(t, \theta, ten)$ is given in equation (8). It depends on both the cumulative job-separation probability and the vesting schedule. It is smaller than one and increasing in tenure before the end of the vesting period and is equal to one afterward. Importantly, this specification captures the fact that vesting matters more for employees who—based on their age and tenure—are more likely to separate from their employer.

$$\Upsilon_e(t, \theta, ten) = 1 - \sum_{j=0}^{T^R - t} \left(\prod_{k=1}^{j-1} \left(1 - \pi^{EU}_{t+k,ten+k} - \pi^{JJ}_{t+k,ten+k}\right) \right) \left(\pi^{EU}_{t+j,ten+j} + \pi^{JJ}_{t+j,ten+j}\right) \left(1 - vst_e(ten + j)\right) \tag{8}$$

Employer contributions adjusted for "vesting risk" ($\widetilde{ec}_e$) are given by equation (9).

$$\widetilde{ec}_{e,t}(\tau_t, \theta, ten_t) = \Upsilon_e(t, \theta, ten_t) \times ec_e(\tau_t, ten_t) \tag{9}$$

On average, in the estimation sample, employer matching contributions are initially vested at 44%, and this share increases over tenure. The average vested share reaches 73% after the end of

[7] This assumption is similar to the approach in Kaplan and Violante (2014). Absent this no-borrowing-in-retirement constraint, older individuals in the model borrow excessively because their high mortality risk reduces their repayment liability.

[8] Workers who separate from their employer after reaching age 55 can make penalty-free withdrawals from their current 401(k) plan. However, this exemption does not extend to funds in 401(k) plans from previous employers or to funds that have been rolled over into an IRA, which are still subject to a 10% early withdrawal penalty until age 59½. In the model, the majority of retirement wealth is represented by assets accumulated in jobs from which separation occurred before age 55.

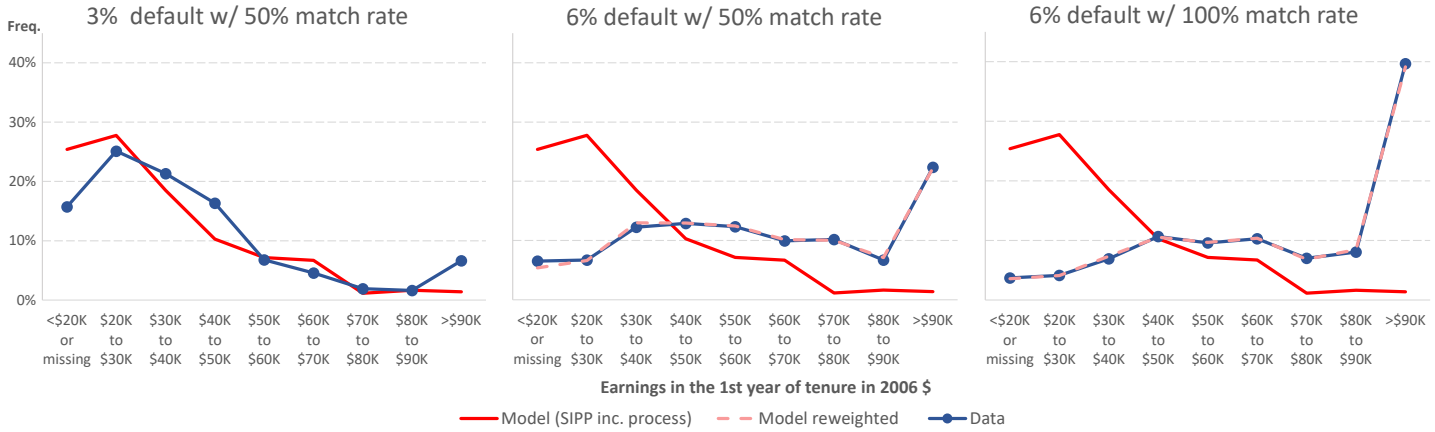
the second year of tenure. I assume that all matching contributions are fully vested starting from the 9th quarter of tenure.

E Estimation with a 6% Default Contribution Rate

E.1 Reweighting approaches

Figure A.10 shows the distribution of labor earnings for employees in their first year of tenure in the model and in the three estimation samples. As shown in the left panel, the baseline income process estimated in nationally representative data from the Survey of Income Programs and Participation (SIPP) generates a distribution of earnings similar to the one observed for auto-enrolled employees in plans adopting a 3% default with a 50% match rate up to 6% of salary.[9] The figure also shows that employees in plans with a 6% default and a 50% match rate, and to a greater extent those with a 6% default and a 100% match rate, have significantly higher earnings and that the income process estimated in the SIPP data offers a poor fit for the distribution of earnings in these two samples with a higher default option. To address this issue, I consider two approaches.

Figure A.10: Income distribution for auto-enrolled workers: Model vs. data



Notes: Each panel corresponds to the probability distribution function of first-year earnings (in 2006 $) accross nine income bins. The "Model (SIPP inc. process)" corresponds to the distribution of earnings in the simulated data based on the income process estimated in SIPP data. The "Model reweighted" series corresponds to the simulated data reweighted to match the empirical earnings distribution as described in Appendix E.1. The "Data" series corresponds to auto-enrolled employees hired in the 12 months following the policy adoption.

In the first approach, I reweight observations in the simulated data to match the empirical

[9] I focus on the distribution of earnings among the auto-enrolled because I observe earnings only for workers making positive contributions to the 401(k) plan. Since participation among the auto-enrolled is above 95%, I have earnings data for over 95% of auto-enrolled workers.

distribution of earnings in the samples with a 6% default and, respectively, a 50% and a 100% match rate. The weight assigned to a simulated observation with earnings in group k is equal to $weight_k = \frac{sh_k^{data}}{sh_k^{sim}}$, where sh_k^{data} and sh_k^{sim} are the shares of observations with income in group k in, respectively, the data and the model simulation. That is, I give more weight to observations with a higher simulated income in computing the simulated moments. Preference parameter estimates using this approach are reported in columns 6 through 8 of Table 4 and the values of sh_k^{data} and sh_k^{sim} are plotted in Figure A.10.

The second approach allows for heterogeneous preferences. I consider two types of agents: (i) standard agents with the baseline preference estimates (i.e., the values estimated in column 1 of Table 4) and the distribution of earnings generated by the income process estimated in SIPP data and (ii) high-earners, who are allowed to have different preferences. I simulated the model twice, once with the baseline preference parameters and once with the high-earners preference parameters. I set the share of high earners to $\alpha^H = 45\%$ and I calculate the simulated moment conditions by applying a weight of $1 - \alpha^H$ on observations simulated under the baseline preferences, and a weight of $\alpha^H weight_k^H$ on observations with earnings in group k simulated under the preferences of the high earners. These weights are chosen such that the reweighted distribution of simulated earnings matches the empirical distribution of earnings in Figure A.10. That is, the reweighted share of simulated observations in income group k (i.e., the weighted sum of sh_k^B, the simulated shares under baseline preferences and sh_k^H, the simulated shares under the preferences of the high earners) is equal to the empirical share of observation in income group k (sh_k^{Data}):

$$sh_k^{Data} = \left(1 - \alpha^H\right) \times sh_k^B + \alpha^H \times weight_k^H \times sh_k^H$$
$$\Rightarrow weight_k^H = \frac{sh_k^{Data} - \left(1-\alpha^H\right) \times sh_k^B}{\alpha^H \times sh_k^H}$$

I confirm in Figure A.10 that the reweighted distributions of simulated earnings match the empirical distributions in the two samples with a 6% default. In Table A.6, I report the preference parameter estimates for high-earners in the two samples with a 6% default contribution rate.

Both approaches yield similar results; workers in plans with a 6% default contribution are estimated to have significantly higher discount factors than those in plans adopting a 3% default. The switching cost estimated under a 6% default is lower in plans with a 100% match rate and larger in plans with a 50% match rate.

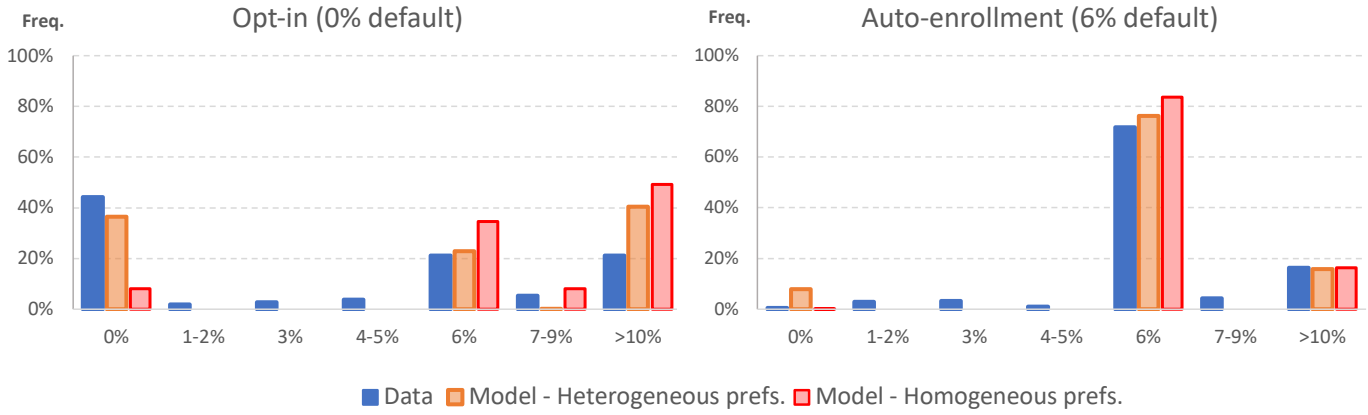
Table A.6: Preference parameter estimates for high-earners

	(1) Switching cost (k)	(2) Quarterly disc. factor (δ)	(3) EIS (σ)
6% default & 50% match rate	\$ 289	0.999	0.586
	(70)	(0.002)	(0.242)
6% default & 100% match rate	\$ 99	0.998	0.761
	(24)	(0.002)	(0.488)

Notes: Preference estimates for high-earners (i.e., 45% of simulated observations) following the reweighting approach outlined in Section E.1. The remaining 55% of observations are simulated under the baseline preference parameter values from column 1 of Table 4 and the baseline earnings distribution. Standard errors adjusted for measurement error in the estimation of the income process are in parentheses.

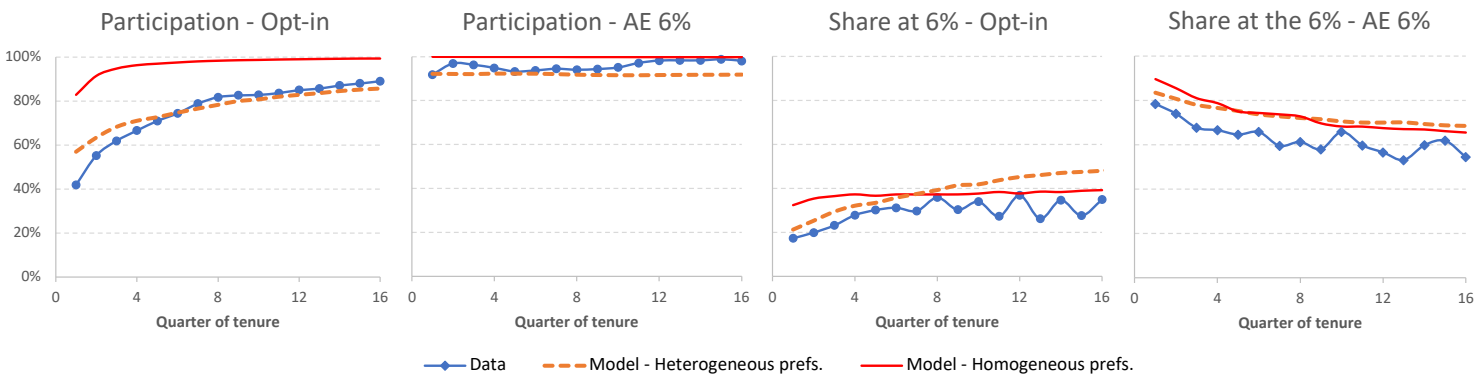
E.2 Model fit: Sample of plans with a 6% default and a 50% match rate up to 6%

Figure A.11: Empirical and simulated distributions of contribution rates in the first 12 months of tenure



Notes: The left (right) panel shows contribution rates of individuals hired in the 12 months before (after) the introduction of auto-enrollment at 6% for new hires in plans with a 50% match rate up to 6% of salary. The model series corresponds to simulations for identical individuals under a 0% and a 6% default.

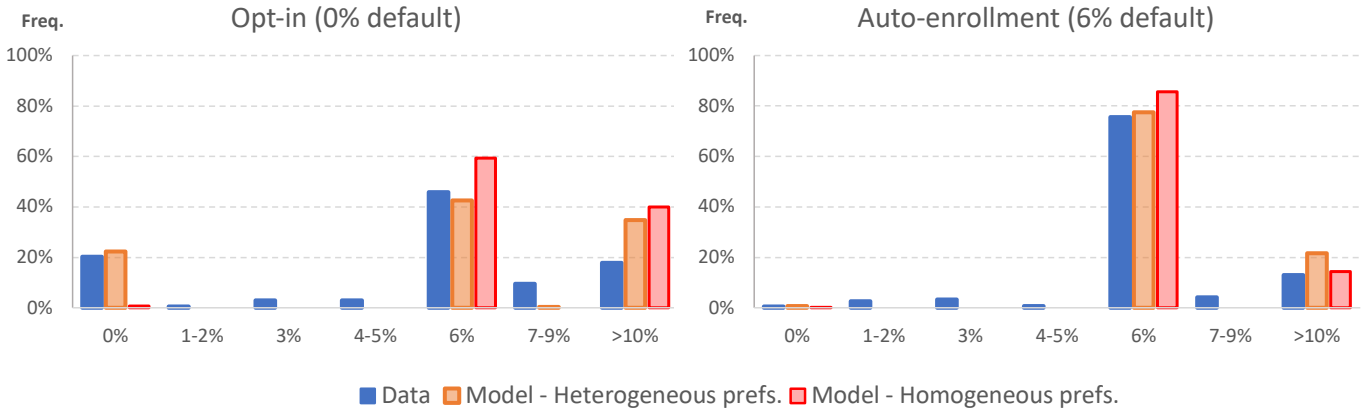
Figure A.12: Empirical and simulated contribution behavior over 4 years of tenure



Notes: The two left (right) panels show the 401(k) participation rate (the share who contribute exactly 6% of salary) for workers hired in the 12 months before and 12 months after the adoption of auto-enrollment at 6% for new hires in plans with a 50% match rate up to 6% of salary. The model series are simulations for identical individuals under a 0% and a 6% default.

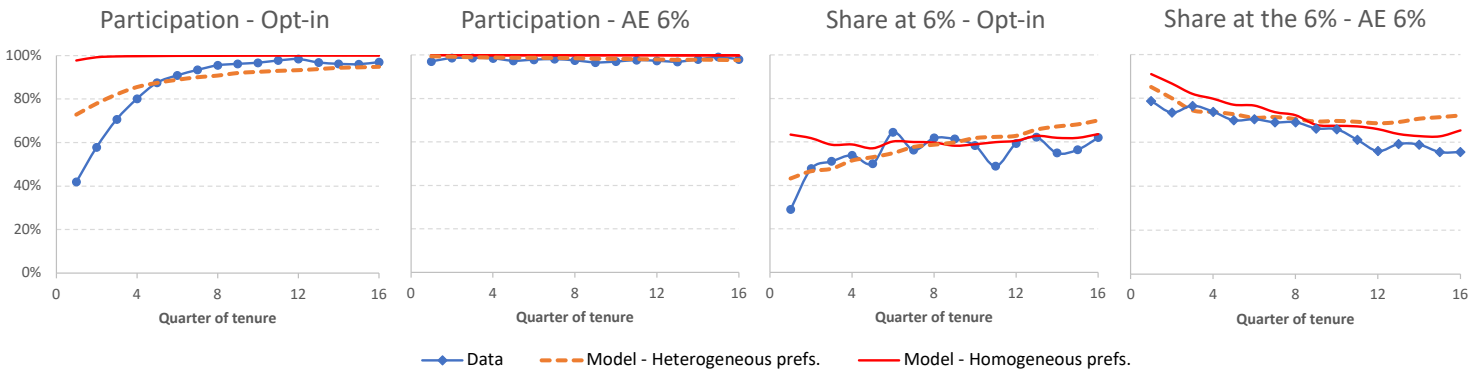
E.3 Model fit: Sample of plans with a 6% default and a 100% match rate up to 6%

Figure A.13: Empirical and simulated distributions of contribution rates in the first 12 months of tenure



Notes: The left (right) panel shows the contribution rates of individuals hired in the 12 months before (after) the introduction of auto-enrollment at 6% for new hires in plans with a 100% match rate up to 6% of salary. The model series corresponds to simulations for identical individuals under a 0% and a 6% default.

Figure A.14: Empirical and simulated contribution behavior over 4 years of tenure



Notes: The two left (right) panels show the 401(k) participation rate (the share who contribute exactly 6% of salary) for workers hired in the 12 months before and 12 months after the adoption of auto-enrollment at 6% for new hires in plans with a 100% match rate up to 6% of salary. The model series are simulations for identical individuals under a 0% and a 6% default.

E.4 Robustness of the preference estimates

In this Appendix, I assess the robustness of the switching cost estimates. I fix either δ or σ and re-estimate the remaining two preference parameters. Estimation results are reported in Table A.7. I estimate a larger switching cost—around $380—when I fix the quarterly discount factor to be lower (i.e., $\delta = 0.96$, implying an annual discount factor of $\delta^4 = 0.85$) or when I fix the elasticity of intertemporal substitution to be lower (i.e., $\sigma = 0.2$ implying a coefficient of relative risk aversion of 5). In contrast, when I fix the EIS to be higher, I estimate a lower switching cost than in the baseline: $210. The model's fit worsens in these specifications relative to the baseline reported in column 1 as reflected by a larger goodness-of-fit measure).

Table A.7: Preference parameter estimates

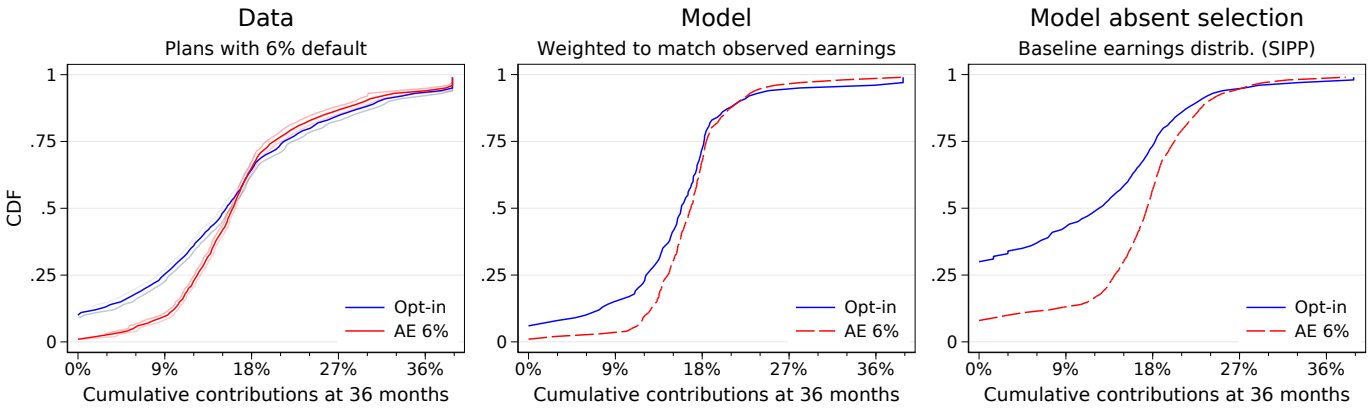
	Auto-enrollment at 3% & 50% match rate			
	(1) Baseline (i.e., col 1 of Table 4)	(2) Fixed $\sigma = 1.25$ (i.e., RRA: $\sigma^{-1} = 0.8$)	(3) Fixed $\sigma = 0.2$ (i.e., RRA: $\sigma^{-1} = 5$)	(4) Fixed $\delta = 0.96$ (i.e., annual freq: δ^4=0.85)
Switching cost (k)	\$ 235	\$ 210	\$ 383	\$ 371
	(31)	(7)	(68)	(48)
Quarterly disc. factor (δ)	0.989	0.990	0.978	0.96
	(0.002)	(0.001)	(0.006)	–
EIS (σ)	0.520	1.25	0.20	0.189
	(0.085)	–	–	(0.007)
Goodness of fit	323	582	684	5,098

Notes: preferences are estimated with data from plans with a 3% default and a 50% match rate up to 6% of salary. Column 1 report the baseline estimate from column 1 of Table 4. Columns 2 and 3 (Column 4) fixe the value of the EIS (discount factor) and estimate the remaining two preference parameters. I use the identity weighting matrix and standard errors adjusted for measurement error in the estimation of the income process are in parentheses. I construct the same goodness-of-fit measure as in Laibson et al., 2024.

E.5 Correcting for selection in the sample of plans with a 6% default

As described in Section 5.7, I use the model to assess the effect of a higher default absent these two sources of selection. The left panel of Figure A.15 corresponds to the observed CDFs in plans with a 6% default. The middle panel, corresponds to model simulations when simulated observations are reweighted to match the distribution of earnings observed in the plans with a 6% default and a 50% match rate. Finally, the right panel corresponds to the model prediction under the baseline income process estimated in SIPP data.

Figure A.15: Heterogeneity in the effect of auto-enrollment at 6%: Data vs Model



Notes: Each line corresponds to the simulated cumulative distribution function (CDF) of cumulative employee 401(k) contributions divided by individuals' annual salary after 36 months of tenure. The left panel corresponds to the empirical distribution in plans adopting a 6% default. The Model series correspond to simulated data using the baseline preference estimates from column 1 of Table 4. In the middle panel, the simulated observations are reweighted to match the empirical distribution of earnings in plans with a 6% default and a 50% match rate following the procedure outlined in Appendix E.1. In the right panel, earnings are simulated based on the income process estimated in SIPP data.

F Model Extensions

F.1 Present-biased preferences

Several empirical studies have documented that measures of time inconsistency correlate with the propensity to stay at the default option in retirement savings plans (Blumenstock et al., 2018; Brown et al., 2016; Brown and Previtero, 2020, Goda et al., 2020). Under present-biased preferences, workers procrastinate on changing their contribution rate because the switching cost is paid in the present while the benefits of retirement saving are in the future (O'Donoghue and Rabin, 1998; Carroll et al., 2009; DellaVigna 2009 and 2018). There are two different approaches for modelling the role of present-biased preferences.

A first view is that contribution choices reflect only the regular discount factor δ because contribution changes are implemented by employers with a delay—usually at the beginning of the next pay period. As noted by Carroll et al. (2009), this implies that a contribution rate choice serves as a commitment to save in the future, starting from the next paycheck. Under this view, introducing present bias acts in a similar way to assuming a higher switching cost: it reduces the probability of switching but does not affect the contribution choice conditional on the employee's making an active decision. In the absence of a source of variation in my data that would allow me to separately

identify this form of present bias from switching costs, I interpret my baseline estimates as capturing the combined effect of present bias and other forms of switching costs.

A second view is that present bias directly affects contribution preferences. Workers choose low contribution rates because they do not want to reduce their consumption in the present and believe (mistakenly) that they will contribute more later. I explore the role of this specification of present bias by introducing in the model an additional discount factor $\beta \leq 1$ between the present and all future periods (Laibson, 1997). I assume that workers are naive about their future selves' present bias. The objective function for a present-biased worker V^{PB} is given by equation (10), where V_{t+1} is the value function in the baseline exponential discounting case.

$$V_t^{PB}(X_t) = \max_{s_t \in S, l_{t+1}} u_a\left(c_t - \mathbb{1}_{(s_t \neq d_t)}.k\right) + \beta\delta\left(1 - m_a\right) \int V_{t+1}\left(X_{t+1}\right) dF\left(\theta_t, emp_t, e_t\right) \tag{10}$$

I calibrate the short-term discount factor $\beta = 0.5$, which is the value estimated by Laibson et al. (2024) in a lifecycle setting, and re-estimate the model's three other preference parameters following the same procedure outlined in Section 5. The results are reported in Table A.8. My estimates of the preference parameters are very close to those estimated by Laibson et al. (2024) in a different setting with different data: I estimate a long-term quarterly discount factor of 0.997 (or 0.988 at annual frequency), which is close to the estimate in Laibson et al. (2024) of 0.99, while my EIS estimate of 0.64 is only slightly lower than the estimate of the inverse of relative risk aversion in Laibson et al. (2024) of 0.77.

Table A.8: Preference parameter estimates – Quasi-hyperbolic time preferences

	(1) Baseline model $\beta = 1$	(2) Naive present bias $\beta = 0.5$
k	\$ 235	\$ 179
	(31)	(38)
δ	0.989	0.997
	(0.002)	(0.001)
σ	0.452	0.640
	(0.085)	(0.069)

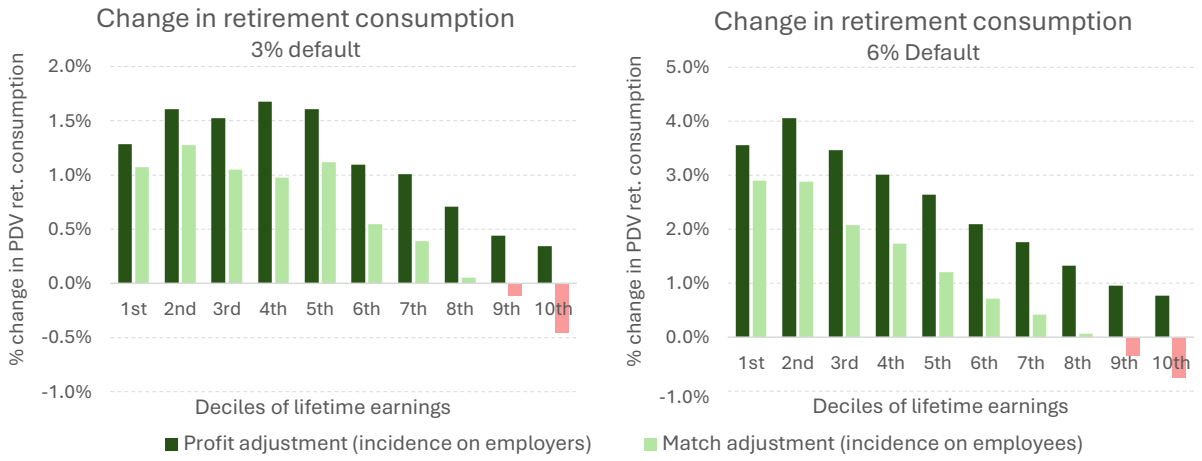
Notes: Column 1 reproduces the baseline estimates from column 1 of Table 4. Column 2 shows preference estimates with naive quasi-hyperbolic discounting and a short term discount factor set equal to $\beta = 0.5$, the value estimated by Laibson et al. (2024).

As shown in the top panel of Figure A.16, the model estimated with quasi-hyperbolic time

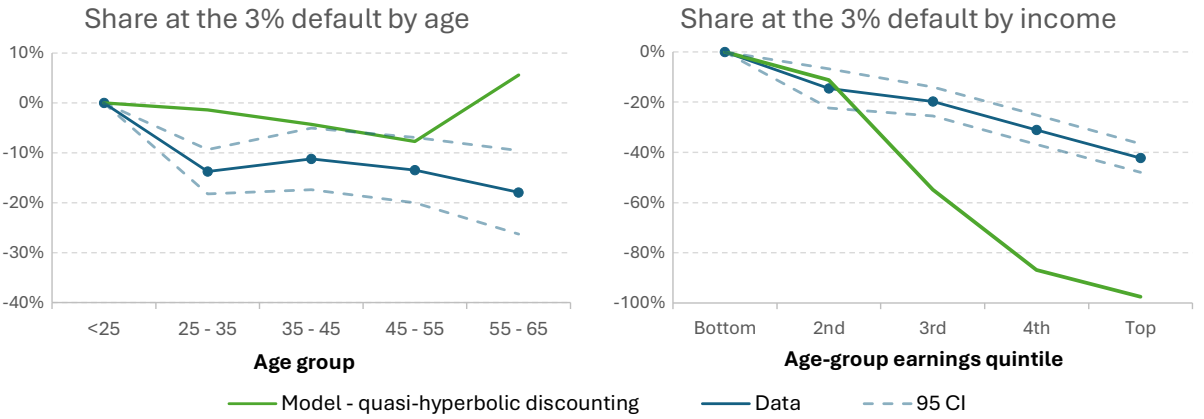
preferences generates a qualitatively similar heterogeneity in the long-term savings gains from an auto-enrollment policy adopted by all employers. However, as shown in the bottom panel, this specification does a worse job capturing the heterogeneity in default effects by age an income.

Figure A.16: Lifetime impact and heterogeneous responses to auto-enrollment - Quasi-hyperbolic time preferences

(a) Effect of a universal auto-enrollment policy on retirement consumption



(b) Heterogeneity in default effects by age and income



Notes: In Panel (a), each point corresponds to the coefficients of the interaction terms of auto-enrollment status (based on a worker's date of hire) with income and age group dummies, estimated according to equation (2). The dependent variable is a dummy equal to 1 if, conditional on participating, a worker observed in her first year of tenure contributes exactly 3% (the auto-enrollment default contribution rate). In Panel (b) , each bar corresponds to the model-predicted percentage change in the present discounted value of consumption between ages 65 and 90 under an auto-enrollment policy adopted by all employers relative to that under an opt-in regime. I use the risk-free interest rate (3%) for discounting retirement consumption. The policies are simulated under the baseline preference estimates from column 2 of Table A.8. Dark-shaded bars correspond to the results based on the assumption that the incidence of the policy falls on employers, while the light-colored bars correspond to the results based on the assumption that employers reduce the 401(k) matching rate to satisfy their budget constraint.

F.2 Proportional switching costs

As an extension of the baseline model, I consider the case where the switching cost is proportional to each worker's labor earnings. For instance, this specification captures the fact that, if changing the DC contribution rate takes time, the opportunity cost of time is higher for highly paid workers. I introduce a switching cost $\tilde{k}$ that is proportional to earnings:

$$u_a\left(c_t - \mathbb{1}_{(s_t \neq d_t)}\tilde{k}.w_t\right)$$

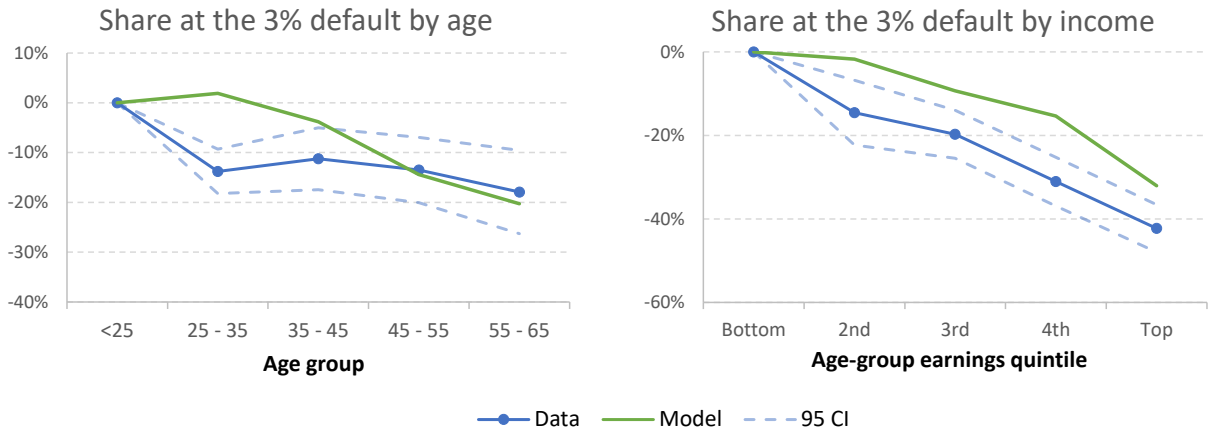
I present estimation results in Table A.9. I estimate $\tilde{k}$ to be equal to 3.16% of quarterly income (i.e., equivalent to $\$311$ at the average earnings level). Under this specification, high-income workers are more likely to stay at the default compared to the baseline model because the switching cost is now larger for them (Figure A.17, right panel). As a consequence, the model with a proportional switching cost generates less correlation between income and the propensity to stay at the default option. While this assumption reduces the correlation between earnings and default behavior, high-income workers are still predicted to deviate from the default contribution rate more frequently than low-income workers because: (i) low-income workers face higher separation probabilities and, as a consequence, expect to stay in the firm for a shorter period, and (ii) high-income individuals have more to gain by contributing to the retirement savings plan because they face a higher marginal tax rate and expect a lower Social Security replacement rate.

Table A.9: Preference parameter estimates with switching costs proportional to earnings

(1) $\tilde{k}$	(2) δ	(3) σ
0.0336	0.987	$\$0.454$
(0.009)	(0.003)	(0.032)

Notes: Columns 1 to 3 present second-stage parameter estimates for the switching cost, the quarterly discount factor, and the elasticity of intertemporal subsitution under the assumption that the switching cost is a constant fraction of earnings.

Figure A.17: Heterogeneity in default effects by age and income – Proportional switching cost



Notes: Each point corresponds to the coefficients of the interaction terms of auto-enrollment status (based on a worker's date of hire) with income and age group dummies, estimated according to equation (2). The dependent variable is a dummy equal to 1 if, conditional on participating, a worker observed in her first year of tenure contributes exactly 3% (the auto-enrollment default contribution rate). The area between the dashed lines corresponds to the 95% confidence intervals of the empirical coefficients. Data source: Administrative 401(k) records from 34 U.S. firms.

F.3 No offset of unemployment benefits

In the baseline specification of the model, early withdrawals from the DC account reduce unemployment benefits. This modeling assumption is consistent with the Unemployment Compensation Amendments of 1976, which require that all retirement income be offset against unemployment compensation. However, there are differences across states in how this offset is implemented. In general, for withdrawals from a 401(k) plan, only the amount contributed by the employer offsets unemployment benefits. For instance, in New Jersey, 50% of retirement income, including *"benefits paid in a lump sum such as 401(k)"*, is substracted from the unemployment benefit if both the claimant and the base-period employer contributed to the pension plan.[10]

However, the rules are complex and vary across states (see Franco (2004) for a discussion of differences across U.S. states). Depending on the specific way the offset is implemented by states' unemployment insurance agencies, unemployed individuals may be able to avoid this offset by first rolling over their 401(k) assets into an individual retirement account before withdrawing those resources. As a robustness check, I re-estimate the model under the assumption that 401(k) withdrawals do not offset unemployment insurance benefits. The estimated parameters, in Table A.10, are little affected by this alternative assumption: the contribution switching cost and the quarterly

[10]Source: https://nj.gov/labor/ui/content/faq.html

discount factor are estimated to be a little smaller ($175 and 0.985 relative to baseline values of, respectively, $235 and 0.989) and the EIS slighly higher (0.801 relative to an estimate of 0.520 in the baseline model).

Table A.10: Preference parameter estimates with no offset of unemployment benefits

(1)	(2)	(3)
k	δ	σ
\$ 175	0.985	\$ 0.801
(20)	(0.001)	(0.071)

Notes: Columns 1 to 3 present second-stage parameter estimates for the switching cost, the quarterly discount factor, and the elasticity of intertemporal subsitution under the assumption of no unemployment offset of early 401(k) withdrawals.

F.4 Discussion of the counterfactual policy analysis

While the model attempts to capture the main aspects of the retirement saving environment, it abstracts from a number of features. In this section, I discuss how the absence of some of these elements could affect the counterfactual policy exercise.

Endogenous labor supply. Employees may be able to undo some of the effect of auto-enrollment by adjusting their future labor supply or shifting the timing of their retirement by a few weeks or months. Introducing these additional margins of adjustment would likely reinforce the main findings of the paper: when employees can adjust both their future contributions and labor supply, the size of the switching cost needed to generate the observed inertia at the default and the long-term effect of auto-enrollment on wealth may become even smaller. However, since I am not able to estimate the causal effect of auto-enrollment on labor supply and retirement decisions in the data, I make the more conservative assumption that auto-enrollment has no effect on labor supply and retirement timing.

Default asset allocation. Under auto-enrollment, the default contribution rate is often combined with a default asset allocation. In this paper, I study the effect of changing the default contribution rate independently from the choice of a default asset allocation. Changing the default asset allocation could have a separate impact on wealth accumulation that is potentially large but beyond the scope of this paper.[11]

[11] Dahlquist et al. (2018) estimate large welfare gains from the implementation of an optimal default asset allocation

General equilibrium interest rate. A national auto-enrollment policy could raise aggregate saving and lead to a reduction in the equilibrium interest rate. Given the estimated value of the elasticity of intertemporal substitution (EIS), this general equilibrium channel would likely reduce the effect of auto-enrollment on individual saving. However, given my finding that auto-enrollment has a very modest impact on long-term asset accumulation and that this effect is concentrated at the bottom of the income distribution (a population that accounts for a small share of aggregate savings), this general equilibrium effect is likely to be very small.

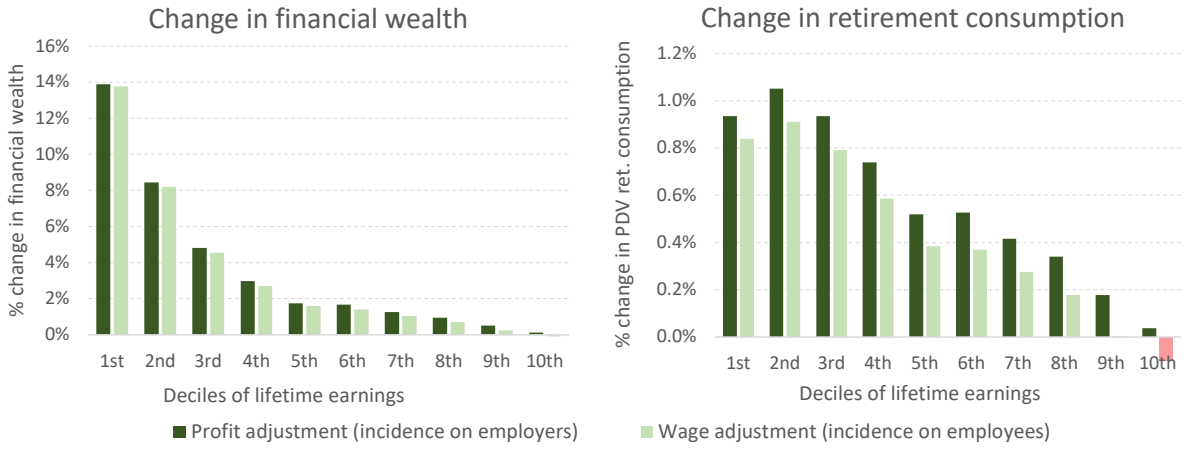
F.5 Alternative incidence assumption: Wage adjustment

Another possibility for satistifying the budget constraint is that employers reduce wages to offset higher matching costs. I assume that all wages are reduced by the same proportional factor to satisfy employers' aggregate budget constraint. Employers adjust the wage level to keep the total compensation bill $W\left(d^{SP}\right) + Mtc\left(d^{SP}\right)$ constant. I show in Figure A.18 simulation results on the effect of a universal auto-enrollment policy adopted by all employers over a lifetime under this incidence assumption.

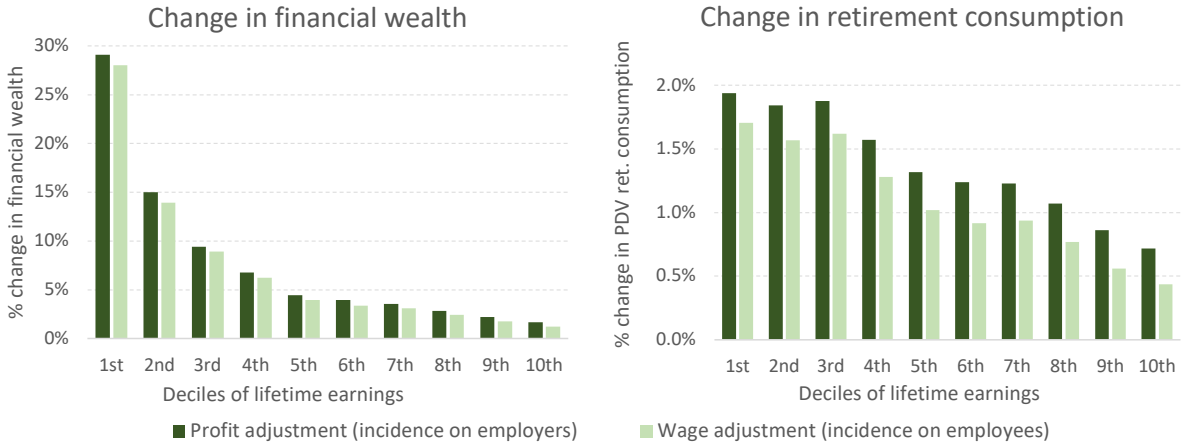
in the context of Swedish pension plans. Choukhmane and de Silva (2023) extend the model to a portfolio choice problem.

Figure A.18: Effect of auto-enrollment on financial wealth at age 65 and on retirement consumption

(a) Universal auto-enrollment policy at 3% over a lifetime



(b) Universal auto-enrollment policy at 6% over a lifetime



Notes: Notes: In the left (right) two panels, each bar corresponds to the model-predicted percentage change in financial wealth at age 65 (present discounted value of consumption between ages 65 and 90) under an auto-enrollment policy adopted by all employers relative to that under an opt-in regime. I use the risk-free interest rate (3%) for discounting retirement consumption. The policies are simulated under the baseline preference estimates from column 1 of Table 4 and the baseline distribution of earnings estimated in SIPP data. Dark-shaded bars correspond to the results based on the assumption that the incidence of the policy falls on employers, while the light-colored bars correspond to the results based on the assumption that employers reduce wages uniformly to satisfy their budget constraint.

F.6 Additional evidence on the lifecycle profile of auto-enrollment savings gains

F.6.1 3% default & match rate adjustment

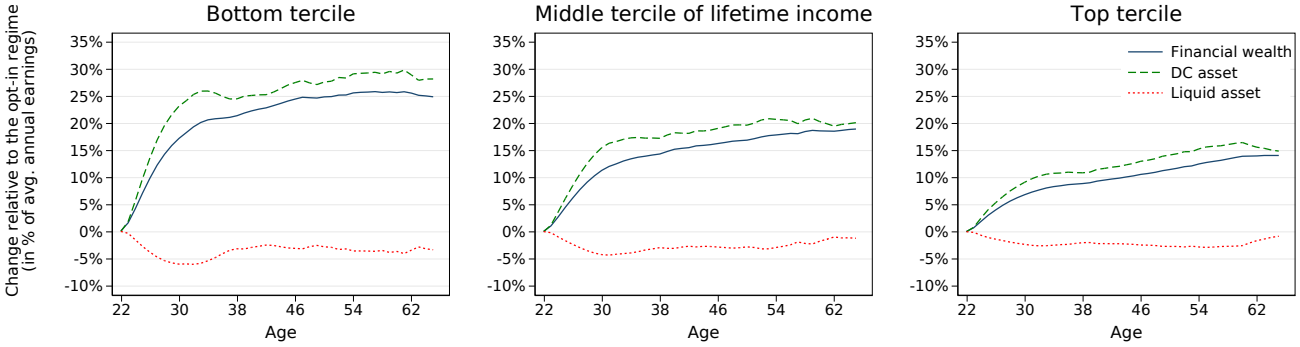
Figure A.19: Effect of auto-enrollment on lifecycle DC savings, liquid assets, and financial wealth



Notes: Each series corresponds to the model-predicted difference between assets under an auto-enrollment policy at 3% and those under an opt-in regime, adopted by all employers. Changes are expressed as a fraction of each simulated individual's average annual earnings between the ages of 22 and 65. The incidence of the policy is assumed to fall on employees, with employers reducing the 401(k) match rate to balance their budget.

F.6.2 6% default & incidence on employers

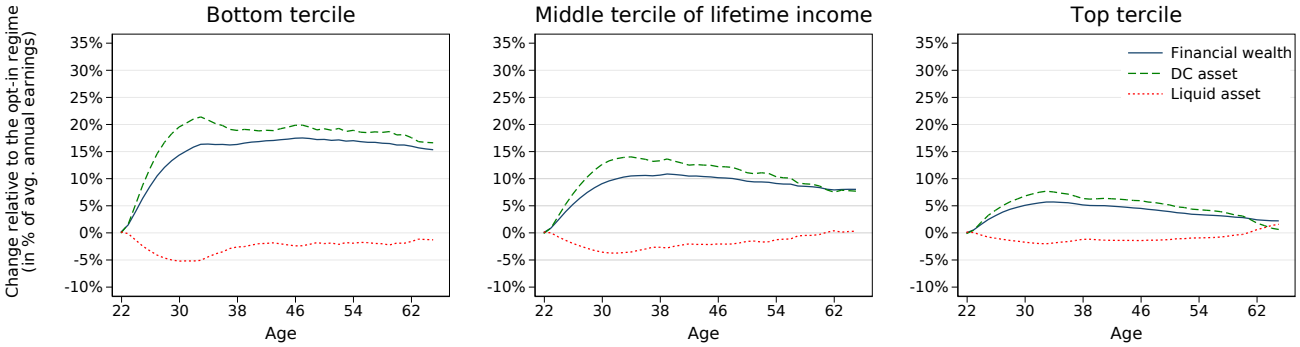
Figure A.20: Effect of auto-enrollment on lifecycle DC savings, liquid assets, and financial wealth



Notes: Each series corresponds to the model-predicted difference between assets under an auto-enrollment policy at 6% and those under an opt-in regime, adopted by all employers. Changes are expressed as a fraction of each simulated individual's average annual earnings between the ages of 22 and 65. The incidence of the policy is assumed to fall on employers, who reduce their profits to cover higher matching costs under auto-enrollment.

F.6.3 6% default & match rate adjustment

Figure A.21: Auto-enrollment's impact on lifecycle DC savings, liquid assets, and financial wealth



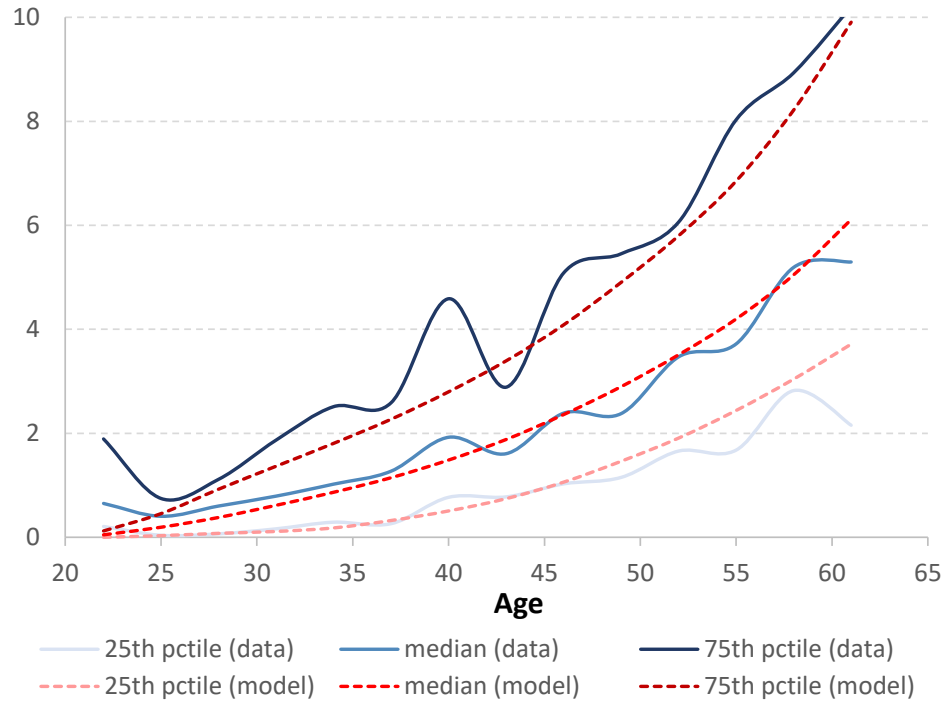
Notes: Each series corresponds to the model-predicted difference between assets under an auto-enrollment policy at 6% and those under an opt-in regime, adopted by all employers. Changes are expressed as a fraction of each simulated individual's average annual earnings between the ages of 22 and 65. The incidence of the policy is assumed to fall on employees, with employers reducing the 401(k) match rate to balance their budget.

G Additional Model Counterfactuals

G.1 Wealth accumulation over the lifecycle

While the model is estimated with 401(k) data only, my baseline estimates are also consistent with the evolution of total wealth accumulation over the lifecycle in the 2016 Survey of Consumer Finances (SCF). Because the wealth measures in the SCF are the household level, they are not directly comparable to the model simulations, which are at the individual level. Instead, I compare the evolution of wealth-to-income ratios over the lifecycle. Total wealth (or total net worth in the SCF) is defined as the sum of all assets net of all outstanding debts but excluding the expected value of future defined benefit and Social Security income. For comparability with the estimation sample, I restrict the SCF sample to households in which either the head or the spouse has any type of account-based pension plan in the current job. The evolution (and heterogeneity) in the observed wealth-to-wage income ratios are reported in Figure A.22 alongside the model predictions. Overall the model predicts well the evolution of wealth accumulation over the lifecycle.

Figure A.22: Ratio of net wealth to earnings by age: Data vs. model



Notes: The data series (straight lines) plot the ratio of total wealth to wage income by age of the head of household. Total wealth corresponds to the sum of all assets net of all outstanding debt. The sample is restricted to households where either the head or spouse has any type of account-based pension plan in the current job. The model series (dashed lines) correspond to the model simulation under an opt-in regime in all jobs. Total wealth in the model series corresponds to the sum of retirement and liquid assets net of unsecured debt liabilities. Data: Survey of Consumer Finances 2016.

H U.K. National Auto-enrollment Policy and Calibration

H.1 U.K. national auto-enrollment policy

The U.K. Pension Act of 2008 requires employers to automatically enroll their employees (with the option to opt-out) in a workplace pension scheme. The law sets the minimum combined employee plus employer default contribution rate at 2% of pensionable pay with a minimum employer contribution of 1%. Each employer was given a staging date based on its number of employees. The staging dates for each employer size are reported in Table A.11. Employers were required to enroll all employees aged between 22 and the stage pension age in a workplace pension plan by the staging date.[12] The auto-enrollment requirement applies to both new hires and nonparticipating seasoned

[12]Employers could apply to postpone the implementation of auto-enrollment by up to 3 months after the staging date. Since I do not observe whether a firm applied to use postponement, I treat the staging date as binding.

employees. I refer the reader to Cribb and Emmerson (2016) for more details on the policy and its rollout.

Table A.11: Auto-enrollment staging dates by employer size

Employer size	Auto-enrollment staging date	Employer size	Auto-enrollment staging date	Employer size	Auto-enrollment staging date
120,000+	October 1, 2012	2,000+	August 1, 2013	61+	August 1, 2014
50,000+	November 1, 2012	1,250+	September 1, 2013	60+	October 1, 2014
30,000+	January 1, 2013	800+	October 1, 2013	59+	November 1, 2014
20,000+	February 1, 2013	500+	November 1, 2013	58+	January 1, 2015
10,000+	March 1, 2013	350+	January 1, 2014	54+	March 1, 2015
6,000+	April 1, 2013	250+	February 1, 2014	50+	April 1, 2015
4,100+	May 1, 2013	160+	April 1, 2014	40+	August 1, 2015
4,000+	June 1, 2013	90+	May 1, 2014	30+	October 1, 2015
3,000+	July 1, 2013	62+	July 1, 2014		

Source: Cribb and Emmerson (2016).

H.2 U.K. calibration

In what follows, I detail the parametrization of the model in the U.K. validation exercise presented in Section 7.2. Other parameters are set equal to their value in the U.S. calibration.

Employer matching formulas. ASHE does not collect data on each employer's matching formula, and I am not aware of any representative evidence on the distribution of employer matching formulas in the U.K. To overcome this data limitation, I use data on the distribution of employee and employer contributions prior to the auto-enrollment rollout to back out each employer's matching formula. Because a cap on matching introduces a kink in the budget set, we should expect the distribution of contribution rates to feature excess mass at the employer matching threshold. Therefore, I hypothesize that the cap on the employer matching contribution is equal to the modal (positive) employee contribution rate in each firm. If this hypothesis is correct, employer contributions should increase in employee contributions up to the modal employee contribution rate in each firm, but not above this hypothesized cap on employer matching. I test and validate this hypothesis in Table A.12 by regressing employer contributions on employee contributions above and below the modal employee contribution rate. Columns 1 and 2 show that, up to the modal employee contribution, each percentage-point increase in contributions raises the employer contribution rate by half a percentage point. Columns 1 and 3 show that employee contributions above the modal contribution rate do not lead to significantly higher employer contributions. These results are consistent with employers offering a 50% match on contributions up to the modal contribution.

In the calibration of the model to the U.K. environment, I group employers into 5 types based

on their employer contribution formula. Motivated by the evidence in Table A.12, I set the match rate equal to 50% for all employers with immediate vesting. I assume that employers have one of five different matching thresholds (cap_e), based on the five most common modal contribution rates observed in the data (and covering more than 80% of employees). The calibrated probability distribution of employer types is reported in Table A.13 and reflects the empirical distribution of modal employee contributions.

U.K. income profile. I estimate the age-income profile using earnings data for U.K. private sector workers aged 22 to 65 using ASHE data from between 1997 and 2016. I use the same regression specification used to estimate the income profile for U.S. workers. The estimation results are reported in Table A.14. For simplicity, I assume that labor market transitions in the U.K. calibration vary by age and tenure but not by income.

Table A.12: Relationship between employer and employee retirement contributions in the U.K.

	Employer contribution in % of pay		
	(1) Full sample	(2) Employee contrib. ≤ modal contrib.	(3) Employee contrib. > modal contrib.
Employee contribution rate			
×dummy (contrib.≤modal)	0.500	0.504	
	(0.14)	(0.14)	
×dummy (contrib.>modal)	0.229		0.155
	(0.158)		(0.123)
Modal contribution rate in the firm			
×dummy (contrib.≤modal)	0.471	0.499	
	(0.306)	(0.333)	
×dummy (contrib.>modal)	0.923		0.812
	(0.408)		(0.407)
R-squared	0.146	0.143	0.111
Observations	21771	17856	3915

Notes: The outcome variable is the employer retirement contribution amount as a percentage of pensionable pay. The modal (positive) employee contribution is computed separately for each employer. I restrict the sample to firms with at least 50 observations in the three years prior to the auto-enrollment staging date. Data source: ASHE waves 2006 to 2016.

Table A.13: Probability distribution of employer types

	Threshold on matching (cap_e)				
	2%	3%	4%	5%	6%
Prob.	0.12	0.18	0.15	0.25	0.30

Notes: This table shows the probability distribution of match thresholds in the U.K. calibration of the model. This distribution is inferred from the empirical distribution of modal contribution rates in the three years prior to the auto-enrollment staging date. Data source: ASHE waves 2006 to 2016.

Table A.14: U.K. age-earnings profile

Age component			
δ_0	δ_1	δ_2	δ_3
5.306	0.255	−0.0043	0.0000219

Notes: Age-earnings profile estimated on a panel of workers continuously employed in the same job. Data source: ASHE waves 1997 to 2016.

Early withdrawals. Early withdrawals from a DC account are not allowed in the U.K. Thus, I impose a 100% tax penalty on early withdrawals $pen_a = 1$.

Income taxation. Individuals' income tax liability is calculated according to the U.K. income tax schedule of 2006 for an individual claiming the personal allowance for singles. The quarterly income tax liability is equal to:

$$tax_t^i = \begin{cases} 0 & if\ 4y^{tax} \leq £4,895 \\ 0.22\left(y^{tax} - \frac{4,895}{4}\right) & if\ £32,400 \geq 4y^{tax} > £4,895 \\ 0.22\left(\frac{32,400-4,895}{4}\right) + 0.40\left(y^{tax} - \frac{32,400}{4}\right) & if\ 4y^{tax} > £32,400 \end{cases}$$

Public pension. The U.K. State Pension provides public pension benefits similar to Social Security in the U.S. The relationship between average lifetime earnings and public pension entitlements is modeled to match the evidence in Figure 1 of O'Dea (2018), expressed in 2006 pounds.

Unemployment benefits. The main unemployment benefit in the U.K. is the Jobseeker's Allowance. In 2006, the quarterly allowance was equal to £ 746 for unemployed individuals older than 24.

H.3 Out-of-sample validation: job switchers in the context of the U.K. national auto-enrollment policy

Fact VII of Section 3 shows that auto-enrollment causes workers to contribute less in their next employer's opt-in retirement savings plan, with no signficant effect when the new job has auto-enrollment. The model predicts a qualitatively similar response.

Model exercise. For each simulated individual, I change the enrollment regime faced in the first two jobs after a randomly drawn quarter t^{exp}. I assume that all employers have an opt-in enrollment regime prior to t^{exp}. I simulate the model under four scenarios: (i) AE to non-AE: auto-enrollment in job number 1 (in period t^{exp} for both new hires and seasoned nonparticipating employees) and opt-in in job number 2, (ii) non-AE to non-AE: opt-in in both jobs, (iii) AE to AE: auto-enrollment in both jobs, and (iv) non-AE to AE: opt-in in job number 1 and auto-enrollment in job number 2. From job number 3 after t^{exp}, all employers auto-enroll their new hires. I compute the difference in contributions in the first year of tenure at job number 2 between scenarios (i) and (ii) to obtain the model counterpart to the coefficient AEtononAE in Table A.3. I compare scenarios (iii) and (iv) to obtain the model counterpart to the coefficient AEtoAE in Table A.3.

Results. The model predicts that auto-enrolling workers in job number 1 causes them to participate 8.4% less in their first year when job number 2 has an opt-in regime. This prediction aligns well with the empirical estimate of 12.8% (*s.e.* 5.5%) lower participation in the next opt-in job (column 1 of Table 3). However, the model-predicted decline in the contribution rate of 0.1 p.p. of salary is smaller than the empirical estimate of a 0.55 p.p.(*s.e.* 0.226) of salary (column 2 of Table 3). Interestingly, the model predictions for both the participation and contribution rates are closer to the point estimates for new hires in jobs who do not expect be auto-enrolled in the next 12 months: a drop in participation of 7.5% (*s.e.* 5.2%) and in the contribution rate of 0.12 p.p. (*s.e.* 0.266) (see Appendix Table A.4, column 6). For employees moving from an auto-enrollment employer to another auto-enrollment employer (i.e., AE to AE relative to non-AE to AE), the model predicts a modest drop in participation of 2.8% and a small drop in contributions of 0.01 p.p. of salary. This is broadly consistent with the empirical finding of no statistically significant difference between previously auto-enrolled and previously non-auto-enrolled workers when the new employer has auto-enrollment (coefficient AE to AE in Table 3).

I Numerical procedure

Discretization. The model has four discrete and five continuous state variables that need to be discretized: labor productivity, tenure, liquid assets, retirement wealth, and average lifetime earnings. Labor productivity is placed on a grid with 7 elements using the method of Tauchen (1986). Tenure is placed on a grid with 9 elements (with all periods past the 9th quarter of tenure treated equally). Liquid assets and retirement wealth are discretized in a way that gives smaller gaps between successive entries on the grid at lower levels and are placed on a grid with 50 elements for each period of retirement and 16 elements for each period of working life. Average earnings are placed on a grid of 7 elements.

Consumption is not placed on a grid, and individuals can choose any feasible consumption level. 401(k) contribution rates are naturally discrete, and in the firms in the estimation sample, contributions are often restricted to below 15% of salary. Thus, I adopt a grid of contribution rates with 16 elements: 0%, 1%,...,15% of salary. Early withdrawals from the DC account are restricted to taking one of ten values. That is, unemployed individuals can withdraw 0%, 1%, 2%, 5%, 10%, 15%, 25%, 50%, 75% or 100% of their DC wealth. In retirement, DC withdrawals can take one of 101 values (0%, 1%,2%,...,99%, 100%).

Interpolation. To evaluate the value function at points not in the discrete subset of points in the discretized state space, I use linear interpolation in multiple dimensions. To limit approximation error despite having fairly coarse grids, I use a method proposed by Carroll (2022). I approximate a quasi-linear transformation of the value function. The transformed value function is closer to linear than the actual one, which implies that linear interpolation is more accurate despite my having few grid points. I verify the accuracy of my numerical procedure by simulating the model—at the parameter values in my baseline estimation—with more grid points and obtain similar results.

Optimization. I solve for the DC contribution and withdrawal choices using a grid search. Given a DC contribution or withdrawal choice, I solve for the liquid asset level (which implies the consumption level) using a golden section search.

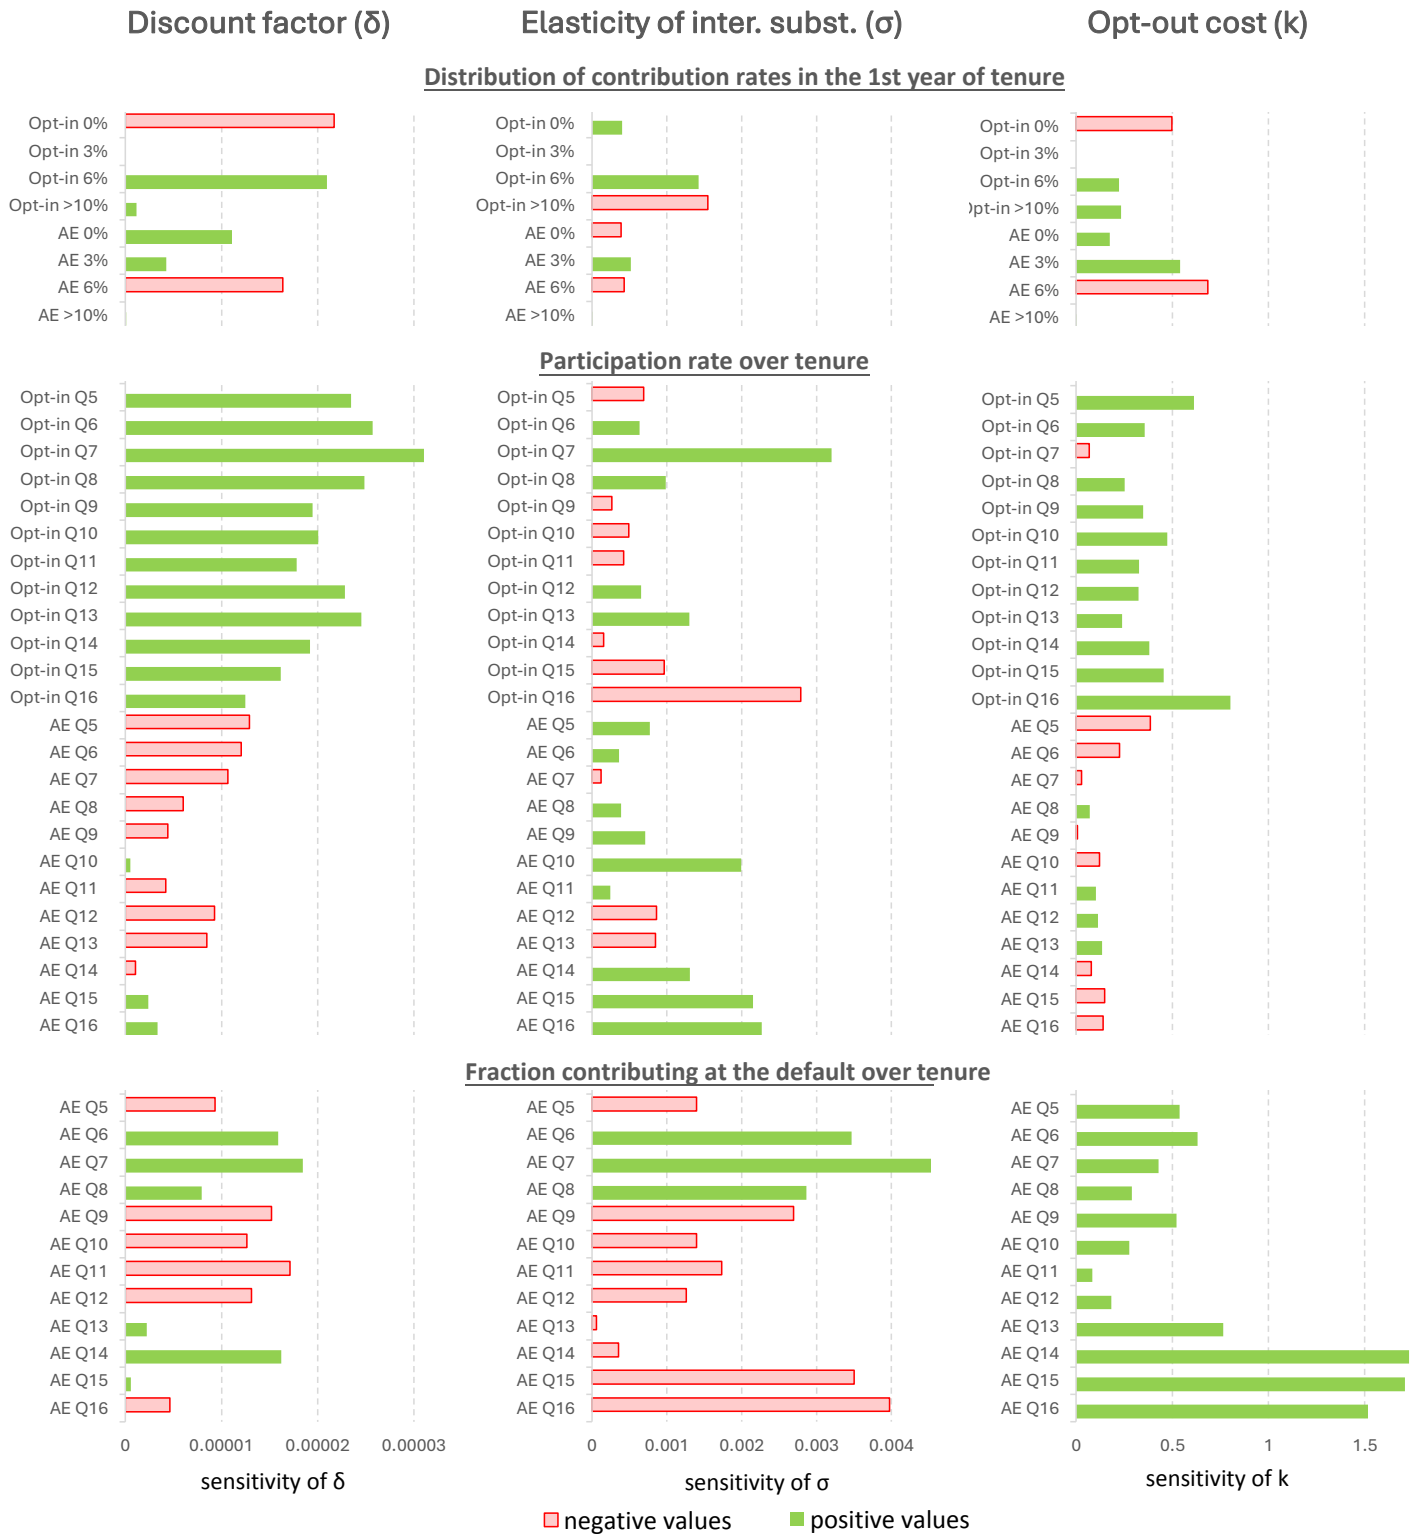
Estimation. I estimate the model in two steps. In the first step I solve the model for 2,744 unique combinations of the model's three parameters for the baseline model specification under both default option (3% and 6%). In the second step, I use the Hooke-Jeeves optimization method to estimate the preference parameters in a region close to the optimum values from the first step. For each combination of preference parameters, I solve for the value function by iterating the problem

from the last period of life (period 272 corresponding to age 90). I then simulate 5,000 individuals in the opt-in group and 5,000 in the auto-enrollment group with identical realization of the stochastic variables. Estimation is computationally intensive. The program was compiled with Intel Fortran, and estimation was is done across 16 simultaneous jobs each parallelized on 42 computer processors using OpenMP on the MIT High Power Computing facility.

J Sensitivity

I plot in Figure A.23 the elements of the sensitivity matrix as defined in Andrews et al. (2017). These values capture how a 1-percentage-point increase in each empirical moment changes the estimated parameters.

Figure A.23: Sensitivity of preference parameters to estimation moments



Notes: This figure shows the sensitivity matrix as defined in Andrews et al. (2017) under the baseline preference estimates from column 1 of Table 4. . Sensitivity values are rescaled to correspond to a 1-percentage-point increase in each moment.

Time preferences. According to the sensitivity matrix plotted in Figure A.23, more contributions at the employer match threshold (6% of salary), and fewer contributions above (≥10%) in the opt-in cohort, would imply to a higher estimate of the elasticity of intertemporal substitution. This result is consistent with the finding in Best et al. (2020) that bunching at—with missing mass above—an interest rate notch can identify the elasticity of intertemporal substitution separately from the discount factor.

Switching cost. The sensitivity measure implies that more bunching at the auto-enrollment default (especially later in employees' tenure) would increase the estimated switching cost. The share of 0% contributions under opt-in is jointly determined by (i) the discount factor (the lower the intertemporal discount factor, the higher is the share contributing 0%) and (ii) the switching cost (0% is the default contribution rate for the opt-in group, and the higher the switching cost, the higher is the share staying at the 0% default). The sensitivity matrix indicates that a higher share at 0% leads to a lower estimate of the discount factor (both in the first year of tenure and later up to four years of tenure). When agents are estimated to be more impatient, a smaller switching cost is required to match the observed bunching at the default. This indirect effect (i.e., a lower discount factor leading to a lower switching cost) dominates the direct effect (i.e., a higher switching cost leading to more mass at the 0% opt-in default), and the sensitivity matrix implies that a higher share at 0% under opt-in leads to a smaller switching cost. Identification is simpler in the case of contributions at 3% under the auto-enrollment regime. The intertemporal discount factor is not as sensitive to the share of contributions exactly at 3%, and thus, in the model, the size of bunching at the 3% auto-enrollment default primarily reflects the size of the switching cost. Both in the first year of tenure and for every quarter 5 to 16 of tenure, a higher share of contributions at 3% under auto-enrollment leads to a higher estimate of the switching cost. A 1 p.p. increase in the share contributing 3% in the first year of tenure leads to a higher switching cost by around 50cts, and a 1 p.p. increase in the share contributing exactly 3% after 4 years (16th quarter of tenure) increases the estimated switching cost by $1.5.

References

Altmann, S., F. Armin, P. Heidhues, and R. Jayaraman (2019). Defaults and donations: Evidence from a field experiment. *Review of Economics and Statistics 101*(5), 808–826.

Andrews, I., M. Gentzkow, and J. M. Shapiro (2017). Measuring the sensitivity of parameter

estimates to estimation moments. *The Quarterly Journal of Economics 132*(4), 1553–1592.

Arnoud, A., T. Choukhmane, J. Colmenares, C. O'Dea, and A. Parvathaneni (2021). The evolution of u.s. firms retirement plan offerings: Evidence from a new panel data set. *Working Paper*.

Bernheim, B. D., A. Fradkin, and I. Popov (2015). The welfare economics of default options in 401 (k) plans. *American Economic Review 105*(9), 2798–2837.

Best, M., J. Cloyne, and E. Ilzetzki (2020). Estimating the elasticity of intertemporal substitution using mortgage notches. *Review of Economic Studies 87*(2), 656–690.

Blumenstock, J., M. Callen, and T. Ghani (2018). Why do defaults affect behavior? experimental evidence from afghanistan. *American Economic Review 108*(10), 2868–2901.

Borella, M., M. De Nardi, and F. Yang (2018). The aggregate implications of gender and marriage. *The Journal of the Economics of Ageing 11*, 6–26.

Brown, J. R., A. M. Farrell, and S. J. Weisbenner (2016). Decision-making approaches and the propensity to default: Evidence and implications. *Journal of Financial Economics 121*(3), 477–495.

Brown, J. R. and A. Previtero (2020). Saving for retirement, annuities and procrastination. *Working Paper*.

Caplin, A. and D. Martin (2017). Defaults and attention: the drop out effect. *Revue économique 5*(5), 747–755.

Carroll, C. D. (2022). Solution methods for microeconomic dynamic stochastic optimization problems. *Working Paper*.

Carroll, G. D., J. J. Choi, D. Laibson, B. C. Madrian, and A. Metrick (2009). Optimal defaults and active decisions. *The quarterly journal of economics 124*(4), 1639–1674.

Choukhmane, T. and T. de Silva (2023). What drives investors portfolio choices? separating risk preferences from frictions. Technical report, Working paper.

Cribb, J. and C. Emmerson (2016). What happens when employers are obliged to nudge? automatic enrolment and pension saving in the uk. *IFS Working Paper*.

Dahlquist, M., O. Setty, and R. Vestman (2018). On the asset allocation of a default pension fund. *The Journal of Finance 73*(4), 1893–1936.

DellaVigna, S. (2009). Psychology and economics: Evidence from the field. *Journal of Economic literature 47*(2), 315–372.

DellaVigna, S. (2018). Structural behavioral economics. In *Handbook of Behavioral Economics: Foundations and Applications*, Volume 1, pp. 613–723.

DellaVigna, S., A. Lindner, B. Reizer, and J. F. Schmieder (2017). Reference-dependent job search: Evidence from hungary. *The Quarterly Journal of Economics 132*(4), 1969–2018.

Franco, C. (2004). Unemployment benefits reduced by pensions and social security: A fact sheet.

Goda, G. S., M. R. Levy, C. F. Manchester, A. Sojourner, and J. Tasoff (2020). Who is a passive saver under opt-in and auto-enrollment? *Journal of Economic Behavior & Organization 173*, 301–321.

Guvenen, F. (2009). An empirical investigation of labor income processes. *Review of Economic dynamics 12*(1), 58–79.

Haggag, K. and G. Paci (2014). Default tips. *American Economic Journal: Applied Economics 6*(3), 1–19.

Kaplan, G. and G. L. Violante (2014). A model of the consumption response to fiscal stimulus payments. *Econometrica 82*(4), 1199–1239.

Laibson, D. (1997). Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics 112*(2), 443–478.

Laibson, D., P. Maxted, S. Chanwook Lee, A. Repetto, and J. Tobacman (2024). Estimating discount functions with consumption choices over the lifecycle. *The Review of Financial Studies*.

Low, H., C. Meghir, and L. Pistaferri (2010). Wage risk and employment risk over the life cycle. *American Economic Review 100*(4), 1432–1467.

Lusardi, A., P.-C. Michaud, and O. S. Mitchell (2017). Optimal financial knowledge and wealth inequality. *Journal of Political Economy 125*(2), 431–477.

Madrian, B. C. and D. F. Shea (2001). The power of suggestion: Inertia in 401 (k) participation and savings behavior. *The Quarterly journal of economics 116*(4), 1149–1187.

Menzio, G., I. A. Telyukova, and L. Visschers (2016). Directed search over the life cycle. *Review of Economic Dynamics 19*, 38–62.

O'Dea, C. (2018). Insurance, efficiency and the design of public pensions. *Working Paper*.

O'Donoghue, T. and M. Rabin (1998). Procrastination in preparing for retirement. In *Behavioral Dimensions of Retirement Economics*.

Samuelson, W. and R. Zeckhauser (1988). Status quo bias in decision making. *Journal of risk and uncertainty 1*, 7–59.

Sun, L. and S. Abraham (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics 225*(2), 175–199.

Tauchen, G. (1986). Finite state markov-chain approximations to univariate and vector autoregressions. *Economics letters 20*(2), 177–181.