

NBER WORKING PAPER SERIES

STATE AUTO-IRA POLICIES AND FIRM BEHAVIOR:
LESSONS FROM ADMINISTRATIVE TAX DATA

Adam Bloomfield
Lucas Goodman
Manita Rao
Sita Slavov

Working Paper 32817
<http://www.nber.org/papers/w32817>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2024, Revised April 2025

We are grateful to Angela Antonelli, Jacob Berman, Christine Cheng, Gopi Shah Goda, Sreeraahul Kancherla, Kim Olson, Paul Organ, Shanthi Ramnath, Melanie Wallskog, and Andrew Whitten, as well as seminar participants at Texas Tech University's Free Market Institute, and attendees of the Treasury Department's Office of Tax Analysis Research Conference, the 117th Annual Conference on Taxation, and the 2024 APPAM Fall Research Conference for helpful input. This research was conducted while Goodman was an employee at the U.S. Department of the Treasury. The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors and do not necessarily reflect the views or the official positions of the U.S. Department of the Treasury, AARP, FDIC, or the U.S. Government. Any taxpayer data used in this research was kept in a secured Treasury or IRS data repository, and all results have been reviewed to ensure that no confidential information is disclosed. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w32817>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2024 by Adam Bloomfield, Lucas Goodman, Manita Rao, and Sita Slavov. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

State Auto-IRA Policies and Firm Behavior: Lessons from Administrative Tax Data
Adam Bloomfield, Lucas Goodman, Manita Rao, and Sita Slavov
NBER Working Paper No. 32817
August 2024, Revised April 2025
JEL No. D14, H75, J26, J32

ABSTRACT

Several states have recently attempted to boost retirement saving by adopting “auto-IRA” policies that require employers not currently offering an employer-sponsored retirement plan (ESRP) to either (1) establish an ESRP or (2) enroll employees in state-facilitated Individual Retirement Accounts (IRAs). We identify the effect of these state policies on firm decisions to offer ESRPs, exploiting the phased rollout of these policy treatments across states and employer size categories. Using U.S. tax microdata, we estimate that about 17% of treated firms have been induced to offer an ESRP by these policies, although there is substantial heterogeneity in these effects across firm and worker characteristics. This effect is large considering that, for employers, establishing and maintaining an ESRP is more costly than utilizing the state-facilitated IRAs. We explore both rational and behavioral explanations for why firms might choose the higher-cost option to comply with state auto-IRA policies.

Adam Bloomfield
U.S. Federal Deposit Insurance
Corporation (FDIC)
550 17th St NW
Washington, DC 20429
adam.bloomfield@gmail.com

Lucas Goodman
Office of Tax Analysis
Department of Treasury
1500 Pennsylvania Avenue NW
Washington, DC 20220
Lucas.Goodman@treasury.gov

Manita Rao
AARP Public Policy Institute
mr Rao@aar p.org

Sita Slavov
Schar School of Policy and Government
George Mason University
3351 Fairfax Drive, MS 3B1 Arlington,
VA 22201
and NBER
sslavov@gmu.edu

I. Introduction

Fringe benefits – largely in the form of employer-sponsored retirement plans (ESRPs) and health insurance – have grown dramatically as a share of employee compensation since the middle of the 20th century (Chen 1981; Wiatrowski 1999; U.S. Bureau of Labor Statistics 2024). A large literature has established that workers place substantial value on nonwage compensation, and that benefits like health insurance and defined benefit pensions affect labor supply and job transitions (e.g., Gruber and Madrian 1994; Garthwaite et al. 2014; Ni and Podgursky 2016; Nyce et al. 2013; Kim 2020). There has also been considerable analysis regarding how fringe benefits and other nonwage job attributes affect measurement of income growth and inequality (e.g., Pierce 2001; Burkhauser, Larrimore, and Simon 2012; Guvenen et al. 2022; Maestas et al. 2023). By contrast, there has been comparatively little research on how employers determine whether to offer fringe benefits to workers. In this paper, we use recent, quasi-experimental variation in state policies to shed light on why employers choose to establish ESRPs.

Most private retirement saving in the U.S. takes place through ESRPs.¹ Over the last fifty years, defined contribution (DC) plans, in which workers contribute to a personal retirement savings account established by their employer, have gradually replaced traditional defined benefit (DB) plans, which pay a retirement benefit according to a formula (U.S. Department of Labor 2023). Federal policy has also incentivized employers to automatically enroll workers in their DC plans (requiring workers to opt out if they do not wish

¹ See Investment Company Institute (2022a). Investment Company Institute (ICI) data on IRAs are drawn from an annual household survey.

to participate), leading to large increases in participation rates. However, a significant proportion of the American workforce works for employers that do not offer ESRPs.² Those who do not have access to an ESRP are disproportionately likely to be young, Black or Hispanic, to have less education and lower incomes, and to be employees of smaller firms (Sabelhaus 2022). While individuals who do not have access to an ESRP may contribute to non-employment based Individual Retirement Accounts (IRAs) – which also provide access to tax-preferred retirement saving – most do not (see, e.g., Investment Company Institute 2022a). In recent years, many states have attempted to boost retirement savings by mandating that all employers offer a way for employees to save for retirement through payroll deduction. Employers can comply with the policy mandate by either offering an ESRP or automatically enrolling employees in an IRA that is facilitated by the state and managed by professional retirement service providers. We refer to these state-facilitated IRAs as state auto-IRA *programs* (or just ‘state programs’), and to a combination of an employer mandate and state auto-IRA program as an auto-IRA *policy*. Employees may opt out of either of these employer mediated savings options.

Starting in 2017, there has been a staggered rollout of these auto-IRA policies across states and firm size categories. In each of the policy expansions that we study, we find clear and substantial increases in the share of affected firms establishing an ESRP immediately upon implementation. We refer to this induced increase in ESRP offerings as the “crowd-in” effect of the policy. We do not find evidence of any offsetting “crowd-out”, i.e. firms

² The administrative tax data used in this paper suggest that 24.5% of U.S. workers who receive a W-2 work for a firm that does not offer an ESRP.

terminating existing ESRPs in favor of utilizing the state auto-IRA program. We estimate that approximately one sixth of firms in the affected employer size ranges were induced to offer an ESRP by the policies. These effects are also substantial relative to the number of firms participating in the state auto-IRA programs directly; we find that ESRP crowd-in accounts for 30 to 45 percent of the total increase in employer coverage, depending on the state. Using a treatment effects framework, we then estimate mean characteristics of “complier” firms induced to offer ESRPs by the policy, as well as their employees and owners. We find that complier firms are more similar to firms that never offer an ESRP than they are to firms that already offer ESRPs.

We analyze our findings in the context of a neoclassical framework where rational employers weigh the costs and benefits of offering ESRPs. ESRPs have substantial costs, as they are typically administered by third parties for a fee and generally require compliance with nondiscrimination rules. Employers do not pay a fee to participate in state auto-IRA programs; however, they do face an administrative burden of signing up for the state’s program and facilitating contributions by employees. On the benefit side, workers value ESRPs because they offer access to tax-advantaged saving. Owners may capture this benefit both directly in their role as an owner-employee (i.e., by participating in the ESRP themselves) and indirectly because their employees value ESRPs. While fully rational workers gain no benefit from the state auto-IRA program – since workers have long been able to contribute to IRAs on their own – “behavioral” workers may value the convenience of automatic contributions, as well as the commitment mechanism they provide. The large crowd-in that we observe empirically requires some combination of high perceived

administrative cost and low perceived benefits of state auto-IRA programs relative to ESRPs. By contrast, if workers value auto-IRAs and ESRPs similarly and the administrative cost of state auto-IRA programs is low, then employers may have an incentive to drop their ESRPs and utilize the less costly state auto-IRA program instead (“crowd out”).³

We then test whether this framework can explain our empirical findings. On the benefit side, we do not find substantial evidence that the crowd-in is driven by perceived low benefits of state auto-IRAs relative to ESRPs. For example, auto-IRAs may have lower benefits than ESRPs because IRAs have lower contribution limits. However, we estimate that the auto-IRA contribution limits typically do not bind, either for employees or owners. On the cost side, while we do not have direct evidence on administrative burden, we find that take-up rates for a federal tax credit for establishing an ESRP – a task that may require several hours of compliance efforts – are very low, suggesting that that paperwork and compliance burdens may weigh heavily on this population.

Finally, we consider how behavioral factors at the firm level might cause a state auto-IRA policy to lead to ESRP crowd-in. In the presence of inertia, for instance, removing the default option of offering no plan may induce employers to revisit their ESRP decision and choose to offer a plan. Furthermore, employers may be responding to marketing that ESRP administrators have undertaken in response to the auto-IRA policies.

There exists a small, recent literature studying state auto-IRA policies. Several papers have studied their direct effect – i.e., the boost in IRA participation brought about by the state

³ Some opponents of auto-IRA policies have cited the potential for crowd-out (e.g., U.S. Chamber of Commerce 2017).

auto-IRA programs (Dao 2024; Quinby et al. 2020; Chalmers et al. 2022). Willingham (2024) finds that employees at adopting firms increase their retirement savings and tend not to make withdrawals even at job separation. Closer to our paper, previous work by Bloomfield et al. (2024) used individual-level data from the Current Population Survey (CPS) and firm-level ESRP filings and found that auto-IRA policies induced ESRP offers.⁴ We add to this literature and provide three substantial contributions beyond Bloomfield et al. (2024). First, our paper confirms and refines the effects identified by Bloomfield et al. (2024). While statistically significant, their CPS-based results are potentially confounded by survey response biases, as household surveys are known to underreport retirement coverage. Thus, these findings may primarily (or entirely) reflect increased worker awareness of existing benefits rather than new ESRP formation. Bloomfield et al. (2024) attempt to address this limitation by examining plan-level disclosure filings, but these data provide limited visibility into firm-level behaviors.⁵ Second, our paper uses administrative tax data that allows for precise treatment assignment (i.e., by firm size and time) and detailed analysis of heterogeneity across firms, workers, and owners. In contrast, the previous study lacked precision in assigning treatment at the firm level and offered limited granularity with regard to firm and worker characteristics. The data precision and granularity in the current study is crucial for understanding the heterogeneous effects of these policies. Third, and most

⁴ Additionally, several commentators have used publicly available or proprietary data to observe increases in ESRP offers in treatment states (Scott 2021, Pardue 2023).

⁵ Firms that do not offer plans or fail to file the required paperwork do not appear in the data. Therefore, the dynamics of firm entry, exit, or non-reporting cannot be disentangled from those of ESRP establishment and termination.

importantly, the comprehensive data in our paper enables us to explore and rule out potential mechanisms driving the employer behavior that we observe.

Our study contributes to several additional streams of research. First, a large body of research has shown that – contrary to the predictions of the rational model – automatic enrollment substantially boosts employee participation in retirement saving accounts in the short run (Madrian and Shea 2001; Choi et al. 2004), although the impact on long-term saving is less clear (Choukmane 2024; Beshears et al. 2022; Derby, Mackie, and Mortenson 2023). However, this literature primarily analyzes worker decisions to participate in ESRPs rather than employer decisions to establish them. We focus on firms’ decisions to establish ESRPs and show that frictions may affect these choices as well.

Second, we contribute to the empirical literature on why firms offer fringe benefits to workers. Previous papers in this literature have considered the role of unions (Freeman 1981), tax policy (Long and Scott 1982, 1984; Turner 1987a, 1987b), worker characteristics (Rhine 1987), economies of scale in benefit provision and search costs (Oyer 2008), and the desire of employers to attract female employees (Liu et al. 2023). Third, our study contributes to the literature on how firms respond to government mandates concerning compensation level and structure. For example, recent papers in the minimum wage literature have examined how minimum wage laws, which regulate monetary compensation, affect nonwage compensation such as health insurance (Clemens et al. 2018; Clemens 2021; Meiselbach and Abraham 2023). Some studies have also examined the impact of the mandates in the Affordable Care Act (ACA), as well as an earlier Massachusetts law, which required some employers to offer health insurance (Kolstad and Kowalski 2016; Lyons 2017;

Abraham et al. 2019). We contribute to both literatures by investigating how state policies affect firms' decisions to offer DC ESRPs, which have become the dominant retirement savings vehicle in the U.S.

Finally, we contribute to a small but growing literature on the role of behavioral economics in explaining firm decision-making (a recent paper is Gertler et al. 2025; for reviews, see Heidhues and Kozzegi 2018; Malmendier 2018). Our results suggest that neoclassical factors cannot fully explain the large crowd-in that we find, leaving behavioral factors such as inertia, salience, and marketing as plausible channels. Thus, these findings are consistent with the hypothesis that owners of small- and medium-sized businesses are susceptible to many of the same behavioral biases that are well-studied in the context of consumers, savers, and employees.

The policy implications of this study are particularly pertinent. As states and the federal government (through the SECURE Act and SECURE 2.0, as well as prospective federal laws under consideration) continue to implement measures designed to boost retirement saving, understanding the impact of state auto-IRA policies on firm and worker behavior is paramount. More broadly, employer mandates are a widespread policy tool that governments use to affect societal goals, so it is important to assess their impact on labor markets. This paper seeks to contribute to a more robust knowledge base that can guide policymakers and employers in their decisions with regard to retirement savings options.

II. Policy Background

Employer-sponsored retirement plans (ESRPs) are a non-wage benefit provided by firms as part of their workers' compensation package.⁶ ESRPs are a tax-advantaged form of saving that may take the form of either a Defined Contribution (DC) or a Defined Benefit (DB) plan.⁷ In a DB plan, employers, and sometimes workers, make mandatory contributions to the plan. Workers receive a retirement benefit, typically in the form of an annuity or lump sum payment, based on a formula that accounts for age, years of service, and salary. In a DC plan, workers make voluntary contributions to an account. Employers may also contribute to the account or match employee contributions. Funds in the account are invested in mutual funds or other securities, and the worker can draw down on the savings to finance consumption during retirement. Over the past several decades, DC plans have grown in popularity, becoming the main type of ESRP offered by the smaller and mid-sized private firms that we focus on in this study (see Department of Labor, 2023). Our analysis in this paper focuses on DC plans, including 401(k)s, 403(b)s, and SIMPLE IRAs.^{8,9} Beyond ESRPs, Individual Retirement Accounts (IRAs) provide most workers with access to tax-advantaged retirement saving even if their employer does not offer a workplace savings vehicle. Contributions to IRAs are voluntary and do not allow an employer match. Though assets held

⁶ Employers may exclude groups of employees (e.g., part-time staff) from participating in their ESRP, although such exclusions may make it harder to meet nondiscrimination tests.

⁷ Tax-advantaged in this context means that no federal, state, or local income taxes are collected on capital gains, dividends, or interest accruing in the plan. In a traditional DC account, withdrawals are taxed as income. In a Roth DC account, employee contributions are taxed as income (but withdrawals are tax-free).

⁸ Savings Incentive Match Plan for Employees (SIMPLE) IRAs are DC ESRPs available to employers with 100 or fewer employees. Both employee and employer contributions can be made to a SIMPLE IRA.

⁹ We do not study DB plan offers in this paper. In principle, a DB plan would satisfy the auto-IRA mandate. However, we do not reliably observe DB plan offer, and we expect that it is uncommon for a firm with fewer than 100 employees to initiate a new DB plan. In 2023, only 5 percent of workers at firms with 1-49 employees had access to a DB plan, and only 8 percent of workers at firms with 50-99 employees had access to a DB plan (Topolski et al. 2024).

within IRAs are substantial in the U.S., most of these funds are the result of “rollovers” from ESRPs (which may take place when a worker leaves a job) rather than direct IRA contributions (Investment Company Institute 2022b).

In this study we focus on policy changes in four states – Oregon, Illinois, California, and Connecticut. Each of these states implemented a policy with two components. The first component is an **auto-IRA program** facilitated by the state and managed by professional retirement service providers. Under these state programs, employers provide the state with a list of their employees, each of whom is automatically enrolled in an IRA managed by a third-party financial institution. Employers facilitate worker savings through payroll deductions, at the level specified by the state program, which are subsequently invested in mutual funds within the IRA. Employees can modify or opt out of these contributions at any time. The second component is an **employer mandate**, under which firms must offer workers a workplace retirement savings option either by participating in the state auto-IRA program or by offering its own ESRP to their employees.¹⁰ Employers who fail to adhere to the employer mandate face the risk of financial penalties, typically charged on a per-employee basis. Throughout this paper, we refer to the combination of both components as a state **auto-IRA policy**.^{11,12}

¹⁰ An ESRP used to satisfy the mandate does not need to feature automatic enrollment.

¹¹ Some states (e.g., Missouri, Massachusetts, and Washington) have experimented with alternative policy structures to encourage retirement saving, such as voluntary state-based ESRP marketplaces and multiple employer plans (MEP). However, employer mandates like those we study in this paper have become by far the most common policy approach across the 19 states that have adopted retirement savings policies.

¹² State-facilitated auto-IRA programs were generally created and open for voluntary enrollment for a period prior to the first mandate deadline. Employers not yet subject to the mandate have the option to enroll employees in the auto-IRA program.

The policy changes we study took effect at various points between 2017 and 2022. In each adopting state, the employer mandate was implemented in waves, with larger firms becoming subject to it before smaller firms. We consider each expansion of a state’s employer mandate as a separate policy “treatment”. Due to the small number of 100+ employee firms that are affected by the policy, we focus on expansions of the policies to firms with fewer than 100 employees. The policy treatments we consider in this paper are summarized in Table 1. They include two expansions in Oregon (to firms with 20-99 and 5-19 employees), two expansions in Illinois (to firms with 25-99 and 16-24 employees), two expansions in California (to firms with 50-99 and 5-49 employees) and one expansion in Connecticut (to firms with 26-99 employees). All policy treatments took place between 2018 and 2022. These treatments encompass all state auto-IRA policies (for firms with fewer than 100 employees) that were implemented in 2022 or earlier, giving us at least two years of post-implementation data for each treatment.¹³

III. Data:

We use administrative tax data to investigate the impact of state auto-IRA policies on firm ESRP offerings. We create an employer-level annual panel from the universe of tax filings from 2012 through 2023. We start with Form 941, a quarterly form which employers use to report (among other things) payroll tax obligations and individual income taxes withheld on

¹³ There were several additional auto-IRA policy implementations that took place in 2023. Additionally, Maryland implemented a version of the auto-IRA policy in 2022 that did not include an employer mandate. We do not study such policy treatments in the main analysis in the paper, and we exclude firms in these states (in affected size categories) from the control group.

behalf of employees. Form 941 indicates the state where the firm is located, as well as counts of employees at a point in time during the quarter. We collapse across quarters within a given year to get employee counts at the annual level. We treat each unique Employer Identification Number (EIN) as a distinct employer; while some large employers have many EINs, this is less common for the small and medium-sized firms that are our primary employers of interest.¹⁴ We drop government employers, who are not affected by the state auto-IRA policies. We take the full population of firms from states with an auto-IRA policy during our analysis window (California, Oregon, Illinois, and Connecticut);¹⁵ for computational tractability, we take a 10% random sample of firms from other states. We retrieve an industrial classification for most firms from their business tax returns. We also retrieve information about offers of health insurance coverage from Forms 1095-B and 1095-C beginning in 2015.

We then link these employers via EIN to Form W-2. We use Form W-2 for two purposes. First, while we cannot directly observe whether a firm offers an ESRP, we can observe (in Box 12) employee contributions to such plans.¹⁶ We code an employer as offering an ESRP if and only if at least one employee makes a contribution.¹⁷ Second, we use Form W-2 to identify all employees of a given firm in a given year (and their wages). We then link

¹⁴ Less commonly, some small- to mid-sized firms contract with Professional Employer Organizations (PEOs), who are typically large payroll administrators; in such cases, the employees of the firm appear on Form 941 of the PEO. Implicitly, such firms are not in our universe.

¹⁵ For auxiliary analyses, we also retrieve a 100% sample of firms in Colorado. Such firms are not included in the main analysis.

¹⁶ Additionally, Box 12 allows us to differentiate between SIMPLE IRAs and other types of DC plans.

¹⁷ To test the sensitivity of this firm ESRP offer definition, we tested the share of firms classified to offer an ESRP if and only if at least 10% of employees make an ESRP contribution. We find that 87.3% of these firms remain classified as offering an ESRP at this much higher threshold.

these employees to other databases by Taxpayer Identification Number (TIN), which is usually the Social Security Number (SSN). In particular, we link to data derived from Social Security Administration (SSA) records to identify date of birth and sex. We link to other tax forms, including Form 1040 (individual income tax return), to retrieve additional characteristics such as geography, the presence of dependents, marital status, and income composition. We impute race and Hispanic ethnicity based on residential zip code, first name, and surname.¹⁸

We also link employers to their natural person owners, when possible. This is more feasible for certain entity types (e.g., S corporations and sole proprietorships) than others (e.g., C corporations and non-taxable entities). Specifically, finding the owner of a sole proprietorship is immediate – the sole proprietorship is linked to an individual tax return. For S corporations and partnerships, we use Schedule K-1 (of Form 1120S and 1065, respectively). While S corporation owners are usually natural persons, that is often not the case for partnerships; we do not attempt to trace through layers of partnership tiers to arrive at indirect owners. For closely held C corporations, we use Schedule G of Form 1120, which lists owners with at least 20% direct ownership share. We retrieve the same information for each owner that we do for each employee.

IV. Methods and Results

¹⁸ Specifically, we use the BIFSG method, described for instance in Voicu (2018). We obtain data on the relationships among race, first name, surname, and zip code from Tzioumis (2018) (race and first name), U.S. Census Bureau (2021) (race and surname), and Manson et al. (2023) (race and zip code). We note that imputing race and ethnicity introduces measurement error and therefore additional uncertainty into our estimates (see Lu et al. 2024). Given the large sample size, we do not expect adjusting for this uncertainty to change the statistical significance of our findings related to race and ethnicity.

a. Summary Statistics

Table 2 provides summary statistics for our sample for the treatment and control groups in the pre- and post-policy periods. For firms in each category, we present the sample size, the share offering at least one ESRP, the average number of employees, the average annual employee contributions to a firm's ESRP(s), and both the average and median annual employee wages. Our sample includes approximately 180,000 distinct firms in control states (recall that we randomly sample firms in control states at a rate of 10%) and 270,000 distinct firms in treated states. Pre-policy, the share of firms offering at least one retirement plan is modestly smaller in the treatment group compared to the control group. Workers at treated firms also earn modestly higher wages and contribute a larger amount to ESRPs on average, although the ratio of average ESRP contributions to wages is quite similar for the treatment and control firms. The average number of workers at treated and control firms remains similar in the post-policy period. Post-policy, the share of firms offering an ESRP is greater in the treatment group compared to control firms. Workers at treated firms continue to earn modestly higher wages than those at control firms. The average level of worker ESRP savings is larger for both treatment and control firms, although the percentage increase is greater in the treatment group.

b. Overall Crowd-In and Crowd-Out

We employ a stacked event study methodology (Cengiz, et al. 2019) to estimate the effect of state auto-IRA policies on firms' decisions to offer an ESRP. For each of the seven policy treatments in Table 1 (indexed by e , with treated state designated by $s^*(e)$), we create

a panel of treatment firms and control firms. Treatment firms are those firms in the relevant size range in the treated state, while control firms are those firms in the same size range in the never-treated states.¹⁹

We consider several dependent variables, including two “flows” and one “stock.” The primary “flow” of interest is a dummy for starting a plan – that is, it equals one if the firm does not offer a plan in the prior year and offers a plan in the current year, and it equals zero in all other cases. That is, “starts plan” measures the unconditional probability of starting to offer an ESRP among firms that existed during the previous year. We then stack all seven panels and run the following regression for firm i (located in state $s(i)$) in policy treatment e at event time k :²⁰

$$start_{iek} = \lambda_{ie} + \mu_{ek} + \sum_{m=-5, m \neq -2}^{k^{max}} \beta_m \times 1(s(i) = s^*(e)) \times 1(k = m) + u_{iek} \quad (1)$$

This regression traces out the mean value of $start_{iek}$ (relative to event time -2) after subtracting out the change in $start$ in same-sized firms in never-treated states. We also consider the opposite flow outcome, $stop_{iek}$, which equals one if and only if the firm offered a plan in the prior year and not in the current year. While studying flow outcomes, we allow the sample of firms to be unbalanced – i.e., firms can enter and exit the sample, either because they start or stop operations entirely, or because they exit or enter the relevant size range for the given policy treatment.²¹ However, we enforce sample balance at the state-year

¹⁹ We assign firms to size buckets based on their lagged employee count from Form 941. To improve the accuracy of treatment assignment, we exclude firms within 10% of the bottom of the size range and within 20% of the top of the size range. The larger margin at the top of the size range is designed to avoid classifying larger firms as “small”, as larger firms are treated earlier in each state.

²⁰ We estimate this regression using the user-written Stata command `reghdfe` (Correia 2017).

²¹ Firms are in the sample at event time k if their lagged firm size (i.e., number of employees at $k - 1$) is within the relevant range for policy treatment e .

level – that is, we keep observations from event times -5 through k^{max} , and we restrict the sample to policy treatments where that entire time period is observed. In order to interpret the results causally, we make the standard parallel trend assumption that the flow outcome would have followed a parallel trend in treatment states versus control states, but for treatment. We cluster our standard errors by firm.

We also consider a the “stock” outcome of having a plan – i.e., $plan_{iek}$, a dummy that equals one if the firm offers a plan in the current year. We modify our estimation method slightly in this case. First, we define whether a firm is in the sample using their firm size at event time -2. Second, we enforce balance in terms of true entry and exit – that is, we drop any firms that start or cease operations entirely during our sample window. And third, we add a treatment-specific linear interaction; this allows us to impose the milder parallel trend assumption in *flows* rather than a stronger parallel trend assumption in *levels*.²² We omit event time -5 in order to identify the linear interaction; this means that the trends are estimated using changes between event times -5 and -2.

Figure 1 presents the main results of estimating equation (1). Panels A and B present results for the “starts plan” outcome, while Panels C and D present results for the “offers plan” outcome. Panels A and C use the full set of seven main policy treatments; this requires us to impose $k^{max} = 1$, as the 2022 treatments cannot extend past event time 1. Panels B and D drop the 2022 policy treatments and allow us to extend to $k^{max} = 2$. In all panels, we see immediate (and precisely estimated) increases in the rates of offering or starting an ESRP

²² For example, if treatment firms started plans at a higher rate prior to treatment, this would violate parallel trends in levels but not necessarily parallel trends in flows.

at event time zero. These increases continue into event time +1, at a slower rate. In Panels B and D, we see that the increase in ESRP offering levels off and approximately stops by event time +2. The treatment effects in Panels B and D (dropping the 2022 policy treatments) are somewhat smaller than in Panels A and C. We note that the 2022 policy treatments – especially the California 5-49 employee treatment – represent a clear majority of the affected firms. Dropping the 2022 policy treatments tends to shift the composition of affected firms toward larger firms, which may have a smaller treatment effect. In Appendix Figures A1-A7, we plot the full event study (using all available years of data) for each of the seven policy treatments separately; in each case, the ESRP offer rate increases noticeably at event time zero. Additionally, in Appendix Figure A8, we plot the stacked event study for the “stops plan” outcome; we do not uncover any economically significant effects on the probability of ending an ESRP.²³ These results are broadly consistent with a substantial “crowd-in” effect – with a small, if any, crowd-out effect – of state auto-IRA policies on firm ESRP offerings.²⁴

Table 3 considers the magnitude of the crowd-in effect in more detail. For each of the seven policy treatments, we compute the number of firms induced to offer an ESRP, which we report in column (1). We compute this in two steps. First, at event times $k = \{-1, 0, 1\}$, we compute the product of (a) the event time k event study coefficient for the “starts plan”

²³ For the “stops plan” outcome, we must drop observations in 2023 and thus reduce the post-period window by one year. The 2023 W-2 data is slightly incomplete, and the “stops plan” outcome is much more sensitive to this incompleteness than other outcomes are.

²⁴ In Appendix Figures A9-A11, we show that these results are robust to (a) changing the definition of “offers plan” from *any* employee making a DC contribution to at least 10 percent of employees making a DC contribution (b) dropping firms that may operate in multiple states, defined by having at least 20% of their workforce living in a state that differs from the headquarters address, and (c) restricting to firms that file an annual tax return (Forms 1120S, 1065, 1120, 990, or Form 1040, Schedule C) with the same EIN.

outcome, estimated solely for that policy treatment, and (b) the number of firms in the relevant size range in the treated state at event time k . Second, we take the sum of this object across event times $k = \{-1, 0, 1\}$. In column (2), we report the number of firms in the size range in each policy treatment at event time -2 that do not offer a plan; in column (3), we report the ratio of column (1) to column (2). We find that, across treatments, state auto-IRA policies induce between approximately 8% (in Illinois 16-24) and 23% (in California, 50-99) of non-offering firms to offer an ESRP.

In Table 4, we compare the magnitude of firms induced to offer ESRPs in Oregon, California, Illinois, and Connecticut (column 1) to the number of firms that are actively participating in the state-facilitated auto-IRA programs (column 2).²⁵ In the absence of crowd-out, the sum of columns (1) and (2) reflect the increase in the number of firms that respond to these state retirement policies by offering an ESRP or participating in the state program, respectively. This increase can be interpreted as the “total effect” of the state policies on firm offerings of retirement savings vehicles. In column (3), we compute the share of firms opting to adhere to the mandate by offering an ESRP. We find that ESRPs account for between 30% (Oregon) and 45% (Connecticut) of the total effect. Thus, an analysis of state

²⁵ We retrieve the latter from publicly available data compiled by the Center for Retirement Initiatives (CRI) at Georgetown University (<https://cri.georgetown.edu/states/state-data/current-year/>) as well as directly from states that post their auto-IRA program statistics to their websites (<https://www.treasurer.ca.gov/calsavers/reports/2024/index.asp>, <https://osc.ct.gov/crsa/reports.html>, https://illinoistreasurer.gov/Individuals/Secure_Choice/Secure_Choice_Performance_Dashboards, and <https://www.oregon.gov/treasury/financial-empowerment/Pages/Oregon-Retirement-Savings-Board.aspx>). These data are only available at the state level (rather than at the policy treatment / firm size level). We measure participation in 2022 in Oregon, Illinois, and Connecticut and 2023 for California; these reflect the most recent years of data prior to the expansion to smaller firms (which we do not study). For disclosure reasons, we are unable to use the tax data to estimate the number of firms participating in state auto-IRA programs.

auto-IRA policies that studies only the direct effects on auto-IRA participation would miss a large share of the overall retirement coverage increases induced by the policy.

The aggregates reported in Table 4, when compared to column (3) in Table 3, imply that many treated employers neither establish a new ESRP nor participate in the state auto-IRA program, despite the mandate to choose one of these two options. To interpret this finding, we note that we define an employer to be participating in the state auto-IRA program if they have provided payroll deductions in the past 90 days. In general, most states also post the number of employers that have (1) registered for their auto-IRA program and (2) added employee data to the system. In Appendix Table 1, we show that these alternate definitions for employer take-up find substantially higher auto-IRA program participation, and thus a much lower share of firms apparently failing to comply with the mandate. Some reasons why employers may appear not to be complying with the state policies could include lack of awareness, conscious noncompliance, time lags between administrative set up steps, or cases where all employees may have opted out. Nevertheless, the magnitude of this wedge remains a puzzle to us and policy experts.²⁶ In any case, our results suggest that state retirement plan mandates result in substantial “crowd-in” (employers establishing new ESRPs instead of utilizing state auto-IRAs) with no meaningful “crowd-out” (employers dropping ESRPs in favor of state auto-IRAs).

²⁶ We have had discussions with individuals who administer the state auto-IRA programs who also consider this gap in compliance to be an unresolved puzzle and an issue that receives substantial attention from state and provider staff. Another possible explanation is that employees may often fail “Know Your Customer” checks — a set of rules intended to reduce money laundering -- which are required to be passed prior to making contributions to the auto-IRA.

c. *Characteristics of firms induced to offer ESRPs*

i. *Approach*

In this section, we estimate the characteristics of “compliers”: firms induced to offer an ESRP as a result of state auto-IRA policies.²⁷ We study the characteristics of compliers for two reasons. First, as a purely descriptive matter, policymakers might be interested in understanding which types of firms are induced into offering ESRPs, and which types of workers gain ESRP coverage, as a result of auto-IRA policies. This type of analysis can contribute to an assessment of the distributional consequences of retirement plan mandates and similar policies. Second, we can compare the characteristics of compliers to characteristics of other firms to shed some light on what is driving the crowd-in of ESRPs that we observe.

We proceed in the spirit of Marbach and Hangartner (2020), modified to our difference-in-differences setting. In particular, let S denote the set of firms that start a plan in the treatment state in the post-period. Using the language of Angrist and Pischke (2009), this set comprises two groups: compliers (C) and always-takers (AT). Always-takers are those who would have started a plan regardless of whether a state implemented a retirement plan mandate; compliers are those who start an ESRP *only* in the state of the world where its state implements a retirement plan mandate. Of course, we cannot observe whether any given firm is a complier or an always-taker. Nevertheless, we can estimate the mean

²⁷ In this section, we use the term “complier” in a narrower sense than we do in other sections. Here, “complier” is used in the econometric sense to refer to firms that create an ESRP in response to the employer mandate of a state’s auto-IRA policy. It does not include firms that comply with the new state laws by sending employee contributions to state-facilitated IRAs.

characteristics for these groups. In particular, for any characteristic X , we can observe directly the expected value of X conditional on starting a plan in the post period in the treated group, $E(X|S)$. This expected value can be decomposed into two components.

$$E(X|S) = E(X|C, S)P(C|S) + E(X|AT, S)(1 - P(C|S)) \quad (2)$$

The first component, $E(X|C, S)P(C|S)$, is the mean value of X among compliers multiplied by the probability of being a complier conditional on starting a plan (i.e., the share of employers starting a plan who are doing so in response to the policy). The second component, $E(X|AT, S)(1 - P(C|S))$, is the expected value of X among always-takers multiplied by the probability of being an always-taker conditional on starting a plan (which is the complement of the probability of being a complier conditional on starting a plan). We can estimate $P(C|S)$ and $E(X|AT, S)$, which can then be used to back out $E(X|C, S)$.

We implement this procedure as follows. First, for any of the policy treatments described in Table 1, we define a “pre year”, a “base year”, and a “post year”. The pre year is five years prior to the policy year, the base year is two years prior, and the post year is one year after. We define the set S to be those firms (in the relevant firm size bucket) that do not offer a plan in the base year but do offer a plan in the post year. We estimate $E(X|S)$ directly for this group of firms.

Next, we estimate $P(C|S)$ using a simple 2-by-2 difference-in-differences approach. We estimate the share of firms starting a plan between event times -2 and 1 in the treatment group (denoted $P(S_{treat}^{post})$) and control group (denoted $P(S_{control}^{post})$). We also estimate the share of firms starting a plan between event times -5 and -2 in the treatment group (denoted $P(S_{treat}^{pre})$) and control group (denoted $P(S_{control}^{pre})$). The difference-in-differences estimate of

the impact of the policy on the probability of starting a plan is then $\left(P(S_{treat}^{post}) - P(S_{treat}^{pre})\right) - \left(P(S_{control}^{post}) - P(S_{control}^{pre})\right)$. We can estimate $P(C|S)$ as the ratio of this difference-in-differences estimate to the directly observed probability of starting a plan in the treatment group in the post-period.

Next, we estimate $E(X|AT, S)$. To do so, we need to identify a group of firms that can definitely be classified as always-takers – i.e., a group of untreated firms that are observed to start an ESRP. This group includes treatment group firms that began to offer an ESRP during the pre-period (denoted S_{treat}^{pre}), as well as control-group firms that began to offer an ESRP during either the pre- or post-period (denoted $S_{control}^{pre}$ and $S_{control}^{post}$ respectively). We then impute $E(X|AT, S)$ assuming parallel trends between control states and treatment states. That is, our estimate for $E(X|AT, S)$ is $E(X|AT, S_{treat}^{pre}) + \left(E(X|AT, S_{control}^{post}) - E(X|AT, S_{control}^{pre})\right)$, where these three terms can be estimated directly from the data using the corresponding sample means.

Finally, to give context to these estimates, we identify a set of “always-offering” (AO) firms and “never-offering” (NO) firms and. The former are those that offer an ESRP in both the base and the post year, while the latter are those firms in the treatment state that do not offer a plan in either the base year or the post year. We note that the NO firms include both firms that participate in the state auto-IRA program, as well as any firms that appear to ignore the mandate. We compute $E(X|NO)$ and $E(X|AO)$ directly using the sample means. We compute all of these objects, $(E(X|C, S), E(X|AT, S), E(X|AO), \text{ and } E(X|NO))$, separately by policy treatment and aggregate across treatments weighted by the number of compliers.

Unless otherwise specified, we measure X lagged by three years. That is, $E(X|S, C)$ reflects the means measured in the base year (-2). Likewise, $E(X|AT, S_{treat}^{pre})$ and $E(X|AT, S_{control}^{pre})$ reflect the means measured in the pre year (-5).

ii. Results

Table 5 presents means for firm-level variables. The first six rows show the sector breakdown. We find that, in general, the compliers have a similar industry mix as never-offerers. However, they are noticeably different from the always-offerers (i.e., firms that offer at both event times -2 and +1) and the always-takers (i.e., firms that would have started an ESRP in the absence of the policy). Relative to the always-offerers and always-takers, compliers are much more likely to be in the leisure and hospitality sectors (NAICS codes 71 and 72) and much less likely to be in the professional services sector (NAICS codes 54, 55, and 56).

We see even larger differences between the firm types when it comes to (lagged) offers of health insurance coverage; 36% of compliers offered health insurance at event time -2, compared to 73% of always-offering firms, 26% of never-offering firms, and 50% of always-takers. These differences suggest a correlation between the decision to offer an ESRP and the decision to offer employer-sponsored health insurance.²⁸ In the final row of this table, we study the *type* of ESRP offered by compliers. The types of ESRPs include SIMPLE IRAs and more standard plans such as 401(k)s.²⁹ A SIMPLE IRA is a special type of

²⁸ We observe health insurance offers beginning in 2015, so we drop policy treatments beginning in 2019 or earlier (when the pre-year would be before 2015) in this row.

²⁹ These outcomes are *not* lagged – i.e., they are measured at event time 1. Furthermore, SEP IRAs are much less common in this population.

ESRP available to small employers that avoids some administrative burden (especially nondiscrimination testing to determine whether the plan favors highly-compensated employees); unlike 401(k) plans, however, SIMPLE IRAs require employer contributions. We find that approximately 22% of compliers (compared to 17% of the always-offerers and 14% of the always-takers) choose to offer SIMPLE IRAs rather than a 401(k)-type plan. Finally, we consider the age and growth rates of affected firms. We find that complier firms are slightly younger than never-offering and (especially) always-offering firms, though somewhat older than always-taker firms. Compliers have a relatively similar growth rate, measured by employment or gross receipts, to never-offering and always-offering firms, though much lower than always-taker firms. That is, firms that start ESRPs in the absence of policy intervention tend to be quickly growing, while firms that start ESRPs due to the policy do not share this trait.

Table 6 presents characteristics of the employees and owners of complier firms; the first four columns report employee means, and the next four columns report owner means. For the latter, we restrict the sample to the firms whose owners can be identified. Among firms that start a plan in treatment states after the policy (that is, in the set S), this restriction causes us to drop about one quarter of firms. As in Table 5, we find that compliers appear to resemble never-offerers along many dimensions but are dissimilar from always-offerers and in some cases always-takers as well. Complier firms' employees are demographically (in terms of age, race, marital status, and sex) very similar to never-offering firms' employees. Workers at always-offerers are slightly older, less likely to be male, more likely to be married, and less likely to be Hispanic. At always-offering firms, 49% of employees have some taxable

interest or dividend income. In comparison, only 36-39% of employees have such income in the other three categories of firms. Complier firms' employees tend to be more highly compensated than employees of never-offering firms; however, they are less highly compensated than employees of those of always-offerer and always-taker firms. At complier firms, 36% of employees did not work for that firm in the prior year, a fraction that is similar to never-offerers; by contrast, only 24% of always-offerer employees were new to their firm this year.

We see similar patterns for owners. Complier firm owners and never-offering firm owners appear fairly similar, but complier owners are less likely to be male, less likely to have investment income, and more likely to be Hispanic relative to always-offering firm owners.

Our findings highlight that, with some notable exceptions, complier firms appear to resemble never-offering firms in terms of industry composition, employee characteristics, and owner characteristics. However, complier firms tend to differ from always-offering firms. For example, always-offering firms tend to be in industries associated with higher-skilled labor, to pay substantially more in wages, and to have more employees and owners who are married while compliers tend to have a higher share of low-wage workers, more female owners, and both owners and employees are more likely to be Hispanic.

V. Explaining crowd-in

Empirically, we observe substantial crowd-in and little (if any) crowd-out of ESRPs, which have higher costs to firms than state auto-IRAs. What factors can explain this outcome? Under the foundation laid by Summers (1989), firms decide whether to offer any

form of non-wage compensation by comparing its benefits and costs. Firm owners benefit from ESRPs both directly in their role as wage-earners – e.g., a firm owner may value his or her own ability to contribute to the firm’s ESRP³⁰ -- and indirectly through the benefits received by their employees. This indirect benefit can take the form of being able to pay lower wages (as in the frictionless model of Summers 1989), or more generally by improving worker retention, recruitment, and job satisfaction.

Meanwhile, the cost of offering an ESRP includes fees paid directly to plan administrators, investment fees, administrative burdens on employers, as well as any costs required to ensure plan compliance with nondiscrimination rules, which prohibit large differences in participation and contributions by low-wage and high-wage employees.³¹ While we are unaware of any representative data on the overall expenses charged by ESRP administrators to small businesses, recent research indicates that small employers can face one-time start-up fees that are as low as \$500, ongoing annual administration costs ranging from \$950 to \$1800, and annual per employee costs ranging from \$72 to \$96 (Chen 2024). These prices indicate that in the first year of establishing an ESRP, an employer with 10 employees would face nominal costs ranging from about \$2,410 to \$3,020 and an employer with 50 employees would face nominal costs ranging from \$5,900 to \$6,250.

³⁰ We estimate that 70% of complier firms are S corporations or C corporations, which are entity types where owners are often W-2 employees. These shares are similar for never-offerer, always-offerer, and always-taker firms as well.

³¹ The cost of complying with non-discrimination rules may include providing a safe harbor employer match (if wages cannot be adjusted downward to cover the cost of the match) and administrative expenses associated with annual nondiscrimination testing. Alternatively, a firm may opt to offer a SIMPLE IRA (which satisfies the ESRP mandate) which tends to have a lower administrative cost; however, SIMPLE IRAs require employer contributions, which may be costly to the employer if frictions prevent offsetting wage reductions.

In the remainder of this section, we use this framework to argue that there are two possible neoclassical explanations for the large crowd-in effects of state auto-IRA policies. The first explanation is that (1) firm owners perceive large administrative burdens associated with state auto-IRAs relative to ESRPs *and* (2) owners or employees place little or no value on auto-IRAs relative to ESRPs (perhaps because they replicate a longstanding savings option, or because of binding contribution limits on IRA contributions). The second explanation is that employees (or owners as employees) attach negative value to state auto-IRAs, perhaps because they are averse to automatic enrollment. (ESRPs satisfy the employer mandates of these state policies regardless of whether they feature automatic enrollment.) We do not find compelling support for either of these explanations. Thus, we are left to consider other behavioral factors that could explain our findings.

a. Explanation 1: State Auto-IRAs have high costs to employers and low benefits to workers

Employers compare the costs and benefits of offering an ESRP relative to the next best alternative. Prior to the policy, the next best alternative is to offer no retirement savings option; after the policy, the next best alternative is utilizing the state auto-IRA program.³² Consider a firm that does not offer an ESRP prior to the auto-IRA policy. Such a firm has decided that the net benefit of offering an ESRP (relative to no plan at all) is negative. After the policy goes into effect, this firm compares the net benefit of starting an ESRP to the net

³² Strictly speaking, firms also have the option post-policy to continue to offer no plan and instead pay a penalty. For the sake of conciseness, we abstract from this part of the choice set.

benefit of utilizing the state auto-IRA program. If employees derive a low or zero benefit from the auto-IRA program, then the administrative burden of facilitating auto-IRA contributions may induce the firm to start an ESRP, especially if it was previously close to the margin of indifference between offering an ESRP and not offering a retirement savings option. Moreover, the administrative cost of utilizing the state auto-IRA program would have to be sufficiently high to induce a large number of firms to start offering an ESRP. To determine whether this explanation is plausible, it is critical to understand the benefits and costs of the state auto-IRA program relative to not offering any retirement savings option at all.

Benefits: In a purely neoclassical model, rational workers would not place any value on the state program, since it merely replicates an already-existing part of their choice set (i.e., saving earnings in an IRA outside of work). However, we find this case unlikely. In particular, we can examine employee participation in IRAs outside of work directly within the complier framework. As we report in Table 8, only 9% of complier firm employees contribute to an IRA at event time -2. This is much less than the average worker participation rate within complier firms' ESRPs at event time 1 (33%), strongly suggesting that most complier firm employees do not view saving in IRAs outside of work to be a perfect substitute for saving in an ESRP through payroll deductions.

In addition, there is strong evidence that workers are not fully rational with regard to retirement savings, and that their choices are heavily influenced by behavioral factors such as cognitive constraints, salience / availability bias, loss aversion, and other phenomena (Beshears et al. 2018). Workers subject to savings inertia likely place some value on the ease of participating in a state auto-IRA program with payroll deductions (which operate similar

to a direct deposit) and automatic investment of contributions. Workers who suffer from present-biased preferences or a lack of financial knowledge – and who are “sophisticated” in their self-awareness of these problems (O’Donoghue and Rabin 1999) – may also value automatic enrollment as a way to commit to saving for retirement. In fact, it is possible that workers value state auto-IRAs similarly to ESRPs with automatic enrollment. In both cases, employees save in a tax-preferred manner through automatic payroll deductions. Furthermore, to the extent that workers have limited financial literacy (van Rooij et al. 2012), they may lack the knowledge to distinguish between ESRPs, state auto-IRAs, and other retirement savings options, which would make the incremental benefit of an ESRP very small relative to the state auto-IRA option. If workers value state auto-IRAs similarly to ESRPs, and if employers’ administrative burden of the auto-IRA program is smaller than the cost of the typical ESRP, then we would expect many firms to drop their ESRPs in favor of the state program.

In this model, therefore, the crowd-in -- and lack of crowd-out -- that we observe suggests that the valuation of an auto-IRA must be low relative to ESRPs. If workers are not fully rational (and therefore value workplace savings options even when they replicate existing choices), then a low auto-IRA value is most plausible for high-income or sophisticated workers. Compared to ESRPs, IRAs (including state auto-IRAs) feature lower contribution limits and do not allow employer matching contributions. Similarly, workers who expect to live a long life, experience high health costs during retirement, or have strong bequest motives may prefer to defer a larger share of wages than auto-IRAs can accommodate. Additionally, state auto-IRA programs make use of Roth IRAs which means

that workers do not have the choice between Roth and traditional tax treatment of their savings as they might under an ESRP. Furthermore, high-income workers (with income above \$228,000 for a married worker in 2023) are not eligible to contribute to a Roth IRA (and thus a state auto-IRA) at all, while there is no such (direct) limitation for individuals contributing to ESRPs. Finally, workers who already contribute to an IRA outside of work are likely to place a higher value on ESRPs than state auto-IRAs, since any outside IRA contributions reduce the effective contribution limit for the state auto-IRA.

However, we find that these income and contribution limitations tend not to be binding for employees: as we report in Table 7, only 2.5% of complier firm employees made contributions to their ESRP at event time +1 in excess of the Roth IRA annual contribution limit and 7.5% of complier firm employees had (lagged) income above the Roth IRA income limits. For owners, the income limitations are more binding relative to employees: 46% had lagged income above the Roth IRA income limits. These owners may place a small value on state auto-IRAs relative to ESRPs to the extent that they are able to capture a larger share of their direct benefit than the indirect benefit accruing to their workers. However, that is inconsistent with the fact that a relatively modest share of complier owners (36%) participated in the ESRP at event time +1 at all – possibly because many owners do not receive a W-2 wage from their firm, and thus are not eligible to contribute to its ESRP.³³ This low participation rate casts doubt on owners' direct benefit playing an outsize role in ESRP

³³ We estimate that, conditional on matching to an owner, we find an owner with a W-2 wage for 71% of complier firms.

adoption – which implies that it is unlikely that auto-IRA valuation is low due to binding income and contribution limits.

Costs: This explanation for crowd-in further requires that business owners perceive the burden of facilitating auto-IRA contributions to be high. While state auto-IRAs are advertised as “free” to employers, we expect that participating in a state’s program has a positive cost to the firm. The employer faces the administrative burden of registering for the program initially, automatically enrolling new employees, and facilitating the payroll deductions (although, unlike ESRPs, state auto-IRAs do not require the payment of any fees or compliance with nondiscrimination rules). Furthermore, while running an ESRP likely has a higher cost to employers than using the state auto-IRA program, it may be easier for firms to outsource management to a third-party ESRP provider -- allowing firm owners to flexibly trade hassle cost for financial cost in the case of ESRPs but not the state auto-IRA program.

While we cannot estimate the perceived burden of using the state program directly, we can observe several indirect pieces of evidence from the tax data. We report these results in Table 8. First, we find that compliers are less likely to use a paid preparer than never-offering firms for their payroll tax return.³⁴ Furthermore, among S corporations, compliers tend to file their income tax returns slightly later in the year relative to never-offering firms.

This evidence is somewhat inconclusive regarding this proposed mechanism. Complier firms appear more likely to prepare and file their own tax returns rather than outsourcing them to paid preparers. Furthermore, the fact that complier firms file their tax

³⁴ Virtually all firms in our data utilize a paid preparer for preparing their annual tax return (e.g., Form 1120S), so there is little variation of interest along that dimension.

returns slightly later suggests that owners may feel overwhelmed by paperwork. It is possible that some of these employers perceive the state auto-IRA option – which adds to the paperwork they need to handle – as administratively burdensome and therefore prefer to outsource employee retirement saving to a third party via an ESRP. However, this evidence is also consistent with an alternative story in which employers who feel comfortable preparing their own tax returns also feel more comfortable doing state auto-IRA paperwork themselves.

Next, we study take-up of the Credit for Small Employer Pension Costs under section 45E. This credit was generally made more generous throughout the sample period; from 2020-2022, the credit was equal to up to 50% of the cost of establishing and administering an ESRP over the first three years. In our data, the average amount of positive credit claimed was \$831 for these years. However, claiming this credit requires about 4.5 hours of recordkeeping and compliance activities as estimated by the IRS in the instructions to the relevant tax form. Studying the 2022 policy treatments, we find that only a tiny share (4%) of compliers claim this credit at any point from 2020-2022, despite the vast majority being eligible to do so.³⁵ (In contrast, around 9 percent of always takers claim the credit.) This suggests that in a context similar to (though not exactly the same as) facilitating worker participation in a state auto-IRA program, perceived or actual administration costs are substantial enough to cause employers to forego meaningful amounts of monetary gain.

Finally, we consider the possibility that some firm owners may have a distaste for auto-IRAs due to ideological feelings about government programs more generally. We proxy

³⁵ This analysis is restricted to firms that file an income tax return, where the credit claim would be found.

for such feelings with the Democratic vote share in the firm's county in presidential elections from 2008 to 2020. The idea behind this proxy is that conservative political attitudes (i.e. voting for the Republican party) may be correlated with a distaste for participating in the government-facilitated auto-IRA program. Yet, we find no economically meaningful differences in the political geography of complier firms relative to the other groups of firms, suggesting that ideology is not a major driver of perceived "costs" of state auto-IRAs.

b. *Explanation 2: Auto-IRAs have negative value to workers*

Until now, we have assumed that the value to workers of state auto-IRA program participation is nonnegative. However, some workers may in fact place negative value on auto-IRAs due to their automatic enrollment feature. In particular, it may be suboptimal for some workers to save for retirement – for example, due to liquidity constraints or high-interest rate debt. Alternatively, even if it is rational to save for retirement, some workers may undervalue saving due to hyperbolic discounting, lack of salience, or psychological frictions. Auto-IRAs lower the utility of these workers' present selves either by inducing them to save for retirement or by imposing administrative and cognitive costs when they choose to opt out. State retirement plan mandates generally do not require ESRPs to feature automatic enrollment. Thus, providing an ESRP may allow firms to adhere to the mandate while sidestepping the need to automatically enroll workers in a workplace retirement savings program. In this scenario, the policy may induce crowd in by employers wishing to avoid automatically enrolling workers in a state-facilitated auto-IRA. However, if this is the

mechanism behind crowd-in, we would expect the new ESRPs not to feature automatic enrollment.

To determine whether this is a plausible explanation for crowd-in, we need to know whether the new ESRPs established in response to the mandates feature automatic enrollment. Unfortunately, we cannot observe automatic enrollment directly in the data available to us.³⁶ Thus, we use two pieces of indirect evidence – which each suggest that complier firms’ plans are just as likely to feature automatic enrollment.

First, we examine the share of employees that contribute to the plan at event time 1. Approximately a third of employees in complier firms (32.7 percent) contributed to an ESRP, which is only 4 percentage points lower than workers at always-taker firms (36.4 percent). This evidence is inconsistent with the plans of complier firms being much less likely to include automatic enrollment. Furthermore, this evidence is also inconsistent with the idea that complier firms’ employees generally do not wish to save for retirement, which is what would motivate offering a plan without automatic enrollment in the first place.

Second, we make use of variation driven by the federal SECURE Act 2.0, which was enacted in 2022. This law requires employers to auto-enroll new hires in their ESRPs beginning in 2025; however, the requirement only applies to plans established in 2023 and later. It would be reasonable to expect a rational, forward-looking employer to consider the auto-enrollment preferences of future employees when choosing between starting an ESRP and enrolling workers in the state auto-IRA program. Therefore, if employers are motivated

³⁶ Larger plans (with 100 or more participants) tend to report automatic enrollment features in the attachments to Form 5500. By contrast, smaller plans are typically required only to file the “short form” version of Form 5500, which does not contain information on automatic enrollment.

to start ESRPs to avoid having their employees be automatically enrolled, we would expect the ESRP crowd-in effect to be smaller in 2023 (holding event time and other factors fixed) than in earlier years.

We test this possibility by comparing the event time zero effect in the Connecticut 5-24 employee and Colorado 5-49 policy treatments (which took effect in 2023) to the event time zero effect in the California 5-49 employee treatment (which took place in 2022). To strengthen comparability, we restrict each two-way test to firms in the same size range as the Connecticut and Colorado policy treatments.³⁷ We report these results in Table 9. For the “starts plan” outcome (columns (1) and (2)), we find that the effect at event time zero is slightly larger in Colorado (11.1 percentage points) than in California (9.6 percentage points in this sample) while the effect is smaller for Connecticut (8.1 percentage points) compared to California (9.9 percentage points in this sample). We find broadly similar patterns using the “offers plan” outcome as well. We conclude from these modest differences and inconsistent signs that the desire to avoid automatic enrollment does not appear to be a major driver of the adoption of ESRPs in response to state auto-IRA policies.

c. Explanation 3: Crowd-in is driven by other behavioral factors.

Non-neoclassical explanations may also drive the crowd-in that we observe. For example, inertia may play a role in firms’ decisions regarding workplace benefits. When a firm is established, its owners might find it optimal to not offer an ESRP, as the new firm may

³⁷ We also exclude firms with fewer than 10 employees from the sample because small businesses are exempt from the 2025 auto-enrollment requirement in Secure 2.0.

have very few (if any) non-owner employees. In addition, new firms are often relatively liquidity and credit constrained in ways that make the pecuniary costs of benefits provision particularly burdensome to the owners of such establishments. Inertia – whether in the form of explicit switching costs, procrastination, or other behavioral factors – may cause such firms to continue not to offer an ESRP even if, as the firm grows, the value of an ESRP begins to exceed its cost. When a state auto-IRA policy comes into effect, it removes the default option (“offer no plan”) from firm owners’ choice set, forcing an active decision and thus inducing such firms to offer an ESRP. That is, the auto-IRA policy does not just induce firms along or just below the margin of indifference to offer ESRPs; it also induces firms that are beyond the margin of indifference, but experiencing inertia, to offer ESRPs.

A related behavioral explanation is marketing. Third party ESRP administrators have responded to these state policies through targeted marketing, designed to convince small business owners to comply with the mandate by offering an ESRP rather than participating in the state auto-IRA program. It is possible that this marketing was particularly successful and effectively altered decision-makers’ perceptions of the costs and benefits of both ESRPs and state-facilitated auto-IRAs.³⁸

Low financial literacy on the part of business owners may contribute both to inertia and susceptibility to marketing. A voluminous literature has documented low, if uneven, financial literacy rates among consumers in the U.S. and around the world (for a review, see

³⁸ For example, a recent Morgan Stanley brief discussing the company’s ESRP services notes that “many states are mandating that employers offer some type of retirement savings plan and workers are looking for job opportunities that offer this type of benefit.” See <https://www.morganstanley.com/atwork/articles/small-business-retirement-plans-sep-simple-ira-401k>.

Lusardi and Mitchell 2011, 2023). Given that the majority of firms, especially small businesses, are controlled by individual owners, extrapolating this evidence in consumer finance to employer decision-making may be plausible.

Finally, we close with one more piece of evidence that could be consistent with either behavioral or neo-classical explanations. In particular, we study early enforcement actions in Illinois aimed at fostering employer compliance with the state's mandate. Specifically, in February of 2023, the Illinois Department of Revenue started issuing Notices of Proposed Assessment to employers who failed to offer an ESRP or join the state auto-IRA program despite being required to do so.³⁹ In addition to notifying employers of their failure to comply with the policy, these notices emphasized that penalties would be avoided if the employer chose to comply with the rule or claim exemption from it within 120 days of receiving the notice. Noncompliant employers from the Illinois 25-99 employee policy treatment would have been in noncompliance for more than three years at the time that these initial notices were received.

In Figure 3, we plot the event study for starting a plan in the Illinois 25-99 employee policy treatment, which took effect in 2019.⁴⁰ We see an increase of about 1 percentage point in starting plans in Illinois in 2023 relative to 2022, which is non-negligible relative to the 3

³⁹ We have information on enforcement activities in Illinois during the study period based on discussions with state administrators as well as publicly available enforcement letters sent to employers. In contrast, we do not have hard data on similar activities in other states. Our understanding, based on discussions with state administrators, is that California has pursued enforcements recently, primarily after our study period.

⁴⁰ One concern is that this 2023 effect may reflect the effect of the 2023 deadline in Illinois for firms with fewer than 15 employees due to mismeasurement of firm size. To mitigate this concern, we restrict to firms with at least 37.5 average quarterly employees (150% of 25) in this exhibit.

and 2 percentage point effects, respectively, in 2019 and 2020 (that is, event times 0 and 1). This suggests that these letters increased ESRP offers.

These letters may have increased ESRP offers through two mechanisms. First, they may have served as simple reminders to those who overlooked the requirement to comply with the state policy by joining the state auto-IRA program or offering an ESRP. Such reminders (or “nudges”) may have helped firm owners overcome behavioral frictions that led to inaction. Second, the letters may have altered employers’ perceptions about the penalties that they could pay – and the probability that those penalties would be imposed – if they failed to comply with the state mandate. That is, to extend the logic of our conceptual framework slightly, firms compare offering an ESRP to the best alternative, which may be either participating in the state auto-IRA program or doing nothing. There may have been some firms whose post-policy (but pre-letter) best alternative was “doing nothing,” as they perceived the expected penalty costs of non-compliance to be small or zero. Upon receiving the letter, the best alternative may have remained “doing nothing”, but at a higher cost – potentially leading ESRPs to become optimal. Alternatively, the letter may have changed the best alternative to “participating in the auto-IRA program” – in which case firms could be induced into offering an ESRP for all the reasons discussed in the prior subsections. We are not able to distinguish between these two hypotheses.

VI. Conclusions

The analysis presented in this paper provides insight into the effects of state auto-IRA policies on the decisions of employers to offer ESRPs. Our findings indicate a substantial

"crowd-in" effect, where a substantial proportion of firms that previously did not offer ESRPs began doing so in response to auto-IRA policies. The event study results demonstrate that firms treated with the policies were more likely to start offering ESRPs compared to similar firms in states without such policies. This effect is consistently observed across multiple states and firm size categories. However, the proportion of firms induced to offer ESRPs varies across states and firm sizes, ranging from approximately 8% to 23% of non-offering firms, depending on the policy treatment. Our findings regarding crowd-in, as well as the impact of enforcement letters, have important policy implications for the large number of states that are in the process of implementing auto-IRA policies. These findings also have implications for current federal policies aimed at increasing worker access to and participation in workplace retirement savings vehicles. More broadly, these results contribute to the evidence base regarding how employer mandates shape the structure of nonwage compensation. While our analysis is conducted at the firm level and focuses on firm behavior, we note that studying the impact on individual savers is also important and a useful area for future research.

Firms induced to offer ESRPs (compliers) tend to resemble never-offering firms rather than always-offering firms. Complier firms are more likely to be in the leisure and hospitality sectors and less likely to be in professional services. Additionally, these firms are generally smaller, offer lower wages, and are less likely to offer health insurance relative to the always-offering firms. Employees at complier firms are younger, more likely to be male, and less likely to have investment income compared to their counterparts at always-offering firms.

Similarly, owners of complier firms are less likely to be male, less likely to have investment income, and more likely to be Hispanic relative to owners of always-offering firms.

Theoretical considerations suggest several possible mechanisms behind the crowd-in effect. Workers at complier firms may find ESRPs more valuable than auto-IRAs due to higher contribution limits, the ability to incorporate employer matching contributions, and the absence of a state requirement for auto-enrollment. Additionally, the perceived administrative burden of state-facilitated auto-IRAs may drive complier firms to opt for ESRPs. However, we do not find strong evidence consistent with these factors driving crowd-in. These findings, in combination with the fact that compliers are quite similar to never-offerers, suggest that the neoclassical explanations that we explore do not explain the degree of crowd-in that we observe. It is therefore possible that other less visible and non-neoclassical factors may be influential in employer decisions to offer ESRPs. These behavioral factors may include inertia, owners' perceptions of ESRPs as complex, and targeted marketing by ESRP administrators that may have also played significant roles in influencing firm decisions. Future research using experimental methods could be used to test some of these non-neoclassical and behavioral factors driving firms' ESRP decisions.

References

- Abraham, Jean M., Anne B. Royalty, and Coleman Drake. 2019. The Impact of Medicaid Expansion on Employer Provision of Health Insurance. *International Journal of Health Economics and Management* 19: 317-340.
- Angrist, Joshua D. and Jorn-Steffen Pischke. *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press, 2009.
- Beshears, John, James J. Choi, David Laibson, Bridgette C. Madrian, and William L. Skimmyhorn. 2022. Borrowing to Save? The Impact of Automatic Enrollment on Debt. *The Journal of Finance* 77: 403-447.
- Beshears, John, James J. Choi, David Laibson, Bridgette C. Madrian. 2018. Behavioral Household Finance. *Handbook of Behavioral Economics: Application and Foundations* 1 1, 177-276, 2018.
- Beshears, John, Matthew Blakstad, James J. Choi, Christopher Firth, John Gathergood, David Laibson, Richard Notley, Jesal D. Sheth, Will Sandbrook, and Neil Stewart. 2024. Does Pension Automatic Enrollment Increase Debt? Evidence from a Large-Scale Natural Experiment. NBER Working Paper no. 32100. February.
- Bloomfield, Adam, Kyung Min Lee, Jay Philbrick and Sita Slavov. 2024. How Do Firms Respond to State Retirement Plan Mandates? *Economic Inquiry* 63(1): 265–288.
- Burkhauser, Richard V., Jeff Larrimore, and Kosali I. Simon, 2012. A "Second Opinion" on the Economic Health of the American Middle Class. *National Tax Journal* 65(1): 7-32.
- Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer. 2019. The Effect of Minimum Wages on Low-Wage Jobs. *The Quarterly Journal of Economics* 134(3): 1405–1454.
- Chalmers, John, Olivia S. Mitchell, Jonathan Reuter, and Mingli Zhong. 2022. Do State-Sponsored Retirement Plans Boost Retirement Saving? *AEA Papers and Proceedings* 112: 142-46.
- Chen, Yung-Ping. 1981. The Growth of Fringe Benefits: Implications for Social Security. *Monthly Labor Review*. November, retrieved from <https://www.bls.gov/opub/mlr/1981/11/art1full.pdf>.
- Chen, Anqi. 2024. Small Business Retirement Plans: How Firms Perceive Benefits and Costs. Issue in Brief 24-7. Center for Retirement Research at Boston College.

- Choi, James J., David Laibson, Brigitte C. Madrian, and Andrew Metrick. 2004. For Better or for Worse: Default Effects and 401(k) Savings Behavior. In *Perspectives on the Economics of Aging* (ed. David Wise). University of Chicago Press.
- Choukmane, Taha. 2024. Default Options and Retirement Saving Dynamics. Retrieved from [Choukhmane-2024-Default-Options-and-Retirement-Saving-Dynamics.pdf](#)
- Clemens, Jeffrey. 2021. How Do Firms Respond to Minimum Wage Increases? Understanding the Relevance of Non-Employment Margins. *The Journal of Economic Perspectives* 35(1): 51-72.
- Clemens, Jeffrey, Lisa B. Khan, and Jonathan Meer. 2018. The Minimum Wage, Fringe Benefits, and Worker Welfare. NBER Working Paper 24635.
- Correia, Sergio. 2017. “Linear Models with High-Dimensional Fixed Effects: An Efficient and Feasible Estimator” Working Paper. <http://scoreia.com/research/hdfe.pdf>
- Dao, N. 2024. Does a Requirement to Offer Retirement Plans Help Low-Income Workers Save for Retirement? Early Evidence from the OregonSaves Program. *Contemporary Economic Policy* 42(3): 524-543.
- Derby, Elena, Kathleen Mackie, and Jacob Mortenson. 2023. Worker and spousal responses to automatic enrollment. *Journal of Public Economics* 223.
- Freeman, Richard B. 1981. The Effect of Unionism on Fringe Benefits. *ILR Review* 34(4): 489-509.
- Garthwaite, Craig, Tal Gross, and Matthew J. Notowidigdo. 2014. Public Health. Insurance, Labor Supply, and Employment Lock. *The Quarterly Journal of Economics* 129(2): 653–96.
- Gertler, Paul, Sean Higgins, Ulrike Malmendier and Waldo Ojeda. 2025. Do Behavioral Frictions Prevent Firms from Adopting Profitable Opportunities? NBER Working Paper no. 33387. January.
- Gruber, Jonathan and Bridgette C. Madrian. 1994. Health Insurance and Job Mobility: The Effects of Public Policy on Job-Lock. *ILR Review* 48(1): 86-102.
- Guvenen, Fatih, Greg Kaplan, Jae Song, and Justin Weidner. 2022. Lifetime Earnings in the United States over Six Decades. *American Economic Journal: Applied Economics* 14(4): 446–79.
- Heidhues, Paul and Koszegi, Botond. Behavioral Industrial Organization. In *Handbook of Behavioral Economics*, Volume 1. Elsevier, 2018.

- Investment Company Institute. 2022a. Ten Important Facts About IRAs. July. Retrieved from <https://www.ici.org/system/files/2022-07/ten-facts-iras.pdf>.
- Investment Company Institute. 2022b. The Role of IRAs in US Households' Saving for Retirement, 2021. ICI Research Perspective. January. Retrieved from <https://www.ici.org/system/files/2022-01/per28-01.pdf>.
- Kim, Dongwoo. 2020. Worker retirement responses to pension incentives: Do they respond to pension wealth? *Journal of Economic Behavior & Organization* 173: 365-385.
- Kolstad Jonathan T., and Amanda E. Kowalski. 2016. Mandate-based health reform and the labor market: Evidence from the Massachusetts reform. *Journal of Health Economics* 47: 81-106.
- Liu, Tim, Christos A Makridis, Paige Ouimet, Elena Simintzi. 2023. The Distribution of Nonwage Benefits: Maternity Benefits and Gender Diversity. *The Review of Financial Studies* 36(1): 194–234.
- Long, James and Frank Scott. 1982. The Income Tax and Nonwage Compensation. *Review of Economics and Statistics* 64(2): 211-19.
- Long, James and Frank Scott. 1984. The Impact of the 1981 Tax Act on Fringe Benefits and Federal Tax Revenues. *National Tax Journal* 37(2): 151-258.
- Lusardi, Annamaria and Olivia S. Mitchell. 2011. Financial Literacy around the World: An Overview. *Journal of Pension Economics and Finance* 10(4): 497–508.
- Lusardi, Annamaria, and Olivia S. Mitchell. 2023. The Importance of Financial Literacy: Opening a New Field. *Journal of Economic Perspectives* 37(4): 137–54.
- Lyons S. 2017. Are employer mandates to offer health insurance effective in reducing subsidized coverage crowd-out of employer-sponsored insurance? *American Journal of Health Economics* 3(3):370–91.
- Lu, Benjamin, Jia Wan, Derek Ouyang, Jacob Goldin, and Daniel E. Ho. 2024. Quantifying the Uncertainty of Imputed Demographic Disparity Estimates: The Dual-Bootstrap. NBER Working Paper no. 32312. April.
- Madrian, Bridgitte and Dennis F. Shea. 2001. The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior. *The Quarterly Journal of Economics* 116(4): 1149-1187.

- Maestas, Nicole, Kathleen J. Mullen, David Powell, Till von Wachter, and Jeffrey B. Wenger. 2023. The Value of Working Conditions in the United States and the Implications for the Structure of Wages. *American Economic Review* 113(7): 2007-47.
- Malmendier, Ulrike. Behavioral Corporate Finance. In *Handbook of Behavioral Economics*, Volume 1. Elsevier, 2018.
- Manson, Steven, Jonathan Schroeder, David Van Riper, Katherine Knowles, Tracy Kugler, Finn Roberts, and Steven Ruggles. 2023. IPUMS National Historical Geographic Information System: Version 18.0 [dataset]. Minneapolis, MN: IPUMS.
- Marbach, Moritz and Dominik Hangartner, D. 2020. Profiling compliers and noncompliers for instrumental-variable analysis. *Political Analysis* 28(3): 435-444.
- Meiselbach, Mark and Jean Abraham. 2023. Do Minimum Wage Laws Affect Employer-Sponsored Insurance Provision? *Journal of Health Economics* 92: 1-17.
- Ni, Shawn and Michael Podgursky. 2016. How Teachers Respond to Pension System Incentives: New Estimates and Policy Applications. *Journal of Labor Economics* 34(4): 1075-1104.
- Nyce, Steven, Sylvester J. Schieber, John B. Shoven, Sita Nataraj Slavov, David A. Wise. 2013. Does Retiree Health Insurance Encourage Early Retirement? *Journal of Public Economics* 104: 40-51.
- O'Donoghue, Ted and Matthew Rabin. 1999. Doing it Now or Later. *American Economic Review* 89(1): 103-124.
- Oyer, Paul. 2008. Salary or Benefits? *Research in Labor Economics* 28: 429-467.
- Pardue, Luke. 2023. State Auto-IRA Mandates Boost 401(k) Adoption, With Largest Gains Among Lower-Income Workers. Gusto. [State Auto-IRA Mandates Boost 401\(k\) Adoption, With Largest Gains Among Lower-Income Workers | Gusto](#)
- Pierce, Brooks. 2001. Compensation Inequality. *The Quarterly Journal of Economics* 116(4): 1493-1525.
- Quinby, Laura D., Alicia H. Munnell, Wenliang Hou, Anek Belbase, and Geoffrey T. Sanzenbacher. 2020. Participation and Pre-Retirement Withdrawals in Oregon's Auto-IRA. *Journal of Retirement* 8(1): 8-21.
- Reuter, Jonathan. 2024. Plan Design and Participant Behavior in Defined Contribution Retirement Plans: Past, Present, and Future. NBER Working Paper no. 32653. July.

- Rhine, Sherrie S. L. 1987. The Determinants of Fringe Benefits: Additional Evidence. *Journal of Risk and Insurance* 54(4): 790-799.
- Sabelhaus, John and Alice Henriques Volz. 2019. Are Disappearing Employer Pensions Contributing to Rising Wealth Inequality? FEDS Notes. February. Retrieved from <https://www.federalreserve.gov/econres/notes/feds-notes/are-disappearing-employer-pensions-contributing-to-rising-wealth-inequality-20190201.html>.
- Sabelhaus, John. 2022. The Current State of U.S. Workplace Retirement Plan Coverage. Pension Research Council Working Paper WP2022-07. The Wharton School, University of Pennsylvania. Retrieved from <https://repository.upenn.edu/server/api/core/bitstreams/53f6523d-7c37-4052-add0-676774a6181a/content>
- Scott, John. 2021. Availability of State Auto-IRAs Appears to Complement Private Market for Retirement Plans. Pew Trusts. [Availability of State Auto-IRAs Appears to Complement Private Market for Retirement Plans | The Pew Charitable Trusts \(pewtrusts.org\)](https://www.pewtrusts.org/en/research-and-analysis/articles/2021/07/27/availability-of-state-auto-iras-appears-to-complement-private-market-for-retirement-plans)
- Summers, Lawrence H. 1989. Some Simple Economics of Mandated Benefits. *American Economic Review* 79(2): 177-183.
- Topolski, John J., Elizabeth A. Myers, and Sylvia L. Bryan. 2024. Worker Participation in Employer-Sponsored Pensions: Data in Brief and Recent Trends. Congressional Research Service Report. Retrieved from <https://crsreports.congress.gov/product/pdf/R/R43439>.
- Turner, Robert W. 1987a. Taxes and the Number of Fringe Benefits Received. *Journal of Public Economics* 33(1): 41-57.
- Turner, Robert W. 1987b. Are Taxes Responsible for the Growth in Fringe Benefits? *National Tax Journal* 40(2): 205-220.
- Tzioumis, Konstantinos, 2018, "Data for: Demographic aspects of first names", <https://doi.org/10.7910/DVN/TYJKEZ>, Harvard Dataverse, V1, Retrieved from <https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/TYJKEZ>.
- U.S. Bureau of Labor Statistics. 2024. Employer Costs for Employee Compensation Summary. June 18. Retrieved from <https://www.bls.gov/news.release/ecec.nr0.htm>.

- U.S. Census Bureau. 2021. Frequently Occurring Surnames from the 2010 Census, File A: Top 1000 Names. Retrieved from <https://www.nhgis.org/documentation/tabular-data>.
- U.S. Chamber of Commerce. 2017. State Auto-IRAs: The Wrong Answer. Retrieved from: <https://www.uschamber.com/assets/archived/images/us-chamber-state-auto-iras-white-paper.pdf>.
- U.S. Department of Labor. 2023. Private Pension Plan Bulletin Historical Tables and Graphs 1975-2021. September. Retrieved from: <https://www.dol.gov/sites/dolgov/files/ebsa/researchers/statistics/retirement-bulletins/private-pension-plan-bulletin-historical-tables-and-graphs.pdf>.
- van Rooij, Maarten C.J., Annamaria Lusardi and Rob J.M. Alessie. 2012. Financial Literacy, Retirement Planning, and Household Wealth. *The Economic Journal* 122: 449-478.
- Voicu, Ioan. 2018. Using First Name Information to Improve Race and Ethnicity Classification. *Statistics and Public Policy*. 5(1): 1-13.
- Wiatrowski, William J. 1999. Tracking Changes in Benefit Costs. *Compensation and Working Conditions* Spring: pp. 32-37. Retrieved from <https://www.bls.gov/opub/mlr/cwc/tracking-changes-in-benefit-costs.pdf>.
- Willingham, Thomas. 2024. Nudging Policies and Retirement Saving: Evidence and Theory from State Auto-IRA Programs. Working Paper. Retrieved from [AutoIRA_ShortTerm_TW_111224.pdf](#)

Table 1: State Auto-IRA policy treatments and employer mandate years

State	Employer Size	Year of Employer Mandate
Oregon	20-99	2018
Oregon	5-19	2019
Illinois	25-99	2019
California	50-99	2021
Illinois	16-24	2022
California	5-49	2022
Connecticut	26-99	2022

Notes: This table reports the seven policy treatments that we study in the main analysis. Employer size refers to the number of employees employed by the firm.

Table 2: Summary Statistics

Treatment and pre/post status	Unweighted number of distinct firms	Share of firms offering at least one ESRP	Average annual employee contribution to an ESRP	Average number of employees at firms	Average gross employee wages	Median gross employee wages
Control firms pre-policy (event time -2)	177,600	38.8%	\$887	17.4	\$30,628	\$23,085
Control firms post-policy (event time +1)	181,800	41.2%	\$1,055	17.9	\$35,448	\$27,273
Treated firms pre-policy (event time -2)	264,900	35.2%	\$963	16.9	\$34,652	\$26,282
Treated firms post-policy (event time +1)	271,800	49.1%	\$1,289	17.5	\$41,207	\$32,368

Notes: This table reports summary statistics for firms in our data. Median gross employee wages are computed at the firm level, and then averaged across firms. All statistics (except number of firms) are computed separately by policy treatment, then weighted across treatment according to the number of treated firms. Firm counts are rounded to the nearest 100. Recall that firms from control states are sampled at a rate of 10%. Source: Authors' calculations from confidential tax data.

Table 3: Relative effects across state policy treatments:

State	Employer Size	Year of Implementation	Firms induced to offer (1)	Firms not offering, pre-policy (2)	Share induced (3)
Oregon	20 - 99	2018	416 (66)	3,170	13.1% (2.1%)
Illinois	25 - 99	2019	887 (98)	6,918	12.8% (1.4%)
Oregon	5 - 19	2019	1,794 (119)	16,164	11.1% (0.7%)
California	50 - 99	2021	1,396 (140)	6,177	22.6% (2.3%)
California	5 - 49	2022	26,490 (476)	165,339	16.0% (0.3%)
Connecticut	26 - 99	2022	319 (47)	1,857	17.2% (2.5%)
Illinois	16 - 25	2022	483 (133)	5,897	8.2% (2.2%)

Notes: This table reports the estimated magnitude of the effect for each policy treatment. Specifically, in column (1), we compute the number of induced firms at event time k to be the event study “starts plan” coefficient for event time k multiplied by the number of firms in the treated state at time k in the relevant size range. Then, we sum across event times -1, 0, and 1 to arrive at the total number of firms induced to offer. In column (2), we compute the number of firms in the treated state in the relevant size range that do not offer a plan at event time -2. Column 3 reports the ratio of column (1) to column (2). Standard errors computed via the delta method are reported in parentheses. Source: Authors’ calculations from confidential tax data.

Table 4: Induced ESRP offers as a share of total induced employer offerings

State	Firms Induced to establish an ESRP	Firms facilitating payroll deductions to the state program	Induced ESRPs as a share of sum
	(1)	(2)	(3)
Oregon	2,211	5,234	30%
California	27,908	37,407	43%
Illinois	1,360	3,044	31%
Connecticut	319	392	45%

Notes: This table compares the estimated magnitude of the ESRP effect in each state to the number of firms that actively participate in the state auto-IRA program by December of the year before a mandate deadline that we do not study in this paper. The estimated magnitude of the ESRP effect corresponds to the amounts reported in Table 3, column (1), aggregated by state. The number of firms participating in the state auto-IRA program is retrieved publicly available data from state program websites and refers to the number of employers submitting payroll deductions in the last 90 days. Column (3) reports the ratio of column (1) to the sum of columns (1) and (2). Source: Authors' calculations from confidential tax data and publicly available data retrieved from state auto-IRA program websites.

Table 5: Means of firm-level variables for Compliers, Never-Offerers, Always-Offerers

Characteristic	Compliers	Never-offerers	Always-offerers	Always-takers
<i><u>Industries</u></i>				
Education/Health	0.146 (0.006)	0.129 (0.001)	0.200 (0.002)	0.203 (0.007)
Goods Producing	0.162 (0.006)	0.172 (0.001)	0.207 (0.002)	0.186 (0.007)
Leisure/Hospitality	0.253 (0.005)	0.278 (0.002)	0.044 (0.001)	0.125 (0.005)
Professional Services	0.102 (0.005)	0.083 (0.001)	0.229 (0.002)	0.173 (0.006)
Trade, Transportation, Utilities	0.207 (0.006)	0.204 (0.001)	0.151 (0.001)	0.157 (0.007)
All other industries	0.131 (0.005)	0.135 (0.001)	0.169 (0.001)	0.157 (0.007)
<i><u>Other benefits</u></i>				
Number of employees	16.561 (0.184)	15.009 (0.148)	18.365 (0.147)	16.880 (0.216)
Offers health insurance	0.358 (0.008)	0.270 (0.002)	0.732 (0.002)	0.507 (0.010)
Offers SIMPLE IRA (at event time 1)	0.217 (0.006)	0.000 (0.000)	0.163 (0.002)	0.138 (0.008)
<i><u>Other firm characteristics</u></i>				
Firm age	10.047 (0.101)	10.797 (0.024)	13.961 (0.027)	9.005 (0.121)
Employment growth (% change)	0.041 (0.013)	0.016 (0.003)	0.084 (0.003)	0.318 (0.016)
Gross receipts growth (% change)	0.071 (0.007)	0.034 (0.002)	0.069 (0.002)	0.230 (0.009)

Notes: This table reports estimated means of various outcomes for four groups of firms. Compliers are those firms that offer an ESRP only when their state implements an auto-IRA policy. Never-offerers are those who do not offer a plan even when their state implements an auto-IRA policy. Always-offerers are those who offer a plan both prior to the policy and after the policy. Always-takers are those who *start* a plan whether or not their state implements an auto-IRA policy. See text for how these objects are calculated. “Education/health” includes NAICS codes 61 and 62. “Goods-producing” includes NAICS codes 11, 21, 23, and 31-33. “Leisure/hospitality” includes NAICS codes 71 and 72. “Professional services” includes NAICS codes 54, 55, and 56. “Trade/transportation” includes NAICS codes 22, 42, 44-45, 48, and 49. All outcomes are measured as of two years prior to implementation, except “Offers SIMPLE IRA”, which is measured one year after implementation. Due to data limitations, the “offers health insurance” outcome uses policy treatments taking place in 2020 or later only. Firm age is the difference between the current year and the first year when they filed Form 941 (with the latter censored at 1999, our first year of data). Employment growth and gross receipts growth is measured as the three-year arc change from event time -5 through event time -2. Bootstrapped standard errors are in parentheses. Source: Authors’ calculations from confidential tax data.

Table 6: Characteristics of employees and owners for Compliers, Never-Offerers, Always-Offerers, and Always-Takers

Characteristic	<u>Employees</u>				<u>Owners</u>			
	Compliers	Never-offerers	Always-offerers	Always-takers	Compliers	Never-offerers	Always-offerers	Always-takers
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Age	38.267 (0.107)	39.013 (0.031)	41.773 (0.026)	37.657 (0.136)	52.660 (0.193)	54.039 (0.045)	55.586 (0.051)	50.807 (0.227)
Male	0.526 (0.004)	0.529 (0.001)	0.510 (0.001)	0.507 (0.005)	0.684 (0.006)	0.692 (0.001)	0.741 (0.002)	0.706 (0.007)
Married	0.344 (0.003)	0.362 (0.001)	0.470 (0.001)	0.373 (0.004)	0.742 (0.006)	0.750 (0.002)	0.812 (0.002)	0.761 (0.007)
Have Dependents	0.314 (0.002)	0.329 (0.001)	0.364 (0.001)	0.345 (0.003)	0.464 (0.007)	0.431 (0.002)	0.436 (0.002)	0.498 (0.009)
Black	0.047 (0.001)	0.047 (0.000)	0.045 (0.000)	0.051 (0.002)	0.043 (0.001)	0.045 (0.000)	0.037 (0.000)	0.044 (0.002)
Hispanic	0.365 (0.003)	0.364 (0.001)	0.268 (0.001)	0.301 (0.003)	0.198 (0.003)	0.191 (0.001)	0.125 (0.001)	0.154 (0.003)
Investment income	0.357 (0.003)	0.350 (0.001)	0.483 (0.001)	0.387 (0.003)	0.801 (0.006)	0.786 (0.002)	0.923 (0.001)	0.841 (0.008)
Log wages	9.800 (0.012)	9.720 (0.003)	10.678 (0.003)	10.136 (0.014)	N/A	N/A	N/A	N/A
New workers this year	0.366 (0.004)	0.355 (0.001)	0.241 (0.001)	0.393 (0.005)	N/A	N/A	N/A	N/A

Notes: This table reports estimated means of various outcomes for compliers, never-offerers, always-offerers, and always-takers. See text and notes to Table 5 for how these objects are calculated. Columns (1)-(4) report these means for the firm's employees, while columns (5)-(8) report these means for the firm's owners. Columns (5)-(8) are restricted to firms where we can identify the firm's natural person owners. Outcomes for race and ethnicity are imputed using a BIFSG algorithm; see text for further details. "Investment income" is a dummy (at the owner or employee level) for having any taxable dividend or interest income. "Log wages" refers to the mean log of employees' annual earnings at the given firm. "New this year" is a dummy (at the employee level) for not being employed by the firm in the prior year. Bootstrapped standard errors are in parentheses. Source: Authors' calculations from confidential tax data.

Table 7: Complier means for participation and contribution constraints

Characteristic	<u>Employees</u>				<u>Owners</u>			
	Compliers (1)	Never- offerers (2)	Always- offerers (3)	Always- takers (4)	Compliers (5)	Never- offerers (6)	Always- offerers (7)	Always- takers (8)
<u>IRA and ESRP participation rate</u>								
Contribute to IRA (-2)	0.089 (0.001)	0.078 (0.000)	0.098 (0.000)	0.107 (0.002)	0.207 (0.006)	0.190 (0.001)	0.113 (0.001)	0.236 (0.008)
Contribute to this ESRP (1)	0.327 (0.004)	0.000 (0.000)	0.451 (0.001)	0.364 (0.005)	0.370 (0.047)	0.001 (0.000)	0.688 (0.021)	0.571 (0.030)
<u>Contribution limits</u>								
Income above Roth IRA limits (-2)	0.074 (0.002)	0.070 (0.000)	0.176 (0.001)	0.119 (0.002)	0.458 (0.008)	0.445 (0.002)	0.717 (0.002)	0.608 (0.010)
Contributions to ESRP above IRA contribution limits (1)	0.024 (0.001)	0.000 (0.000)	0.156 (0.001)	0.071 (0.002)	0.118 (0.021)	0.000 (0.000)	0.554 (0.022)	0.345 (0.015)

Notes: This table reports estimated means of various outcomes for compliers, never-offerers, always-offerers, and always-takers. See text and notes to Table 5 for how these objects are calculated. Columns (1)-(4) report these means for the firm's employees, while columns (5)-(8) report these means for the firm's owners. Columns (5)-(8) are restricted to firms where we can identify the firm's natural person owners. Contributions to any IRA and having income above Roth IRA limits are measured in event time -2. "Contribute to this ESRP" and contributions above IRA contribution limits are measured at event time 1. Source: Authors' calculations from confidential tax data.

Table 8: Complier means for outcomes related to burden, hassle costs, and ideology

Characteristic	Compliers	Never-offerers	Always-offerers	Always-takers
<i>Tax filing behavior</i>				
Has Form 1120S preparer	0.130 (0.006)	0.242 (0.001)	0.115 (0.001)	0.154 (0.007)
Has Form 941 preparer	0.973 (0.004)	0.960 (0.001)	0.980 (0.001)	0.949 (0.005)
Average Form 1120S filing time (days)	187.726 (1.920)	179.030 (0.449)	180.620 (0.554)	166.182 (2.233)
<i>Take-up of section 45E credit</i>				
Receives section 45E credit	0.029 (0.004)	0.002 (0.000)	0.008 (0.000)	0.081 (0.005)
<i>Proxied political ideology</i>				
County Democratic vote share	0.635 (0.002)	0.636 (0.000)	0.640 (0.001)	0.645 (0.003)


Notes: This table reports estimated means of various outcomes for compliers, never-offerers, always-offerers, and always-takers. See text and notes to Table 5 for how these objects are calculated. The outcomes “Has Form 1120S preparer” and “Average Form 1120S filing time (days)” are restricted to S corporations; the latter is computed relative to January 1 of the year following the tax year in question. “Receives section 45E credit” outcome is restricted to 2022 policy treatments and represents a dummy for the firm claiming the credit at any point between 2020 and 2022. The county Democratic vote share is the share of the two-party vote for the Democratic presidential candidate, aggregated between the 2008, 2012, 2016, 2020 general elections. Source: Authors’ calculations from confidential tax data.

Table 9: Testing automatic enrollment: Comparing 2022 and 2023 policy treatments at event time zero

	Starts plan		Offers plan	
	(1)	(2)	(3)	(4)
2022 policy treatment	California	California	California	California
2023 policy treatment	Colorado	Connecticut	Colorado	Connecticut
Firm size	5-49	5-24	5-49	5-24
2022 effect	0.096 (0.001)	0.099 (0.002)	0.122 (0.002)	0.121 (0.003)
2023 effect	0.111 (0.003)	0.081 (0.005)	0.130 (0.004)	0.077 (0.007)


Notes: This table compares the event time 0 treatment estimates for the 2022 California policy treatment (or some subset thereof) to those from 2023 policy treatments in Colorado and Connecticut. In each case, the sample is restricted to firms within the same size range (5-49 when comparing to Colorado and 5-24 when comparing to Connecticut). Source: Authors' calculations from confidential tax data.

Figure 1: Staggered event study results



Notes: This figure reports the main stacked event study results for ESRP offer, estimated using Equation (1). The “stock” variable (“offers plan”) takes on a value of 1 if the firm offers an ESRP to any of its employees and 0 otherwise. The “flow” variable (“starts plan”) takes on a value of 1 if a firm did not offer an ESRP in the previous year and does offer an ESRP in the current year; it takes on a value of zero otherwise. In Panels A and C, we study all seven main policy treatments, with $k^{max} = 1$. In Panels B and D, we drop policy treatments implemented in 2022, allowing us to increase k^{max} to 2. Standard errors are clustered by firm. Source: Authors’ calculations from confidential tax data.

Figure 2: Case study for Illinois 25-99 employee policy treatment: the role of early enforcement actions



Notes: This figure reports the event study for the “starts plan” outcome for the Illinois 25-99 employee policy treatment. The sample is restricted to firms with at least 37.5 average quarterly employees in the prior year. Standard errors are clustered by firm. Source: Authors’ calculations from confidential tax data.


Appendix

Table A1: Induced ESRP offers as a share of total induced employer offerings using alternative employer participation metrics

State	Firms Induced to establish an ESRP (1)	Firms facilitating payroll deductions to the state program (2)	Induced ESRPs as a share of sum (3)	Firms that have added employee data to the state program (4)	Induced ESRPs as a share of sum (5)	Firms that have registered for the state program (6)	Induced ESRPs as a share of sum (7)
OR	2,211	5,234	30%	16,059	12%	17,671	11%
CA	27,908	37,407	43%	119,970	19%	131,250	18%
IL	1,360	3,044	31%	7,211	16%	9,312	13%
CT	319	392	45%	N/A	N/A	876	27%


Notes: This table compares the estimated magnitude of the ESRP effect in each state to the number of firms participating in the state auto-IRA program by December of the year before a mandate deadline that we do not study in this paper. The estimated magnitude of the ESRP effect corresponds to the amounts reported in Table 3, column (1), aggregated by state. The number of firms participating in the state auto-IRA program is retrieved publicly available data from state program websites. Several metrics for employer participation are considered when calculating the share of firms induced to offer an ESRP. For example, for the metric of firms facilitating payroll deductions in the last 90 days, Column (3) reports the ratio of column (1) to the sum of columns (1) and (2). Source: Authors' calculations from confidential tax data and publicly available data retrieved from state auto-IRA program websites.

Figure A1: Event studies for Oregon, 20-99 employee policy treatment




Notes: This figure plots the event studies from Equation (1), restricted to the Oregon 20-99 employee policy treatment, using all years of available data. See the notes to Figure 1 for further details. Source: Authors' calculations from confidential tax data.

Figure A2: Event studies for Oregon, 5-19 employee policy treatment




Notes: This figure plots the event studies from Equation (1), restricted to the Oregon 5-19 employee policy treatment, using all years of available data. See the notes to Figure 1 for further details. Source: Authors' calculations from confidential tax data.

Figure A3: Event studies for Illinois, 25-99 employee policy treatment




Notes: This figure plots the event studies from Equation (1), restricted to the Illinois 25-99 employee policy treatment, using all years of available data. See the notes to Figure 1 for further details. Source: Authors' calculations from confidential tax data.

Figure A4: Event studies for California, 50-99 employee policy treatment



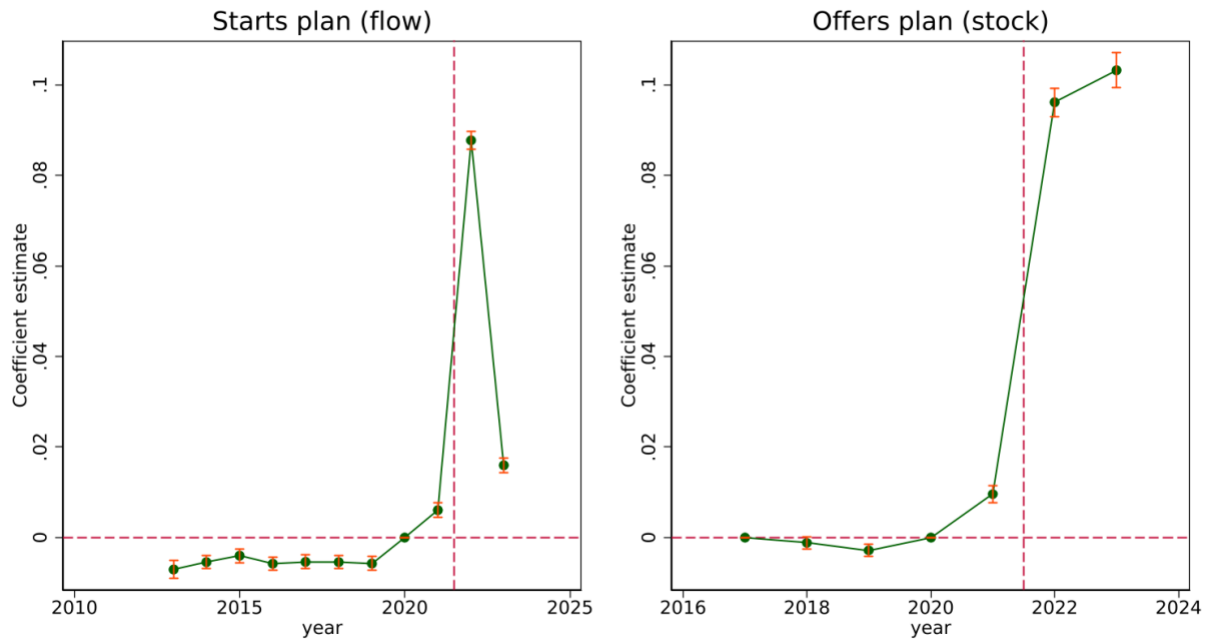
Notes: This figure plots the event studies from Equation (1), restricted to the California 50-99 employee policy treatment, using all years of available data. See the notes to Figure 1 for further details. Source: Authors' calculations from confidential tax data.

Figure A5: Event studies for Illinois, 16-24 employee policy treatment




Notes: This figure plots the event studies from Equation (1), restricted to the Illinois 16-24 employee policy treatment, using all years of available data. See the notes to Figure 1 for further details. Source: Authors' calculations from confidential tax data.

Figure A6: Event studies for California, 5-49 employee policy treatment




Notes: This figure plots the event studies from Equation (1), restricted to the California 5-49 employee policy treatment, using all years of available data. See the notes to Figure 1 for further details. Source: Authors' calculations from confidential tax data.

Figure A7: Event studies for Connecticut, 26-99 employee policy treatment




Notes: This figure plots the event studies from Equation (1), restricted to the Connecticut 26-99 employee policy treatment, using all years of available data. See the notes to Figure 1 for further details. Source: Authors' calculations from confidential tax data.

Figure A8: Stacked event study for “stops plan” outcome



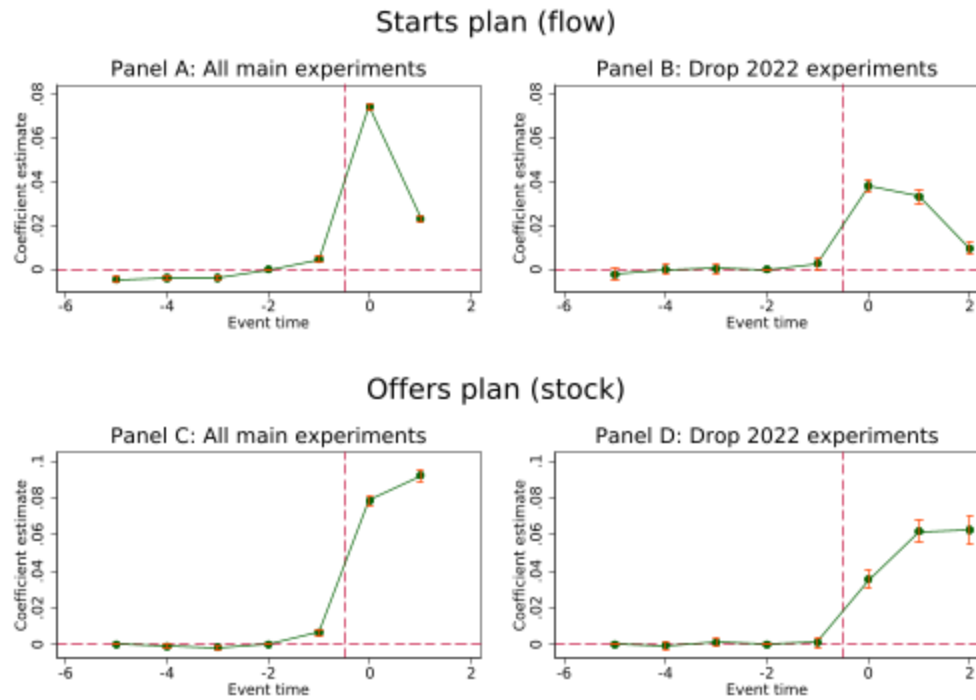
Notes: This figure plots stacked event studies from Equation (1), where the dependent variable is “stops plan”, which equals one in year t if and only if the firm offers a plan at time $t - 1$ and not at time t . We drop data involving 2023, and thus reduce k^{max} by one relative to Figure 1. Panel A uses all main policy treatments, while Panel B drops treatments implemented in 2022. See notes to Figure 1 for further details. Source: Authors’ calculations from confidential tax data.

Figure A9: Stacked event study (stock), using 10% threshold for defining ESRP offer




Notes: This figure plots stacked event studies from Equation (1), where the dependent variable is “offers plan” (stock), which equals one in year t if and only if at least 10% of current employees make a DC contribution in the current year. Panel A uses all main policy treatments, while Panel B drops treatments implemented in 2022. See notes to Figure 1 for further details. Source: Authors’ calculations from confidential tax data.

Figure A10: Staggered event study results, dropping possible multi-state firms



Notes: This figure reports the main stacked event study results for ESRP offer, estimated using Equation (1), but dropping firms that are proxied to be multi-state. We define a firm to be multi-state if at least 20 percent of its employees live in a state that differs from the state of the headquarters of the firm. See notes to Figure 1. Source: Authors' calculations from confidential tax data.

Figure A11: Staggered event study results, keeping only firms that file annual tax return with same EIN



Notes: This figure reports the main stacked event study results for ESRP offer, estimated using Equation (1) for the “stock” outcome, but restricting to firms who file (in event time -2) an annual tax return with the same EIN as their Form 941. We include Forms 1120S, 1065, 1120, 990, and Form 1040, Schedule C as “annual tax returns.” See notes to Figure 1. Source: Authors’ calculations from confidential tax data.