

NBER WORKING PAPER SERIES

WHY DO EMPLOYERS ESTABLISH RETIREMENT SAVINGS PLANS? EVIDENCE
FROM STATE “AUTO-IRA” POLICIES

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

We are grateful to Kim Olson, Christine Cheng, and Angela Antonelli for helpful input regarding policy details. 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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.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.

Why Do Employers Establish Retirement Savings Plans? Evidence from State “Auto-IRA” Policies

Adam Bloomfield, Lucas Goodman, Manita Rao, and Sita Slavov

NBER Working Paper No. 32817

August 2024

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 retirement plan mandates on firm decisions to offer ESRPs, treating the gradual rollout of these policies across states and employer size categories as a series of “experiments.” Using U.S. tax microdata, we estimate that at least 30,000 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 approach to complying with 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
manitara@usc.edu

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 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. Those who do not have access to an ESRP are

¹ See Investment Company Institute (2022a).

disproportionately likely to be racial and ethnic minorities, to have less education and lower incomes, and to be employees of smaller firms (Sabelhaus 2023). 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 auto-IRA *programs*, and to a combination of an employer mandate and 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. Across all seven policy implementations (“experiments”) 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” (firms terminating existing ESRPs in favor of utilizing the state auto-IRA program). We estimate that approximately 30,000 firms were induced to offer an ESRP by the policies, which corresponds to approximately one sixth of all firms in policy states in the affected firm size ranges. These effects are also substantial relative to the number of firms participating in the auto-IRA programs directly; we find that ESRP crowd-in accounts for between 27% and 45% 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 resemble firms that never offer an ESRP along several dimensions, including workforce and owner race, gender, wages, and sector. By contrast, similar-sized firms that already offer ESRPs tend to have much higher wages, are much more likely to offer employer-sponsored health insurance, and their workforce is less likely to be Hispanic.

We analyze our findings in the context of a neoclassical framework where rational employers weigh the costs and benefits of ESRPs relative to an outside option. The auto-IRA policy changes the outside option from not offering a plan to participating in the auto-IRA program, and thus changes the relative costs and benefits. Workers may derive a larger benefit from ESRPs than auto-IRAs, as the former have looser contribution limits and generally have more flexibility. Owners may receive this benefit both directly in their role as an owner-employee (i.e., by participating in the ESRP themselves) and indirectly by capturing some of the increased worker surplus. At the same time, ESRPs are generally costlier to administer than auto-IRA programs. Auto-IRA programs are free for employers, though the employer faces some administrative burden in signing up for the auto-IRA program and facilitating contributions by employees. In contrast, ESRPs, which are typically administered by third parties for a fee, can have substantial costs to employers. The large crowd-in that we observe empirically requires some combination of high perceived cost and low perceived benefits of auto-IRA programs relative to ESRPs.

However, we do not find substantial evidence that the crowd-in is driven by perceived low benefits of auto-IRAs relative to ESRPs. For example, we estimate that the auto-IRA contribution limits typically did not bind, either for employees or owners. It is possible that some employers may find auto-IRA participation burdensome, which could explain crowd-in. 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 requires 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 might cause an 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 the direct effect – i.e., the boost in IRA participation brought about by the auto-IRA programs (Dao 2024; Quinby et al. 2020; Chalmers et al. 2022). Closer to our paper, previous work by Bloomfield et al. (2023) used individual-level data from the Current Population Survey and firm-level filings reporting the status of their ESRPs (Form 5500) and found that auto-IRA policies induced ESRP offers.² We contribute to this literature by using rich administrative microdata updated through 2023. These data offer us a window into more recent expansions of auto-IRA policies, while also allowing for improved treatment assignment (i.e., using firm-size)

² Additionally, several commentators have used publicly available or proprietary data to observe increases in ESRP offers in treatment states (Scott 2021, Pardue 2023).

and greatly increased statistical precision. Additionally, these data permit a more granular exploration of the responses by firms and workers across various dimensions such as employer size, industry, worker, and owner demographics. This analysis both (1) provides a clearer picture of which firms and employees are most affected by the policy (2) allows us to explore the behavioral channels that may be driving employer responses to the policy.

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, although the impact on long-term saving is less clear (Madrian and Shea 2001; Choi et al. 2004; Choukmane 2021; Beshears, Choi, and Laibson 2021). However, this literature primarily analyzes worker decisions to participate in ESRPs rather than employer decisions to establish them. We contribute to this literature by focusing on firms’ decisions to establish ESRPs and show that frictions may affect these choices as well.

Second, we contribute to the empirical literature focused on the reasons 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 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 (for reviews, see Heidhues and Kozzegi 2018; Malmendier 2018). Our results suggest that it is unlikely that neoclassical factors can fully explain the large crowd-in that we find, leaving behavioral factors such as inertia, salience, and marketing as plausible candidates. 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 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, employers, and advocates in the pursuit of enhancing retirement security for all working Americans.

II. Policy Background

Employer-sponsored retirement plans (ESRPs) are a non-wage benefit provided by firms as part of their workers' compensation package. 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.³ 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.⁴ Beyond ESRPs, Individual Retirement Accounts (IRAs) provide most workers with access to tax-advantaged retirement saving independent of employment. Contributions to IRAs are voluntary and do not allow an employer match. Though assets held 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).

³ 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. DB pension payments to retirees are taxed as income. 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).

⁴ 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.

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 an auto-IRA program, 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 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 retirement plan mandate** (“employer mandate”), under which every firm must offer a workplace retirement savings option either by participating in the auto-IRA program or by offering its own ESRP to their employees. 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 policies as an **auto-IRA policy**.⁵

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 “experiment”. 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 experiments we consider in this paper are summarized in Table 1. These

⁵ 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.

experiments encompass all 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 experiment.

As the rollout of auto-IRA policies is ongoing, some states expanded or implemented auto-IRA policies in 2023. These experiments include 2023 expansions in Illinois (to firms with 5-15 employees) and Connecticut (to firms with 5-25 employees). In addition, Colorado implemented a new auto-IRA policy applying to firms with 5 or more employees in 2023, while Virginia and Maine began implementing auto-IRA programs in 2023 as well. These 2023 experiments are not included in our main analysis, and we exclude Colorado, Virginia, and Maine firms from the control groups, as those states began implementing auto-IRA programs in 2023. Finally, we exclude Maryland firms from the control group; Maryland implemented an auto-IRA policy in 2022, but the employer mandate is not enforced by imposing financial penalties for noncompliance. Instead, it is incentivized by giving employers a credit for a state plan reporting fee.⁶

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 from Form 941, a quarterly form which employers use to

⁶ The financial incentives associated with Maryland's program are also smaller than those in other states. Maryland's program offers participating firms a "credit" of \$300. In contrast, in California, covered employers who fail to comply within 90 days of receiving a noncompliance notice pay a penalty of \$250 per employee. There is an additional penalty of \$500 per employee if the employer continues to be noncompliant after 180 days of receiving the notice.

report (among other things) payroll tax obligations and individual income taxes withheld on 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.⁷ 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);⁸ for computational tractability, we take a 10% random sample of firms from other states. We are able to 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. We code an employer as offering an ESRP if and only if at least one employee makes a contribution.⁹ Second, we use Form W-2 to identify all employees of a given firm in a given year (and their wages). We then link 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)

⁷ 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.

⁸ For auxiliary analyses, we also retrieve a 100% sample of firms in Colorado. Such firms are not included in the main analysis.

⁹ Additionally, Box 12 allows us to differentiate between SIMPLE IRAs and other types of DC plans.

records to identify date of birth and gender. 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.¹⁰

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

a. Overall Crowd-In and Crowd-Out

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

¹⁰ 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).

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.¹¹ We then stack all seven panels and run the following regression for firm i (located in state $s(i)$) in experiment e at event time k :¹²

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

This regression traces out the mean value of y (relative to event time -2) after subtracting out the change in y in same-sized firms in never-treated states. We cluster our standard errors by firm. While the sample of firms is unbalanced (i.e., firms enter and exit), the sample is balanced at the state-year level – that is, we keep observations from event times -5 through k^{max} , and we restrict the sample to experiments where that entire time period is observed.

We consider several dependent variables, including one “stock” and two “flows.” 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” variables are dummies for “starting” or “stopping” a plan. The 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. That is, “starts plan” measures the unconditional probability of starting to offer an ESRP among firms that existed during the previous year. The variable “stops plan” takes on a value of 1 if a firm makes the opposite transition, switching from offering an ESRP in the previous year to not

¹¹ 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.

¹² We estimate this regression using the user-written Stata command `reghdfe` (Correia 2017).

offering an ESRP in the current year; it takes on a value of 0 otherwise. That is, “stops plan” measures the unconditional probability of terminating an ESRP among all existing firms.

Our preferred measure is the “starts plan” outcome, as it reduces the role that composition changes can play in driving the estimate. For example, consider the California 50-99 employee experiment, which was implemented in 2021. In 2019 (i.e., event time -2, the omitted year), firm size is measured in 2018. In 2021 (i.e., event time zero), firm size is measured in 2020. Because California experienced a worse shock due to COVID-19 than the nation as a whole,¹³ firms likely shrank more in California than in control states, meaning that smaller or more marginal firms – which might be less likely to offer ESRPs – would have dropped out of the sample more in California than in control states. This compositional effect could drive the “any offer” effect upward in this experiment. By contrast, prior to the policy changes, there is virtually no relationship between firm size and starting a plan (if anything, the slope is slightly negative), largely because starting a plan is a relatively rare outcome in the baseline. As a result, the bias from compositional changes is very small for the “starts plan” outcome. Nevertheless, we show both outcomes for completeness.

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

¹³ For example, in August 2020, California’s unemployment rate was 3 percentage points higher than the national average.

precisely estimated) increases in the rates of offering or starting an ESRP 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 experiments) are somewhat smaller than in Panels A and C. We note that the 2022 experiments – especially the California 5-49 employee experiment – represent a clear majority of the affected firms. Dropping such experiments 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 experiments separately; in each case, the ESRP offer rate increases noticeably at event time zero.¹⁴ 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.¹⁵ These results are broadly consistent with a substantial “crowd-in” effect – with a small, if any, crowd-out effect – of auto-IRA policies on firm ESRP offerings.

Table 2 considers the magnitude of this effect in more detail. For each of the seven experiments, 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” outcome, estimated solely for that experiment, and (b) the number of firms in the relevant size range in the treated

¹⁴ We also note that, as hypothesized above, we find a modestly larger effect for “offers plan” than “starts plan” for the California 50-99 employee experiment.

¹⁵ 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.

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 experiment 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 experiments, 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 3, we compare the magnitude of firms induced to offer ESRPs in Oregon, California, and Illinois (column 1) to the number of firms that are actively participating in the state-facilitated auto-IRA programs at the end of 2023 (column 2).¹⁶ 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 auto-IRA 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. In Oregon and Illinois, ESRPs account for about 27% of the total effect, while ESRPs account for 45% of the effect in California. Furthermore, these ratios in Oregon and Illinois may be underestimates. In these states, the denominator likely includes firms responding to additional waves of mandates that took place in 2023 for smaller firms, while the numerator does not include any induced effects for these 2023 experiments. In sum, we infer that an analysis of state auto-IRA policies that studies only the direct effects on

¹⁶ We retrieve the latter from publicly available data compiled by the Center for Retirement Initiatives at Georgetown University. These data are available at the state level (rather than at the experiment level) and are not available for Connecticut. For disclosure reasons, we are unable to use the tax data to estimate the number of firms participating in auto-IRA programs.

auto-IRA participation would miss a large share of the overall retirement coverage increases induced by the policy.

b. 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 the auto-IRA policies.¹⁷ 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.

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 characteristics for these groups. In particular, for any

¹⁷ 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 an auto-IRA mandate. It does not include firms that comply with the new state laws by sending employee contributions to state-facilitated IRAs.

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.

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

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 experiments 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 $(P(S_{treat}^{post}) - P(S_{treat}^{pre})) -$

$(P(S_{control}^{post}) - P(S_{control}^{pre}))$. 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 $E(X|AT, S_{treat}^{pre}) + (E(X|AT, S_{control}^{post}) - E(X|AT, S_{control}^{pre}))$, 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 “never-offering” (*NO*) firms and “always-offering” (*AO*) firms. The former are those firms in the treatment state that do not offer a plan in either the base year or the post year, while the latter are those that do offer a plan in both the base and the post year. We compute $E(X|NO)$ and $E(X|AO)$ directly using the sample means. We compute all of these objects, ($E(X|C, S)$, $E(X|AT, S)$, $E(X|AO)$, and $E(X|NO)$), separately by experiment and aggregate across experiments 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 4 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.¹⁸ 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.¹⁹ A SIMPLE IRA is a special type of ESRP available to small employers that avoids some administrative burden (especially nondiscrimination testing to

¹⁸ We observe health insurance offers beginning in 2015, so we drop experiments beginning in 2019 or earlier (when the pre-year would be before 2015) in this row.

¹⁹ These outcomes are *not* lagged – i.e., they are measured at event time 1. Furthermore, SEP IRAs are much less common in this population.

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.

Table 5 presents characteristics of the employees and owners of complier firms; the first four columns refer to employee means, and the next four report mean characteristics of owners. 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 4, 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-offerer firms' employees. 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 offerer 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-offerer 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-offerer firm owners.

Our findings highlight that, in general, complier firms appear to resemble never-offering firms, both in terms of industry composition, employee characteristics, and owner characteristics, while they 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**

a. *Theory*

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). The crowd-in that we observe implies that some firms prefer offering no plan to offering an ESRP, but also prefer offering an ESRP over an auto-IRA. In this subsection, we consider how these two preference orderings could coexist simultaneously in a model where firms are fully rational.

Utility maximization by firm owners implies that they decide to offer an ESRP when the value of doing so exceeds the cost. The value of offering an ESRP comprises two parts. First, employees value ESRPs – perhaps because they provide access to tax-advantaged saving, or

because they make retirement saving administratively easier. This value that employees receive flows to employers indirectly; under the foundation laid by Summers (1989), which assumes a frictionless labor market, workplace benefits such as ESRPs allow employers to reduce wages paid, thus increasing firm profits directly. In the presence of frictions (such as downward nominal wage rigidity), ESRPs could increase firm profits through other mechanisms instead, such as by improving worker retention, recruitment, and job satisfaction. Second, firm owners may benefit from ESRPs 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 ESRP.²⁰

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.²¹ 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

²⁰ 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.

²¹ 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.

about \$2410 to \$3020 and an employer with 50 employees would face nominal costs ranging from \$5900 to \$6250.

We illustrate employers' decisions using a two-dimensional plane, depicted in Figure 2. The x axis reflects the cost to the firm of providing an ESRP (relative to a non-ESRP outside option), while the y axis reflects the combined benefit that the firm owner and workers receive from ESRP availability (relative to the same outside option). The outside option is either offering no retirement savings vehicle (the baseline case) or participating in the auto-IRA program (which applies to firms in treated states after the auto-IRA policy is implemented). Points on the plane formed by (x, y) pairs that lie above the 45-degree line, along which the benefit to workers of an ESRP equals the cost to the firm, represent counterfactuals in which the firm will offer an ESRP. Conversely, cost-benefit pairs that correspond to points that lie below the 45-degree line, represent counterfactuals in which the firm will not offer an ESRP.

i. Pre-policy

Prior to the auto-IRA policy, the relevant outside option is not offering a retirement savings option of any sort. The neoclassical model predicts that firms that lie above the diagonal will choose to offer ESRPs. For these firms, offering an ESRP has a cost that is lower than the benefit to workers (and the owner). By analogous logic, firms that lie below the diagonal do not offer ESRPs. We consider a firm that does not offer an ESRP pre-policy. Such a firm is situated below the diagonal, as indicated by the point marked "Pre-Policy" in Figure 2.

ii. Post-policy

The auto-IRA policy removes the outside option of “not offering a plan” and changes it to “offering an auto-IRA.”²² Although auto-IRAs are advertised as “free” to employers, we expect that an auto-IRA program has a positive, though likely small, 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. (Unlike ESRPs, however, auto-IRAs do not require the payment of any fees or compliance with nondiscrimination rules.) Therefore, the employer will move left (possibly only slightly) in Figure 2 after the policy.

Additionally, the auto-IRA program may have value (relative to an ESRP) of either sign to employees; if positive, the firm moves down the vertical line in Figure 2, while if negative, the firm moves up the vertical line. There are several cases to consider. First, in a purely neoclassical model, rational workers would not place any value on the auto-IRA program, since it merely replicates an already-existing part of their choice set (i.e., saving earnings in an IRA outside of work). This case, in which there is no vertical change in the (x, y) -plane for the firm, corresponds to the point marked “Post-Policy (rational)” in Figure 2. In this depiction, the leftward shift is sufficient to induce the firm to offer an ESRP post-policy. But more generally, whether a firm is induced to offer an ESRP would depend on the magnitude of that leftward shift relative to how far the firm was initially from the margin of indifference.

The fully rational model imposes constraints on the auto-IRA policy’s effect size. If firms face only a small administrative cost from the auto-IRA program, then we would observe a large increase in the probability of offering an ESRP only if many firms were somehow located along

²² 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.

or just below the 45-degree line (relative to much further below). Given that many factors likely influence the value of ESRPs for each worker, and that worker characteristics vary substantially across firms, there is no reason to expect a large share of firms to be located along or just below the margin of indifference pre-policy implementation.

Second, and more realistically, 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 an auto-IRA 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 plausible that workers value 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, auto-IRAs, and other retirement savings, which would make the incremental benefit of an ESRP very small relative to the auto-IRA option. This case is depicted as a movement down to the x axis of Figure 2, as indicated by the point marked “Post-Policy (ESRP value \approx auto-IRA value)”. In this scenario, even firms that were initially above the 45-degree line (and therefore offered ESRPs) prior to the policy would move down to the x axis of Figure 2.

On the other hand, some high-income or sophisticated workers may place substantially more value on an ESRP than an auto-IRA, meaning that the downward movement of a firm in Figure 2 would be muted. Compared to ESRPs, IRAs (including 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, auto-IRA plans make use of Roth IRAs which means that workers do not have the choice between Roth and traditional 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 an 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 auto-IRAs.

Third, the value of an auto-IRA to some workers may be negative. 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 causes the firm to move upward because the value of the

outside option is now negative rather than zero, as illustrated by the point “Post-Policy (negative auto-IRA value)” in Figure 2. The firm depicted in the figure is induced to offer an ESRP. But more generally, the choice to offer an ESRP depends on the magnitude of the leftward and upward shifts relative to the initial position. Additionally, in this scenario, if a firm establishes an ESRP to avoid automatically enrolling workers in an auto-IRA, then we would expect the new ESRP not to feature automatic enrollment.

More generally, depending on the value that workers place on ESRPs relative to auto-IRAs, the firm depicted in Figure 2 may end up at any point along the vertical dotted line. The model predicts crowd-in only if the firm ends up on the portion of the dotted line that lies above the 45-degree line.

Finally, we note that crowd-out (i.e., firms stopping offering ESRPs) is a theoretical possibility among firms that start above the 45-degree line. This would occur if the downward move is large (i.e., because its employees value an auto-IRA to nearly the same extent as an ESRP) and its leftward move is small (i.e., because it is nearly costless to offer an auto-IRA). Empirically, we do not observe a meaningful amount of crowd-out, so such shifts are unlikely to be frequent in practice.

b. Empirical evidence

Empirically, we observe substantial crowd-in and little (if any) crowd-out of ESRPs, which have higher costs to workers and firms than auto-IRAs. Under the neoclassical framework above, this somewhat counterintuitive observation could be caused by a combination of two factors. First, owners or employees may value ESRPs to a much higher extent than auto-IRAs –

perhaps because of binding contribution limits in IRAs, because of pre-existing IRA participation, or because ESRPs do not require automatic enrollment – causing an upward, or at most small downward, shift in Figure 2. Second, employers may perceive that participating in the auto-IRA program has a high administrative cost relative to the cost of running an ESRP (i.e., a large leftward shift in Figure 2). In this section, we consider the evidence in favor of or against these hypotheses. We conclude with evidence showing that certain behavioral factors outside this framework may be important as well.

i. Possibility 1: Auto-IRAs provide negative value to workers that prefer not to save.

It is possible that auto-IRAs, which are designed with positive recurring employee savings as the default, provide negative value to a firm’s employees who prefer not to save for retirement. These workers may incur cognitive costs in opting out of the default or disutility from failing to do so. In this case, the firm may respond to these state policies by establishing an ESRP that does not 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.²³ 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 (33.5 percent) contributed to an ESRP,

²³ 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.

which is only 4 percentage points lower than workers at always-taker firms (37.5 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 auto-IRA program. Therefore, if employers are motivated 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 4-49 experiments (which took effect in 2023) to the event time zero effect in the California 5-49 employee experiment (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 experiments.²⁴ We report these results in Table 6. 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

²⁴ 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.

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 auto-IRA policies.

ii. Possibility 2: Auto-IRA participation offers little value to employees because of pre-existing IRA participation.

Auto-IRAs should have zero value to fully rational workers and may have zero value to less than rational workers if such employees are already saving in IRAs outside of work. We can examine employee participation in IRAs outside of work directly within the complier framework. As we report in Table 7, 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 complier employees do not view saving in IRAs outside of work to be a perfect substitute for saving in an ESRP through payroll deductions. Thus, unless owners are fully unable to capture benefits to employees, we find this possibility implausible.

iii. Possibility 3: Auto-IRAs have little value to employees because of IRA contribution and income limits.

Employees or owners may place substantially less value on an auto-IRA relative to an ESRP if they find the Roth IRA limitations to be highly binding, which would imply only a small

downward shift in Figure 2; this, in combination with a large perceived burden of auto-IRA participation (i.e., a large leftward shift) could explain crowd-in.

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 auto-IRAs relative to ESRPs, and therefore eschew auto-IRAs in favor of 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 a relatively small share of complier owners (26%) participated in the ESRP at event time +1 at all – possibly because many owners did not receive a W-2 wage from their firm, and thus are not eligible to contribute to its ESRP. This low participation rate casts doubt on owners' direct benefit playing an outside role in ESRP adoption.

iv. Possibility 4: Auto-IRAs are perceived to impose large administrative burdens on employers.

It may also be the case that business owners perceive the burden of facilitating auto-IRA contributions to be high. While running an ESRP likely has a higher cost to employers than using the auto-IRA program, it may be easier to outsource this cost to a third party ESRP provider. While we cannot estimate the perceived burden of using the auto-IRA 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, both for payroll tax returns and for income tax returns (with the latter restricted to S corporations). 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. It is possible that some of these employers perceive the 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 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 (2022 is the most recent year when this is observed), 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 experiments, 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.²⁵ (In contrast, around 9 percent of always takers claim the

²⁵ This analysis is restricted to firms that file an income tax return, where the credit claim would be found.

credit.) This suggests that in a context similar to (though not exactly the same as) administering an auto-IRA, perceived or actual administration costs are substantial enough to cause employers to forego meaningful amounts of 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 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 auto-IRAs.

v. Possibility 5: 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 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. That is, in the framework of Figure 2, some share of firms above the diagonal pre-policy were

nonetheless not offering an ESRP. When an 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.

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 auto-IRA. It is possible that this marketing was particularly successful and effectively altered decision-makers’ perceptions of the costs and benefits of both ESRPs and auto-IRAs.²⁶

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 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 auto-IRA program despite

²⁶ 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>.

being required to do so. 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 experiment 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 experiment, which took effect in 2019.²⁷ 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 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 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 Figure 2 slightly, firms compare offering an ESRP to the best alternative, which may be either participating in the auto-IRA policy 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.

²⁷ 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.

Upon receiving the letter, the best alternative may have remained “doing nothing”, but at a higher cost – in which case firms would move left in Figure 2, potentially leading ESRPs to become optimal. Alternatively, the letter may have changed the best alternative to “offering an auto-IRA” – 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 experiment. 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.

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 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. 2023. How Do Firms Respond to State Retirement Plan Mandates? NBER Working Paper no. 31398. June.
- 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.
- 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. 2021. Default Options and Retirement Saving Dynamics. Retrieved from <https://tahachoukmane.com/wp-content/uploads/2021/04/Choukmane-2021-Default-options.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://scorreia.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*.
- 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.
- 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.
- Lusardi, Annamaria and Olivia S. Mitchell. 2011. Financial Literacy around the World: An Overview. *Journal of Pension Economics and Finance* 10(4), 497–508.
doi:10.1017/S1474747211000448
- 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. 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
- Madrian, Bridgitte and Dennis F. Shea. 2001. The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior. *Quarterly Journal of Economics* 116(4): 1149-1187.
- 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.
- 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.
- 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. 2023. 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 \(pewtrusts.org\)](#)
- Summers, Lawrence H. 1989. Some Simple Economics of Mandated Benefits. *American Economic Review* 79(2): 177-183.

- 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. 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.
- 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>.

Table 1: Auto-IRA Policy Experiments and Implementation Years

State	Employer Size	Year of Implementation
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 experiments that we study in the main analysis. Employer size refers to the number of employees employed by the firm.

Table 2: Relative effects across experiments and states:

State	Employer Size	Year of Implementation	Firms induced to offer	Firms not offering, pre-policy	Share induced
			(1)	(2)	(3)
Oregon	20-99	2018	416	3170	13.1%
Oregon	5-19	2019	1795	16164	11.1%
Illinois	25-99	2019	883	6918	12.8%
California	50-99	2021	1395	6176	22.6%
Illinois	16-24	2022	477	5895	8.1%
California	5-49	2022	26513	165315	16.0%
Connecticut	26-99	2022	321	1856	17.3%

Notes: This table reports the estimated magnitude of the effect for each experiment. 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). Source: Authors’ calculations from confidential tax data.

Table 3: Induced ESRP offers as a share of new plan offerings

State	Induced ESRP	Firms participating in Auto-IRA	ESRP as share of sum
	(1)	(2)	(3)
Oregon	2211	6000	27%
California	27908	34000	45%
Illinois	1360	3700	27%

Notes: This table compares the estimated magnitude of the ESRP effect in each state to the number of firms actively participating in the auto-IRA program by May of 2023. The estimated magnitude of the ESRP effect corresponds to the amounts reported in Table 2, column (1), aggregated by state. The number of firms participating in the auto-IRA program is retrieved publicly available data compiled by the Center for Retirement Initiatives at Georgetown University and refers to the number of employers actively submitting payroll deductions. Column (3) reports the ratio of column (1) to the sum of columns (1) and (2). Source: Authors' calculations from confidential tax data and data compiled by the Center for Retirement Initiatives at Georgetown University.

Table 4: Means of firm-level variables for Compliers, Never-Offerers, Always-Offerers

Characteristic	Compliers	Never-offerers	Always-offerers	Always-takers
	<i>Industries</i>			
Education/Health	0.144 (0.006)	0.124 (0.001)	0.199 (0.002)	0.203 (0.007)
Goods-Producing	0.165 (0.006)	0.177 (0.001)	0.203 (0.002)	0.181 (0.007)
Leisure/Hospitality	0.249 (0.005)	0.266 (0.002)	0.043 (0.001)	0.127 (0.005)
Professional Services	0.102 (0.005)	0.089 (0.001)	0.233 (0.002)	0.176 (0.007)
Trade, Transportation, Utilities	0.206 (0.006)	0.204 (0.001)	0.151 (0.001)	0.154 (0.007)
All other industries	0.134 (0.005)	0.140 (0.001)	0.171 (0.002)	0.158 (0.006)
	<i>Other characteristics</i>			
Offers health insurance	0.356 (0.008)	0.264 (0.002)	0.728 (0.002)	0.498 (0.009)
Offers SIMPLE IRA (at event time 1)	0.220 (0.006)	0.000 (0.000)	0.166 (0.002)	0.142 (0.008)

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 experiments taking place in 2020 or later only. Bootstrapped standard errors are in parentheses. Source: Authors’ calculations from confidential tax data.

Table 5: Characteristics of employees and owners for Compliers, Never-Offerers, Always-Offerers, and Always-Takers

Characteristic	<i>Employees</i>				<i>Owners</i>			
	Compliers (1)	Never- offerers (2)	Always- offerers (3)	Always- takers (4)	Compliers (5)	Never- offerers (6)	Always- offerers (7)	Always- takers (8)
Age	38.37 (0.11)	39.30 (0.03)	41.82 (0.03)	37.63 (0.13)	52.69 (0.35)	54.02 (0.09)	55.56 (0.19)	50.63 (0.36)
Male	0.527 (0.004)	0.533 (0.001)	0.509 (0.001)	0.507 (0.005)	0.686 (0.014)	0.692 (0.002)	0.740 (0.008)	0.703 (0.011)
Married	0.347 (0.003)	0.368 (0.001)	0.471 (0.001)	0.372 (0.004)	0.742 (0.014)	0.749 (0.005)	0.811 (0.003)	0.758 (0.011)
Have Dependents	0.315 (0.002)	0.331 (0.001)	0.362 (0.001)	0.345 (0.003)	0.462 (0.012)	0.430 (0.007)	0.435 (0.005)	0.501 (0.014)
Black	0.046 (0.001)	0.047 (0.000)	0.045 (0.000)	0.052 (0.002)	0.043 (0.002)	0.045 (0.001)	0.037 (0.001)	0.045 (0.003)
Hispanic	0.365 (0.003)	0.366 (0.001)	0.267 (0.001)	0.300 (0.003)	0.199 (0.026)	0.194 (0.024)	0.126 (0.018)	0.154 (0.023)
Investment income	0.360 (0.003)	0.355 (0.001)	0.486 (0.001)	0.388 (0.003)	0.801 (0.013)	0.782 (0.013)	0.922 (0.011)	0.837 (0.016)
Log wages	9.814 (0.012)	9.732 (0.003)	10.679 (0.003)	10.125 (0.014)	N/A	N/A	N/A	N/A
New this year	0.362 (0.004)	0.352 (0.001)	0.240 (0.001)	0.396 (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 4 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 6: Testing automatic enrollment: Comparing 2022 and 2023 experiments at event time zero

	<u>Starts plan</u>		<u>Offers plan</u>	
	(1)	(2)	(3)	(4)
2022 experiment	California	California	California	California
2023 experiment	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)	0.131 (0.004)	0.078 (0.007)

Notes: This table compares the event time 0 treatment estimates for the 2022 California experiment (or some subset thereof) to those from 2023 experiments 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.

Table 7: Complier means for participation, contribution constraints

Characteristic	<i>Employees</i>				<i>Owners</i>			
	Compliers (1)	Never- offerers (2)	Always- offerers (3)	Always- takers (4)	Compliers (5)	Never- offerers (6)	Always- offerers (7)	Always- takers (8)
<i>IRA and ESRP participation rate</i>								
Contribute to IRA (-2)	0.090 (0.002)	0.079 (0.000)	0.098 (0.000)	0.107 (0.002)	0.206 (0.012)	0.188 (0.008)	0.113 (0.009)	0.238 (0.015)
Contribute to this ESRP (1)	0.335 (0.004)	0.000 (0.000)	0.461 (0.001)	0.375 (0.005)	0.260 (0.026)	0.000 (0.000)	0.624 (0.023)	0.504 (0.030)
<i>Contribution limits</i>								
Income above Roth IRA limits (-2)	0.075 (0.002)	0.072 (0.000)	0.177 (0.001)	0.118 (0.002)	0.455 (0.041)	0.435 (0.028)	0.712 (0.017)	0.600 (0.013)
Contributions to ESRP above IRA contribution limits (1)	0.025 (0.002)	0.000 (0.000)	0.163 (0.001)	0.076 (0.002)	0.071 (0.020)	0.000 (0.000)	0.491 (0.018)	0.323 (0.025)

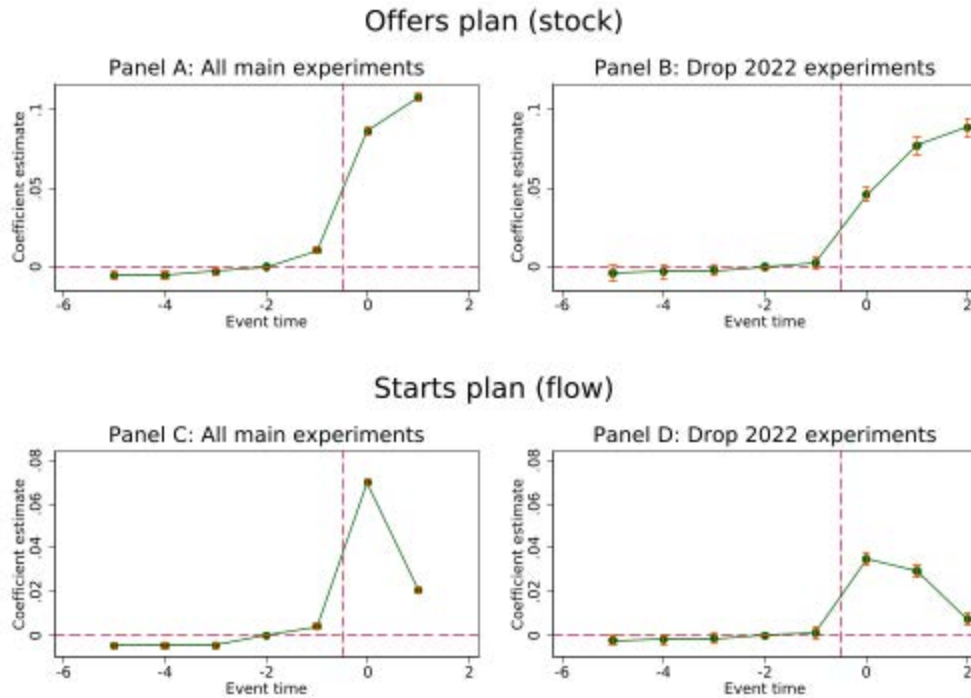
Notes: This table reports estimated means of various outcomes for compliers, never-offerers, always-offerers, and always-takers. See text and notes to Table 4 for how these objects are calculated. Columns (1)-(4) report these means for the firm's employees, while columns (5)-(8) report these means for 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 and hassle costs

Characteristic	Compliers	Never-offerers	Always-offerers	Always-takers
<i>Tax filing behavior</i>				
Has Form 941 preparer	0.503 (0.007)	0.546 (0.002)	0.502 (0.002)	0.527 (0.009)
Has Form 1120S preparer	0.135 (0.009)	0.256 (0.003)	0.130 (0.002)	0.168 (0.010)
Average Form 1120S filing time (days)	151.9 (2.2)	146.7 (0.6)	148.8 (0.6)	142.2 (2.3)
<i>Take-up of section 45E credit</i>				
Receives section 45E credit	0.035 (0.006)	0.002 (0.000)	0.009 (0.001)	0.093 (0.007)
<i>Proxied political ideology</i>				
County Democratic vote share	0.636 (0.002)	0.638 (0.000)	0.641 (0.001)	0.645 (0.002)


Notes: This table reports estimated means of various outcomes for compliers, never-offerers, always-offerers, and always-takers. See text and notes to Table 4 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 experiments 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.

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 experiments, with $k^{max} = 1$. In Panels B and D, we drop experiments 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: Stylized depiction of ESRP decision



Notes: This figure depicts a stylized analysis of the costs and benefits of offering an ESRP, relative to an alternative. Prior to the auto-IRA policy, the alternative is not offering a plan. After the auto-IRA policy, the alternative is participating in the auto-IRA program. See text for a discussion of each of the points on this figure.

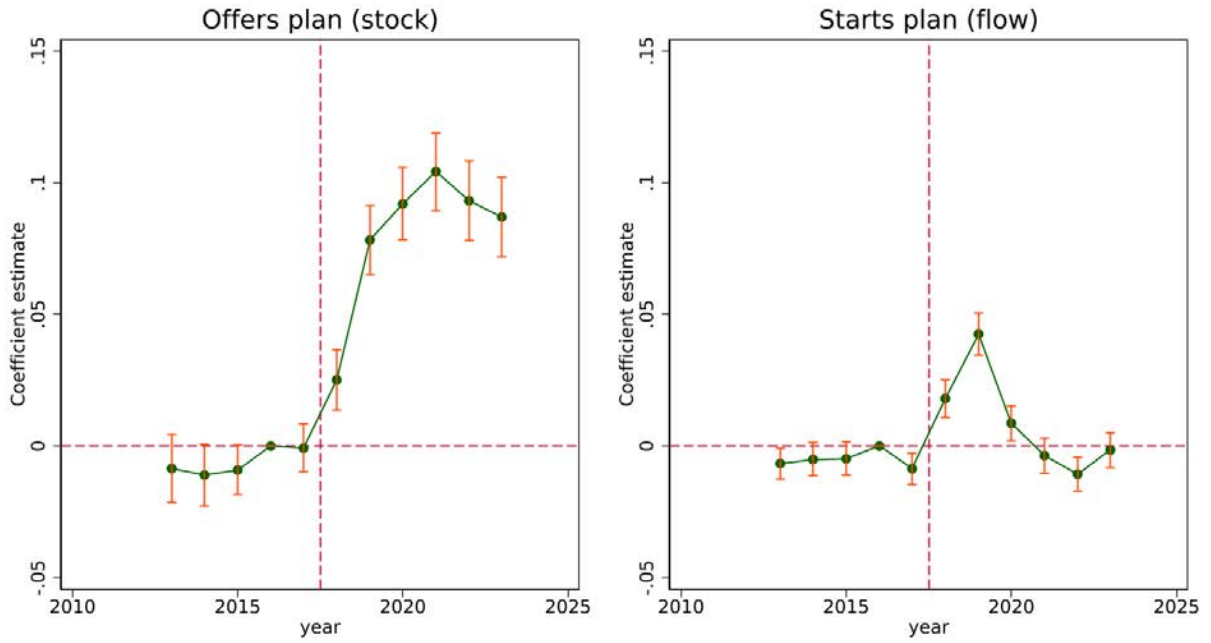
Figure 3: Case study for Illinois 25-99 employee experiment: the role of early enforcement actions



Notes: This figure reports the event study for the “starts plan” outcome for the Illinois 25-99 employee experiment. 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 Figures

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




Notes: This figure plots the event studies from Equation (1), restricted to the Oregon 20-99 employee experiment, 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 experiment




Notes: This figure plots the event studies from Equation (1), restricted to the Oregon 5-19 employee experiment, 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 experiment




Notes: This figure plots the event studies from Equation (1), restricted to the Illinois 25-99 employee experiment, 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 experiment



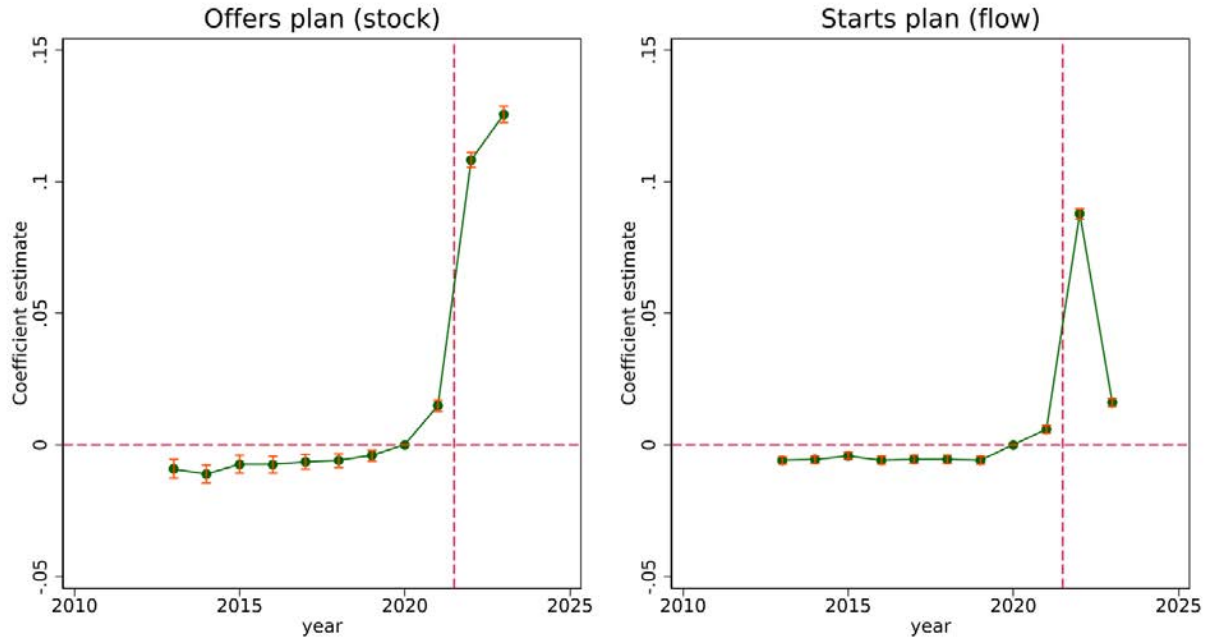
Notes: This figure plots the event studies from Equation (1), restricted to the California 50-99 employee experiment, using all years of available data. See the notes to Figure 1 for further details. Source: Authors' calculations from confidential tax data.

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




Notes: This figure plots the event studies from Equation (1), restricted to the Illinois 16-24 employee experiment, 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 experiment



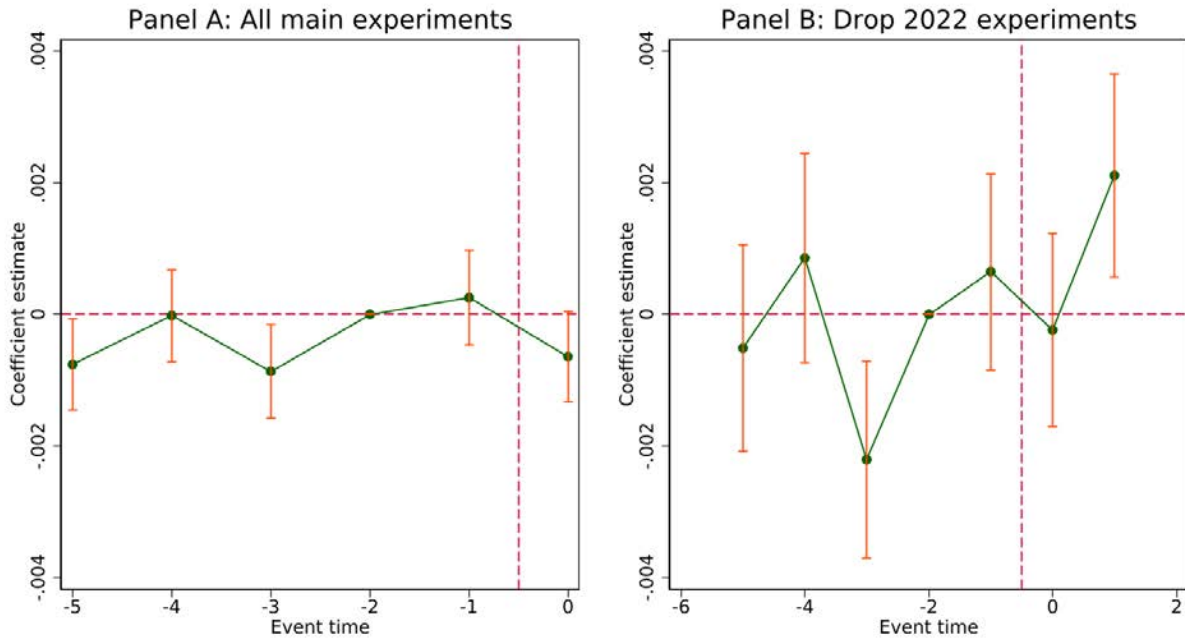
Notes: This figure plots the event studies from Equation (1), restricted to the California 5-49 employee experiment, using all years of available data. See the notes to Figure 1 for further details. Source: Authors' calculations from confidential tax data.

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



Notes: This figure plots the event studies from Equation (1), restricted to the Connecticut 26-99 employee experiment, 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 experiments, while Panel B drops experiments implemented in 2022. See notes to Figure 1 for further details. Source: Authors’ calculations from confidential tax data.