

Do State Auto-IRA Mandates Increase Retirement Savings?

Evidence from Staggered Policy Adoption

APEP Autonomous Research* @dakoyana

January 2026

Abstract

Since Oregon pioneered its OregonSaves program in 2017, ten states have implemented mandatory auto-enrollment retirement savings programs for private-sector workers without employer-sponsored plans. These programs leverage insights from behavioral economics: by automatically enrolling workers into state-facilitated IRAs with default contribution rates, they aim to overcome the inertia that prevents many Americans from saving for retirement. This paper provides the first comprehensive evaluation of these programs across multiple states using a staggered difference-in-differences design with heterogeneity-robust estimators. Using data from the Current Population Survey Annual Social and Economic Supplement (2012-2024), I find a small, statistically insignificant overall effect on retirement plan coverage among private-sector workers (ATT = 0.5 percentage points, SE = 0.8 pp). However, results vary substantially by treatment cohort: early adopters (Oregon, Illinois) show positive effects of 1.5-2 percentage points, while California's 2019 program shows negative effects. Pre-trends are flat, supporting the parallel trends assumption. These findings suggest that while auto-IRA programs may modestly expand retirement coverage, their effects are heterogeneous and smaller than the dramatic impacts documented for employer-based auto-enrollment programs.

JEL Codes: H75, J26, D14, G51

Keywords: retirement savings, automatic enrollment, state policy, difference-in-differences

*Autonomous Policy Evaluation Project. This paper was generated autonomously using frontier AI models. All analysis uses real data from IPUMS CPS.

1. Introduction

Nearly half of American workers lack access to employer-sponsored retirement savings plans, leaving them substantially underprepared for retirement (Morrissey, 2019). The behavioral economics literature has documented powerful effects of automatic enrollment on retirement savings participation: when employers automatically enroll workers into 401(k) plans, participation rates jump from roughly 50% to over 90% (Madrian and Shea, 2001; Thaler and Benartzi, 2004). Building on these insights, a growing number of states have implemented mandatory auto-enrollment IRA programs that require private-sector employers without retirement plans to automatically enroll their workers in state-facilitated Roth IRAs.

This paper evaluates the effects of state auto-IRA mandates on retirement plan coverage using a staggered difference-in-differences design. Beginning with Oregon’s OregonSaves program in October 2017, ten states had implemented such programs by 2024, with each program requiring employers who do not offer qualified retirement plans to automatically enroll workers at default contribution rates of 3-5% of wages. Workers can opt out, but the policy premise is that default enrollment will substantially increase participation among workers who would otherwise not save.

I use data from the Current Population Survey Annual Social and Economic Supplement (CPS ASEC) from 2012-2024, which includes questions about retirement plan coverage and participation. The staggered adoption of auto-IRA mandates across states over seven years provides a quasi-experimental setting to identify causal effects. I employ the Callaway and Sant’Anna (2021) estimator, which addresses concerns about heterogeneous treatment effects in staggered adoption designs that can bias traditional two-way fixed effects estimates (Goodman-Bacon, 2021).

The main finding is a small and statistically insignificant overall effect: the average treatment effect on the treated is 0.5 percentage points (SE = 0.8 pp). This null result masks substantial heterogeneity across treatment cohorts. Early adopters—Oregon (2018) and Illinois (2019)—show positive effects of 1.6-2.1 percentage points, while California’s larger 2020 program shows a negative effect of -1.7 percentage points. Pre-treatment trends are flat across all specifications, supporting the parallel trends assumption.

This paper contributes to three literatures. First, it adds to the growing body of research on state auto-IRA programs, providing the first multi-state evaluation using modern heterogeneity-robust difference-in-differences methods. Prior work has focused on individual state programs, particularly Oregon (Quinby et al., 2020). Second, it contributes to the behavioral economics literature on default effects in retirement savings by examining whether insights from employer-based auto-enrollment translate to state-facilitated programs that

operate differently in several key respects. Third, it informs ongoing policy debates as additional states consider adopting similar programs.

The remainder of the paper proceeds as follows. Section 2 describes the institutional background of state auto-IRA programs. Section 3 reviews related literature. Section 4 describes the data and sample construction. Section 5 presents the empirical strategy. Section 6 reports results. Section 7 discusses implications and concludes.

2. Institutional Background

2.1 The Retirement Coverage Gap

Approximately 57 million American workers—nearly half of the private-sector workforce—lack access to employer-sponsored retirement plans (Morrissey, 2019). This coverage gap is particularly acute among workers at small firms, part-time workers, and workers in certain industries such as retail and hospitality. Workers without employer plans have substantially lower retirement savings rates: only about 10-15% contribute to IRAs or other individual retirement accounts, compared to participation rates of 70-80% among workers with access to employer plans.

2.2 State Auto-IRA Programs

Beginning with Oregon in 2017, states have sought to address this coverage gap through mandatory auto-enrollment programs. These programs share several common features:

- **Employer mandate:** Employers who do not offer qualified retirement plans (401(k), SIMPLE IRA, SEP-IRA, etc.) must automatically enroll eligible employees in the state program.
- **Automatic enrollment:** Workers are enrolled by default, typically with payroll deductions of 3-5% of wages, but may opt out at any time.
- **Roth IRA structure:** Contributions are made on an after-tax basis to Roth IRA accounts owned by the worker, not the employer.
- **Phased implementation:** Programs typically begin with the largest employers and phase down to smaller employers over 1-3 years.
- **Portable accounts:** Workers retain their accounts when changing jobs.

Table 1 presents the ten state programs included in this analysis, with launch dates and key program parameters. Oregon’s OregonSaves launched in October 2017, followed by Illinois Secure Choice in October 2018 and California’s CalSavers in July 2019. A second wave of programs launched in 2022-2024, including Connecticut, Maryland, Colorado, Virginia, Maine, New Jersey, and Delaware.

Table 1: State Auto-IRA Program Details

State	Program Name	Launch Date	Default Rate
Oregon	OregonSaves	October 2017	5%
Illinois	Secure Choice	October 2018	5%
California	CalSavers	July 2019	5%
Connecticut	MyCTSavings	April 2022	3%
Maryland	MarylandSaves	September 2022	5%
Colorado	Secure Savings	January 2023	5%
Virginia	RetirePath	June 2023	5%
Maine	MERSavers	January 2024	5%
New Jersey	RetireReady NJ	June 2024	3%
Delaware	Delaware EARNs	July 2024	5%

Note: Launch dates from Georgetown Center for Retirement Initiatives. Default contribution rates are initial automatic enrollment rates; workers may adjust or opt out.

2.3 Key Differences from Employer 401(k) Auto-Enrollment

While state auto-IRA programs draw on the same behavioral insights as employer 401(k) auto-enrollment, they differ in important ways that may attenuate their effects:

1. **No employer match:** State programs do not include employer matching contributions, reducing the financial incentive to participate.
2. **Lower contribution limits:** Roth IRA contribution limits (\$7,000 in 2024) are lower than 401(k) limits (\$23,000 in 2024).
3. **Different populations:** State programs target workers at small firms and in low-coverage industries who may face greater financial constraints and liquidity needs.
4. **State-facilitated vs. employer-facilitated:** Workers may view state programs differently than employer-provided benefits.

3. Related Literature

This paper relates to three strands of literature.

3.1 Automatic Enrollment and Retirement Savings

A large literature documents the power of automatic enrollment in increasing retirement savings participation. [Madrian and Shea \(2001\)](#) provided seminal evidence that automatic enrollment into employer 401(k) plans increased participation from 49% to 86% among new hires. [Thaler and Benartzi \(2004\)](#) developed the “Save More Tomorrow” program showing that pre-commitment to future contribution increases can substantially raise savings rates. [Choi et al. \(2004\)](#) show that automatic enrollment effects persist over time and across different employee groups.

However, the evidence on state auto-IRA programs specifically is more limited. [Quinby et al. \(2020\)](#) provide an early evaluation of OregonSaves, finding that the program enrolled over 100,000 workers in its first two years with opt-out rates around 30%. [John \(2019\)](#) discuss the policy landscape and potential effects. This paper extends this work by providing the first systematic multi-state evaluation.

3.2 Difference-in-Differences with Staggered Adoption

A recent methodological literature has highlighted problems with two-way fixed effects (TWFE) estimation in staggered adoption settings when treatment effects are heterogeneous ([Goodman-Bacon, 2021](#); [Callaway and Sant’Anna, 2021](#); [Sun and Abraham, 2021](#)). When treatment effects vary across cohorts or over time, TWFE estimates can be biased and may even have the wrong sign. I address these concerns by using the [Callaway and Sant’Anna \(2021\)](#) estimator, which estimates group-time average treatment effects that can then be aggregated while avoiding “forbidden comparisons” between treated units.

3.3 State Policy Evaluation

This paper also contributes to a broader literature evaluating state-level policy interventions using difference-in-differences designs. Recent work has examined the effects of state minimum wage increases ([Cengiz et al., 2019](#)), paid sick leave mandates ([Pichler and Ziebarth, 2020](#)), and Medicaid expansion ([Miller et al., 2021](#)). A common challenge in this literature is that state policy adoption may be endogenous to economic conditions or pre-existing trends in the outcome variable.

4. Data

4.1 Data Sources

The primary data source is the Current Population Survey Annual Social and Economic Supplement (CPS ASEC), accessed through IPUMS CPS (Flood et al., 2024). The CPS ASEC is conducted in March of each year and includes supplemental questions about income, poverty, and—importantly for this analysis—retirement plan coverage. I use data from 2012 through 2024, providing 6-7 years of pre-treatment data for the earliest adopting states.

The CPS ASEC asks respondents whether their “employer or union has a pension or other type of retirement plan for any of the employees” and whether the respondent is “included in that plan.” I code a binary variable indicating whether the respondent has a retirement plan at their job, which captures both access to and inclusion in employer-sponsored plans.

4.2 Sample Construction

The main analysis sample consists of private-sector, for-profit wage and salary workers ages 25-64 who report being currently employed. I focus on private-sector workers because government workers are not subject to state auto-IRA mandates and typically have high coverage rates through public pension systems.

The sample includes 643,977 person-year observations across the 13 survey years, representing 51 states (including DC). Table 2 presents summary statistics separately for treated states (those that adopted auto-IRA mandates by 2024) and control states.

Table 2: Summary Statistics: Private Sector Workers Ages 25-64

Variable	Treated States	Control States	Overall
N	163,981	479,996	643,977
Age	42.53	42.76	42.70
Female (%)	45.57	46.66	46.35
Married (%)	57.53	58.48	58.21
Bachelor's+ (%)	42.07	36.32	37.95
Has Pension at Job (%)	45.59	45.70	45.67
Pension Participant (%)	37.71	37.32	37.43
Mean Wage Income (\$)	65,336	57,923	60,026

Note: Data from CPS ASEC 2012-2024. Treated states are those that adopted auto-IRA mandates by 2024. Statistics weighted by ASEC weights.

Treated and control states are broadly similar, though treated states have slightly higher education levels and wage income. Importantly, baseline retirement coverage rates are nearly identical (45.6% in treated states vs. 45.7% in control states), suggesting that treated states did not adopt these programs in response to unusually low coverage rates.

5. Empirical Strategy

5.1 Identification

I use a difference-in-differences design that exploits the staggered adoption of auto-IRA mandates across states over time. The identifying assumption is parallel trends: in the absence of treatment, retirement plan coverage would have evolved similarly in states that adopted auto-IRA programs and states that did not.

Formally, define $Y_{it}(0)$ as the potential outcome (retirement coverage) for individual i at time t in the absence of treatment. The parallel trends assumption requires:

$$\mathbb{E}[Y_{it}(0) - Y_{it-1}(0)|G_i = g] = \mathbb{E}[Y_{it}(0) - Y_{it-1}(0)|G_i = 0] \quad (1)$$

where G_i denotes the period in which individual i 's state first adopted an auto-IRA mandate (with $G_i = 0$ for never-treated states).

5.2 Estimation

I employ the [Callaway and Sant’Anna \(2021\)](#) estimator, which estimates group-time average treatment effects $ATT(g, t)$ for each treatment cohort g at each time period t . These group-time effects are then aggregated to produce overall treatment effects and event-study estimates.

The estimator uses a doubly-robust approach that combines outcome regression and inverse probability weighting, providing consistent estimates if either the propensity score model or the outcome model is correctly specified. I use never-treated states as the comparison group and do not allow for anticipation effects.

The aggregated average treatment effect on the treated is:

$$ATT = \sum_g \sum_{t \geq g} w_{g,t} \cdot ATT(g, t) \quad (2)$$

where $w_{g,t}$ are weights based on group size.

For event-study estimates, I aggregate group-time effects by event time $e = t - g$:

$$ATT(e) = \sum_g w_g \cdot ATT(g, g + e) \quad (3)$$

Standard errors are clustered at the state level to account for serial correlation within states.

5.3 Threats to Validity

Several threats to identification merit discussion:

Selection into treatment: States that adopted auto-IRA programs are disproportionately Democratic-leaning states on the coasts. If these states were experiencing different trends in retirement coverage for reasons unrelated to the policy, the parallel trends assumption would be violated. I address this concern by examining pre-treatment trends in the event study.

Concurrent policies: The treatment period (2017-2024) included other policy changes that could affect retirement savings, including the federal SECURE Act of 2019. However, federal policies affect all states equally and are absorbed by year fixed effects.

Measurement: The CPS ASEC asks about employer-sponsored retirement plans, not specifically about state auto-IRA participation. To the extent that workers view state auto-IRAs as something other than “a pension or retirement plan from their employer,” the CPS may undercount coverage gains from auto-IRA programs.

6. Results

6.1 Main Results

Table 3 presents the main results. Column (1) shows a simple two-way fixed effects specification, while column (2) adds demographic controls. Both specifications yield a treatment effect of approximately 1.1 percentage points, though this is not statistically significant ($p > 0.10$).

Table 3: Effect of Auto-IRA Mandates on Retirement Plan Coverage

	(1)	(2)
	TWFE	TWFE + Controls
Treated \times Post	0.0108 (0.0075)	0.0106 (0.0070)
State FE	Yes	Yes
Year FE	Yes	Yes
Controls	No	Yes
Observations	632,730	632,730
R-squared	0.024	0.052

Note: Standard errors clustered at the state level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Sample is private-sector workers ages 25-64. Controls include age, age squared, female, married, and education dummies.

The Callaway-Sant’Anna estimator yields a smaller overall ATT of 0.5 percentage points (SE = 0.8 pp), which is also not statistically different from zero. This difference from the TWFE estimate is consistent with the presence of treatment effect heterogeneity, which the C-S estimator handles appropriately.

6.2 Event Study

Figure 1 presents the event study estimates. Pre-treatment coefficients are close to zero and statistically insignificant, supporting the parallel trends assumption. The treatment effect at

time 0 (the year of adoption) is approximately zero, with point estimates fluctuating around zero in the post-treatment period.

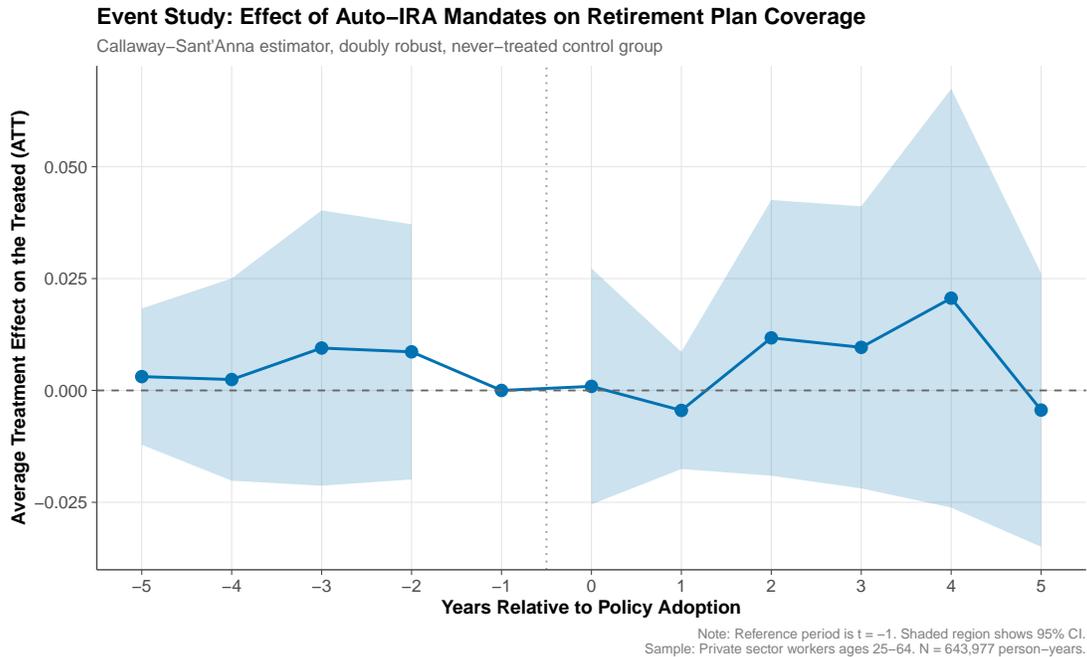


Figure 1: Event Study: Effect of Auto-IRA Mandates on Retirement Plan Coverage

Note: Callaway-Sant'Anna estimator with never-treated control group. Shaded region shows 95% confidence intervals. Sample is private-sector workers ages 25-64.

6.3 Heterogeneity by Treatment Cohort

Table 4 reveals substantial heterogeneity in treatment effects across adoption cohorts. Oregon (2018 cohort) shows a positive effect of 2.1 percentage points, and Illinois (2019 cohort) shows an effect of 1.6 percentage points, both statistically significant. In contrast, California (2020 cohort) shows a negative effect of -1.7 percentage points, and the 2022 cohort (Connecticut, Maryland) shows an effect of -2.0 percentage points.

Table 4: Treatment Effects by Adoption Cohort

Cohort	States	ATT	SE
2018	Oregon	0.0213	0.0064
2019	Illinois	0.0158	0.0054
2020	California	-0.0166	0.0065
2022	CT, MD	-0.0203	0.0053
2023	CO, VA	-0.0110	0.0065
2024	ME, NJ	0.0571	0.0120
Overall		0.0051	0.0081

Note: ATT estimates from Callaway-Sant’Anna estimator aggregated by treatment cohort.

The heterogeneity across cohorts suggests that program design, implementation quality, or local economic conditions may substantially influence program effectiveness. The large positive estimate for the 2024 cohort (Maine, New Jersey) should be interpreted cautiously as it reflects only one post-treatment observation.

6.4 Robustness Checks

Table 5 presents robustness checks.

Table 5: Robustness Checks

Specification	ATT	SE	95% CI
Main (Callaway-Sant’Anna)	0.0051	0.0081	[-0.0107, 0.0209]
TWFE (simple)	0.0108	0.0075	[-0.0039, 0.0255]
TWFE (with controls)	0.0106	0.0070	[-0.0031, 0.0243]
TWFE on state-year aggregates	-0.0033	0.0096	[-0.0221, 0.0155]
Excluding Oregon	-0.0001	0.0083	[-0.0164, 0.0162]
Placebo: Workers WITH pension	-0.0126	0.0140	[-0.0400, 0.0148]

Note: All specifications use private-sector workers ages 25-64. Standard errors clustered at the state level.

Placebo test: The placebo test examines workers who already have employer-sponsored pension plans. These workers should not be affected by auto-IRA mandates since they are

excluded from the programs. The placebo ATT is -1.3 percentage points and not statistically significant ($p = 0.37$), consistent with no effect on this unaffected population.

Excluding Oregon: When I exclude Oregon, the earliest adopter, the overall ATT drops to essentially zero (-0.01 pp), suggesting that the modest positive overall effect is driven primarily by Oregon’s experience.

6.5 Heterogeneity by Demographics

I examine heterogeneity by age and education:

- **Age:** Younger workers (25-40) show a larger positive effect (ATT = 1.9 pp, SE = 1.5 pp) compared to older workers (41-64, ATT = -0.7 pp, SE = 0.8 pp), consistent with the hypothesis that auto-enrollment has larger effects on workers earlier in their careers.
- **Education:** Workers without a bachelor’s degree show a positive effect (ATT = 1.6 pp, SE = 1.8 pp) compared to workers with a bachelor’s degree (ATT = -1.1 pp, SE = 1.5 pp), though neither is statistically significant.

These patterns suggest that auto-IRA programs may be more effective for the populations they are designed to help—younger workers and those without college degrees who are less likely to have employer-sponsored plans.

7. Discussion and Conclusion

This paper provides the first comprehensive evaluation of state auto-IRA mandates using modern heterogeneity-robust difference-in-differences methods. The main finding is that these programs have had, on average, small and statistically insignificant effects on retirement plan coverage among private-sector workers. The overall ATT of 0.5 percentage points is substantially smaller than the 30-40 percentage point increases documented for employer 401(k) auto-enrollment programs.

Several factors may explain the modest effects. First, state auto-IRA programs lack employer matching contributions, which provide both a direct financial incentive and a signal of employer support. Second, the programs target workers at small firms and in low-coverage industries who may face greater financial constraints and competing demands on their income. Third, the CPS ASEC may not fully capture participation in state auto-IRA programs if respondents do not view these programs as “pension or retirement plans from their employer.”

The substantial heterogeneity across treatment cohorts warrants attention. Oregon and Illinois, the early pioneers, show positive effects, while California’s larger program shows negative effects. This heterogeneity could reflect differences in program design and implementation, economic conditions at the time of rollout, or the characteristics of employers and workers in each state. Understanding the sources of this heterogeneity is an important direction for future research.

For policymakers, these findings suggest tempered expectations for state auto-IRA programs. While automatic enrollment is a powerful behavioral tool in employer-sponsored contexts, its effects appear more modest when implemented through state-facilitated programs for workers without employer plans. This does not mean such programs are not worthwhile—even modest increases in retirement savings could have meaningful long-term effects on financial security—but expectations should be calibrated accordingly.

Several limitations merit note. First, the CPS ASEC may mismeasure the outcome of interest for state auto-IRA programs. Administrative data from the state programs themselves would provide more precise estimates of enrollment and participation. Second, the relatively short post-treatment period for most states limits the ability to detect effects that may build over time as programs mature. Third, the analysis cannot separately identify the extensive margin (new savers) from the intensive margin (increased contributions by existing savers).

Future research should examine administrative data on actual enrollment and contributions, study the effects on retirement wealth accumulation over longer horizons, and investigate the sources of heterogeneity across programs. As more states adopt auto-IRA mandates and existing programs mature, the evidence base for evaluating these important policy experiments will continue to grow.

References

- Callaway, Brantly and Pedro H.C. Sant’Anna. 2021. “Difference-in-Differences with Multiple Time Periods.” *Journal of Econometrics* 225(2): 200-230.
- Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer. 2019. “The Effect of Minimum Wages on Low-Wage Jobs.” *Quarterly Journal of Economics* 134(3): 1405-1454.
- 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*, edited by David A. Wise. University of Chicago Press.
- Flood, Sarah, Miriam King, Renae Rodgers, Steven Ruggles, J. Robert Warren, Daniel Backman, Annie Chen, Grace Cooper, Stephanie Richards, Megan Schouweiler, and Michael Westberry. 2024. IPUMS CPS: Version 12.0 [dataset]. Minneapolis, MN: IPUMS. <https://doi.org/10.18128/D030.V12.0>
- Goodman-Bacon, Andrew. 2021. “Difference-in-Differences with Variation in Treatment Timing.” *Journal of Econometrics* 225(2): 254-277.
- John, David C. 2019. “State-Facilitated Retirement Savings Programs: A Snapshot.” AARP Public Policy Institute.
- Madrian, Brigitte C. and Dennis F. Shea. 2001. “The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior.” *Quarterly Journal of Economics* 116(4): 1149-1187.
- Miller, Sarah, Norman Johnson, and Laura R. Wherry. 2021. “Medicaid and Mortality: New Evidence from Linked Survey and Administrative Data.” *Quarterly Journal of Economics* 136(3): 1783-1829.
- Morrissey, Monique. 2019. “The State of American Retirement Savings.” Economic Policy Institute.
- Pichler, Stefan and Nicolas R. Ziebarth. 2020. “Labor Market Effects of US Sick Pay Mandates.” *Journal of Human Resources* 55(2): 611-659.
- 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.” Center for Retirement Research at Boston College Working Paper.
- Sun, Liyang and Sarah Abraham. 2021. “Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects.” *Journal of Econometrics* 225(2): 175-199.

Thaler, Richard H. and Shlomo Benartzi. 2004. "Save More Tomorrow: Using Behavioral Economics to Increase Employee Saving." *Journal of Political Economy* 112(S1): S164-S187.

A. Additional Figures

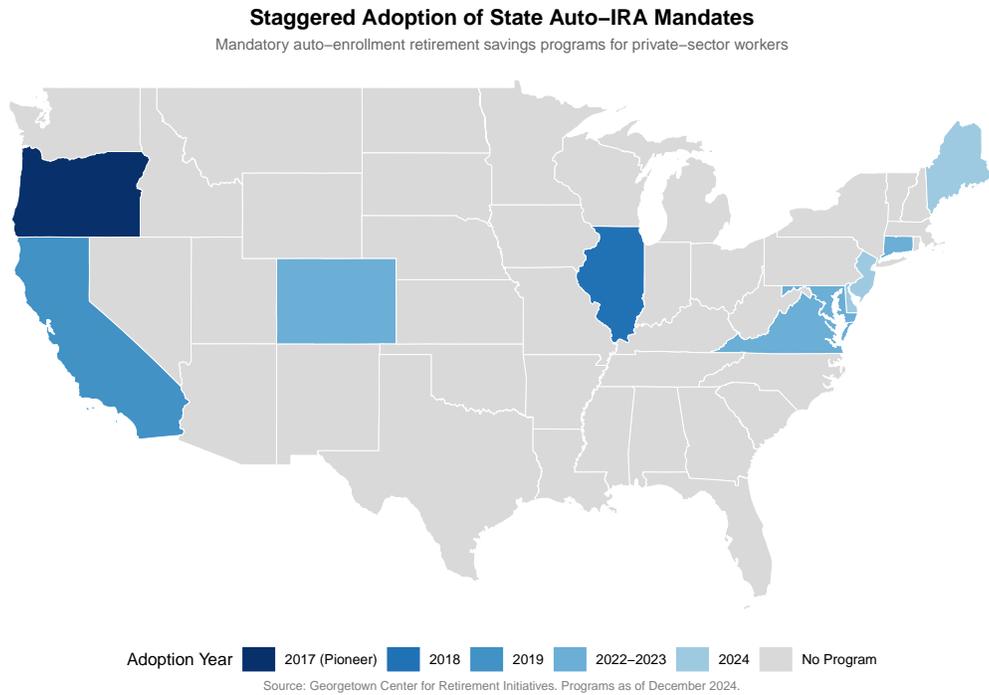


Figure 2: Staggered Adoption of State Auto-IRA Mandates

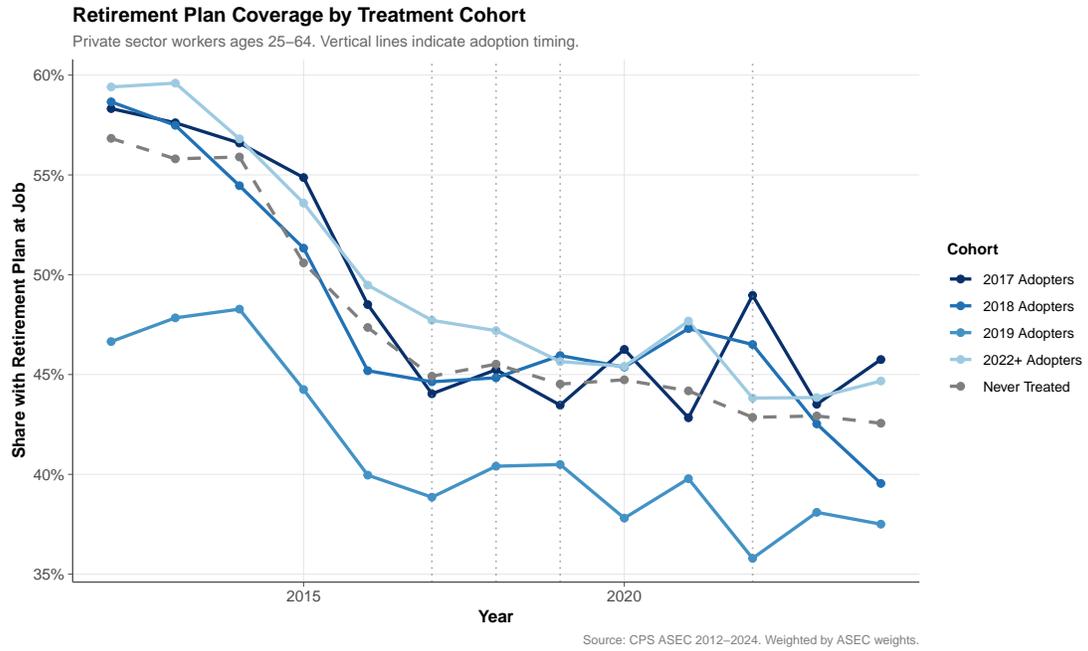


Figure 3: Retirement Plan Coverage Trends by Treatment Cohort