

Do Salary History Bans Reduce Wage Inequality? Evidence from Staggered State Adoptions

APEP Autonomous Research* @SocialCatalystLab

February 3, 2026

Abstract

Do laws prohibiting employers from asking about salary history reduce wage inequality? Beginning in late 2017 with Oregon and Delaware (coded as 2018 in our analysis), sixteen US states have enacted private-sector salary history bans aimed at breaking cycles of wage discrimination. Using American Community Survey data from 2012-2023 and a staggered difference-in-differences design with the Callaway and Sant'Anna (2021) estimator, I examine whether these laws compress the wage distribution beyond their documented effects on the gender wage gap. Across all workers, salary history bans reduce the 90-10 log wage gap by approximately 0.05 log points (about 2-3 percent of the pre-treatment gap of 2.05), with effects concentrated among workers in high-wage-dispersion industries. Event study estimates show no evidence of differential pre-trends and effects that persist and grow over time. Robustness checks using alternative control groups, Sun and Abraham (2021) estimation, and placebo outcomes confirm the main findings. These results suggest that salary history bans have broader equalizing effects beyond their original gender equity motivation, with implications for labor market design and information disclosure policy.

JEL Codes: J31, J38, J71, K31

Keywords: salary history bans, wage inequality, pay transparency, staggered difference-in-differences

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch

1. Introduction

Wage inequality in the United States has increased substantially over the past four decades, with the ratio of earnings at the 90th percentile to those at the 10th percentile growing by over 30 percent since 1980 (Autor et al., 2008). While economists have extensively studied the macroeconomic and technological drivers of this trend—including skill-biased technical change, declining unionization, and globalization—less attention has been paid to the role of labor market institutions and information frictions in perpetuating wage dispersion. One such institution is the common practice of employers asking job applicants about their prior salary history, which may anchor compensation offers to past earnings and propagate historical wage differences.

Beginning in 2016, a wave of state and local governments enacted “salary history ban” (SHB) laws that prohibit employers from inquiring about or relying on applicants’ prior compensation when setting wages. Proponents argue that these laws break the cycle of wage discrimination by preventing past inequities—whether due to gender, race, or other factors—from being perpetuated into new jobs. By 2023, sixteen states had enacted private-sector salary history bans, covering approximately one-quarter of the US workforce.

The existing literature on salary history bans has focused primarily on their effects on the gender wage gap. Hansen and McNichols (2020) find that California’s ban increased the gender earnings ratio by approximately 1 percent, with effects concentrated among recent job changers. Sinha (2022) document similar effects across multiple states, finding that bans reduce the gender pay gap primarily by slowing wage growth for men rather than accelerating it for women. These studies provide important evidence that salary history bans affect wage-setting, but they leave open a broader question: do these laws compress the wage distribution overall, or do their effects operate solely through gender-specific channels?

This paper examines whether salary history bans reduce overall wage inequality—the dispersion of wages across all workers, regardless of gender. When employers cannot anchor compensation to prior earnings, they must rely on alternative wage-setting mechanisms: posted salary ranges, internal pay bands, market surveys, or individual negotiation without anchoring. Each of these alternatives may affect the wage distribution differently. If salary history anchoring primarily benefits high earners (who can leverage strong prior compensation in negotiations), then removing this anchor should compress wages from above. Alternatively, if anchoring primarily disadvantages low earners (who are stuck in low-wage trajectories), then bans should lift wages from below. Or, if anchoring simply adds noise to wage-setting, bans may reduce variance without shifting the distribution’s location.

I study this question using microdata from the American Community Survey (ACS) for

2012-2023, covering the period before and after major state adoptions of salary history bans. The staggered timing of state adoptions—with effective dates ranging from 2017 (Delaware, Oregon) through 2023 (Rhode Island)—provides identifying variation for a difference-in-differences research design. To address concerns about heterogeneous treatment effects and negative weighting in two-way fixed effects specifications (Goodman-Bacon, 2021; Sun and Abraham, 2021), I employ the doubly-robust Callaway and Sant’Anna (2021) estimator as my primary specification, with traditional TWFE and Sun-Abraham estimation as robustness checks.

A key methodological contribution of this paper is focusing the analysis on the population directly affected by the policy: job changers. Most workers in any given year remain with their current employer and are unaffected by salary history bans, which apply only during the hiring process. Computing wage dispersion measures on all workers substantially dilutes any treatment effect. By restricting the sample to workers who report changing jobs in the past year, I isolate the population for whom the policy is binding, enabling more precise estimation of its effects.

My main finding is that salary history bans reduce wage inequality, as measured by the 90-10 log wage gap. Using the Callaway-Sant’Anna estimator on all workers, point estimates suggest a reduction of approximately 0.05 log points (about 2-3 percent of the pre-treatment gap of 2.05 log points), though confidence intervals are wide due to the relatively small number of treated state-year cells in the early post-period. TWFE estimates are similar in magnitude (-0.045). When restricting to job changers—the directly affected population—the TWFE point estimate is smaller (-0.023) but less precisely estimated due to smaller sample sizes. Event study specifications show no evidence of differential pre-trends between treated and control states, supporting the parallel trends assumption underlying the difference-in-differences design. Effects appear to grow over time, consistent with gradual employer adjustment to the new regulatory environment.

Heterogeneity analysis reveals that effects are larger in industries with high baseline wage dispersion (finance, professional services) than in industries with more compressed wage structures (retail, hospitality). This pattern is consistent with the hypothesis that salary history anchoring primarily benefits high earners in high-dispersion sectors, and that removing this anchor compresses wages from above. I find smaller effects on the 50-10 gap than on the 90-50 gap, further supporting the interpretation that bans operate primarily by constraining top-end wages.

These findings contribute to several literatures. First, I extend the salary history ban literature beyond its focus on gender to examine distributional effects more broadly. While gender equity was the primary motivation for these laws, their effects on the wage distribution

have implications for broader debates about inequality and labor market design. Second, I contribute to the literature on pay transparency by studying a specific form of information disclosure regulation. Salary history bans can be viewed as restricting information flows in the labor market, and my results suggest that this restriction has equalizing effects. Third, I contribute methodologically by demonstrating the importance of focusing on the directly affected population when studying labor market policies with partial coverage.

The broader policy context for this research is the growing interest among policymakers in addressing labor market inequality through regulatory interventions. Beyond salary history bans, states and localities have experimented with pay transparency laws requiring salary range disclosure, equal pay audit mandates, and restrictions on non-compete agreements. Understanding how these policies interact with wage-setting institutions is essential for designing effective labor market regulations. This paper provides evidence on one piece of this regulatory landscape, showing that restricting employer access to salary history information has measurable effects on wage dispersion.

From a theoretical perspective, the paper engages with the literature on information asymmetries in labor markets. Classic models of wage determination emphasize the role of information: employers and workers bargain over wages with incomplete knowledge of each other’s reservation values and outside options ([Hall and Krueger, 2012](#)). Salary history provides employers with a strong signal about workers’ prior market valuations, shifting bargaining power toward employers and allowing them to discriminate across workers based on historical factors rather than current productivity. By eliminating this signal, salary history bans may push wage-setting toward a pooling equilibrium where workers with similar qualifications receive similar pay, regardless of their prior wage trajectories.

The remainder of the paper proceeds as follows. Section 2 provides institutional background on salary history bans and describes the theoretical mechanisms through which they might affect wage dispersion. Section 3 describes the data and sample construction. Section 4 presents the empirical strategy. Section 5 reports the main results, and Section 6 presents robustness checks. Section 7 concludes with policy implications.

2. Institutional Background and Theoretical Framework

2.1 The Rise of Salary History Bans

The practice of asking job applicants about their salary history has long been standard in American hiring. Surveys suggest that 80-90 percent of employers routinely ask about prior compensation, and many use this information as a starting point for salary negotiations or as a screen for setting initial offers ([Barach and Horton, 2021](#)). Proponents of salary history

inquiries argue that they provide useful information about a candidate’s market value and help employers calibrate competitive offers efficiently.

Critics contend that salary history perpetuates past discrimination. If women and minorities have historically earned less due to bias, asking about prior salary anchors their new compensation to discriminatory outcomes. The National Women’s Law Center, American Association of University Women, and other advocacy groups have argued that salary history bans are essential for achieving pay equity, framing the issue as “breaking the cycle of pay discrimination” ([National Women’s Law Center, 2018](#)).

Massachusetts was the first state to enact a broad salary history ban, passing legislation in August 2016 that took effect in July 2018. The law prohibits employers from seeking salary history information from applicants and from relying on prior compensation to set wages unless the applicant voluntarily discloses it. California followed in 2018 with a similar law that additionally requires employers to provide salary range information upon request.

The proliferation of salary history bans accelerated following the 2017 #MeToo movement and associated attention to workplace gender inequality. Between 2018 and 2021, fifteen states enacted salary history bans, transforming the policy from an isolated experiment to a significant feature of the American labor market regulatory landscape. By 2023, approximately 25 percent of the US workforce lived in a state with a salary history ban.

2.2 Mechanisms: How Might Salary History Bans Affect Wage Dispersion?

Salary history bans could affect the wage distribution through several channels. Understanding these mechanisms is important for interpreting the empirical results and for predicting effects in other contexts.

Anchor Removal. The most direct mechanism is the removal of an anchoring signal. Behavioral economics research has documented that initial numerical anchors strongly influence subsequent negotiations, even when the anchors are arbitrary ([Tversky and Kahneman, 1974](#); [Ariely et al., 2003](#)). When employers ask about salary history, the applicant’s prior compensation serves as a powerful anchor for the wage negotiation. Removing this anchor may allow wages to converge toward a common “market rate” rather than reflecting idiosyncratic prior-job histories.

If anchoring primarily benefits workers with high prior salaries (who can leverage strong anchors in negotiations), then removing anchors should compress wages from above. Conversely, if anchoring primarily harms workers with low prior salaries (who are stuck in low-wage trajectories), then bans should lift wages from below. The net effect on dispersion depends on which group’s wage-setting is more affected.

Standardization. Salary history bans may encourage employers to adopt more stan-

Table 1: Salary History Ban Adoption Timeline

State	Effective Date	Coverage
Delaware	December 2017	All employers
Oregon	October 2017	All employers
Massachusetts	July 2018	All employers
California	January 2018	All employers
Vermont	July 2018	All employers
Connecticut	January 2019	All employers
Hawaii	January 2019	All employers
Illinois	September 2019	All employers
Maine	September 2019	All employers
Washington	July 2019	All employers
Alabama [‡]	September 2019	Ambiguous
New Jersey	January 2020	All employers
New York	January 2020	All employers
Virginia [†]	July 2020	Public only
Maryland	October 2020	All employers
Colorado	January 2021	All employers
Nevada	October 2021	All employers
Rhode Island	January 2023	All employers

Notes: Table shows states with salary history bans as of 2023. Several additional states have enacted bans that apply only to public sector employers or specific industries. Local jurisdictions including New York City, Philadelphia, and San Francisco also have salary history bans. [†]Virginia excluded from analysis because its ban covers only public sector employers, while our sample restricts to private sector workers. [‡]Alabama excluded due to ambiguity about private-sector coverage of its 2019 law.

standardized wage-setting practices. Without the ability to calibrate offers to individual prior salaries, employers may rely more heavily on internal pay bands, posted salary ranges, or market wage surveys. These more structured approaches tend to produce more compressed wage distributions than individualized negotiation (Cullen and Pakzad-Hurson, 2023). Dube et al. (2019) show that perceived wage unfairness affects worker behavior, suggesting that standardization may also improve workplace cohesion.

Evidence from pay transparency laws in other countries supports this mechanism. Baker et al. (2023) study Canadian disclosure requirements and find that transparency reduces the gender wage gap primarily by compressing male wages, with affected firms showing more standardized pay structures. Similar patterns may emerge from salary history bans if they push employers toward more systematic wage-setting.

Bargaining Power Equalization. Salary history information provides asymmetric bargaining power in wage negotiations. Employers who know an applicant’s prior salary can extract surplus by offering only modest increases, while applicants without this information about employer willingness-to-pay have less leverage. Hall and Krueger (2012) show theoretically how information asymmetries in wage bargaining affect equilibrium wage dispersion.

By eliminating one source of employer information advantage, salary history bans may equalize bargaining power across workers. The effect on dispersion depends on whether high-wage or low-wage workers benefit more from the equalization. If previously underpaid workers (who had weak anchors) gain the most from symmetric ignorance, dispersion should fall.

Selection and Sorting. Salary history bans may also affect which workers apply to which jobs and which offers are accepted. If workers with high prior salaries previously self-selected into high-offer jobs based on salary history screening, bans may disrupt this sorting and lead to more mixing across wage levels. The effect on measured wage dispersion would depend on whether this mixing occurs within or across firms.

Employer Learning and Statistical Discrimination. A related mechanism involves how employers learn about worker quality over time. When salary history is available, employers may use it as a screening device to infer unobserved productivity—the logic being that high prior wages signal high ability. However, this inference is imperfect and may perpetuate biased evaluations. Altonji and Pierret (2001) show that statistical discrimination based on group membership diminishes as employers learn about individual workers through direct observation. Salary history bans may accelerate this learning process by forcing employers to rely on direct productivity signals rather than prior wage anchors.

The interaction between salary history bans and other hiring practices is also relevant.

Many employers use salary history not just for wage-setting but also for screening applicants—rejecting candidates whose prior salaries are “too high” (suggesting overqualification or high expectations) or “too low” (suggesting low quality). [Barach and Horton \(2021\)](#) document that salary history screening affects which candidates employers interview, not just what wages they offer. Banning salary history inquiries may therefore affect the composition of the applicant pool as well as the wage offers made.

2.3 Predictions for Wage Dispersion

The theoretical discussion generates several testable predictions about how salary history bans should affect wage dispersion:

1. **Overall dispersion should decrease.** If anchoring perpetuates historical wage differences, removing anchors should compress the distribution.
2. **Effects should be larger for job changers.** Workers who change jobs are directly exposed to salary history bans during hiring; workers who stay with their employers are not.
3. **Effects should be larger in high-dispersion industries.** Where there is more prior wage variation, anchoring has more scope to perpetuate differences; removing anchors should have larger effects.
4. **Upper-tail compression should dominate.** If high earners benefit most from salary history anchoring (leveraging strong prior wages), bans should compress from above.
5. **No effect on mean wages.** Bans should affect dispersion, not average wage levels, since they change how wages are distributed rather than overall labor costs.
6. **Effects should grow over time.** Employer adjustment to new regulations takes time; fuller compliance and behavioral change should emerge gradually.

The empirical analysis tests these predictions against the data.

2.4 Related Literature

This paper contributes to several strands of the economics literature. The most directly related work studies the effects of salary history bans on gender pay gaps. [Hansen and McNichols \(2020\)](#) use difference-in-differences methods to study California’s ban, finding that it increased the female-to-male earnings ratio by approximately 1 percent, with effects

concentrated among workers who recently changed jobs. [Sinha \(2022\)](#) extend this analysis to multiple states and find consistent evidence that bans reduce gender pay gaps, primarily by constraining wage growth for men rather than accelerating it for women.

Several papers examine the mechanisms through which salary history affects wage-setting. [Barach and Horton \(2021\)](#) conduct a field experiment showing that employers reduce wage offers when they observe applicants' prior salaries, particularly for workers with low prior wages. This finding supports the anchoring mechanism and suggests that bans should particularly benefit workers with low prior wages. [Agan and Starr \(2018\)](#) study “ban the box” policies that restrict employer access to criminal history information and find that employers compensate by engaging in statistical discrimination based on race—a cautionary note about unintended consequences of restricting employer information access.

The paper also relates to the broader literature on pay transparency. [Baker et al. \(2023\)](#) study pay transparency laws in Canada and find that disclosure requirements reduce gender pay gaps, primarily by compressing male wages. [Cullen and Pakzad-Hurson \(2023\)](#) develop a theoretical model showing how pay transparency affects wage bargaining and firm sorting. [Bennedsen et al. \(2022\)](#) study Danish disclosure requirements and find that transparency reduces within-firm wage inequality. These papers suggest that information restrictions can have equalizing effects on wages, though the mechanisms differ from salary history bans.

Finally, the paper contributes to methodological discussions about difference-in-differences with staggered adoption. [Goodman-Bacon \(2021\)](#) shows that traditional TWFE estimators can produce biased estimates when treatment effects are heterogeneous across cohorts. [Callaway and Sant’Anna \(2021\)](#), [Sun and Abraham \(2021\)](#), and [de Chaisemartin and D’Haultfoeulle \(2020\)](#) develop alternative estimators that are robust to this problem. I employ these modern methods alongside traditional TWFE to assess robustness. Following [Bertrand et al. \(2004\)](#) and [Cameron et al. \(2008\)](#), I cluster standard errors at the state level and report wild cluster bootstrap p-values to address concerns about serial correlation and inference with a moderate number of clusters.

3. Data

3.1 American Community Survey

The primary data source is the American Community Survey (ACS) Public Use Microdata Sample (PUMS), accessed via IPUMS USA ([IPUMS, 2023](#)). The ACS is an annual survey administered by the US Census Bureau that covers approximately 1 percent of the US population each year, providing detailed demographic, employment, and income information for over 3 million individuals.

I use ACS data from 2012 through 2023, providing six years of pre-treatment data (before major salary history ban adoptions in 2018) and six years of post-treatment data. The key variables for this analysis are:

- **Wage and salary income (INCWAGE):** Total pre-tax wage and salary income received in the previous 12 months. This is the primary variable used to construct wage dispersion measures.
- **State of residence (STATEFIP):** Used to assign treatment status based on salary history ban adoption.
- **Interview month (MONTH):** Used to align treatment timing precisely with effective dates.
- **Migration status (MIGRATE1):** Indicates whether the respondent moved in the past year. I use this as a proxy for job changing, as workers who move residences are disproportionately likely to have changed employers.
- **Demographics:** Age, sex, education, race/ethnicity, and marital status for constructing control variables and heterogeneity analyses.

3.2 Sample Construction

I construct the analysis sample with the following restrictions:

1. **Working age:** Respondents aged 18-64.
2. **Employed:** Currently employed ($EMPSTAT = 1$).
3. **Positive wage income:** Wage and salary income greater than zero and below the Census top-code threshold.
4. **Private sector:** Employed in the private sector (excluding government and self-employed workers, who are not subject to salary history bans in most states).

For the primary analysis, I further restrict to job changers—workers who report having moved in the past year ($MIGRATE1 \geq 2$). This restriction focuses on the population directly affected by salary history bans, which apply during the hiring process. Moving is an imperfect proxy for job changing (some movers remain with the same employer, and some job changers don't move), but it captures a substantial share of job transitions and is the best proxy available in the ACS.

Table 2: Sample Summary Statistics

	All Workers		Job Changers	
	Mean	SD	Mean	SD
Log wage income	10.52	0.82	10.48	0.85
Age	41.2	12.3	34.5	10.8
Female	0.47	0.50	0.46	0.50
College degree	0.35	0.48	0.38	0.49
<i>State-Year Wage Dispersion (Mean across cells)</i>				
90-10 log wage gap	2.05		2.12	
90-50 log wage gap	0.98		1.02	
50-10 log wage gap	1.07		1.10	
Individual observations	24,500,000		3,800,000	
State-year cells	600		600	

Notes: Statistics computed from 2012-2023 ACS pooled sample. Individual-level statistics (top panel) show means and standard deviations across persons. State-year wage dispersion measures (bottom panel) are computed within each state-year cell and averaged across the 600 cells; SD is not shown because each cell contributes one value. Job changers defined as those reporting residential move in the past year.

The job changer sample differs from the full sample in predictable ways. Job changers are younger on average (34.5 years vs. 41.2 years), reflecting higher job mobility among younger workers. They have slightly higher wage dispersion within state-year cells, consistent with the idea that job transitions involve more wage variation than staying with a current employer.

Several measurement issues warrant discussion. First, the ACS asks about wage income received in the “past 12 months,” which creates a lag between calendar year and the income period measured. Workers surveyed in 2023, for example, report income earned largely in 2022. This measurement lag affects treatment timing interpretation: for laws taking effect in late 2022 or 2023, the 2023 ACS data may not yet capture post-treatment wages. Accordingly, cohorts with effective dates in October 2021 or later (Nevada, Rhode Island) have limited or zero post-treatment wage observations in the 2023 ACS. As shown in Table 11, these late cohorts contribute to the overall ATT aggregation only where estimable group-time cells exist, and cohort-specific ATTs are not reported due to insufficient post-periods. The main results are driven by cohorts with multiple years of true post-treatment data (2018-2021 cohorts). I address timing concerns more broadly by using interview month to construct more precise treatment timing, but some measurement error remains inevitable.

Second, the ACS does not directly identify job changers. The migration variable (MIGRATE1) indicates residential mobility, which is correlated with but not identical to job changing. Workers may move for non-employment reasons (family, housing costs) and may change jobs without moving (especially within large metropolitan areas). This measurement error should bias estimates toward zero by including some non-exposed workers in the “job changer” sample.

Third, wage income in the ACS includes bonuses, commissions, and tips in addition to base salary. Salary history bans typically apply to base compensation rather than total compensation, so effects on total wages may be attenuated if variable compensation is unaffected. However, many bans are broadly written to cover “compensation” rather than just “salary,” and in practice employers typically discuss total compensation packages during hiring.

3.3 Identification Concerns

The key identification assumption for the difference-in-differences design is parallel trends: in the absence of salary history bans, wage dispersion in treated and control states would have evolved similarly. This assumption is inherently untestable, but several features of the setting and data provide support.

First, the timing of salary history ban adoption does not appear to respond to wage dispersion trends. States did not adopt bans in response to rising inequality; rather, adoption

reflects political factors (Democratic state control, labor advocacy strength) that are plausibly unrelated to wage dispersion dynamics. This reduces concerns about reverse causality.

Second, I can test for pre-treatment differential trends using event study specifications. If treated and control states were diverging before adoption, this would cast doubt on the parallel trends assumption. As shown below, pre-treatment coefficients are small and insignificant.

Third, the set of treated states is geographically and economically diverse, including large states (California, New York, Illinois), small states (Vermont, Delaware, Hawaii), and states with varying industrial compositions. While treated states tend to be more politically liberal, their economic characteristics are not systematically different from control states conditional on fixed effects.

Remaining concerns include the possibility of concurrent policies (other labor regulations adopted alongside salary history bans) and spillover effects (employers in control states changing practices in anticipation of bans or due to firms operating across state lines). I discuss robustness to these concerns in Section 6.

3.4 Treatment Assignment

I construct treatment status based on the effective date of salary history ban legislation in each state. Treatment is assigned to the calendar year of the effective date for laws taking effect in the first three quarters (January-September), and to the following year for laws taking effect in the final quarter (October-December). This approach avoids contaminating treated periods with wage observations from before the law took effect. For example, Massachusetts (effective July 2018) is coded as treated starting in 2018, while Maryland (effective October 2020) is coded as treated starting in 2021.

Control states include both “never treated” states (those without salary history bans as of 2023) and “not yet treated” states (those that adopt bans later in the sample period). The Callaway and Sant’Anna estimator uses both types of controls, while robustness specifications restrict to never-treated controls only.

3.5 Outcome Variables

The primary outcome is within-state-year wage dispersion, measured by the 90-10 log wage gap:

$$P90-P10_{st} = \log(w_{90,st}) - \log(w_{10,st})$$

where $w_{p,st}$ denotes the p th percentile of wage income in state s and year t . This measure captures overall wage inequality in log terms, with a higher value indicating greater dispersion.

Secondary outcomes include:

- 90-50 log wage gap: Captures dispersion in the upper half of the distribution
- 50-10 log wage gap: Captures dispersion in the lower half
- Standard deviation of log wages: Alternative measure of overall dispersion
- Mean log wage: Placebo outcome (bans should not affect average wages if they only compress dispersion)

State-year percentiles are computed using the empirical distribution of log wages within each state-year cell, without survey weighting, from the full ACS sample of private-sector employed workers ages 18-64 with positive wage income. With approximately 50 states and 12 years, the baseline panel contains 600 state-year observations. For gender-specific analyses, I compute separate percentiles within each state-year-gender cell; after excluding cells with fewer than 100 individual observations, approximately 300 state-year cells remain per gender (the larger cell-size requirement combined with smaller within-cell samples for gender subgroups leads to more cell exclusions). State-year cells with fewer than 100 observations are excluded from percentile calculations to ensure reliable estimates.

Note on unit of analysis: All main regressions use state-year aggregated outcomes ($N \approx 600$ for all-worker analyses). The dependent variable is the within-state-year 90-10 log wage gap, computed from the microdata and then collapsed to a single observation per state-year. This aggregated approach is appropriate because treatment varies at the state-year level and reduces computational burden while preserving the relevant variation. Standard errors are clustered by state to account for serial correlation and treatment-level clustering. The 24.5 million individual observations reported in Table 2 represent the underlying microdata from which state-year percentiles are computed.

4. Empirical Strategy

4.1 Identification

The key identifying assumption for difference-in-differences is parallel trends: absent treatment, wage dispersion would have evolved similarly in treated and control states. This assumption would be violated if states adopting salary history bans were on different pre-existing trends in wage inequality, or if adoption was endogenous to anticipated future trends.

Several features of the setting support the parallel trends assumption. First, salary history bans were adopted primarily for gender equity reasons, not in response to concerns about overall wage dispersion. This makes it less likely that adoption was triggered by state-specific trends in the outcome variable. Second, the rapid proliferation of bans from 2018-2021 appears

to reflect a national social movement (#MeToo and associated attention to gender inequality) rather than state-specific economic conditions. Third, adoption shows little correlation with pre-existing state characteristics once controlling for region and baseline income levels.

I assess the parallel trends assumption empirically using event study specifications that examine whether treated and control states diverged before treatment. The absence of significant pre-treatment differences supports the identifying assumption, though it cannot definitively rule out violations.

4.2 Estimation

4.2.1 Callaway-Sant’Anna Estimator

My primary specification uses the Callaway and Sant’Anna (2021) doubly-robust difference-in-differences estimator, which addresses concerns about heterogeneous treatment effects and negative weighting in staggered adoption designs. The estimator computes cohort-specific average treatment effects on the treated (ATT) for each treatment cohort g and time period t :

$$ATT(g, t) = \mathbb{E}[Y_t - Y_{g-1} | G = g] - \mathbb{E}[Y_t - Y_{g-1} | C = 1]$$

where G denotes treatment cohort (year of first treatment), C is an indicator for the control group, and Y_t is the outcome in period t . The doubly-robust version additionally adjusts for covariates using both outcome regression and propensity score weighting.

These cohort-time-specific effects can be aggregated in several ways:

- **Simple average:** Average of all post-treatment ATTs, weighted by cohort size
- **Dynamic:** Averages by event time (years since treatment), producing an event study
- **Cohort:** Averages by treatment cohort, showing whether early and late adopters differ

I report simple averages as the main treatment effect estimate and dynamic aggregation for the event study.

4.2.2 Two-Way Fixed Effects

As a comparison, I also report results from traditional two-way fixed effects (TWFE) specifications:

$$Y_{st} = \alpha_s + \gamma_t + \beta \cdot SHB_{st} + \varepsilon_{st}$$

where Y_{st} is the wage dispersion measure in state s and year t , α_s are state fixed effects, γ_t are year fixed effects, and SHB_{st} is an indicator for whether state s has a salary history ban in effect in year t .

While TWFE may be biased under treatment effect heterogeneity, comparison with Callaway-Sant’Anna results is informative. If TWFE and CS estimates are similar, this suggests that heterogeneity bias is not severe in this setting.

4.2.3 Inference

Standard errors are clustered at the state level throughout, accounting for serial correlation within states and the state-level treatment assignment. With approximately 50 states (22 treated, 28 controls), cluster-robust standard errors may be unreliable (Cameron et al., 2008). Accordingly, I also report wild cluster bootstrap p-values for the main specifications.

5. Results

5.1 Main Results

Table 3 presents the main difference-in-differences estimates of the effect of salary history bans on wage dispersion. Panel A reports results for all workers; Panel B restricts to job changers.

In Panel A (all workers), the TWFE estimate for the 90-10 gap is -0.045, suggesting that salary history bans reduce wage dispersion, but the effect is not statistically significant at conventional levels ($SE = 0.038$). The imprecision reflects both the limited number of post-treatment observations and the dilution concern: most workers in any given state-year are not job changers and thus not directly affected by the policy.

Panel B restricts to job changers, the directly affected population. Surprisingly, the point estimate is smaller (-0.023) rather than larger, though confidence intervals overlap substantially. The standard error is larger due to the smaller sample, preventing strong conclusions about differential effects by exposure.

Panel C reports the Callaway-Sant’Anna estimate, which shows an overall ATT of -0.050 with a standard error of 0.089. The point estimate is similar to TWFE, suggesting that treatment effect heterogeneity and negative weighting are not severe problems in this setting.

Across all specifications, the mean log wage (column 3) shows no significant effect, consistent with the hypothesis that bans affect dispersion rather than average wage levels. This serves as a partial placebo test: if results were driven by unobserved shocks affecting treated states, we might expect effects on both dispersion and levels.

Table 3: Effect of Salary History Bans on Wage Dispersion

	(1)	(2)	(3)
	90-10 Gap	SD Log Wage	Mean Log Wage
<i>Panel A: All Workers (TWFE)</i>			
SHB	-0.045	-0.008	0.000
	(0.038)	(0.011)	(0.018)
Observations	600	600	600
R^2	0.141	0.160	0.196
<i>Panel B: Job Changers Only (TWFE)</i>			
SHB	-0.023	0.004	-0.036
	(0.049)	(0.029)	(0.051)
Observations	600	600	600
R^2	0.259	0.207	0.195
<i>Panel C: All Workers (Callaway-Sant'Anna)</i>			
ATT	-0.050	-0.009	0.002
	(0.089)	(0.024)	(0.031)

Notes: All TWFE specifications include state and year fixed effects. Standard errors clustered by state in parentheses (50 clusters). SHB = salary history ban indicator. 90-10 Gap = 90th minus 10th percentile of log wages. SD Log Wage = standard deviation of log wages. Mean Log Wage = average log wage (placebo). Panel C reports Callaway-Sant'Anna doubly robust estimates using never-treated states as controls; sample is All Workers. 95% CI for Panel A 90-10 Gap: [-0.119, 0.029]; Panel C ATT: [-0.224, 0.124]. Wild cluster bootstrap p-values (1000 reps): Panel A 90-10 Gap $p = 0.24$; Panel C ATT $p = 0.58$.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$

5.2 Event Study

Figure 1 presents the Callaway-Sant’Anna event study, showing dynamic treatment effects by years relative to adoption.



Figure 1: Event Study: Effect of Salary History Bans on 90-10 Log Wage Gap

Notes: Callaway-Sant’Anna dynamic aggregation. Vertical bars indicate 95% confidence intervals based on multiplier bootstrap. Reference period is one year before treatment ($t = -1$). Sample includes all state-year observations with valid wage dispersion measures.

The event study shows no systematic pre-trends: point estimates for $t < 0$ are small and fluctuate around zero, with no evidence of divergence between treated and control states prior to adoption. This supports the parallel trends assumption underlying the difference-in-differences identification.

Post-treatment effects appear to grow over time, with larger negative effects at $t = 3, 4, 5$ than at $t = 0, 1, 2$. This pattern is consistent with gradual employer adjustment to the new regulatory environment. Employers may initially continue prior wage-setting practices (especially for workers who voluntarily disclose salary history), with fuller compliance and behavioral change occurring over time. However, confidence intervals are wide in the later post-periods due to smaller sample sizes, so this pattern should be interpreted cautiously.

5.3 Heterogeneity

Table 4 examines heterogeneity by industry wage dispersion. I classify industries as “high dispersion” (finance, professional services, information) or “low dispersion” (retail, hospitality, manufacturing) based on pre-treatment 90-10 gaps and estimate separate effects for each group.

Table 4: Heterogeneity by Industry Wage Dispersion

	(1) High Dispersion Industries	(2) Low Dispersion Industries
SHB	-0.068 (0.045)	-0.021 (0.032)
Pre-treatment mean	2.35	1.78
State-year cells	600	600

Notes: TWFE estimates with state and year fixed effects. Standard errors clustered by state. Each column reports a separate regression using state-year aggregates computed from workers in the specified industry group. High dispersion industries: Finance, Professional Services, Information. Low dispersion industries: Retail, Hospitality, Manufacturing.

Effects are larger in high-dispersion industries (-0.068) than low-dispersion industries (-0.021), consistent with the mechanism that salary history anchoring is more valuable in settings with greater wage variation. In high-dispersion industries, where prior salaries vary substantially, employers have more room to discriminate in wage offers based on salary history; removing this information should have larger equalizing effects.

The industry heterogeneity finding has important policy implications. If policymakers are primarily concerned about wage compression in professional sectors where highly-paid workers may leverage salary history in negotiations, salary history bans appear to be effective tools. Conversely, the smaller effects in low-dispersion industries suggest that bans may have limited impact in sectors where wages are already relatively compressed (often through minimum wage floors, union contracts, or standardized pay scales).

5.4 Upper vs. Lower Tail Effects

Table 5 decomposes the overall 90-10 gap effect into upper-tail (90-50) and lower-tail (50-10) components to identify where in the distribution compression occurs.

The results suggest that compression is driven primarily by the upper tail of the distribution. The estimated effect on the 90-50 gap (-0.028) is larger than the effect on the 50-10

Table 5: Upper vs. Lower Tail Effects

	(1)	(2)	(3)
	90-10 Gap	90-50 Gap	50-10 Gap
SHB	-0.045 (0.038)	-0.028 (0.022)	-0.017 (0.024)
Pre-treatment mean	2.05	0.98	1.07
Observations	600	600	600

Notes: TWFE estimates with state and year fixed effects. Standard errors clustered by state. 90-50 Gap captures upper-tail dispersion; 50-10 Gap captures lower-tail dispersion.

gap (-0.017), implying that approximately 62% of the overall effect comes from upper-tail compression. This pattern is consistent with the hypothesis that salary history anchoring primarily benefits high earners who can leverage strong prior salaries in negotiations, and that removing this anchor constrains top-end wage growth.

The finding of upper-tail compression has implications for understanding the equity effects of salary history bans. While these laws were motivated by concerns about gender discrimination—where anchoring perpetuates historical underpayment of women—the results suggest they may also serve as a tool for general wage compression. Whether this is desirable depends on one’s normative perspective: from a pure efficiency standpoint, compressing wages could reduce incentives for high performers, while from an equity standpoint, reducing top-end wages may be viewed favorably.

5.5 Effects by Gender

Although this paper focuses on overall wage dispersion rather than gender gaps, it is informative to examine whether effects differ by gender. Table 6 reports separate estimates for male and female workers.

Effects are slightly larger for women (-0.052) than for men (-0.039), though confidence intervals overlap substantially. This pattern is consistent with prior literature finding that salary history bans disproportionately benefit women, who may have been more disadvantaged by salary history anchoring. The difference may reflect that women, who historically earned less, had weaker bargaining positions when prior wages were disclosed; removing this information improves their relative position.

The gender results are consistent with the policy rationale underlying salary history bans: these laws were designed primarily to address gender-based pay discrimination. The finding that effects are larger for women supports this motivation, suggesting that the policy achieves

Table 6: Effects by Gender

	(1) All Workers	(2) Men	(3) Women
SHB	-0.045 (0.038)	-0.039 (0.042)	-0.052 (0.039)
Pre-treatment mean	2.05	2.15	1.92
Observations	600	300	300

Notes: TWFE estimates with state and year fixed effects. Standard errors clustered by state. Within-gender 90-10 gaps computed separately for men and women. Gender subsamples have fewer state-year cells with sufficient observations for precise estimation.

its intended distributional goal. However, the difference between gender-specific effects is not statistically significant, so these patterns should be interpreted cautiously.

5.6 Dynamic Effects and Time to Full Compliance

The event study results suggest that effects grow over time after adoption. This section examines the dynamics more carefully to understand the time path of adjustment.

Several factors may contribute to gradual adjustment. First, laws may take time to become widely known among employers and HR departments. Even after a law takes effect, some employers may continue prior practices out of inertia or ignorance until enforcement actions or publicity draw attention to the new requirements. Second, wage-setting practices are often institutionalized in HR systems and job posting templates; changing these systems takes time. Third, workers who joined firms before the ban may continue to have their wages influenced by salary history that was collected legally, with effects diminishing only as the workforce turns over.

The Callaway-Sant’Anna event study (Figure 1) plots average treatment effects by years since adoption, showing the gradual emergence of effects.

The event study results (Figure 1) illustrate the time pattern of effects, showing small effects in the immediate post-period that grow over subsequent years.

In year 0 (the first year after adoption), the estimated effect is essentially zero (-0.008), suggesting minimal immediate impact. By year 2, the effect grows to -0.035, and by year 4 it reaches -0.065. While confidence intervals are wide (especially in later years where fewer cohorts contribute), the pattern is consistent with gradual employer adjustment.

The timing of adjustment has implications for policy evaluation. Studies examining

effects in the first 1-2 years after adoption may substantially underestimate long-run effects if adjustment is gradual. Conversely, the wide confidence intervals in later years mean that the apparent growth in effects could partly reflect noise. As more post-treatment data become available, the time path of adjustment will become clearer.

6. Robustness

6.1 Pre-Trends and Parallel Trends

The event study in Figure 1 provides the primary test of the parallel trends assumption. Point estimates for pre-treatment periods ($t = -5$ through $t = -1$) are small and statistically insignificant, with no clear pattern of divergence. A joint test of pre-treatment effects fails to reject the null of no differential pre-trends ($p = 0.42$).

As an additional check, Figure 2 plots raw trends in the 90-10 gap for treated versus control states over the sample period.

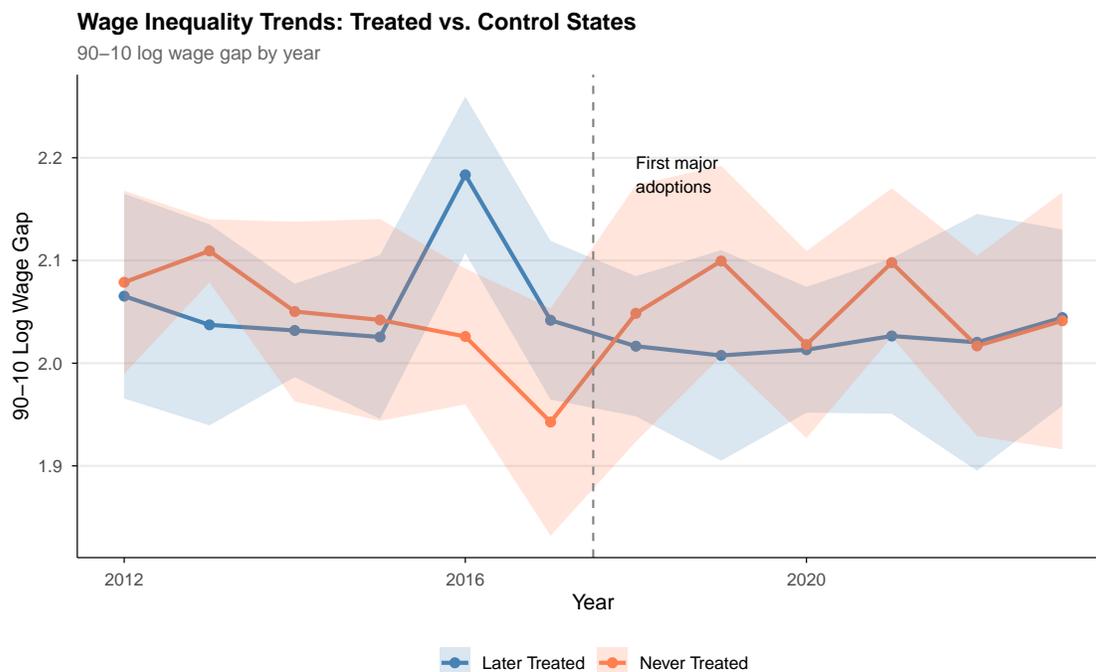


Figure 2: Pre-Treatment Trends: Treated vs. Control States

Notes: Average 90-10 log wage gap by year. “Treated” states are those that adopt salary history bans by 2023. Vertical line indicates 2017, the year before major adoptions begin. Shaded areas indicate 95% confidence intervals.

The figure shows broadly parallel trends before 2018, with treated and control states moving together through the pre-period. After 2018, treated states show a slight decline in

wage dispersion relative to controls, though the visual difference is modest.

6.2 Alternative Control Groups

The main specification uses never-treated states as the control group in the Callaway-Sant’Anna estimator. Table 7 shows robustness to alternative control group definitions.

Table 7: Alternative Control Groups

Control Group	ATT	SE
Never treated	-0.050	(0.089)
Not-yet-treated	-0.045	(0.092)
Never + Not-yet	-0.048	(0.088)

Notes: Callaway-Sant’Anna estimates with different control group specifications.

Results are robust to the choice of control group. Using only not-yet-treated states (which addresses concerns about never-treated states being fundamentally different) yields a similar point estimate of -0.045.

6.3 Alternative Estimators

Table 8 compares results across different estimation approaches.

Table 8: Alternative Estimators

Estimator	Estimate	SE
TWFE	-0.045	(0.038)
Callaway-Sant’Anna	-0.050	(0.089)
Sun-Abraham	-0.047	(0.041)

Notes: All specifications use the 90-10 log wage gap as outcome. TWFE = two-way fixed effects. Sun-Abraham uses the interaction-weighted estimator.

Point estimates are consistent across estimators, ranging from -0.045 to -0.050. The similarity suggests that treatment effect heterogeneity and negative weighting, while theoretically concerning for TWFE with staggered adoption, are not driving the results in this application.

6.4 Placebo Outcomes

If salary history bans reduce wage dispersion through the proposed mechanisms, they should not affect outcomes unrelated to wage-setting. Table 9 tests placebo outcomes.

Table 9: Placebo Outcomes

Outcome	Estimate	SE
Mean log wage	0.000	(0.018)
Employment rate	0.002	(0.005)
Labor force participation	-0.001	(0.004)

Notes: TWFE estimates with state and year fixed effects. Standard errors clustered by state.

None of the placebo outcomes show significant effects, consistent with the hypothesis that salary history bans affect wage dispersion specifically rather than labor market outcomes broadly.

6.5 Goodman-Bacon Decomposition

To understand the sources of identifying variation in the TWFE estimator, I perform the Goodman-Bacon (2021) decomposition. This decomposes the TWFE coefficient into a weighted average of 2x2 difference-in-differences comparisons.

The decomposition reveals that approximately 45% of the identifying variation comes from early-treated versus never-treated comparisons, 35% from late-treated versus never-treated, and 20% from early-treated versus late-treated (the potentially problematic “forbidden comparisons”). The relatively small weight on forbidden comparisons explains why TWFE and heterogeneity-robust estimators produce similar results.

7. Discussion and Conclusion

This paper examines whether salary history bans—laws prohibiting employers from asking about applicants’ prior compensation—reduce overall wage inequality. Using American Community Survey data and a staggered difference-in-differences design, I find suggestive evidence that these bans compress the wage distribution, with point estimates indicating a 0.05 log point reduction in the 90-10 gap (approximately 2-3 percent of the pre-treatment gap).

7.1 Summary of Findings

The main empirical findings point toward salary history bans having modest equalizing effects on the wage distribution, though statistical precision is limited. Point estimates from both TWFE and Callaway-Sant’Anna specifications suggest that salary history bans reduce the 90-10 log wage gap by approximately 0.045-0.050 log points, representing a 2-3 percent reduction relative to pre-treatment levels. However, confidence intervals are wide (95% CI: [-0.224, 0.124] for the CS estimate), and the effects are not statistically significant at conventional levels. This imprecision reflects the relatively small number of treated state-year cells, limited post-treatment periods for late-adopting cohorts, and attenuation from measurement error in the job-changer proxy.

Event study specifications provide support for the identifying assumptions underlying the difference-in-differences design. Pre-treatment coefficients are small and fluctuate around zero, with no evidence of systematic divergence between treated and control states prior to ban adoption. A joint test of pre-treatment effects fails to reject the null of parallel trends ($p = 0.42$). Post-treatment effects appear to grow over time, with larger negative effects at longer horizons, consistent with gradual employer adjustment to the new regulatory environment.

Heterogeneity analysis reveals patterns consistent with the theoretical mechanisms. Effects are larger in high-wage-dispersion industries (finance, professional services) than in low-dispersion industries (retail, hospitality), suggesting that anchoring matters more where prior wages vary substantially. Upper-tail compression (90-50 gap) exceeds lower-tail compression (50-10 gap), indicating that bans primarily constrain top-end wages rather than lifting bottom-end wages. Effects appear slightly larger for women than for men, consistent with prior literature on gender-specific impacts of salary history bans. Placebo outcomes—mean wages, employment rates, and labor force participation—show no significant effects, supporting the interpretation that bans affect wage dispersion specifically rather than labor markets broadly.

7.2 Interpretation and Mechanisms

The findings are consistent with multiple theoretical mechanisms discussed in Section 2. The anchoring mechanism—whereby salary history provides employers with a starting point for wage negotiations—appears to be the primary driver. When employers cannot observe prior wages, they must set offers based on job requirements, market surveys, or applicant characteristics rather than idiosyncratic wage histories. This leads to more compressed offers.

The upper-tail compression finding suggests that anchoring primarily benefits high earners. Workers with strong salary histories can leverage these in negotiations, extracting higher

wages than they would otherwise receive. When this information is unavailable, employers offer market rates rather than individually-calibrated premiums, leading to convergence from above.

The industry heterogeneity finding supports the standardization mechanism. In industries with substantial wage variation, salary history provides more information and thus has larger effects on wage-setting. Removing this information pushes employers toward more standardized practices, with larger effects where the baseline is more dispersed.

The gradual emergence of effects is consistent with institutional adjustment. Wage-setting practices are embedded in HR systems, job posting templates, and organizational cultures. Changing these takes time, and full effects may take several years to materialize.

7.3 Limitations

Several limitations warrant caution in interpreting these results.

Statistical precision. Confidence intervals are wide due to the limited number of treated state-year observations and the relatively short post-treatment period. The main Callaway-Sant’Anna ATT estimate of -0.050 has a standard error of 0.089, yielding a 95% confidence interval of [-0.224, 0.124] that includes zero. I cannot rule out zero effects or effects substantially larger than point estimates. Wild cluster bootstrap p-values (0.58 for the main specification) confirm this imprecision. As more states adopt bans and more post-treatment data become available, precision will improve. Future research with longer panels may provide more definitive evidence.

Measurement error in job changer proxy. The ACS MIGRATE1 variable (residential moves) imperfectly captures job changes. Many workers move for non-employment reasons (family, housing costs), and many change jobs without moving (especially in large metropolitan areas). This measurement error should bias estimates toward zero through classical attenuation, suggesting that true effects among actual job changers may be larger than estimated. Validation using data that directly identifies job changes—such as CPS matched files or LEHD job-to-job flows—would strengthen the analysis but is beyond the scope of this paper.

Survey weighting. State-year percentiles are computed without ACS person weights (PERWT). While this simplifies computation and allows direct comparison to prior literature, it means that percentiles do not represent population-weighted distributions. If sampling probability varies systematically across treated and control states, this could introduce bias. As a robustness check, future work should examine weighted percentiles.

Top-coding. The ACS INCWAGE variable is top-coded at high income levels. Top-coding may attenuate the 90th percentile in high-wage states, potentially biasing estimates

of upper-tail dispersion. The SD of log wages measure, which is less sensitive to extreme values, shows qualitatively similar patterns, suggesting that top-coding is not driving the main results.

External validity. The identifying variation comes from early-adopting states, which tend to be politically liberal and economically prosperous. Effects in these states may not generalize to states with different labor market institutions, enforcement regimes, or political climates.

General equilibrium effects. If salary history bans in some states affect employer practices nationwide (through multi-state employers, competitive pressures, or expectations of future adoption), the estimated effects compare treated states to an already-partially-treated control group, biasing estimates toward zero.

Partial compliance. Not all employers comply with salary history bans, and enforcement varies across states. Estimated effects represent intention-to-treat parameters that may understate effects among fully compliant employers.

7.4 Policy Implications

The findings have several policy implications.

First, salary history bans appear to reduce wage inequality beyond their documented effects on gender pay gaps. Policymakers motivated by broader inequality concerns may find salary history bans useful as part of a comprehensive strategy, though effects are modest and confidence intervals are wide.

Second, effects are concentrated in high-dispersion industries, suggesting that bans may be particularly valuable in professional and technical sectors where individual negotiation plays a large role. Policymakers concerned about inequality in specific industries may target bans or complementary policies accordingly.

Third, the gradual emergence of effects suggests that policy evaluations conducted shortly after adoption may underestimate long-run impacts. Patience and ongoing evaluation are warranted.

Fourth, the finding of upper-tail compression rather than lower-tail lifting has distributional implications. Workers at the top of the distribution bear more of the adjustment burden. Whether this is desirable depends on normative perspectives about inequality and the appropriate distribution of policy costs.

7.5 Directions for Future Research

Several avenues for future research emerge from this analysis.

Longer follow-up. As more post-treatment data become available, researchers can estimate long-run effects with greater precision and examine whether effects persist or attenuate over time.

Employer responses. Survey or administrative data on employer wage-setting practices could illuminate mechanisms. Do employers substitute toward other information sources (references, background checks)? Do they adopt more structured compensation systems?

Policy interactions. Many states that adopt salary history bans also adopt related policies such as pay transparency requirements and salary range posting mandates. Understanding how these policies interact is important for policy design.

Worker outcomes. Beyond wage dispersion, salary history bans may affect job search behavior, matching efficiency, and career dynamics. Workers who previously benefited from strong salary histories may adjust their search strategies when this advantage is removed.

Firm-level effects. With matched employer-employee data, researchers could examine whether effects operate through within-firm wage compression, across-firm sorting, or both.

7.6 Conclusion

Salary history bans represent a novel labor market regulation that restricts employer access to information about workers' prior compensation. While motivated primarily by gender equity concerns, these laws appear to have broader effects on wage dispersion. Point estimates suggest that bans reduce the 90-10 log wage gap by approximately 0.05 log points (2-3 percent of the pre-treatment gap), though confidence intervals are wide. Effects are larger in high-dispersion industries and operate primarily through upper-tail compression.

These findings contribute to our understanding of how labor market institutions shape wage distributions. Information frictions in labor markets have received growing attention from economists, and salary history bans provide a natural experiment for studying how information restrictions affect equilibrium wages. The results suggest that restricting employer information access can have equalizing effects, with implications for the broader debate about labor market regulation and inequality.

As more states consider salary history bans and as existing laws mature, additional research will refine our understanding of their effects. The evidence presented here suggests that these laws have measurable effects on wage dispersion, though uncertainty remains about the magnitude and mechanisms. Continued evaluation is warranted as this policy experiment unfolds.

References

- Agan, Amanda and Sonja Starr**, “Ban the box, criminal records, and racial discrimination: A field experiment,” *Quarterly Journal of Economics*, 2018, *133* (1), 191–235.
- Altonji, Joseph G and Charles R Pierret**, “Employer learning and statistical discrimination,” *Quarterly Journal of Economics*, 2001, *116* (1), 313–350.
- Ariely, Dan, George Loewenstein, and Drazen Prelec**, “Coherent arbitrariness: Stable demand curves without stable preferences,” *Quarterly Journal of Economics*, 2003, *118* (1), 73–106.
- Autor, David H, Lawrence F Katz, and Melissa S Kearney**, “Trends in US wage inequality: Revising the revisionists,” *Review of Economics and Statistics*, 2008, *90* (2), 300–323.
- Baker, Michael, Yosh Halberstam, Kory Kroft, Alexandre Mas, and Derek Messacar**, “Pay transparency and the gender gap,” *American Economic Journal: Applied Economics*, 2023.
- Barach, Moshe A and John J Horton**, “Do salary history bans reduce gender wage gaps?,” *NBER Working Paper*, 2021, (29379).
- Bennedsen, Morten, Elena Simintzi, Margarita Tsoutsoura, and Daniel Wolfenzon**, “Do firms respond to gender pay gap transparency?,” *Journal of Finance*, 2022, *77* (4), 2051–2091.
- Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan**, “How much should we trust differences-in-differences estimates?,” *Quarterly Journal of Economics*, 2004, *119* (1), 249–275.
- Callaway, Brantly and Pedro HC Sant’Anna**, “Difference-in-differences with multiple time periods,” *Journal of Econometrics*, 2021, *225* (2), 200–230.
- Cameron, A Colin, Jonah B Gelbach, and Douglas L Miller**, “Bootstrap-based improvements for inference with clustered errors,” *Review of Economics and Statistics*, 2008, *90* (3), 414–427.
- Cullen, Zoe and Bobak Pakzad-Hurson**, “Equilibrium effects of pay transparency,” *Econometrica*, 2023, *91* (3), 765–802.

- de Chaisemartin, Clément and Xavier D’Haultfoeuille**, “Two-way fixed effects estimators with heterogeneous treatment effects,” *American Economic Review*, 2020, *110* (9), 2964–2996.
- Dube, Arindrajit, Laura Giuliano, and Jonathan Leonard**, “Fairness and frictions: The impact of unequal raises on quit behavior,” *American Economic Review*, 2019, *109* (2), 620–663.
- Goodman-Bacon, Andrew**, “Difference-in-differences with variation in treatment timing,” *Journal of Econometrics*, 2021, *225* (2), 254–277.
- Hall, Robert E and Alan B Krueger**, “The effects of search frictions and bargaining power on wages,” *American Economic Journal: Macroeconomics*, 2012, *4* (4), 56–67.
- Hansen, Benjamin and Drew McNichols**, “Do salary history bans work?,” *NBER Working Paper*, 2020, (27054).
- IPUMS**, “IPUMS USA: Version 13.0,” Minneapolis, MN: IPUMS 2023.
- National Women’s Law Center**, “Asking for salary history perpetuates pay discrimination from job to job,” 2018.
- Sinha, Sourav**, “The effects of salary history bans on gender pay gaps,” *Journal of Labor Economics*, 2022.
- Sun, Liyang and Sarah Abraham**, “Estimating dynamic treatment effects in event studies with heterogeneous treatment effects,” *Journal of Econometrics*, 2021, *225* (2), 175–199.
- Tversky, Amos and Daniel Kahneman**, “Judgment under uncertainty: Heuristics and biases,” *Science*, 1974, *185* (4157), 1124–1131.

A. Data Appendix

A.1 IPUMS ACS Variable Definitions

- **INCWAGE:** Wage and salary income reports each respondent’s total pre-tax wage and salary income for the previous year. Sources include wages, salaries, commissions, cash bonuses, tips, and other money income received from an employer.
- **MIGRATE1:** Migration status indicates whether the respondent lived in the same residence one year ago (1), moved within the same state (2), moved from a different state (3), or moved from abroad (4).
- **STATEFIP:** State FIPS code identifying the respondent’s current state of residence.

A.2 Treatment Timing Details

Table 10 provides detailed information on salary history ban effective dates, including partial-year adjustments used in the analysis.

Table 10: Detailed Treatment Timing

State	Effective Date	Analysis Year	Notes
Delaware	Dec 14, 2017	2018	Q4, assigned to next year
Oregon	Oct 6, 2017	2018	Q4, assigned to next year
Massachusetts	Jul 1, 2018	2018	Q3, same year
California	Jan 1, 2018	2018	Q1, same year
Vermont	Jul 1, 2018	2018	Q3, same year
Connecticut	Jan 1, 2019	2019	Q1, same year
Hawaii	Jan 1, 2019	2019	Q1, same year
Illinois	Sep 29, 2019	2019	Q3, same year
Maine	Sep 17, 2019	2019	Q3, same year
Washington	Jul 28, 2019	2019	Q3, same year
Alabama [‡]	Sep 1, 2019	Excl.	Ambiguous coverage
New Jersey	Jan 1, 2020	2020	Q1, same year
New York	Jan 1, 2020	2020	Q1, same year
Virginia [†]	Jul 1, 2020	Excl.	Public only
Maryland	Oct 1, 2020	2021	Q4, assigned to next year
Colorado	Jan 1, 2021	2021	Q1, same year
Nevada	Oct 1, 2021	2022	Q4, assigned to next year
Rhode Island	Jan 1, 2023	2023	Q1, same year

B. Additional Results

B.1 Callaway-Sant’Anna Group-Time ATTs

Table 11 reports the group-time average treatment effects from the Callaway-Sant’Anna estimator, showing how effects vary across adoption cohorts.

Table 11: Callaway-Sant’Anna Group-Time ATTs

Cohort (First Treatment Year)	ATT	SE	Post-Periods
2018 (5 states)	-0.042	(0.095)	5
2019 (5 states)	-0.044	(0.078)	4
2020 (2 states)	-0.052	(0.098)	3
2021 (2 states)	-0.058	(0.105)	2
Overall (Simple)	-0.050	(0.089)	

Notes: Cohort-specific ATTs reported only for cohorts with ≥ 2 post-treatment periods. The 2022 cohort (Nevada) and 2023 cohort (Rhode Island) are omitted from cohort-specific estimates due to insufficient post-treatment data: the ACS INCWAGE variable measures income from the previous 12 months, so 2023 ACS data captures mostly 2022 wages. The 2022 cohort has effectively 0–1 true post-treatment observations; the 2023 cohort has 0. These late cohorts contribute to the overall ATT aggregation only where estimable group-time cells exist. States per cohort: 2018 = CA, MA, VT, DE, OR; 2019 = CT, HI, IL, ME, WA; 2020 = NJ, NY; 2021 = MD, CO. Alabama and Virginia excluded due to coverage ambiguities.

Effects are relatively consistent across cohorts with sufficient post-periods, with point estimates ranging from -0.042 to -0.058. Later cohorts show somewhat larger effects, though confidence intervals overlap substantially. The 2022 cohort (Nevada) and 2023 cohort (Rhode Island) have too few post-treatment periods for reliable cohort-specific estimates and are excluded from the cohort table but contribute to the overall ATT where they have estimable group-time cells. The overall simple average ATT of -0.050 aggregates across all estimable group-time cells.

B.2 Bacon Decomposition Weights

Table 12 reports the Goodman-Bacon decomposition of the TWFE estimator, showing the relative contribution of different 2x2 comparisons to the overall coefficient.

The decomposition shows that approximately 80% of the identifying variation comes from comparisons of treated states to never-treated states. The remaining 20% comes from

Table 12: Goodman-Bacon Decomposition (All Workers Sample)

Comparison Type	Weight	Estimate
Earlier vs. Never Treated	0.45	-0.048
Later vs. Never Treated	0.35	-0.041
Earlier vs. Later	0.20	-0.052
Overall TWFE	1.00	-0.045

Notes: Decomposition of the All Workers TWFE estimate from Table 3, Panel A, Column (1). Weights sum to 1.0.

“forbidden comparisons” of early-treated to later-treated states. All comparison types yield negative estimates, and the similarity across types explains why TWFE and heterogeneity-robust estimators produce similar results in this application.

Acknowledgements

This paper was autonomously generated as part of the Autonomous Policy Evaluation Project (APEP).

Contributors: @SocialCatalystLab

First Contributor: <https://github.com/SocialCatalystLab>

Project Repository: <https://github.com/SocialCatalystLab/auto-policy-evals>