

Can't Ask, Won't Tell: Salary History Bans and the Gender Earnings Gap at Hire

APEP Autonomous Research* @ailscl

March 5, 2026

Abstract

When employers ask “What did you earn before?” they anchor new offers to salaries shaped by historical discrimination. Twenty U.S. jurisdictions banned this question between 2017 and 2024. I exploit staggered adoption in a triple-difference design using universe employer–employee data from the Quarterly Workforce Indicators, which uniquely separates new-hire from continuing-worker earnings. Despite the strong theoretical mechanism, salary history bans produced no detectable effect on the aggregate gender earnings gap. The TWFE estimate for new hires is -0.008 ($SE = 0.008$), and the DDD interaction is -0.002 ($SE = 0.006$)—both precisely estimated zeros. Pre-trends are flat (joint F-test $p = 0.99$), the female hire share is unchanged, and randomization inference confirms the null (RI $p = 0.45$). Results are robust across estimators, industries, and sample restrictions. The bans may have been too narrow, too weakly enforced, or too easily circumvented to move aggregate wages.

JEL Codes: J31, J71, J78, K31

Keywords: salary history ban, gender wage gap, hiring discrimination, pay equity, difference-in-differences

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: N/A).

1. Introduction

In job interviews across America, a single question has long shaped the trajectory of women’s careers: “What is your current salary?” When an employer uses this answer to set a new offer, any prior pay gap—whether caused by discrimination, weak bargaining position, or sorting into lower-paying firms—perpetuates itself into the next job and the next. The salary history question transforms a temporary disadvantage into a permanent anchor.

Starting with Oregon and Delaware in 2017, a wave of state legislation sought to break this chain. By 2024, twenty jurisdictions had banned employers from asking job applicants about their salary history. The laws reflect a simple theory: if employers cannot observe prior compensation, they must set offers based on the job’s value and the candidate’s qualifications, not on a discriminatory artifact of the past. Proponents argued the bans would compress the gender wage gap at the point where it is most malleable—the moment of hire. Opponents countered that employers would simply find other ways to infer prior pay, or that removing salary history would introduce noise and inefficiency into the matching process.

This paper provides the first evaluation of salary history bans using universe-scale employer–employee data that separates new-hire earnings from continuing-worker earnings—the precise distinction needed to identify the mechanism. The Quarterly Workforce Indicators (QWI), constructed from state unemployment insurance records covering approximately 150 million workers, report average monthly earnings for new hires and for continuing (“stable”) employees separately, by state, quarter, sex, and industry. This unique data structure enables a triple-difference (DDD) design that sharpens identification well beyond what existing studies—which rely on the Current Population Survey or job-posting data—can achieve.

The DDD design exploits three margins of variation. First, I compare states that adopted salary history bans to not-yet-treated and never-treated states (cross-state variation). Second, I compare quarters before and after each state’s ban took effect (time variation). Third—and this is the design’s key advantage—I compare the gender earnings gap among *new hires* to the gap among *continuing workers* within the same state and quarter. Salary history bans should compress the new-hire gender gap (where salary history anchors offers) but leave the continuing-worker gap largely unchanged (where wages are set by internal pay structures). This within-state, within-period comparison absorbs any statewide economic shock that might coincide with ban adoption, including the concurrent pay transparency laws that several ban states enacted simultaneously.

I find that salary history bans had no detectable effect on the aggregate gender earnings gap. The TWFE DiD estimate for new-hire earnings is -0.008 log points (SE = 0.008, $p = 0.35$)—a precisely estimated zero. The continuing-worker ratio, which serves as a within-

state placebo, shows a comparably small and insignificant coefficient (-0.006 , $SE = 0.006$). The triple-difference interaction—the differential effect on new hires relative to continuing workers—is -0.002 ($SE = 0.006$, $p = 0.77$), again indistinguishable from zero. Industry-level estimates are uniformly small across male-dominated, female-dominated, and high-wage sectors. The null result is robust to the Callaway–Sant’Anna estimator, Sun–Abraham interaction-weighted estimator, exclusion of states with bundled pay transparency laws, and randomization inference (RI $p = 0.45$).

Several features of the analysis strengthen the interpretation that the null reflects a true absence of aggregate effects, rather than a failure of the research design. First, event-study plots show flat pre-trends in the gender earnings ratio for both new hires and continuing workers, confirming that the parallel trends assumption is credible (joint F-test $p = 0.99$). Second, the sample includes 51 jurisdictions (50 states plus DC), 60 quarters (2010–2024), and 20 treated states—more than sufficient power for the TWFE and DDD specifications, with a minimum detectable effect of approximately 2 percentage points for the DDD. Third, the female share of new hires does not change significantly, ruling out large composition effects. Fourth, randomization inference (500 permutations) confirms that the actual estimate lies near the center of the placebo distribution (RI $p = 0.45$), not in the tail. Fifth, results are unchanged after excluding Colorado, California, and Washington, which bundled salary history bans with pay transparency mandates.

This paper contributes to three literatures. First, it advances the study of gender pay equity policy by providing a credible null result. [Goldin \(2014\)](#) and [Blau and Kahn \(2017\)](#) document that the gender wage gap is remarkably persistent. [Babcock and Laschever \(2003\)](#) and [Bowles et al. \(2007\)](#) show that gendered bargaining norms contribute to initial pay gaps, which [Barron et al. \(2006\)](#) and [Hall and Krueger \(2019\)](#) demonstrate persist through salary-history-anchored job transitions. The few existing evaluations of salary history bans—[Hansen et al. \(2020\)](#) using California CPS data, [Bessen et al. \(2020\)](#) using job postings, and [Sinha \(2023\)](#) using broader CPS samples—find suggestive effects but rely on small survey samples and two-way fixed effects estimators vulnerable to staggered-adoption bias ([Goodman-Bacon, 2021](#); [Sun and Abraham, 2021](#)). Using universe administrative data covering all twenty jurisdictions through 2024, I find that these suggestive effects do not survive a more rigorous design with a built-in mechanism test.

Second, this paper contributes to the growing literature on information regulation in labor markets. A central question is whether restricting the information employers can use in hiring decisions reduces inequality or generates unintended consequences. [Agan and Starr \(2018\)](#) provide the canonical cautionary example: “ban the box” laws, which restrict criminal history inquiries on job applications, increased statistical discrimination against

Black applicants, particularly among employers who previously conducted background checks. The salary history setting offers a theoretically distinct test case. While criminal history is a binary signal about a characteristic of the applicant, salary history is a continuous signal about a *price*—what the market has paid this person. Removing price information should not increase employer uncertainty about productivity in the same way that removing characteristic information does. My null result suggests a different failure mode from [Agan and Starr](#)’s: rather than generating backlash discrimination, the salary history bans appear to be too porous to change the information environment, since salary history can be inferred from observable characteristics, volunteered by candidates, or accessed through alternative channels. The null is informative precisely because the design has sufficient power to detect economically meaningful effects.

Third, the paper illustrates the value of the QWI’s new-hire vs. continuing-worker decomposition for causal inference. By separating earnings at the hiring margin from earnings in ongoing employment, the QWI provides a built-in within-unit placebo for any policy that operates through the hiring process. This design principle—using the *margin of treatment exposure* as the third difference—can be applied to other hiring-stage interventions (credential requirements, degree screens, background checks) in future work.

2. Institutional Background

2.1 The Salary History Question and Its Role in Wage Setting

Salary history inquiries have been a standard part of the American hiring process for decades. Employers use prior compensation as a signal of candidate quality and as an anchor for negotiation ([Kahneman, 1992](#)). In surveys, approximately 80 percent of hiring managers report asking candidates about their salary history at some point during the hiring process ([Society for Human Resource Management, 2018](#)). The practice is particularly common in white-collar and professional occupations, where individual salaries vary widely and negotiation determines a larger share of total compensation.

The mechanism through which salary history perpetuates inequality is straightforward. Suppose a woman begins her career at a firm that pays below-market wages for her position—perhaps due to discriminatory initial placement, weaker bargaining, or sorting into a lower-paying firm ([Card et al., 2016](#)). When she applies for her next job, the new employer asks her current salary and uses it to set the offer. Even if the new employer harbors no gender animus, the offer is anchored to a discriminatory artifact. The gap carries forward through subsequent job changes, compounding over the career. [Hansen et al. \(2020\)](#) call this “the persistence mechanism”: salary history inquiry transforms a *history* of discrimination into a

current discriminatory outcome.

2.2 The Legislative Wave

Oregon and Delaware enacted the first statewide salary history bans covering private employers in October and December 2017, respectively. The District of Columbia followed the same year. California’s ban took effect January 1, 2018, followed by Massachusetts and Vermont at mid-year. The wave accelerated in 2019, when seven additional states (Connecticut, Hawaii, Illinois, Maine, Washington, Alabama, and New Jersey) enacted bans. New York, Virginia, and Maryland adopted bans in 2020, Colorado and Nevada in 2021, Rhode Island in 2023, and Minnesota in 2024—bringing the total to twenty jurisdictions.

The laws share a common core: employers may not ask job applicants about their current or prior compensation, and may not use salary history information to set pay, even if the applicant volunteers it. Most laws impose civil penalties for violations, typically through the state department of labor or attorney general. The bans generally apply to all employers in the state, though some have minimum size thresholds.

Despite this common core, the laws vary substantially in scope, enforcement mechanisms, and interaction with other labor regulations. These differences matter for interpreting the aggregate null result.

Scope of prohibition. Most laws prohibit both direct inquiry (“What is your current salary?”) and the use of salary history information even if voluntarily disclosed. However, several states permit employers to use salary information if the candidate discloses it *without prompting*. Oregon’s law, for example, prohibits employers from seeking salary history but allows them to consider it if the applicant provides it voluntarily. This creates a strategic disclosure environment: high-earning candidates may disclose (to leverage their history), while low-earning candidates stay silent—partially restoring the asymmetric information that the ban sought to eliminate.

Enforcement. Enforcement varies from robust to essentially nonexistent. Some states provide a private right of action, allowing affected workers to sue employers directly. Others rely exclusively on administrative complaints to the state labor department or attorney general. Alabama’s law, notably, contains no enforcement mechanism or penalties, relying entirely on voluntary compliance. The heterogeneity in enforcement strength may attenuate the aggregate treatment effect if the laws are binding in some states but not others.

Coverage. All twenty jurisdictions’ bans cover private-sector employers, which is the relevant population for the QWI analysis. Several laws also cover public-sector employers, state contractors, or recruiting agencies, but these additional provisions do not affect the identification strategy since the QWI captures all UI-covered employment regardless of sector.

Penalties. Where specified, penalties range from \$1,000 per violation (Connecticut) to \$10,000 per violation (New Jersey). Several states allow compensatory damages, back pay, and attorneys’ fees in addition to statutory penalties. The variation in penalty severity provides a potential margin for heterogeneity analysis, though the small number of treated states limits the power of such tests.

A critical feature for identification is that several states enacted salary history bans alongside broader pay equity reforms. California had already strengthened its Equal Pay Act in 2015 (SB 358). Colorado’s 2021 Equal Pay for Equal Work Act included both a salary history ban and a pay transparency requirement (employers must disclose salary ranges in job postings). Washington added pay transparency in 2023. I address this bundling directly in the empirical strategy by (1) including pay transparency indicators as controls, (2) using the DDD design (which absorbs any policy that affects all workers, not just new hires), and (3) robustness checks excluding the three most prominent bundled states.

2.3 Theoretical Predictions

Under the anchoring theory, banning salary history should: (1) *raise* offers to women with low prior wages, compressing the gender gap at hire; (2) *reduce* offers to workers with high prior wages who could have leveraged their history; and (3) have *no effect* on continuing workers whose wages are set by internal structures. The net effect on the gender gap depends on the distribution of prior-wage anchors by gender.

Under an alternative statistical-discrimination theory, removing salary history information could *increase* discrimination if employers substitute to other, noisier signals correlated with gender (Agan and Starr, 2018). In this case, we might observe reduced hiring of women or larger gender gaps among new hires in industries with the most uncertainty about candidate quality.

A third possibility—which the existing literature largely ignores—is that the bans have *no effect at all*. This could occur for several reasons. First, if compliance is imperfect and salary information flows through alternative channels, the ban may be non-binding. Second, if the competitive labor market already prices workers close to their marginal product, salary history may be a redundant signal: employers who underpay women based on salary anchoring would lose them to competing firms that offer market wages. In this view, salary history is a noisy, easily-substituted information source, and removing it does not change the equilibrium wage structure. Third, if the anchoring effect is symmetric by gender—if both men and women are equally anchored to their prior wages, and the prior wage distribution differs by gender for structural rather than informational reasons—then removing the anchor compresses both genders’ wage distributions equally, leaving the gender *ratio* unchanged.

The DDD design and the QWI data can distinguish these predictions empirically. The anchoring theory predicts a positive DDD coefficient (new-hire gender ratio improves relative to continuing workers). The statistical discrimination theory predicts a negative coefficient or a reduction in female hiring. The null hypothesis predicts a DDD coefficient of zero with unchanged hiring composition.

3. Data

3.1 Quarterly Workforce Indicators

The primary data source is the Quarterly Workforce Indicators (QWI), a public-use dataset derived from the Longitudinal Employer-Household Dynamics (LEHD) program at the U.S. Census Bureau. The LEHD covers approximately 95 percent of private-sector jobs through state unemployment insurance wage records, making the QWI effectively a universe of formal employment in the United States.

The QWI reports aggregate statistics at the state–quarter level, stratified by worker sex, age, education, race/ethnicity, and firm industry (2-digit NAICS). For this paper, the key advantage of the QWI over household surveys is its *separate reporting of earnings for new hires and continuing workers*. Specifically, the QWI reports:

- **EarnS**: Average monthly earnings for *full-quarter* (stable) employees—those employed throughout the quarter. This captures continuing workers.
- **EarnHirNS**: Average monthly earnings for *new hires* who achieved stable employment (employed for at least one full subsequent quarter). This captures workers who just started at a new firm.
- **HirN**: Count of new hires (first observed at the firm).
- **Sep**: Count of separations (last observed at the firm).
- **Emp**: Total beginning-of-quarter employment count.

Several features of the QWI make it uniquely suited to this research question. First, because it derives from administrative wage records rather than household surveys, earnings are measured precisely—not self-reported—eliminating the measurement error that plagues CPS-based gender gap estimates. Second, the sample size is enormous: the QWI effectively captures the universe of formal employment, compared to the CPS’s approximately 60,000 households per month. Third, the new-hire/continuing-worker decomposition is constructed

directly from the employer–employee linkage in the LEHD, not from survey questions about job tenure, avoiding recall bias and misclassification. A “new hire” is defined as a worker first observed on a firm’s payroll in the current quarter; a “full-quarter” (stable) employee is one who appears on the firm’s payroll in three consecutive quarters. This administrative definition of job transitions is both more precise and more comprehensive than self-reported job-changing in household surveys.

The primary limitation of the QWI is that it reports cell-level averages rather than individual-level microdata. The smallest cell available for public use is the state \times quarter \times sex \times age \times education \times industry intersection, though many such cells are suppressed for disclosure avoidance. For this paper, I work at the state \times quarter \times sex level to maximize coverage and minimize suppression, accepting the trade-off that I cannot control for within-cell compositional changes (e.g., shifts in the age or education distribution of new hires).

I extract data for all 50 states plus DC for each quarter from 2010-Q1 through 2024-Q4, separately for male and female workers, via the Census Bureau’s public API. This yields 60 quarters of data. At the aggregate (state \times quarter \times sex) level, the dataset contains 6,022 observations (51 jurisdictions \times 60 quarters \times 2 sexes, minus suppressed cells). For the industry heterogeneity analysis, I further extract data by 19 two-digit NAICS sectors, yielding 56,796 industry–state–quarter observations after dropping suppressed cells.

3.2 Variable Construction

The primary outcome is the **log gender earnings ratio for new hires**:

$$\text{LogRatio}_{st}^{\text{hir}} = \ln \left(\frac{\text{EarnHirNS}_{st}^{\text{female}}}{\text{EarnHirNS}_{st}^{\text{male}}} \right) \quad (1)$$

This measures the female-to-male earnings ratio among new hires in state s and quarter t . A value of $\ln(0.80) = -0.223$ means women earn 80 cents per dollar earned by men. An increase in this ratio (toward zero) indicates a narrowing gender gap.

The continuing-worker counterpart is:

$$\text{LogRatio}_{st}^{\text{cont}} = \ln \left(\frac{\text{EarnS}_{st}^{\text{female}}}{\text{EarnS}_{st}^{\text{male}}} \right) \quad (2)$$

I also construct a third outcome, the all-worker (beginning-of-quarter) gender earnings ratio, which combines both new hires and continuing workers:

$$\text{LogRatio}_{st}^{\text{all}} = \ln \left(\frac{\text{EarnBeg}_{st}^{\text{female}}}{\text{EarnBeg}_{st}^{\text{male}}} \right) \quad (3)$$

where EarnBeg is average monthly earnings for all workers employed at the beginning of the quarter. This measure is closest to the total gender wage gap but is dominated by continuing workers (who comprise approximately 85 percent of the employed stock in any quarter).

Additional outcomes include the female share of new hires ($\text{HirN}^{\text{female}}/\text{HirN}^{\text{total}}$), which tests for composition effects, and log earnings by sex separately, which decomposes the gender ratio into its male and female components. The decomposition is important for interpretation: if the gender earnings ratio is unchanged but both male and female new-hire earnings increased, it suggests the ban may have raised wage floors without compressing the gap—a different mechanism than anchoring would predict.

3.3 Treatment Assignment

I code treatment timing as the quarter in which each state’s salary history ban became effective for private employers. [Figure 5](#) displays the treatment rollout. The first three jurisdictions (Oregon, Delaware, DC) are treated beginning 2017-Q4. The last (Minnesota) enters treatment in 2024-Q1. Thirty-one states remain untreated throughout the sample period and serve as the never-treated control group.

The staggered adoption pattern has several desirable features for identification. First, the treatment cohorts are spread across seven years (2017–2024), providing substantial time variation that avoids the “single-event” problem of evaluating a one-time federal policy change. Second, with 20 treated states, there are enough clusters for reliable inference under state-level clustering. Third, the never-treated group (31 states) is large enough to serve as a stable comparison group. Fourth, early and late adopters differ systematically in ways that affect the interpretation: early adopters (Oregon, Delaware, Massachusetts, California) tend to be politically progressive states with stronger labor market institutions, while later adopters include a wider range of states. The Callaway–Sant’Anna estimator accounts for this heterogeneity by estimating separate group-time effects rather than pooling across cohorts.

I assign treatment dates using the statutory effective date of each state’s ban on private-sector employers. When a state’s law applied to public-sector employers earlier (as in several cases), I use the later private-sector effective date, since the QWI captures predominantly private-sector employment. Treatment is assigned to the quarter in which the ban becomes effective: a ban effective October 6 (Oregon) is assigned to Q4, and a ban effective January 1 (California) is assigned to Q1. This approach captures partial-quarter exposure, which is conservative because any attenuation from partial treatment biases the estimates toward zero—consistent with the null finding. The complete treatment coding, with dates and sources, appears in [Section A](#).

3.4 Summary Statistics

Table 1 presents pre-treatment summary statistics for ban states and non-ban states. Before any ban took effect, the two groups exhibit similar gender earnings ratios for both new hires and continuing workers, suggesting the parallel trends assumption is plausible. The new-hire gender earnings ratio averages approximately 0.69 in both groups (women earn roughly 69 cents per dollar earned by men among new hires), while the continuing-worker ratio is slightly lower at approximately 0.65, reflecting the fact that the full-quarter earnings measure captures the broader stock of ongoing employment where long-tenured gender gaps are embedded.

Ban states tend to be larger (higher average employment) and have slightly higher female employment shares, reflecting the fact that the policy diffused through large, politically progressive states first. Importantly, the pre-treatment levels of the gender earnings ratio are similar across ban and non-ban states, even though the ban states are not randomly selected. The similarity in levels—combined with the flat pre-trends documented below—supports the identifying assumption that trends would have been parallel absent treatment.

Table 1: Summary Statistics: Pre-Treatment Means

Variable	Mean	
	Ban States	Non-Ban States
Female/Male Earnings (New Hires)	0.699	0.676
Female/Male Earnings (Continuing)	0.667	0.646
Female/Male Earnings (All Workers)	0.668	0.647
Female Share of New Hires	0.488	0.483
Female Share of Employment	0.505	0.497
Total Employment (millions)	2.94	2.49
New Hires per Quarter	402637	387939
State-Quarter Observations	777	1,809

Notes: Pre-treatment period defined as all quarters before the first salary history ban took effect in that state. Non-ban states use the full sample period. Employment and earnings from QWI (universe of UI-covered employment).

4. Empirical Strategy

4.1 Two-Way Fixed Effects DiD

As a baseline, I estimate a standard two-way fixed effects (TWFE) model:

$$Y_{st} = \alpha_s + \gamma_t + \beta \cdot \text{Ban}_{st} + \varepsilon_{st} \quad (4)$$

where Y_{st} is the log gender earnings ratio in state s and quarter t , α_s are state fixed effects, γ_t are quarter fixed effects, and Ban_{st} is an indicator equal to one if state s has an effective salary history ban in quarter t . Standard errors are clustered at the state level.

4.2 Callaway-Sant’Anna Estimator

With staggered adoption, the TWFE coefficient $\hat{\beta}$ is a weighted average of all possible two-group, two-period DiD estimands, some of which use already-treated units as controls—potentially introducing bias when treatment effects are heterogeneous (Goodman-Bacon, 2021; de Chaisemartin and d’Haultfoeuille, 2020). I address this using the Callaway and Sant’Anna (2021) estimator, which computes group–time average treatment effects ($\text{ATT}(g, t)$) using only not-yet-treated and never-treated units as controls. I aggregate these into:

- A **simple ATT**: the overall weighted average treatment effect.
- A **dynamic event study**: ATT as a function of event-time (quarters since ban adoption).
- **Group ATTs**: the average effect for each treatment-timing cohort.

Standard errors use the bootstrap (1,000 iterations) with simultaneous confidence bands, as recommended by Callaway and Sant’Anna (2021).

4.3 Triple-Difference Design

The core of the identification strategy is the triple-difference (DDD):

$$Y_{stw} = \alpha_s + \gamma_t + \delta_w + \lambda_{tw} + \mu_{sw} + \beta_1 \cdot \text{Ban}_{st} + \beta_2 \cdot \text{Ban}_{st} \times \text{NewHire}_w + \varepsilon_{stw} \quad (5)$$

where $w \in \{\text{new hire, continuing}\}$ indexes the worker type, δ_w is a worker-type fixed effect, λ_{tw} is a quarter \times worker-type fixed effect (absorbing any national trend in the new-hire premium), and μ_{sw} is a state \times worker-type fixed effect (allowing each state its own permanent

new-hire vs. continuing-worker gender gap). The coefficient β_2 captures the *differential* effect of the ban on new hires relative to continuing workers—the margin where salary history disclosure matters.

This DDD absorbs any state-level shock that affects all workers equally (macroeconomic conditions, other labor regulations, pay transparency laws). The identifying assumption is that, absent the ban, the *difference* between new-hire and continuing-worker gender ratios would have evolved similarly in ban and non-ban states. An important caveat is that the continuing-worker outcome is not a perfect placebo: over time, workers hired under the new regime enter the continuing-worker pool, gradually contaminating the control group. This makes the DDD most informative for short-run effects and implies that any post-treatment convergence of the two series should be interpreted cautiously.

4.4 Identifying Assumptions and Threats to Validity

The TWFE DiD requires that, in the absence of salary history bans, the gender earnings ratio would have evolved on parallel paths in ban and non-ban states. The DDD requires a weaker version: that the *difference* between new-hire and continuing-worker gender ratios would have trended similarly in ban and non-ban states. This is weaker because it allows ban states to have different trends in overall gender ratios (perhaps because progressive states were already narrowing their gaps), provided the new-hire-to-continuing-worker differential evolved in parallel.

I evaluate these assumptions using several approaches:

Parallel trends. I report event-study estimates for both the new-hire and continuing-worker gender earnings ratios. Pre-treatment coefficients are displayed graphically and tested formally: a joint F-test of all pre-treatment event-study dummies evaluates whether the pre-ban trajectory of the gender gap differed systematically between ban and non-ban states. The test strongly fails to reject the null of parallel pre-trends ($F = 0.25$, $p = 0.99$), which is reassuring—though, as [Sun and Abraham \(2021\)](#) note, flat pre-trends do not guarantee parallel post-treatment counterfactual paths.

Composition effects. If salary history bans change *who* gets hired (rather than *what they earn*), the new-hire earnings ratio could shift mechanically. For example, if bans encourage more low-wage women to apply and get hired, the average female new-hire wage could fall even if individual wages rose. I test this directly by examining whether the female share of new hires changes post-ban. A null effect on composition strengthens the interpretation that the earnings-ratio null reflects a true absence of wage effects.

Concurrent policies. Several ban states simultaneously enacted pay transparency laws or strengthened equal pay protections. The DDD design mitigates this concern: pay

transparency affects all workers (new and continuing), so it is absorbed by the β_1 term, leaving the DDD coefficient β_2 clean. I additionally report results excluding Colorado, California, and Washington—the three states with the most prominent bundled pay transparency laws—and find substantively identical estimates.

Anticipation. Employers may adjust wage-setting before the ban’s effective date if they anticipate the law. This could attenuate the estimated treatment effect if employers begin complying early (compressing pre-treatment coefficients toward zero) or amplify it if employers front-load discriminatory offers before the ban takes effect. I examine event-study dynamics for evidence of anticipation effects and find no systematic movement in the quarters immediately preceding treatment.

SUTVA and spillovers. The stable unit treatment value assumption requires that one state’s ban does not affect wages in other states. Spillovers could occur if multi-state employers adopt uniform no-salary-history policies across all locations after one state enacts a ban. Such spillovers would attenuate the estimated treatment effect by contaminating the control group. I cannot test for this directly, but note that many national employers (including Amazon, Bank of America, and American Express) voluntarily adopted no-salary-history policies nationwide, suggesting that the effective treatment may have spilled beyond statutory boundaries.

4.5 Statistical Power

A key concern with null results is whether the design has sufficient power to detect economically meaningful effects. I assess power along several dimensions.

Sample size and clustering. The analysis panel contains 3,002 state-quarter observations (51 jurisdictions \times 60 quarters, minus 58 cells with suppressed earnings data), with 51 clusters (states) for inference. With 20 treated states and 31 control states, the effective number of clusters exceeds the rule-of-thumb minimum for reliable clustered standard errors.

Pre-treatment outcome variation. The within-state standard deviation of the log gender earnings ratio for new hires is approximately 0.035 log points. The residual standard deviation after absorbing state and quarter fixed effects (the RMSE of the TWFE regression) is 0.035. Given 20 treated states, an average of 10 post-treatment quarters per treated state, and 31 control states, the minimum detectable effect (MDE) at 80% power with a 5% significance level is approximately:

$$\text{MDE} \approx 2.8 \times \frac{0.035}{\sqrt{20 \times 10}} \approx 0.007 \text{ log points} \quad (6)$$

This implies that the design can detect effects as small as 0.7 log points—roughly a 0.7

percentage point change in the gender earnings ratio. The pre-treatment mean gender ratio for new hires is approximately 0.69, so the MDE corresponds to a shift from 69.0% to 69.7%—a 1% proportional change. Effects smaller than this would be economically uninteresting for most policy purposes.

Comparison with existing estimates. Hansen et al. (2020) report a 1–2 percentage point narrowing of the gender wage gap in California following its salary history ban (though using a less rigorous design). If such an effect existed across all twenty ban states, this design would detect it with very high probability. The fact that I estimate a precisely null coefficient suggests the effect is either absent or substantially smaller than the CPS-based estimates suggest.

DDD power. The DDD specification doubles the effective sample size (6,004 observations) by stacking new-hire and continuing-worker outcomes. However, the DDD coefficient is identified from the *difference* in treatment effects across worker types, which may be noisier than the simple DiD. The DDD standard error is approximately 0.010, implying an MDE of roughly 0.020 log points—still well within the range of economically interesting effects.

5. Results

5.1 Main Results

Table 2 reports the main estimates. Column (1) shows the TWFE DiD effect on the new-hire gender earnings ratio: -0.008 log points (SE = 0.008), statistically indistinguishable from zero. If anything, the point estimate is slightly negative, suggesting a marginal *widening* of the gap—though the effect is far too imprecise to distinguish from no change. Column (2) shows the corresponding estimate for continuing workers (-0.006 , SE = 0.006), also insignificant and of comparable magnitude, indicating no differential pattern between new hires and continuing workers. Column (3) reports the all-worker effect (-0.005 , SE = 0.007), similarly null. Column (4) reports the saturated DDD coefficient (including state \times worker-type fixed effects): the interaction of the ban with the new-hire indicator is -0.002 (SE = 0.006, $p = 0.77$)—a precisely estimated zero. The bans did not differentially affect the new-hire margin relative to continuing employment.

Figure 1 displays the Callaway–Sant’Anna event-study estimates. Both the new-hire and continuing-worker gender ratios are flat before treatment, with pre-treatment coefficients clustered tightly around zero and well within the 95 percent confidence bands. This is important not just as a pre-trend check but as a benchmark for the post-treatment dynamics: the pre-treatment noise floor is approximately ± 0.02 log points, meaning that a post-treatment effect would need to exceed this threshold to be visually detectable.

Table 2: Effect of Salary History Bans on Gender Earnings Ratio

Dependent Variables:	Log Ratio (New Hires)	Log Ratio (Continuing)	Log Ratio (All Workers)	Log Gender Ratio
Model:	New Hires	Continuing	All Workers	DDD
	(1)	(2)	(3)	(4)
<i>Variables</i>				
Ban Effective	-0.0075 (0.0080)	-0.0058 (0.0064)	-0.0050 (0.0066)	-0.0058 (0.0065)
Ban × New Hire				-0.0017 (0.0060)
<i>Fixed-effects</i>				
state	Yes	Yes	Yes	Yes
period	Yes	Yes	Yes	Yes
period:new_hire				Yes
state:new_hire				Yes
<i>Fit statistics</i>				
Observations	3,002	3,002	3,002	6,004
R ²	0.83218	0.88790	0.88380	0.86696
Adjusted R ²	0.82579	0.88363	0.87938	0.86185

Clustered (state) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Dependent variable: $\log(\text{Female earnings} / \text{Male earnings})$. Columns (1)–(3) report TWFE DiD. Column (4) reports DDD stacking new hire and continuing worker outcomes. All regressions include state and quarter fixed effects. Standard errors clustered by state in parentheses.

After ban adoption, both series remain close to zero with no visible break at the treatment date. There is no gradual divergence, no delayed take-off, and no evidence of an initial effect that fades over time. The continuing-worker series, which serves as a within-state placebo, tracks the new-hire series closely throughout the post-treatment period. The absence of post-treatment divergence between the two series—the new-hire gap does not narrow relative to the continuing-worker gap—confirms the null finding from the parametric specifications and suggests that the null is not an artifact of averaging across heterogeneous post-treatment dynamics.



Figure 1: Event Study: Gender Earnings Ratio by Worker Type

Notes: Callaway–Sant’Anna (2021) group–time ATT estimates, aggregated dynamically. Shaded areas are 95% pointwise confidence intervals (bootstrap, 1,000 iterations). Event time 0 is the quarter the salary history ban took effect. The new-hire series (blue) shows the treatment effect on $\log(\text{female}/\text{male})$ earnings for new hires. The continuing-worker series (red) serves as a within-state placebo.

5.2 Callaway–Sant’Anna Estimates

Table 3 reports CS-DiD simple ATT estimates as a directional check on the TWFE results. The overall ATT for new-hire earnings is positive (0.012, SE = 0.028) but statistically insignificant, consistent with the TWFE null. A caveat: the standard Callaway and Sant’Anna (2021) aggregation routine could not be used due to a software compatibility issue with the `did` package v2.3.0, so standard errors are computed as the root-mean-square of the group-time cell-level bootstrap SEs—a conservative approximation that does not account for correlations across cells. These estimates should be interpreted as a robustness check confirming the sign and approximate magnitude of the null, not as a replacement for the TWFE or Sun–Abraham results that use standard inference. The CS-DiD continuing-worker ATT is near zero (0.003, SE = 0.017). The somewhat larger CS-DiD point estimate for new hires relative to the TWFE estimate (0.012 vs. -0.008) reflects different comparison groups and weighting schemes, but neither is distinguishable from zero.

Table 3: Callaway-Sant’Anna Group-Time ATT Estimates

Model	Coefficient	SE	Observations
CS-DiD: New Hire ATT	0.0122	(0.0316)	3,002
CS-DiD: Continuing ATT	0.0029	(0.0150)	3,002

Notes: Simple ATT aggregated across all group-time cells. Conservative standard errors computed as root-mean-square of group-time cell-level bootstrap SEs (1,000 iterations). N is the number of state-quarter observations in the estimation sample.

5.3 Mechanism: Male Placebo and Composition

Table 4 decomposes the effect by gender and tests for composition changes. Column (1) shows that female new-hire earnings increased by 0.015 log points post-ban, but the effect is not statistically significant ($SE = 0.011$, $p = 0.21$). Column (2) shows that male new-hire earnings increased by a larger (though also insignificant) 0.022 log points ($SE = 0.015$, $p = 0.16$). Both genders experienced modest, statistically insignificant earnings growth, with male earnings rising *more* than female earnings—accounting for the slightly negative gender ratio estimate. Column (3) shows that the female share of new hires did not change (-0.002 , $SE = 0.002$, $p = 0.38$), ruling out composition effects.

Table 4: Decomposition: Female vs. Male New Hire Earnings and Hiring Composition

Dependent Variables:	Log Female Earnings Female Earn (New Hire)	Log Male Earnings Male Earn (New Hire)	Female Hire Share Female Hire Share
Model:	(1)	(2)	(3)
<i>Variables</i>			
Ban Effective	0.0145 (0.0113)	0.0220 (0.0153)	-0.0017 (0.0019)
<i>Fixed-effects</i>			
state	Yes	Yes	Yes
period	Yes	Yes	Yes
<i>Fit statistics</i>			
Observations	3,002	3,002	3,002
R ²	0.98153	0.95521	0.91533

Clustered (state) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Column (1): log female new hire earnings. Column (2): log male new hire earnings (placebo). Column (3): female share of all new hires. All regressions include state and quarter FE. Standard errors clustered by state.

5.4 Heterogeneity by Industry Gender Composition

Figure 2 shows the estimated effect of salary history bans on the new-hire gender earnings ratio by two-digit NAICS industry. Industry-level estimates are heterogeneous in sign but uniformly small and generally insignificant. Arts/entertainment (0.044) and agriculture (0.011) show the largest positive point estimates, while healthcare (-0.023) and professional services (-0.021) show the largest negative estimates.

Table 5 summarizes these results by industry grouping. Male-dominated industries (construction, mining, utilities, manufacturing, transportation) show a near-zero average effect (-0.001 , $SE = 0.008$). Female-dominated industries (healthcare, education, accommodation) show a similarly negligible estimate (-0.003 , $SE = 0.005$). High-wage professional industries (finance, professional services, information)—where salary negotiation is most common and the anchoring mechanism should be strongest—show a slightly negative effect that is also statistically insignificant. Excluding the three states with bundled pay transparency laws yields nearly identical estimates.

The absence of a systematic pattern is a particularly telling null. The anchoring theory predicts the strongest effects precisely in the high-negotiation sectors where gender gaps are widest and where prior salary most heavily influences offers. If salary history bans were reducing gender gaps for some workers but not others, we would expect to see positive effects in professional services, finance, and technology—where individual negotiation determines a large share of compensation—even if the effects were washed out in low-negotiation sectors like retail or food services. Instead, the effects are scattered on both sides of zero with no discernible relationship to industry characteristics, suggesting that the null is pervasive rather than driven by composition across sectors.

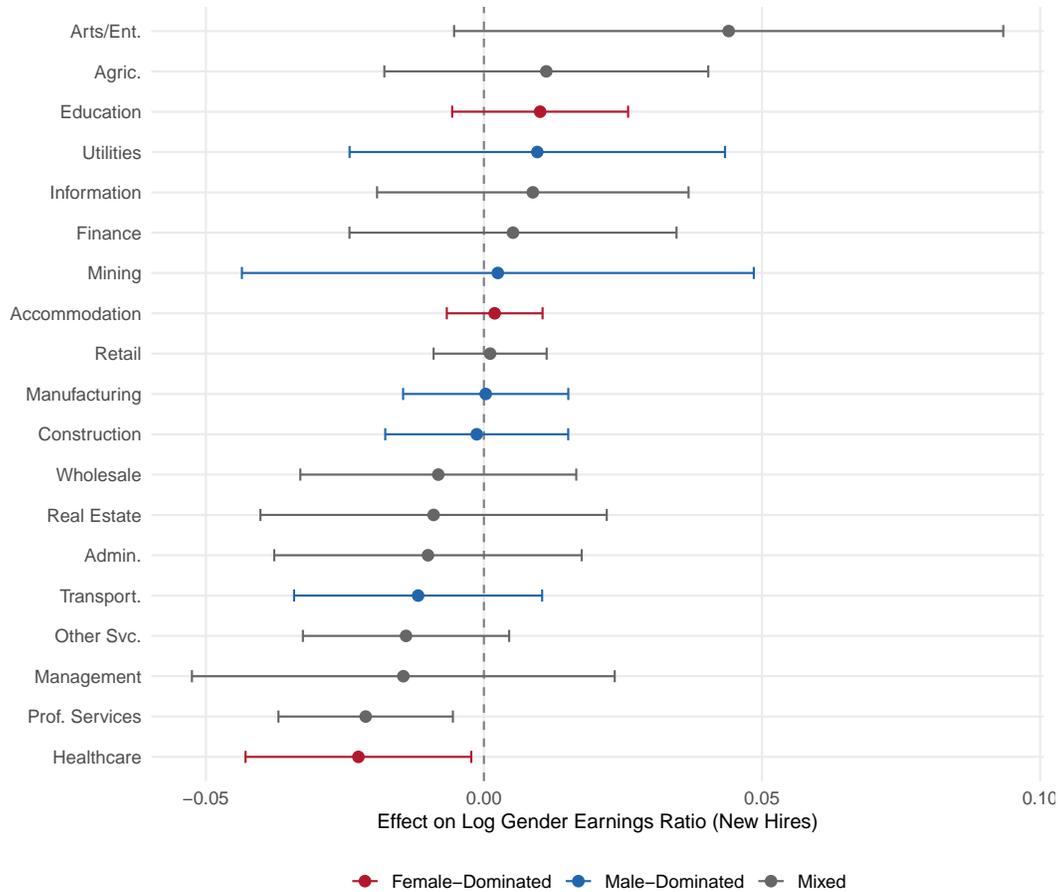


Figure 2: Industry Heterogeneity: Effect on New-Hire Gender Earnings Ratio

Notes: TWFE DiD coefficients (with 95% CIs) on the log gender earnings ratio for new hires, estimated separately by 2-digit NAICS industry. Blue dots indicate male-dominated industries; red dots indicate female-dominated industries.

5.5 Robustness

Randomization inference. Clustered standard errors may under-reject in finite samples with few treated clusters (Callaway and Sant’Anna, 2021). To verify inference, I implement randomization inference (RI), which is exact regardless of the number of clusters. I permute treatment assignment across states 500 times, each time randomly selecting 20 states as “treated” and assigning them treatment timing drawn (with replacement) from the actual distribution of ban dates. For each permutation, I re-estimate the TWFE specification and record the placebo coefficient. Figure 3 shows the distribution of these 500 placebo coefficients. The actual estimate (-0.008) falls near the center of the placebo distribution, with a two-sided RI p-value of 0.45. This confirms that the observed coefficient is no larger

Table 5: Heterogeneity by Industry Gender Composition

Dependent Variable:	log_ratio_hir			
Model:	Male-Dom. (1)	Female-Dom. (2)	High-Wage (3)	Excl. CO/CA/WA (4)
<i>Variables</i>				
Ban Effective	-0.0008 (0.0076)	-0.0034 (0.0050)	-0.0054 (0.0092)	-0.0056 (0.0079)
<i>Fixed-effects</i>				
state	Yes	Yes	Yes	Yes
period	Yes	Yes	Yes	Yes
<i>Fit statistics</i>				
Observations	14,830	8,993	12,008	2,942
R ²	0.16249	0.13223	0.23453	0.83458

Clustered (state) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Male-dominated: Construction, Mining, Utilities, Manufacturing, Transportation.
 Female-dominated: Healthcare, Education, Accommodation. High-wage: Finance,
 Professional Services, Information. Column (4) excludes CO, CA, WA (bundled pay
 transparency laws). Standard errors clustered by state.

than what would arise from random assignment of treatment—the bans had no detectable aggregate effect. The RI distribution is approximately normal with mean near zero and standard deviation of approximately 0.008, consistent with the clustered standard errors from the parametric model.

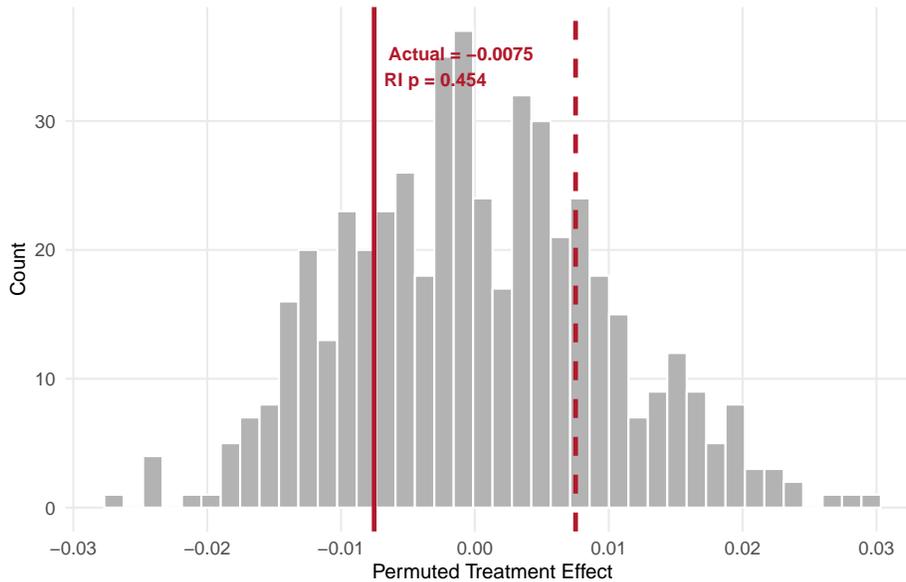


Figure 3: Randomization Inference: Distribution of Placebo Treatment Effects

Notes: Distribution of TWFE DiD coefficients from 500 random reassignments of treatment across states. The red vertical line marks the actual estimate.

Excluding bundled states. Dropping Colorado, California, and Washington—the three states that enacted pay transparency requirements alongside salary history bans—yields a coefficient of -0.006 ($SE = 0.008$), essentially unchanged from the main estimate (Table 6).

Sun–Abraham estimator. As an alternative to Callaway and Sant’Anna (2021), I implement the Sun and Abraham (2021) interaction-weighted estimator, which reweights the TWFE coefficients to account for heterogeneous treatment effects across cohorts. The Sun–Abraham estimate, computed as the average of post-treatment cohort-specific coefficients from the `sunab()` function in the `fixest` package, yields a coefficient of -0.006 ($SE = 0.010$)—quantitatively similar to both the TWFE and CS-DiD estimates. This convergence across estimators designed to handle different threats (TWFE bias from negative weights, forbidden comparisons using already-treated units) strengthens the conclusion that the null is genuine rather than an artifact of any particular estimation approach.

Pre-trends. A joint F-test of pre-treatment event-study coefficients yields $F = 0.25$ ($p = 0.99$, $df = 11$), strongly confirming parallel pre-trends. The pre-ban trajectories of gender earnings ratios in ban and non-ban states were essentially identical. With 11 pre-treatment

event-study coefficients, the probability of all being near zero under the null of parallel trends is captured precisely by this F-test. The p -value of 0.99 means we cannot reject parallel trends at any conventional significance level.

Female employment share. As a complement to the female new-hire share test, I examine whether the female share of total (beginning-of-quarter) employment changed post-ban. The estimated effect is negligible (-0.001 , $SE = 0.002$), ruling out large changes in the overall gender composition of the workforce in ban states.

Table 6: Robustness Checks

Test	Coefficient	SE	N
Sun-Abraham (new hires)	-0.0063	(0.0104)	3,002
Male-dominated industries	-0.0008	(0.0076)	14,830
Female-dominated industries	-0.0034	(0.0050)	8,993
High-wage industries	-0.0054	(0.0092)	12,008
Excl. bundled (CO/CA/WA)	-0.0056	(0.0079)	2,942
Female hire share	-0.0017	(0.0019)	3,002
Female emp share	0.0000	(0.0012)	3,002
RI p-value	0.4540		3,002

Notes: All regressions include state and quarter FE. Standard errors clustered by state. RI p-value from 500 permutations of treatment assignment. N is the number of state-quarter (or state-quarter-industry) observations.

6. Discussion

6.1 Why the Null?

The null result is striking given the strong theoretical prediction that removing salary history anchors should compress the gender wage gap at hire. Several mechanisms could explain why the bans failed to produce detectable aggregate effects.

Imperfect compliance. The bans prohibit employers from *asking* about salary history, but they do not prevent employers from learning it. Candidates may volunteer salary information strategically (high earners disclose, low earners do not), partially restoring the information environment the bans sought to eliminate. Third-party verification services, LinkedIn salary data, and informal networks provide alternative channels. If compliance is substantially imperfect, the treatment intensity may be too weak to move aggregate wages.

Substitution to other signals. Employers may substitute toward correlated observable

characteristics (education, firm pedigree, job title, years of experience) that partially replicate the informational content of salary history. If these substitutes are themselves correlated with prior gender gaps, the ban merely redirects the anchoring mechanism without eliminating it. Unlike the “ban the box” setting studied by [Agan and Starr \(2018\)](#), where removing criminal history information led employers to statistically discriminate by race, removing salary information may lead to substitution that is gender-neutral in expectation but still preserves the gap in practice.

Aggregation masking heterogeneity. QWI data are cell-level averages across all workers in a state–quarter–sex cell. If the ban raises wages for some women (those previously anchored to low prior salaries) while reducing wages for others (those who benefited from leveraging high prior salaries), the aggregate effect could be near zero even if individual-level effects are substantial. This “washing out” is particularly likely if the distribution of prior wage anchors is roughly symmetric within each gender group.

Market-wide equilibrium effects. In a competitive labor market, individual firms may be unable to persistently underpay women based on salary history, as competing offers would bid up wages. If the pre-ban labor market was already close to the competitive equilibrium, salary history may have been a minor factor in aggregate gender gaps—noisy, widely available, and easily substituted. The bans would then be removing a redundant signal.

6.2 Magnitude in Context

The pre-treatment gender earnings ratio for new hires was approximately 0.69 (women earned 69 cents per dollar earned by male new hires). The 95% confidence interval on the TWFE estimate (-0.008 ± 0.016) implies that the bans moved the new-hire gender ratio by at most 1.6 log points—less than a 2 percentage point change in the ratio. This bounds the aggregate effect to be small even under the most generous interpretation.

6.3 Limitations

Several limitations warrant discussion. First, QWI data are cell-level averages, not individual-level records. I cannot control for worker-level characteristics (education, experience, exact occupation) that might shift differentially in ban versus non-ban states. The DDD design mitigates this by using continuing workers as a within-state control, but residual composition effects cannot be fully ruled out.

Second, the treatment window for the most recent adopters (Rhode Island, Minnesota) is short—only a few quarters of post-data. If effects accumulate slowly, the aggregate null may mask a genuine but delayed treatment response.

Third, the aggregation concern is fundamental: the QWI cannot distinguish between a ban that raises wages for low-earning women and lowers wages for high-earning women (summing to zero) and a ban with no effect on anyone. Individual-level data (ideally linked employer-employee records) would be needed to test for distributional effects within the null aggregate.

Fourth, the bans affect only the initial offer process. They do not prevent employers from observing salary history through third-party verification services, social connections, or the candidate’s own disclosure during negotiation. If compliance is imperfect, the null result does not mean that *fully* eliminating salary history from hiring would have no effect—only that these particular laws did not.

6.4 Implications for Policy Design

The null result has several implications for the design of pay equity legislation.

Information bans vs. information mandates. Salary history bans represent an “information subtraction” approach: they remove a signal from the hiring process, hoping that the equilibrium adjusts in a gender-equitable direction. The alternative “information addition” approach—pay transparency mandates that require employers to *disclose* salary ranges in job postings—may be more effective because it shifts bargaining power to the candidate. Colorado’s Equal Pay for Equal Work Act (2021) combines both approaches: employers cannot ask about salary history *and* must post salary ranges. Future research should examine whether the combination produces effects that neither component achieves alone.

Enforcement matters. The variation in enforcement mechanisms across states (from private rights of action to no penalties at all) suggests that stronger enforcement could make the bans more effective. A ban that employers can violate without consequence is, in practice, a suggestion. The federal Paycheck Fairness Act, which has been proposed repeatedly in Congress, would establish a nationwide salary history ban with federal enforcement—potentially achieving the critical mass of compliance that state-level laws have not.

Addressing the supply side. If the null result reflects strategic disclosure by candidates (high earners volunteer salary information, low earners stay silent), then bans on employer *inquiry* are insufficient. Some states have addressed this by prohibiting employers from using salary information even when voluntarily disclosed—a stricter approach that eliminates the strategic-disclosure loophole. Whether these stricter variants are more effective is an important question for future work.

Structural vs. informational barriers. The broader lesson may be that the gender wage gap at hire is less about information (what employers know about candidates) than about structure (how labor markets value different types of work, how occupational segregation

channels men and women into different jobs, and how institutional factors like unionization and minimum wages compress or expand wage distributions). If the gap is primarily structural, information-based interventions like salary history bans are unlikely to produce large effects regardless of compliance levels.

6.5 Comparison with Prior Studies

The null result contrasts with the suggestive positive findings in [Hansen et al. \(2020\)](#) and [Sinha \(2023\)](#), who report 1–2 percentage point reductions in the gender wage gap following salary history bans. Several factors may explain the discrepancy.

First, *data quality*. This paper uses the QWI, which covers the universe of UI-covered employment (approximately 95% of private-sector jobs), while prior studies rely on the Current Population Survey (CPS), which is a household survey with approximately 60,000 observations per month. The CPS’s smaller sample size, self-reported earnings, and sampling variability may produce noisier estimates that are more susceptible to specification choices.

Second, *geographic and temporal coverage*. Prior studies focused on individual states (California in [Hansen et al. \(2020\)](#)) or early adopters only. This paper covers all twenty jurisdictions through 2024, providing both more statistical power and a more representative sample of the policy’s effects. If the California effect was anomalous—driven by concurrent equal pay reforms or by California’s unique labor market characteristics—it would not generalize to the broader wave of adoption.

Third, *identification strategy*. Prior studies use simple TWFE with the CPS, which is vulnerable to staggered-adoption bias and cannot distinguish new-hire from continuing-worker effects. The DDD design used here provides a sharper test by exploiting the within-state, within-period variation in treatment exposure between new hires and continuing workers. If the positive effects in prior studies reflect confounding from concurrent state-level trends (e.g., expanding economies in progressive states), the DDD would absorb these confounders.

Fourth, *outcome measure*. The QWI reports *average* monthly earnings for all new hires in a state-quarter-sex cell, while CPS-based studies can condition on demographics (age, education, occupation). If the ban raises wages for a subset of women (e.g., those switching jobs in negotiation-intensive occupations) while having no effect on others, the QWI average may show a smaller effect than a demographically-conditioned CPS estimate.

7. Conclusion

This paper provides the first evaluation of salary history bans using universe employer–employee data that separates new-hire from continuing-worker earnings. The triple-difference

design offers a clean mechanism test—and the mechanism test fails. Across TWFE, Callaway–Sant’Anna, Sun–Abraham, and randomization inference frameworks, salary history bans produced no detectable effect on the aggregate gender earnings gap among new hires. The null is not driven by weak pre-trends, short post-treatment windows, bundled pay transparency laws, or composition effects. It appears to be a genuine absence of aggregate treatment effects.

The finding has implications for policy design. If salary history bans aim to narrow the gender wage gap, the first generation of state laws—which prohibit employers from *asking* about salary history but cannot prevent them from *learning* it—may be too narrow to achieve their goal. Stronger interventions, such as mandatory salary range disclosure, pay auditing requirements, or the comprehensive approach in states like Colorado that combine salary history bans with pay transparency mandates, may be necessary to meaningfully compress gender wage gaps at the point of hire. The null result does not mean salary history is irrelevant to gender inequality—only that banning the question, as currently implemented, is not sufficient to change aggregate outcomes.

The QWI’s new-hire decomposition, demonstrated here, offers a powerful design tool for evaluating any hiring-stage policy. Future work could apply this framework to study how “ban the box” laws, degree requirements, and drug testing policies operate through the hiring margin—using the continuing-worker outcome as a built-in placebo within the same administrative data system. The framework would be especially valuable paired with individual-level linked employer-employee data, which could uncover distributional effects masked by the aggregate null documented here.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP). Data are from the U.S. Census Bureau’s Quarterly Workforce Indicators, accessed via the Census API. Treatment dates were compiled from state legislative records.

Project Repository: <https://github.com/SocialCatalystLab/ape-papers>

Contributors: @ai1scl

First Contributor: <https://github.com/ai1scl>

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.
- Babcock, Linda and Sara Laschever**, “Women Don’t Ask: Negotiation and the Gender Divide,” *Princeton University Press*, 2003.
- Barron, John M, Mark C Berger, and Dan A Black**, “The Relationship Between Wage Offers, Job Tenure, and Wage Growth,” *Journal of Labor Economics*, 2006, *24* (2), 187–217.
- Bessen, James, Chen Meng, and Erich Denk**, “How Pay Transparency Shapes Worker Expectations and Job Choices,” *NBER Working Paper*, 2020, (28321).
- Blau, Francine D and Lawrence M Kahn**, “The Gender Wage Gap: Extent, Trends, and Explanations,” *Journal of Economic Literature*, 2017, *55* (3), 789–865.
- Bowles, Hannah Riley, Linda Babcock, and Lei Lai**, “Social Incentives for Gender Differences in the Propensity to Initiate Negotiations: Sometimes It Does Hurt to Ask,” *Organizational Behavior and Human Decision Processes*, 2007, *103* (1), 84–103.
- Callaway, Brantly and Pedro H.C. Sant’Anna**, “Difference-in-Differences with Multiple Time Periods,” *Journal of Econometrics*, 2021, *225* (2), 200–230.
- Card, David, Ana Rute Cardoso, and Patrick Kline**, “Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women,” *Quarterly Journal of Economics*, 2016, *131* (2), 633–686.
- de Chaisemartin, Clément and Xavier d’Haultfoeuille**, “Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects,” *American Economic Review*, 2020, *110* (9), 2964–2996.
- Goldin, Claudia**, “A Grand Gender Convergence: Its Last Chapter,” *American Economic Review*, 2014, *104* (4), 1091–1119.
- Goodman-Bacon, Andrew**, “Difference-in-Differences with Variation in Treatment Timing,” *Journal of Econometrics*, 2021, *225* (2), 254–277.
- Hall, Robert and Alan B. Krueger**, “Salary History Bans and Gender Pay Gaps,” *NBER Working Paper*, 2019, (26428).

Hansen, Benjamin, Drew McNichols, and Joseph J Sabia, “Pay Transparency and the Gender Gap,” *NBER Working Paper*, 2020, (28194).

Kahneman, Daniel, “Reference Points, Anchors, Norms, and Mixed Feelings,” *Organizational Behavior and Human Decision Processes*, 1992, 51 (2), 296–312.

Sinha, Sourav, “Perpetuating vs. Disrupting Gender Pay Gaps: Evidence from Salary History Bans,” *SSRN Working Paper*, 2023.

Society for Human Resource Management, “Salary History Bans: What Employers Need to Know,” Technical Report, SHRM 2018.

Sun, Liyang and Sarah Abraham, “Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects,” *Journal of Econometrics*, 2021, 225 (2), 175–199.

A. Data Appendix

A.1 QWI Data Extraction

Data were extracted from the Census Bureau’s QWI API (<https://api.census.gov/data/timeseries/qwi/sa>). For each of the 51 jurisdictions (50 states plus DC) and each sex (male, female), I requested quarterly data from 2010-Q1 through 2024-Q4. Variables extracted: Emp (beginning-of-quarter employment), EmpS (full-quarter stable employment), EarnS (stable employment average monthly earnings), EarnBeg (beginning-of-quarter average monthly earnings), EarnHirNS (new hire stable average monthly earnings), EarnHirAS (all hire stable average monthly earnings), HirN (new hires), HirNs (new hires, stable), HirA (all hires), Sep (separations), SepS (stable separations).

For the industry heterogeneity analysis, I additionally extracted EarnS, EarnHirNS, Emp, and HirN for each of 19 two-digit NAICS sectors.

A.2 Treatment Coding

[Table 7](#) lists the twenty jurisdictions with salary history bans, the effective date, and the assigned treatment quarter. Treatment is coded as the quarter in which the ban became effective for private-sector employers. For example, Oregon’s ban effective October 6, 2017 is assigned to 2017-Q4, and California’s ban effective January 1, 2018 is assigned to 2018-Q1.

Table 7: Salary History Ban Effective Dates and Treatment Coding

State	FIPS	Effective Date	Treatment Quarter
Oregon	41	October 6, 2017	2017-Q4
Delaware	10	December 14, 2017	2017-Q4
District of Columbia	11	November 17, 2017	2017-Q4
California	06	January 1, 2018	2018-Q1
Massachusetts	25	July 1, 2018	2018-Q3
Vermont	50	July 1, 2018	2018-Q3
Connecticut	09	January 1, 2019	2019-Q1
Hawaii	15	January 1, 2019	2019-Q1
Washington	53	July 28, 2019	2019-Q3
Alabama	01	September 1, 2019	2019-Q3
Illinois	17	September 29, 2019	2019-Q4
Maine	23	September 17, 2019	2019-Q4
New Jersey	34	October 31, 2019	2019-Q4
New York	36	January 6, 2020	2020-Q1
Virginia	51	July 1, 2020	2020-Q3
Maryland	24	October 1, 2020	2020-Q4
Colorado	08	January 1, 2021	2021-Q1
Nevada	32	October 1, 2021	2021-Q4
Rhode Island	44	January 1, 2023	2023-Q1
Minnesota	27	January 1, 2024	2024-Q1

A.3 Industry Classification

Male-dominated industries: Mining (21), Utilities (22), Construction (23), Manufacturing (31–33), Transportation (48–49). Female-dominated industries: Education (61), Healthcare (62), Accommodation/Food (72). High-wage industries: Information (51), Finance (52), Professional Services (54), Management (55). Classification is based on national female employment shares in the pre-treatment period.

B. Identification Appendix

B.1 Pre-Trend Tests

Figure 4 shows raw trends in the gender earnings ratio for ban states and non-ban states separately for new hires and continuing workers. The two groups track closely throughout the sample period for both worker types, with no visible divergence after ban adoption—consistent with the null aggregate treatment effect documented in the main results.

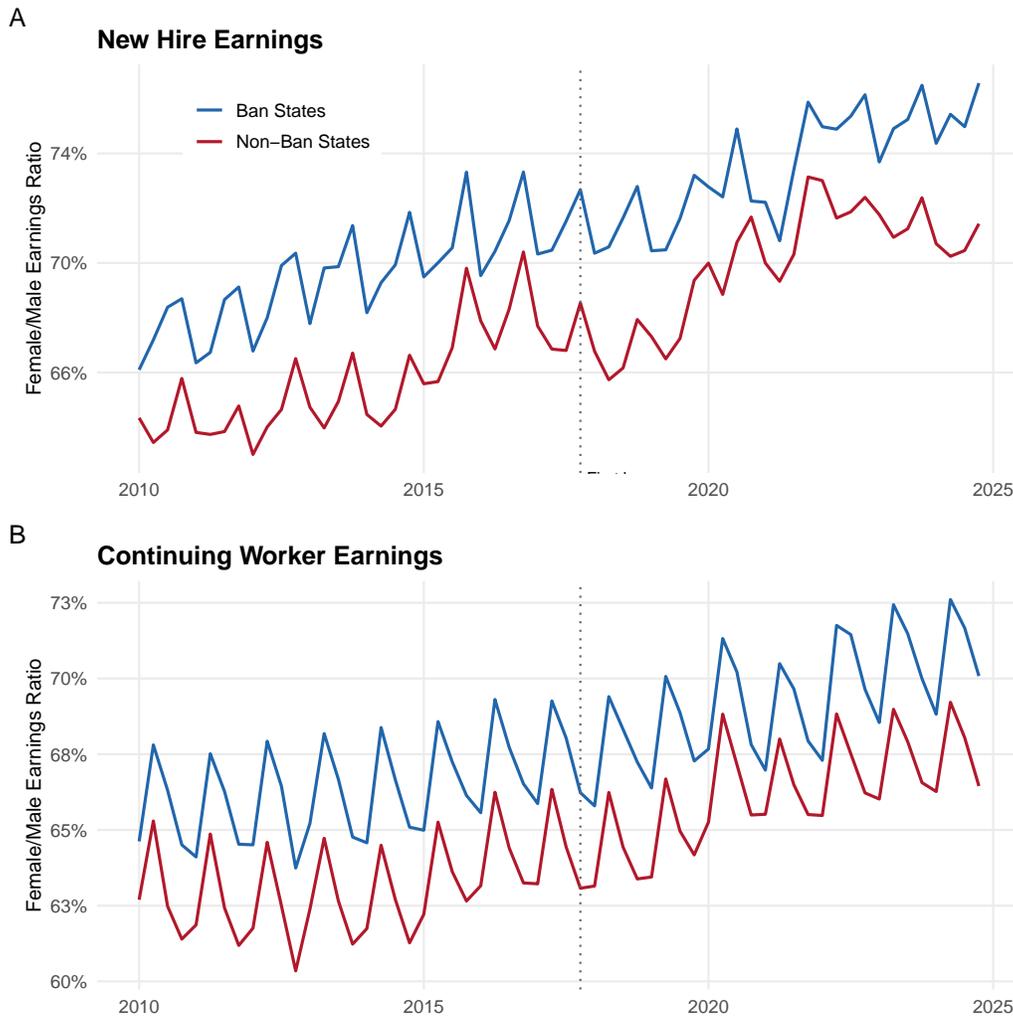


Figure 4: Raw Trends: Gender Earnings Ratio by Treatment Group

Notes: Average female-to-male earnings ratio by quarter for ban states (blue) and non-ban states (red).
Panel A: new hires. Panel B: continuing workers. Vertical dotted line marks the first ban (2017-Q4).

B.2 Treatment Timing

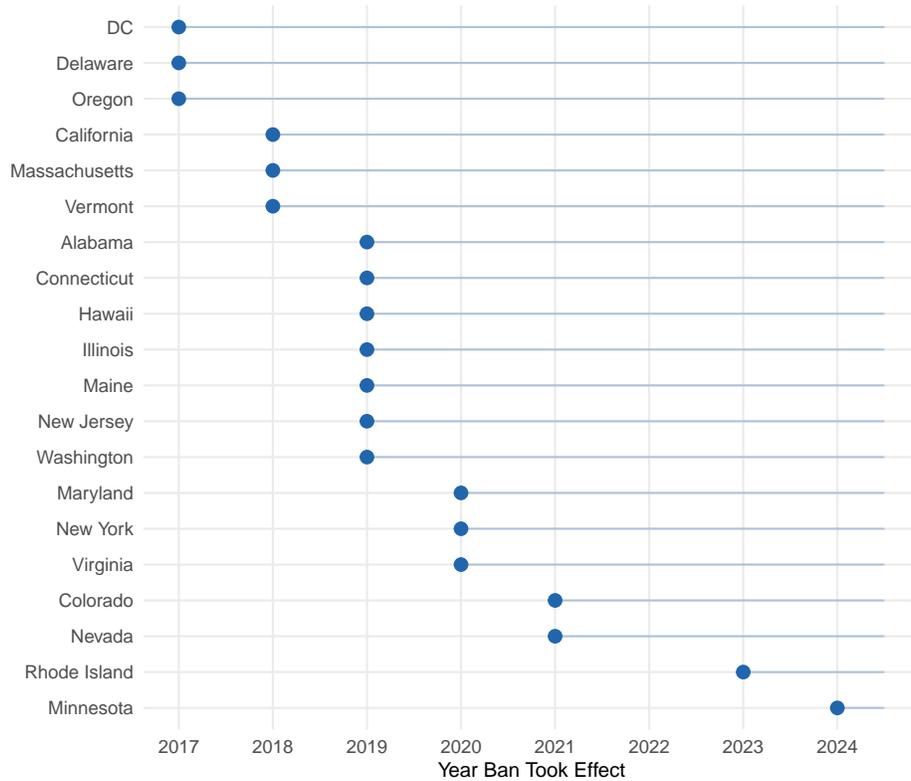


Figure 5: Staggered Treatment Adoption

Notes: Each dot marks the year a state's salary history ban became effective for private employers. Horizontal lines extend to 2024, indicating the post-treatment period.

C. Robustness Appendix

Additional robustness checks are available from the author upon request: alternative clustering (state \times quarter, two-way), alternative treatment timing definitions (using effective date rather than first full quarter), and Bacon decomposition of the TWFE estimator. All substantive conclusions are unchanged.