The Hidden Wage Floor: How Salary History Bans Reshape the Gender Earnings Gap

APEP Autonomous Research* @olafdrw

March 13, 2026

Abstract

A woman who earned less yesterday will earn less tomorrow—not because she is less productive, but because employers anchor offers to prior pay. Between 2017 and 2023, sixteen U.S. states banned employers from asking job applicants about salary history. Using Quarterly Workforce Indicators at the state-industry-quarter level and a Callaway–Sant'Anna staggered difference-in-differences design, I find that salary history bans narrowed the gender gap in new-hire earnings by 2.3 log points (SE = 0.73 log points), a moderate-negative standardized effect of −0.12 SD. The effect grows over time, reaching 4.1 log points by three years post-adoption. Healthcare and professional services show the strongest convergence. I find no evidence that bans triggered statistical discrimination against Black workers: the Black–White earnings gap, if anything, narrowed slightly. Standard two-way fixed effects yield a near-zero estimate, underscoring the importance of heterogeneity-robust methods.

JEL Codes: J31, J71, J78, K31

Keywords: salary history bans, gender pay gap, information asymmetry, staggered DiD, labor market discrimination

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 20m).

1. Introduction

In 2023, a newly hired woman in the United States earned 83 cents for every dollar earned by a newly hired man in the same industry and state. Part of this gap reflects occupational sorting, experience, and preferences. But part of it is a ratchet: employers use salary history to set offers, and past discrimination perpetuates itself through the sequence of wages a worker accumulates over her career (Agan et al., 2022). Each underpayment becomes the anchor for the next.

Between 2017 and 2023, sixteen states attempted to break this ratchet by prohibiting employers from asking job applicants about prior compensation—a policy known as a salary history ban. The logic is straightforward: if employers cannot observe past wages, they cannot anchor to them, and offers should reflect current productivity rather than historical disadvantage. But the prediction is not unambiguous. Doleac and Hansen (2020) demonstrated that "ban-the-box" laws—which removed criminal history from job applications—backfired by inducing statistical discrimination against Black men. Could salary history bans similarly harm the workers they intend to help?

This paper provides the first industry-level decomposition of salary history ban effects using administrative employment data. I combine the Census Bureau's Quarterly Workforce Indicators (QWI)—which uniquely provide new-hire earnings, hiring flows, and separation rates at the state-industry-quarter level—with a Callaway–Sant'Anna staggered difference-in-differences design exploiting seven distinct adoption cohorts between 2017 and 2023.

The main finding is that salary history bans narrowed the gender gap in new-hire earnings by 2.3 log points ($p < 0.01$). To put this in context, the average pre-ban gender gap was roughly 27 log points across industries; the bans closed approximately 8.5 percent of this gap. The effect is not immediate—it builds over time, reaching 4.1 log points by three years post-adoption, consistent with gradual employer adjustment and compositional turnover in the workforce.

Three mechanism tests sharpen the interpretation. First, I test whether effects concentrate in industries where the pre-ban gender gap was largest—the signature prediction of the anchoring hypothesis. Healthcare (NAICS 62) and management of companies (NAICS 55) show the strongest narrowing (approximately 2.9 log points each), consistent with industries where wage negotiation plays a larger role. Second, I test for Doleac and Hansen (2020)-style statistical discrimination: if employers substitute group-level priors for individual salary history, Black workers' hiring rates or earnings should deteriorate. I find no such evidence. The Black–White new-hire earnings gap narrowed by 0.8 log points (not statistically significant), and Black hiring rates were unchanged. Third, a government-sector placebo (NAICS 92—

where public employers are typically exempt from private-employer bans) shows a near-zero effect, supporting the identification.

A critical methodological finding is that standard two-way fixed effects (TWFE) yields a near-zero estimate (-0.001, $p = 0.82$) for the same outcome. The discrepancy arises because TWFE is contaminated by negative weights under staggered adoption with heterogeneous treatment effects (Goodman-Bacon, 2021; de Chaisemartin and D'Haultfœuille, 2020). The Callaway–Sant'Anna estimator, which computes group-time average treatment effects using only clean comparisons with never-treated units, recovers a significant and economically meaningful effect that TWFE masks entirely.

This paper contributes to three literatures. First, it advances the growing body of work on salary history bans (Agan et al., 2022; Barach and Horton, 2021; Sinha, 2024; Hansen and Guidice, 2020), all of which rely on survey data, tax records, or audit studies. The QWI's administrative coverage of the near-universe of private-sector employment, combined with its unique decomposition into new-hire earnings, hiring rates, and separations, enables the first industry-level analysis of where and how these bans operate. Second, it contributes to the literature on information removal in labor markets (Doleac and Hansen, 2020; Agan and Starr, 2018; Bartik and Nelson, 2018), demonstrating that—unlike criminal history bans—salary history bans do not trigger detectable statistical discrimination against minority workers. Third, it provides a clean demonstration of the practical consequences of estimator choice in staggered designs (Callaway and Sant'Anna, 2021; Sun and Abraham, 2021; Roth et al., 2023), where a policy-relevant effect is entirely invisible to the workhorse TWFE specification.

The remainder of the paper proceeds as follows. Section 2 describes the institutional background. Section 3 presents the data. Section 4 details the empirical strategy. Section 5 reports results. Section 6 discusses implications.

2. Institutional Background

2.1 Salary History Bans in the United States

Salary history bans prohibit employers from requesting, requiring, or relying on an applicant's prior compensation in setting wages. The typical statute bars the employer from asking about current or past salary during the application process, though many allow voluntary disclosure by the applicant. Some statutes extend to benefit information; others apply only to base salary.

Massachusetts was the first state to pass a salary history ban (signed August 2016, effective July 2018), following a campaign led by the American Association of University Women. Delaware and Oregon adopted bans effective in late 2017. A wave of adoptions followed:

California, Massachusetts, and Vermont in 2018; six states (Connecticut, Hawaii, Illinois, Washington, Maine, Alabama) in 2019; New Jersey and Maryland in 2020; Colorado and Nevada in 2021; and Rhode Island in 2023. By the end of 2023, sixteen states had enacted bans covering private employers.

The staggered adoption pattern reflects the politics of labor market regulation. Early adopters tended to be politically liberal states (California, Massachusetts, Oregon), while later adopters included states with more mixed political leanings (Colorado, Nevada). Critically, the timing of adoption was driven by state legislative calendars and advocacy cycles—not by contemporaneous changes in the gender pay gap within those states. This supports the parallel trends assumption underlying the difference-in-differences design.

2.2 Why Salary History Matters for Gender Earnings

The theoretical case for salary history bans rests on the wage-anchoring mechanism. Employers use prior salary as a signal of worker productivity and as a starting point for negotiation (Barach and Horton, 2021). If women's wages are depressed at any point in their career—whether due to discrimination, occupational segregation, or time out of the labor force for caregiving—the depressed wage propagates forward through subsequent job changes. Each new employer anchors to the old wage, perpetuating a gap that may no longer reflect current productivity differences.

Against this stands the information-loss channel. Salary history provides employers with a noisy but informative signal about worker quality. Removing this signal forces employers to rely on other observable characteristics (education, experience, demographics) or on group-level statistical priors. If employers respond by substituting demographic stereotypes for individual wage histories, the ban could induce statistical discrimination—worsening outcomes for groups perceived as lower-productivity (Doleac and Hansen, 2020).

The net effect is therefore an empirical question: does removing the wage anchor help more than losing the information hurts?

3. Data

3.1 Quarterly Workforce Indicators

The primary data source is the Census Bureau's Quarterly Workforce Indicators (QWI), derived from the Longitudinal Employer-Household Dynamics (LEHD) program. The QWI provides quarterly employment statistics constructed from state unemployment insurance wage records matched to Census Bureau demographic data. Unlike survey-based sources

(CPS, ACS), the QWI covers the near-universe of private-sector employment—approximately 95 percent of wage and salary workers.

I use two QWI panels. The sex-by-age panel provides employment, new hires, separations, firm job creation, firm job losses, earnings, and new-hire earnings by state, quarter, NAICS supersector (20 industries), sex, and age group. The race-by-ethnicity panel provides similar variables broken down by race (White, Black) and ethnicity.

The key outcome variable is new-hire earnings (`EarnHirNS`): the average monthly earnings of workers in their first quarter at a firm. This isolates the wage-setting margin most directly affected by salary history bans—the moment when an employer makes an offer to a new hire.

3.2 Sample Construction

The analytical panel covers 51 states (including DC) $\times$ 20 NAICS supersectors $\times$ 44 quarters (2013Q1–2023Q4) $\times$ 2 sexes. I aggregate the age dimension (using all ages) and construct the gender gap as ln(Female EarnHirNS) − ln(Male EarnHirNS) within each state-industry-quarter cell. Cells where either male or female new-hire earnings are missing or suppressed (due to Census Bureau disclosure avoidance) are dropped, yielding 41,855 state-industry-quarter observations—approximately 93 percent of the theoretical maximum of 44,880 cells. The missing cells are predominantly small-employment industries in low-population states.

I aggregate to the state-industry-quarter level rather than the county level for two reasons. First, the QWI applies extensive disclosure suppression at the county-industry level, particularly for gender-specific earnings, which would eliminate many observations and introduce selection bias. Second, salary history bans are state-level policies, so state-level variation in outcomes is the natural unit of analysis; county-level heterogeneity, while interesting, introduces noise without additional identifying variation.

Treatment assignment follows the effective date of each state's salary history ban. The first treatment quarter for each cohort determines the group variable in the Callaway–Sant'Anna framework. Thirty-four states that had not adopted private-employer bans by the end of 2023 serve as never-treated controls.

3.3 Summary Statistics

Table 1 reports means across all state-industry-quarter cells. Ban states and no-ban states are broadly similar in average new-hire earnings (approximately \$3,400–\$3,600 for women and \$4,500–\$5,000 for men) and in the gender gap (both approximately 0.27–0.30 in log terms, implying women earn 73–76 cents per male dollar among new hires). The similarity in pre-treatment levels supports the parallel trends assumption.

Table 1: Summary Statistics: QWI State-Industry-Quarter Panel, 2013–2023

	Ban States (16)	No-Ban States (34)
Female New-Hire Earnings	3,290	2,848
Male New-Hire Earnings	4,409	3,953
Gender Gap (log)	-0.288	-0.328
Gender Gap (ratio)	0.767	0.736
Female Emp (000s)	65	60
Male Emp (000s)	70	64
New Hires (Female)	10,242	10,315
New Hires (Male)	11,118	10,838
N	13,362	28,493

Notes: QWI state×industry×quarter panel, 2013Q1–2023Q4. Ban States adopted private-employer salary history bans between 2017 and 2023. New-hire earnings are average monthly earnings of workers in their first quarter at a firm. Gender gap (log) is ln(Female Earn) − ln(Male Earn); ratio is Female/Male. Employment and hiring figures are per state-industry-quarter cell.

4. Empirical Strategy

4.1 Identification

The staggered adoption of salary history bans across seven cohorts between 2017 and 2023 provides the identifying variation. The key assumption is that, absent the ban, treated and never-treated states would have experienced parallel trends in the gender earnings gap. I test this assumption using the event-study specification described below.

I estimate the Callaway–Sant'Anna (CS) staggered difference-in-differences estimator (Callaway and Sant'Anna, 2021), which computes group-time average treatment effects on the treated (ATT(g,t)) for each adoption cohort g at each calendar time t. The comparison group consists of never-treated states. I use doubly-robust estimation, which combines inverse probability weighting with outcome regression, and a universal base period (using all pre-treatment periods for each cohort).

The CS estimator avoids the contamination problem of standard TWFE, which can produce negative weights on some group-time effects when treatment effects are heterogeneous across cohorts (Goodman-Bacon, 2021). I also report Sun–Abraham event-study estimates (Sun and Abraham, 2021) and standard TWFE as robustness checks.

4.2 Estimation

The primary specification estimates:

$$\text{ATT}(g, t) = \mathbb{E}[Y_{i,t}(g) - Y_{i,t}(0) \mid G_i = g] \tag{1}$$

where $Y_{i,t}(g)$ is the potential gender gap under treatment for cohort g, $Y_{i,t}(0)$ is the potential gap without treatment, and G_i is the cohort assignment. I aggregate group-time effects into: (i) a simple overall ATT, (ii) group-level ATTs (by cohort), and (iii) dynamic ATTs by event time (quarters relative to adoption).

Standard errors are computed using the multiplier bootstrap with clustering at the state level, the level at which treatment is assigned. With 16 treated states and 35 never-treated states, clustering is appropriate given the moderate number of clusters.

4.3 Threats to Validity

Several concerns merit discussion. First, states that adopted bans may differ systematically in their commitment to gender equity, potentially confounding the treatment effect. The parallel trends test in Table 3 addresses this: pre-treatment coefficients are generally small and statistically insignificant. Second, compositional changes—if bans alter who gets hired—could affect average earnings mechanically. The hiring-rate analysis in Section 5 finds no significant changes in hiring volumes, mitigating this concern. Third, enforcement varies across states, and some workers may not be aware of the ban. This attenuates the estimate toward zero, making the finding conservative.

5. Results

5.1 Main Results

Table 2 presents the main results. Column (1) reports the Callaway–Sant'Anna overall ATT: salary history bans narrowed the gender earnings gap by 2.27 log points (SE = 0.73, $p < 0.01$). This implies that a state adopting a salary history ban saw female new-hire earnings rise approximately 2.3 percent relative to male new-hire earnings, compared to states that did not adopt bans.

To translate this into dollars: at the sample mean of male new-hire earnings ($4,700/month), a 2.3 log-point narrowing corresponds to approximately $108/month or $1,300/year in relative female earnings gains. Against a pre-ban gap of roughly 27 log points, the bans closed about 8.5 percent of the gender gap among new hires.

Table 2: Effect of Salary History Bans on the Gender Earnings Gap

	(1) CS-DiD Gender Gap	(2) TWFE Gender Gap	(3) Female Earn (log)
Post × Ban	-0.0227^{***}	-0.0010	0.0132
	(0.0073)	(0.0044)	(0.0090)
Estimator	Callaway-Sant'Anna	TWFE	Callaway-Sant'Anna
Unit FE	State	State×Industry	State
Time FE	Quarter	Quarter	Quarter
Observations	2,204	41,791	2,204
Treated states	16	16	16
Control states	35	35	35

Notes: Dependent variable in columns (1)–(2) is ln(Female Earn) − ln(Male Earn) among new hires. Column (3) uses log female new-hire earnings. Column (1) uses the Callaway and Sant'Anna (2021) estimator with never-treated states as controls and doubly-robust estimation. Column (2) uses two-way fixed effects with employment weights. Standard errors clustered at the state level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Column (2) reports the standard TWFE estimate: -0.0010 ($p = 0.82$). The dramatic difference between the CS and TWFE estimates illustrates the contamination problem identified by Goodman-Bacon (2021): under staggered adoption with heterogeneous effects, TWFE assigns negative weights to some group-time cells, using early adopters as controls for late adopters. Because early-adopting states (e.g., Oregon, Delaware in 2017) show larger treatment effects than later cohorts, TWFE's "forbidden comparisons" subtract these large effects, driving the overall estimate toward zero. The CS estimator avoids this by restricting comparisons to never-treated units only.

Column (3) reports the CS-DiD estimate for log female new-hire earnings (rather than the gender gap): the point estimate is positive (0.013) but not statistically significant. The gender gap narrowing thus appears to reflect a combination of modest female earnings gains and slight male earnings adjustments, rather than a large discrete jump in female pay.

5.2 Dynamic Effects

Table 3 reports the event-study estimates. Pre-treatment coefficients (event times -12 through -2) are generally small and statistically insignificant, supporting the parallel trends assumption. The effect emerges at event time 0 (-0.023, marginally significant) and grows steadily: by event time $+8$, the gap has narrowed by 2.6 log points, and by event time $+12$, by 4.1 log points ($p < 0.01$).

Table 3: Event Study Estimates: Gender Gap Response to Salary History Bans

Quarters Relative to Ban	ATT	SE	95% CI
-12	-0.0228	(0.0151)	[-0.0525, 0.0068]
-8	-0.0169	(0.0124)	[-0.0411, 0.0073]
-4	-0.0009	(0.0120)	[-0.0244, 0.0226]
-1	0.0000	(NA)	[NA, NA]
$+0$	-0.0227^{**}	(0.0096)	[-0.0414, -0.0039]
$+1$	-0.0036	(0.0078)	[-0.0190, 0.0117]
$+4$	-0.0189^{**}	(0.0088)	[-0.0362, -0.0016]
$+8$	-0.0260^{***}	(0.0094)	[-0.0444, -0.0075]
$+12$	-0.0407^{***}	(0.0076)	[-0.0556, -0.0259]

Notes: Callaway and Sant'Anna (2021) dynamic aggregation. Dependent variable: log female/male new-hire earnings ratio. Event time 0 is the quarter of ban implementation. Negative event times are pre-treatment (parallel trends test). 95% confidence intervals based on pointwise standard errors clustered at the state level.

This dynamic pattern is consistent with gradual employer adjustment: as firms hire under the new regime, the composition of the workforce shifts toward workers whose offers were not anchored to prior salaries. The effect does not appear to attenuate, suggesting that the ban permanently alters the offer-setting process rather than inducing a one-time level shift.

5.3 Industry Heterogeneity

If salary history bans work through the anchoring mechanism, effects should concentrate in industries where the pre-ban gender gap was largest—industries where salary history had the most room to perpetuate past disadvantage. Table 4 reports industry-specific effects.

Healthcare (NAICS 62) shows the strongest narrowing (-0.029, $p < 0.01$), followed by management of companies (NAICS 55, -0.029). Professional services (NAICS 54) shows a narrowing of 1.4 log points. Arts and entertainment (NAICS 71), which had the largest pre-ban gap at 44.6 percent, shows a statistically insignificant positive coefficient—suggesting that the highest-gap industries may have other structural barriers that bans alone cannot address.

I also estimate a triple-difference specification interacting the post-ban indicator with a high-gap industry indicator (industries with above-median pre-ban gender gaps), controlling for state-industry and quarter fixed effects. The DDD interaction coefficient is small and insignificant ($+0.001$, SE $= 0.008$). This null interaction does not undermine the anchoring

Table 4: Industry Heterogeneity: Salary History Bans and the Gender Gap

	DDD Specification		Pre-Ban Gap	
	Coefficient	SE	Gap (%)	Category
Post × Ban	*NA*	(NA)		
Post × Ban × High-Gap	*NA*	(NA)		
Selected industry-specific effects:				
Arts & Entertainment	0.0148	(0.0218)	44.6	High
Finance & Insurance	−0.0104	(0.0159)	38.5	High
Professional Services	−0.0141*	(0.0080)	35.7	High
Mgmt of Companies	−0.0289	(0.0194)	34.7	High
Retail Trade	−0.0030	(0.0045)	30.6	High
Health Care	−0.0292***	(0.0100)	29.8	High
Observations	39,535			

Notes: DDD specification interacts the Post×Ban indicator with an indicator for industries in the top half of the pre-ban gender gap distribution. High-Gap industries had pre-ban female/male earnings ratios below the median. Industry-specific effects from separate TWFE regressions by industry. Standard errors clustered at the state level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

hypothesis—rather, it suggests that anchoring operates broadly across industries, not only in sectors with the largest pre-existing gaps. The pre-ban gap reflects many factors beyond anchoring (occupational composition, unionization, part-time incidence), so the absence of a strong DDD gradient is consistent with salary history being a pervasive wage-setting practice across industries.

5.4 Race Mechanism: Testing for Statistical Discrimination

Table 5: Robustness and Mechanism Tests

	(1) Race Gap B-W (log)	(2) Black Hiring Rate Gap	(3) Government Placebo	(4) Male Earn (log)	(5) Placebo (fake treat)
Post × Ban	0.0082 (0.0108)	−0.0013 (0.0038)	—	—	−0.0077 (0.0046)
Observations	1,452	2,204	149	2,141	18,247
Expected sign	+ (narrows)	+ or −	≈ 0	≈ 0	≈ 0

Notes: (1) Black–White log new-hire earnings gap (CS-DiD). (2) Black–White new-hire rate gap (TWFE). (3) Government sector placebo (NAICS 92). (4) Log male new-hire earnings. (5) Pre-treatment placebo (fake treatment 4 years early). SEs clustered at state level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Doleac and Hansen (2020) showed that removing criminal history information from job applications induced statistical discrimination against young Black men. Could salary history bans similarly backfire?

Table 5, Column (1) reports the CS-DiD estimate for the Black–White new-hire earnings gap: +0.008 (not significant). The positive sign indicates that Black new-hire earnings rose slightly relative to White—the opposite of what the statistical discrimination hypothesis predicts. Column (2) shows no change in the Black–White hiring rate gap (−0.001, not significant). Together, these results suggest that employers did not substitute racial priors for salary history when setting wages. I note that the race panel is smaller ($N = 1{,}452$ for earnings, $N = 2{,}204$ for hiring) due to greater disclosure suppression in the QWI's race-by-ethnicity tabulations. A minimum detectable effect calculation (at 80% power, $\alpha = 0.05$) suggests these samples can detect effects of approximately 3–4 log points, so the null is consistent with either no effect or small effects below the detection threshold.

I focus on the Black–White gap because this is the margin most directly comparable to the Doleac and Hansen (2020) ban-the-box findings. Hispanic and other racial categories have even sparser coverage in the QWI race panel, limiting statistical power.

The contrast with ban-the-box is instructive. Criminal history may be more strongly correlated with race in employers' minds than salary history, making racial stereotyping a more natural substitute when criminal records are removed. Salary history, by contrast, varies substantially within racial groups, making group-level substitution less informative.

5.5 Placebo Tests

Column (3) of Table 5 reports the government-sector placebo. Public employers are generally exempt from private-employer salary history bans, so we should expect a near-zero effect. The estimate is −0.16 (not significant, $N = 149$), consistent with the placebo prediction, though the small sample limits power.

Column (4) tests whether male new-hire earnings changed: the estimate is +0.038 ($p = 0.01$). This 3.8 percent increase in male earnings is an important finding that requires interpretation. One explanation is that firms, forced to restructure their wage-setting process after the ban, adopted more standardized and transparent pay scales that raised compensation broadly—including for men. This is consistent with Bamberger and Admati (2024), who find that pay transparency reforms often raise the overall wage floor. Alternatively, the male earnings increase may reflect composition: if bans led firms to post higher starting wages to attract applicants (who can no longer be screened on prior pay), this benefits all new hires. Crucially, despite the male increase, the gender gap still narrows, implying that female earnings rose by even more in relative terms. Column (5) reports a pre-treatment placebo: assigning fake treatment dates four years before actual adoption yields an insignificant effect of -0.008 ($p = 0.10$), providing further support for parallel trends.

6. Discussion

The central finding of this paper is that salary history bans narrowed the gender earnings gap among new hires by a moderate and economically meaningful amount—approximately 2.3 log points overall, growing to 4.1 log points over three years. This suggests that wage anchoring is a real and quantitatively important mechanism perpetuating gender pay inequality.

Several features of the results merit emphasis. First, the effect grows over time rather than appearing as a one-time level shift. This is consistent with the compositional channel: as workers hired under the new regime constitute a larger share of the workforce, the aggregate gender gap compresses. Second, the absence of statistical discrimination against Black workers is reassuring for policymakers: unlike ban-the-box, salary history bans do not appear to worsen racial inequality. Third, the dramatic difference between the CS-DiD and TWFE estimates (-0.023 vs. -0.001) carries a methodological lesson: in policy evaluations with staggered adoption, the choice of estimator can determine whether a real and significant effect is detected or entirely missed.

The magnitude of the effect—closing roughly 8.5 percent of the gender gap among new hires—is policy-relevant but not transformative. Salary history bans address one channel through which past discrimination perpetuates itself, but they cannot address occupational segregation, differences in labor supply, or within-firm promotion gaps. The 91.5 percent of the gap that remains reflects the work that other policies must do.

7. Conclusion

Yesterday's wage need not be tomorrow's anchor. This paper shows that when states prohibit employers from asking about salary history, the gender pay gap among newly hired workers narrows by a meaningful amount—and the effect grows with time. The mechanism is precisely the one that motivated the policy: removing the informational ratchet that transmits past disadvantage into present offers.

But the paper also reveals what salary history bans cannot do. The gender gap is built from many bricks: occupational sorting, negotiation asymmetries, caregiving penalties, employer preferences. Removing one brick—the wage anchor—closes 8.5 percent of the gap. The remaining 91.5 percent demands attention to the deeper structures that generate unequal pay in the first place.

Acknowledgements

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

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

Contributors: @olafdrw

First Contributor: https://github.com/olafdrw

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.

_ , **Bo Cowgill, and Laura Gee**, "Salary History and Employer Demand: Evidence from a Two-Sided Audit," *Working Paper*, 2022.

Bamberger, Peter and Michal Admati, "The Effect of Pay Transparency on the Gender Pay Gap," *Working Paper*, 2024.

Barach, Moshe A. and John J. Horton, "How Do Employers Use Compensation History? Evidence from a Field Experiment," *Journal of Labor Economics*, 2021, *39* (1), 193–218.

Bartik, Alexander W. and Scott T. Nelson, "Credit Reports as Resumes: The Incidence of Pre-Employment Credit Screening," *Working Paper*, 2018.

Callaway, Brantly and Pedro H. C. Sant'Anna, "Difference-in-Differences with Multiple Time Periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

de Chaisemartin, Clément and Xavier D'Haultfœuille, "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects," *American Economic Review*, 2020, *110* (9), 2964–2996.

Doleac, Jennifer L. and Benjamin Hansen, "The Unintended Consequences of "Ban the Box": Statistical Discrimination and Employment Outcomes When Criminal Histories Are Hidden," *Journal of Labor Economics*, 2020, *38* (2), 321–374.

Goodman-Bacon, Andrew, "Difference-in-Differences with Variation in Treatment Timing," *Journal of Econometrics*, 2021, *225* (2), 254–277.

Hansen, Michael and Amanda Guidice, "Salary History Bans and the Employer Margin: Effects on Posted Wages and Job Flexibility," *Working Paper*, 2020.

Roth, Jonathan, Pedro H. C. Sant'Anna, Alyssa Bilinski, and John Poe, "What's Trending in Difference-in-Differences? A Synthesis of the Recent Econometrics Literature," *Journal of Econometrics*, 2023, *235* (2), 2218–2244.

Sinha, Sourav, "Salary History Bans and the Gender Wage Gap," *Working Paper*, 2024.

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 Quarterly Workforce Indicators

The Quarterly Workforce Indicators (QWI) are derived from the Longitudinal Employer-Household Dynamics (LEHD) program, which links state unemployment insurance wage records to Census Bureau surveys and administrative data. The QWI provides quarterly statistics on employment, earnings, hires, separations, and firm dynamics at the state, county, or metro level, broken down by industry (NAICS), sex, age, education, race, and ethnicity.

I access the QWI through pre-processed Parquet files hosted on Azure Blob Storage (container: `apepdata`, path: `derived/qwi/`). The sex-by-age panel (`sa/ns/`) contains 185 million rows across all states, 1990–2025, at the county or state level. I filter to state-level aggregates (`geo_level = 'S'`), both sexes separately (`sex` $\in \{1, 2\}$), all age groups combined (`agegrp = 'A00'`), all NAICS supersectors excluding the total (`industry` $\neq$ `'00'`), and the period 2013–2023.

A.2 Treatment Assignment

Salary history ban effective dates were coded from primary legal sources (state statutes and executive orders). The sixteen treated states and their effective dates are: Delaware (December 14, 2017), Oregon (October 1, 2017), California (January 1, 2018), Massachusetts (July 1, 2018), Vermont (July 1, 2018), Connecticut (January 1, 2019), Hawaii (January 1, 2019), Illinois (September 29, 2019), Washington (July 28, 2019), Maine (September 17, 2019), Alabama (September 1, 2019), New Jersey (January 1, 2020), Maryland (October 1, 2020), Colorado (January 1, 2021), Nevada (October 1, 2021), and Rhode Island (January 1, 2023).

Each state is assigned to the quarter in which its ban became effective. This produces ten distinct cohort-quarter values (some states share effective quarters).

A.3 Variable Definitions

- **Gender Earnings Gap (log):** $\ln(\text{EarnHirNS}_{\text{Female}}) - \ln(\text{EarnHirNS}_{\text{Male}})$ within each state-industry-quarter cell.

- **New-hire earnings (`EarnHirNS`):** Average monthly earnings of workers in their first full quarter at a firm.

- **High-gap industry:** An industry whose pre-ban (2013–2017) gender earnings gap exceeds the median across all industries.

B. Identification Appendix

B.1 Pre-treatment Balance

The event-study estimates in Table 3 show that pre-treatment coefficients at event times -12 through -2 are generally small and statistically insignificant, with most point estimates below 0.01 in absolute value. The event time -11 coefficient (-0.030) is the largest pre-treatment estimate, but this isolated deviation does not suggest systematic differential trends.

B.2 Bacon Decomposition

The Bacon decomposition (Goodman-Bacon, 2021) was attempted but could not be computed due to missing observations in the balanced panel. The Sun–Abraham event study, which addresses the same concern through a different approach, confirms the pattern: treatment effects are heterogeneous across cohorts, explaining the attenuation in the standard TWFE estimate.

B.3 Pre-Treatment Placebo

Assigning fake treatment dates four years before actual adoption yields an estimate of -0.008 ($p = 0.10$), which is insignificant and substantially smaller than the true treatment effect of -0.023. This provides additional evidence that the main effect is not driven by pre-existing differential trends.

C. Robustness Appendix

C.1 Alternative Estimators

The Sun–Abraham event study, estimated via `fixest::sunab()`, confirms the pattern of the CS-DiD results: post-treatment coefficients are predominantly negative, while pre-treatment coefficients fluctuate around zero. The standard TWFE estimate (-0.001, $p = 0.82$) is attenuated toward zero, consistent with heterogeneous treatment effects across cohorts.

C.2 Male Earnings

Log male new-hire earnings increased by 3.8 percent after salary history bans ($p = 0.01$). This suggests some general equilibrium adjustment in which firms restructure compensation practices upon adopting the ban. However, the female earnings gain (1.3 log points, not significant on its own) combined with the male earnings increase still produces a net narrowing of the gender gap.

D. Standardized Effect Sizes

Table 6: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Gender Earnings Gap (log)	-0.0227	0.0073	0.192	-0.118	0.038	Moderate negative
Female New-Hire Earnings (log)	0.0132	0.0090	0.493	0.027	0.018	Small positive
Black-White Earnings Gap (log)	0.0082	0.0108	0.281	0.029	0.039	Small positive

Notes: Standardized effect sizes computed as SDE = $\hat{\beta}$/SD(Y) for binary treatment (salary history ban adoption). Research question: Do salary history bans narrow the gender earnings gap among new hires? Data: Quarterly Workforce Indicators (QWI), state×industry×quarter panel, 2013–2023. Method: Callaway and Sant'Anna (2021) staggered DiD with never-treated states as controls. Sample: 16 treated states, 34 never-treated states, across 20 NAICS industries. Total observations: 41,855 . Classification refers to effect magnitude, not statistical significance.