Do Salary-Range Disclosure Laws Raise Earnings and Reduce Pay Inequality?
Evidence from the United States

Master's Thesis

Begyere Bartholomeo

Supervisor

Henry Antonio Redondo Gómez

University of Alicante

September 2026

Abstract

Salary-range disclosure laws reduce information asymmetry by giving workers information about compensation before they apply for a job or negotiate an offer. This information may help workers assess job opportunities more accurately, although the benefits may differ across workers because mobility, bargaining opportunities, and access to outside options are not the same for everyone.

This thesis examines whether state salary-range disclosure laws in the United States raise real annual earnings and whether the gains are broadly shared across workers. Using approximately 19 million employed adults from the 2010–2024 American Community Survey and variation in the timing of state salary-posting mandates, the analysis finds that disclosure is associated with approximately 0.99% higher real annual wage and salary earnings.

The distributional results, however, do not show a systematic narrowing of the main gender and racial or ethnic earnings differentials. The results therefore suggest that better wage information can be associated with higher average earnings without necessarily reducing existing demographic earnings inequality.

Keywords: pay transparency; salary-range disclosure; annual earnings; earnings inequality; American Community Survey.

1. Introduction

Information is central to wage determination, yet workers and employers do not enter the employment relationship with equal knowledge. Employers usually know the compensation budget for a vacancy, the wages paid for comparable work, and the scope for negotiation. Applicants often observe only a fragment of the relevant wage distribution and may misperceive what they could earn elsewhere. Jäger et al. (2024) showed that workers can hold systematically inaccurate beliefs about outside wages, particularly in low-paying firms. When outside options are poorly understood, search may be directed toward less productive matches and bargaining may reproduce existing earnings differences. This informational asymmetry motivates recent laws requiring employers to disclose salary ranges before hiring negotiations are completed.

The policy question extends beyond whether disclosure raises average earnings. New information may be valuable only when workers can act on it. Mobility constraints, occupational segregation, discrimination, professional networks, and unequal access to credible outside offers can shape the capacity to search, bargain, or leave an employer. Salary disclosure may therefore raise earnings broadly, concentrate gains among workers with stronger outside options, or alter gender and racial or ethnic earnings differentials. A change in a differential is itself incomplete evidence because it can result from faster growth for the lower-earning group or slower growth for the reference group. The relevant distributional analysis must consequently examine average earnings, demographic differentials, and the earnings of the constituent groups together.

Existing research establishes that pay information can affect wage setting, but it does not fully resolve how the gains from statewide vacancy-level disclosure mandates are distributed across workers. Evidence from firm-level reporting and public-sector disclosure has largely focused on gender pay differences, while theoretical work shows that information about coworkers' pay and information about pay across employers can operate through different search and bargaining channels (Bennedsen et al., 2022; Cullen and Pakzad-Hurson, 2023; Cullen, 2024). Arnold, Quach and Taska (2025) provide the closest evidence on US salary-posting mandates, documenting increased salary disclosure and higher wages. Building on this evidence, the central question of this thesis is whether statewide salary-range disclosure laws raise real annual earnings and, importantly, whether those gains are broadly shared across workers.

The empirical analysis has three parts. First, it estimates the average earnings response after implementation. Second, it examines raw and residualized men-women and White-minority earnings differentials together with group-specific earnings. This makes it possible to separate changes in earnings gaps from general changes in earnings. Third, the analysis considers differences by educational attainment.

The study uses approximately 19 million individual observations from the 2010-2024 American Community Survey obtained through IPUMS USA. The sample consists of employed adults aged 18-64 with positive wage and salary income and valid person weights. Nominal annual earnings are converted to 2024 dollars. Individual records are used to construct population-weighted state-year outcomes and state-year demographic and education-group outcomes. The ACS provides large and consistently defined state samples, but it is a repeated cross-section rather than a panel of workers or firms. Accordingly, the estimates identify changes in state-level outcome distributions around policy adoption; they do not follow the same worker through a negotiation or observe employer compliance directly.

Identification exploits the staggered implementation of salary-range disclosure requirements across states. The analysis uses the Callaway and Sant'Anna (2021) staggered difference-in-differences estimator, which accounts for states introducing the law in different years by comparing each group of adopting states with states that are untreated in the corresponding year. This design avoids using already-treated states as controls for later adopters when treatment effects may differ across cohorts or over time. Event-study estimates assess dynamics before and after implementation, while a conventional two-way fixed-effects model is retained as a benchmark. Inference is clustered at the state level and supplemented by wild-cluster bootstrap tests because only a small number of jurisdictions are observed after treatment by 2024. Alternative treatment timing, weighting, placebo, and leave-one-treated-state-out exercises examine the sensitivity of the findings.

The baseline estimates indicate that the introduction of salary-range disclosure laws is associated with an increase of approximately 0.99% in real annual earnings. This positive earnings effect is generally maintained across the main sensitivity specifications, although the distributional effects are less uniform. In particular, the estimated effects on the raw men-women, White-Black, White-Hispanic, and White-Asian earnings differentials are not statistically significant at the 5% level. Accounting for observable worker characteristics produces a somewhat different pattern. The residualized men-women differential increases significantly following implementation, while the corresponding racial and ethnic differentials remain statistically insignificant. The group-specific estimates similarly point to positive earnings responses for some demographic groups but provide little evidence of a general narrowing of existing earnings differentials. The results also reveal differences across educational groups. Rather than increasing with educational attainment, the estimated effects are concentrated among workers with a high-school diploma or GED, for whom the effect is positive and statistically significant. By contrast, the estimated effect for workers with a bachelor's degree or higher is not statistically significant.

This thesis adds to the literature in three ways. First, it studies a recent labor-market information policy using nationally representative microdata and a method that allows for staggered policy adoption. Second, it considers the distribution of the estimated earnings effect using raw gaps, residualized gaps, and group-specific outcomes. Third, it examines differences by education.

The remainder of the thesis is organized as follows. Chapter Two describes the development and design of salary-range disclosure laws, explains the economic mechanisms through which they may affect earnings, and situates the analysis within the related evidence. Chapter Three presents the ACS data, outcome construction, treatment definitions, empirical models, identifying assumptions, and sensitivity analyses. Chapter Four develops the empirical results from the average earnings response to demographic incidence and educational heterogeneity. Chapter Five concludes and discusses the interpretation, limitations, and policy implications of the evidence.

2. Institutional background and related evidence

2.1 Salary-range disclosure laws

US salary-range disclosure laws developed from a broader history of equal-pay and pay-information regulation. The Equal Pay Act of 1963 prohibited sex-based wage discrimination for equal work under specified conditions, but it did not require employers to disclose compensation before hiring (US Equal Employment Opportunity Commission, 1963). Executive Order 13665 later protected employees and applicants of covered federal contractors from retaliation for discussing, disclosing, or asking about compensation (US Department of Labor, 2015). These measures improved legal protection around pay, but they did not generally tell applicants what an employer expected to pay for a particular vacancy.

The recent state laws examined in this thesis move transparency to the recruitment stage. Colorado became the first state in the analytical sample to require compensation information in job postings, effective January 1, 2021. California and Washington followed on January 1, 2023, and New York followed in September 2023. Hawaii adopted a statewide requirement in January 2024, the District of Columbia in June 2024, and Maryland in October 2024. Other states introduced similar requirements after the analytical period. The staggered timing of these reforms provides the policy variation used in the empirical analysis.

The laws share a common intervention but differ in their design. California and Washington generally apply their posting requirements to employers with at least 15 employees, while New York uses a lower employer-size threshold. Some states require information on benefits or other compensation in addition to the salary range. Coverage of remote positions, internal opportunities, enforcement procedures, and penalties also differs across jurisdictions (California Department of Industrial Relations, 2023; Washington State Department of Labor & Industries, 2023; New York State Department of Labor, 2023; Colorado Department of Labor and Employment, 2024). These differences may affect employer coverage and compliance.

Despite this variation, the central feature of the laws is the same. Employers must reveal compensation information earlier in the recruitment process. Applicants can use this information before applying or negotiating. Incumbent workers may also compare their current earnings with the ranges offered for similar jobs. Salary-range disclosure differs from salary-history bans, protections for pay discussion, and aggregate gender-pay reporting. The policy studied here is primarily a pre-employment information intervention.

Figure 2.1 summarizes the evolution of US pay-transparency policy and the timing of statewide job-posting mandates. It shows the federal legal foundations and the later clustering of state adoption after Colorado. The figure also distinguishes effective dates from enactment dates, which is important because late-year implementation creates only partial exposure within an ACS survey year.

Figure 2.1. Evolution and diffusion of US pay-transparency policy

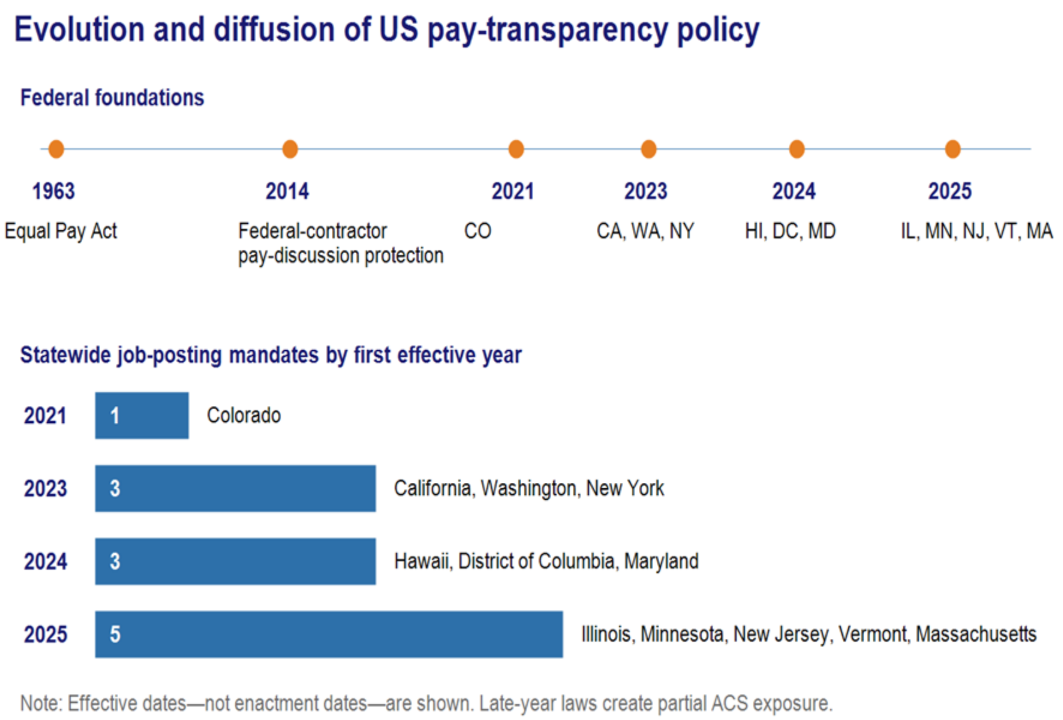

Note: *The figure reports effective dates rather than enactment dates. CO = Colorado; CA = California; WA = Washington; NY = New York; HI = Hawaii; DC = District of Columbia; MD = Maryland; IL =*

Illinois; MN = Minnesota; NJ = New Jersey; VT = Vermont; MA = Massachusetts. Laws effective after 2024 are outside the main 2010–2024 analytical treatment window.

The timing highlights the staggered nature of policy adoption. Colorado is the only treated jurisdiction for a relatively long post-treatment period in the analytical sample. The 2023 and 2024 adopters contribute shorter post-treatment histories. This feature is important for the event-study analysis because later event times are supported mainly by the earliest treatment cohort.

2.2 Economic mechanisms

Salary-range disclosure can reduce search and bargaining frictions by revealing the compensation attached to a vacancy before a worker accepts an offer. In standard search-and-matching models, workers compare the value of available matches with outside opportunities, so better information can change where workers search, which offers they pursue, and the wage at which a match is acceptable (Pissarides, 2000). Posted ranges can also provide a reference point during negotiation and allow incumbent workers to assess external opportunities more accurately.

The value of that information depends on the quality of a worker's outside options and on whether those options are correctly perceived. Jäger et al. (2024) document substantial misperceptions of outside wages and show that information about comparable workers' pay changes beliefs, intended job search, and wage-negotiation behavior. Workers with transferable skills, stronger networks, greater mobility, or access to many potential employers may therefore be better positioned to act on disclosed ranges. Workers facing limited labor demand, occupational segregation, discrimination, or mobility constraints may benefit less. This mechanism allows average earnings to rise even when gender or racial and ethnic earnings differentials do not narrow.

Employers may also adjust wage-setting behavior when compensation becomes more visible. Greater transparency can constrain arbitrary differences in offers to similar workers, but it can also induce firms to standardize offers or bargain differently. Cullen and Pakzad-Hurson (2023) show theoretically that some forms of transparency can reduce wage dispersion while weakening workers' bargaining position. Related evidence from Biasi and Sarsons (2022) shows that when Wisconsin school districts gained greater flexibility to negotiate individual pay, the gender wage gap widened among teachers with comparable credentials, partly because women were less likely to negotiate. These results illustrate why the distributional consequences of wage-setting institutions need not mirror their average effects.

These mechanisms motivate a sequence of empirical questions. The analysis first establishes whether salary-range disclosure changes average real annual earnings. It then asks whether the estimated gains are shared across demographic groups by examining raw and residualized earnings differentials and group-specific earnings. Finally, it studies heterogeneity by educational attainment as one observable dimension related to occupation, mobility, and potential outside opportunities.

Evidence from other pay transparency policies further suggests that the effects of disclosure depend on the design and coverage of the policy. In Denmark, Bennedsen et al. (2022) find that firm-level gender-pay reporting reduced the gender wage gap, with much of the reduction arising from slower wage growth among male employees. Similarly, Baker et al. (2023) show that public-sector salary disclosure in Canada contributed to a reduction in the gender pay gap within universities. While these studies provide evidence that greater pay transparency can influence wage outcomes, the policies they examine differ from salary-range disclosure laws in an important respect. They focus on the disclosure of realized or aggregated employee pay, whereas the laws considered in this study require employers to disclose expected salary ranges for job vacancies. The mechanisms through which these policies operate, as well as the workers directly exposed to them, may therefore differ.

The international policy landscape has also moved toward greater pay transparency. In the European Union, the Pay Transparency Directive requires employers to provide applicants with information on initial pay or the relevant pay range (European Parliament and Council of the European Union, 2023). The Directive also introduces stronger reporting and enforcement requirements. The International Labor Organization (2022) stresses the importance of clear policy coverage and effective implementation.

In the United States, wage information has historically been limited in online job postings (Batra et al., 2023). This provides an important basis for policies that require employers to disclose salary ranges in job advertisements. Recent evidence also shows that pay disclosure can affect workers' job-search decisions. These responses may differ across workers (Feng, 2024). Thus, salary-range disclosure may affect labor-market outcomes through both improved access to wage information and changes in job-search behavior.

2.3 Earnings inequality and the policy context

Salary-range disclosure laws emerged in a labor market with persistent demographic earnings differences. These differences form part of the policy context because workers cannot easily identify or challenge unequal compensation when employers control information about available pay. The descriptive figures below show how the main demographic earnings gaps evolved over 2010-2024. They provide context for the reforms.

Both figures use the 2010-2024 IPUMS American Community Survey analytical sample. Weighted mean log real annual earnings are calculated for each demographic group within the state and year. The state-level gaps are then aggregated using population weights. The measures are unadjusted. They reflect differences in education, occupation, hours, industry, labor-force attachment, migration, and other characteristics, as well as possible differences in wage rates.

Figure 2.2, Panel A, shows the raw men-women annual-earnings differential. The gap remained close to 0.35-0.36 in log-earnings units during much of the first half of the sample. It then declined more steadily and reached approximately 0.32 in 2024.

Figure 2.2. Raw demographic annual-earnings gaps, 2010–2024

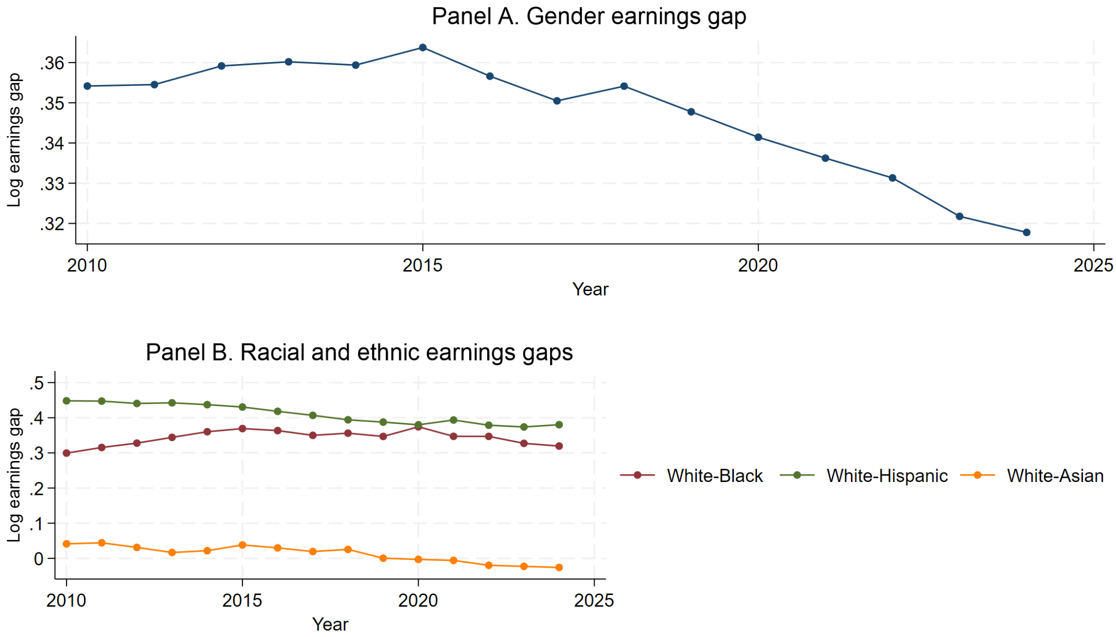

Note: Panels use population-weighted state-year mean log real annual wage and salary earnings from the 2010-2024 IPUMS ACS.

Panel A reports men minus women. Panel B reports non-Hispanic White minus non-Hispanic Black, Hispanic, and non-Hispanic Asian workers. Positive values indicate higher mean earnings for the reference group.

The decline in the gender gap began before Colorado's disclosure mandate took effect in 2021. Figure 2.2, Panel A, documents that the decline preceded Colorado's mandate and therefore does not identify a causal policy effect. The empirical design separates this underlying movement from changes associated with policy adoption.

Figure 2.2, Panel B, extends the descriptive analysis to racial and ethnic earnings differentials. The White-Hispanic differential is the largest over most of the period, declining from approximately 0.45 in 2010 to 0.38 in 2024. The White-Black differential first rises and later falls, ending near 0.32. The White-Asian differential is smaller and becomes negative near the end of the period.

The three racial and ethnic differentials do not follow the same path. This is important for the empirical analysis. A policy can raise average earnings without narrowing every demographic gap. A gap can also narrow because the lower-paid group gains or because the reference group grows more slowly. The institutional evidence motivates three empirical questions. The first is whether salary-range disclosure is associated with higher real annual earnings. The second is whether the policy changes gender and racial earnings differentials. The third is whether the earnings response differs across education groups.

3. Data and empirical strategy

3.1 American community survey and sample construction

The empirical analysis uses repeated cross-sections from the 2010-2024 American Community Survey (ACS), obtained through IPUMS USA (2026). The ACS provides large annual samples of the US resident population and contains consistent information on state of residence, annual wage and salary income, age, sex, race, Hispanic origin, educational attainment, occupation, industry, and person weights. Its large state samples make it suitable for constructing annual earnings outcomes and demographic earnings differentials at the state-year level.

The ACS does not follow the same workers, employers or firms over time. Each survey year is a new cross-section. The analysis describes changes in state-level earnings distributions around policy implementation rather than changes experienced by the same individuals. The absence of employer identifiers also prevents direct measurement of firm compliance, within-firm salary bands or employer-specific wage adjustment.

The analytical sample consists of adults aged 18-64 who were employed at the survey date, reported positive wage and salary income during the preceding 12 months and had a positive ACS person weight. Wage and salary income is retained as the earnings concept because the disclosure laws principally regulate employer-employee relationships; self-employment income is not substituted when wage income is absent. Nominal earnings are converted into 2024 dollars using the Consumer Price Index and transformed using the natural logarithm.[1]

$$y_{ist} = \ln\left(\text{Annual wage and salary income}_{ist} \times \frac{CPI_{2024}}{CPI_t}\right) \tag{1}$$

Where i indexes individuals, s indexes states, and t indexes survey years. The dependent variable y is the logarithm of real annual wage and salary earnings.

After applying the restrictions, the sample contains 18,979,949 person records distributed across 765 state-year cells. Person weights are used in every aggregation so that each cell represents the relevant resident working population rather than the unweighted survey sample.

3.2 Outcome and variable construction

The principal outcome is the person-weighted state-year mean of log real annual wage and salary earnings:

$$\overline{y}_{st} = \frac{\sum_{i\in(s,t)} \omega_{ist} y_{ist}}{\sum_{i\in(s,t)} \omega_{ist}} \tag{2}$$

Where ω_{ist} denotes the ACS person weight.[2] Distributional outcomes are constructed as differences between weighted group means within each state and year.

The gender annual-earnings differential is:[3]

[1]Expressing earnings in 2024 dollars removes changes in the general price level over time, so the outcome reflects real rather than nominal annual earnings.

[2]ACS person weights are used so that the state-year estimates represent the relevant resident working population rather than the raw survey sample.

[3]The raw demographic differentials are differences in weighted mean log real annual earnings. A positive value indicates higher mean earnings for the reference group, while a negative value indicates higher mean earnings for the comparison group.

$$G_{st}^{MF} = \overline{y}_{st}^{M} - \overline{y}_{st}^{F} \tag{3}$$

The White-Black differential is:

$$G_{st}^{WB} = \overline{y}_{st}^{W} - \overline{y}_{st}^{B} \tag{4}$$

The corresponding White-Hispanic and White-Asian outcomes are:

$$G_{st}^{WH} = \overline{y}_{st}^{W} - \overline{y}_{st}^{H} \tag{5}$$

$$G_{st}^{WA} = \overline{y}_{st}^{W} - \overline{y}_{st}^{A} \tag{6}$$

White, Black and Asian workers are restricted to non-Hispanic respondents. Hispanic workers are classified by Hispanic origin irrespective of race, producing mutually exclusive analytical groups. Positive values in Equations (4)–(6) denote higher mean earnings among non-Hispanic White workers; negative values denote higher mean earnings among the comparison group.

A differential may change because either constituent group changes. The analysis also constructs separate state-year mean outcomes for men, women, non-Hispanic White workers, non-Hispanic Black workers, Hispanic workers and non-Hispanic Asian workers. These outcomes distinguish convergence caused by gains among a lower-paid group from convergence caused by weaker earnings among the reference group.

To distinguish raw earnings differences from differences associated with observed worker characteristics, the analysis also constructs Mincer-style residualized earnings gaps.[4] Log real annual earnings are regressed on age, age squared, educational attainment, occupation, usual weekly hours, and calendar-year indicators. Sex and race or ethnicity is excluded because they define the demographic comparisons of interest:

$$y_{ist} = \alpha + \beta_1 Age_{ist} + \beta_2 Age_{ist}^2 + \gamma' Educ_{ist} + \delta' Occ_{ist} + \theta Hours_{ist} + \lambda_t + \varepsilon_{ist} \tag{7}$$

The residualized differential for groups A and B is the difference between their person-weighted mean residuals within state s and year t:

$$RG_{st}^{A-B} = \overline{\varepsilon}_{Ast} - \overline{\varepsilon}_{Bst} \tag{8}$$

These residualized outcomes are used as robustness measures for the men-women, White-Black, White-Hispanic, and White-Asian differentials. They show whether the estimated patterns remain after accounting for observable differences in age, education, occupation, hours worked, and common year effects. They do not remove unobserved worker or employer characteristics.

Educational attainment is divided into four categories: less than high school; high school diploma or GED; some college or an associate degree; and bachelor's degree or higher. For every state-year-education cell, the analysis calculates mean log real annual earnings and demographic and employment-composition

[4]Residualized earnings gaps measure differences that remain after accounting for the observed characteristics included in the earnings regression. They should not be interpreted as removing differences associated with unobserved worker or employer characteristics.

measures. Occupations and industries are ranked using pre-policy mean earnings and divided into low-, middle- and high-paying tertiles, preventing post-treatment outcomes from defining the sorting measures.

3.3 Pre-policy earnings, dispersion and sorting

Pre-policy statistics are calculated over 2010-2020, ending before Colorado implemented the first statewide mandate in 2021. The descriptive framework examines two possible sources of heterogeneous effects. Differences in initial earnings dispersion may affect the scope for compensation adjustment, and apparent education effects may reflect sorting into occupations and industries with different wage-setting practices

Table 3.1. Pre-policy annual earnings, dispersion, and raw differentials by education

Education	**Mean log earnings**	**SD**	**P90-P10**	**Men-women**	**White-Black**
Less than high school	10.059	1.004	2.254	0.462	0.179
High school	10.338	1.002	2.272	0.395	0.247
college	10.451	1.041	2.442	0.387	0.190
Bachelors	11.134	0.973	2.178	0.430	0.226

Notes: Statistics are calculated with ACS person weights for 2010–2020. P90–P10 is the difference between the 90th and 10th percentiles of log real annual wage and salary earnings. Gender and racial differentials use the group definitions in Section 3.2. Source: Calculations based on IPUMS USA.

The descriptive statistics document substantial differences in earnings levels across education groups but do not show uniquely greater dispersion among graduates. Workers with a bachelor's degree or higher have the highest mean log earnings but the smallest P90-P10 range. The some-college or associate-degree group has the widest percentile range. The raw gender differential is largest among workers without a high-school qualification, while the largest White-Black differential appears among high-school/GED workers. These patterns provide descriptive context for the heterogeneity analysis; they do not constitute estimates of the policy effect.

3.4 Empirical strategy

The empirical analysis exploits differences in the timing of state salary-range disclosure mandates. States introduced qualifying requirements in different years, while most jurisdictions had not implemented a comparable statewide mandate by the end of the analytical period. This staggered adoption permits comparisons between changes in adopting states and contemporaneous changes in states that had not yet adopted or never adopted the policy.

The analysis proceeds in four stages. Conventional two-way fixed-effects specifications provide transparent benchmarks. The principal estimates then use the Callaway and Sant'Anna (2021) staggered difference-in-differences design that avoids treating already-exposed states as controls for later adopters. Event-study estimates describe adjustment before and after implementation. Finally, the analysis evaluates identifying assumptions using alternative treatment definitions, placebo assignments, leave-one-state-out estimates and inference procedures designed for a limited number of treated state clusters.

3.4.1 Conventional difference-in-differences

The conventional difference-in-differences model provides a transparent benchmark for the direction and magnitude of the average association. It compares the change in outcomes after 2020 among states that adopt a qualifying mandate by 2024 with the contemporaneous change among states that remain untreated throughout the sample. Let $Adopter_s$ equal one for eventual adopters and let $post2020_t$ equal one from 2021 onward. Equation (9) implements this comparison while controlling for permanent state differences, common year shocks, and observed state-year composition.

$$Y_{st} = \alpha_s + \lambda_t + \beta \left(\text{Adopter}_s \times \text{Post2020}_t\right) + X_{st}^{'}\gamma + \varepsilon_{st} \tag{9}$$

In Equation (9), Y denotes the state-year outcome, and denote state and year fixed effects, X contains observed state-year covariates, and is the error term. The coefficient is the difference between the average pre-to-post change among eventual adopters and the corresponding change among never-treated states. A causal interpretation requires conditional parallel trends, no anticipation, no cross-state spillovers, correct classification of exposure, and either homogeneous treatment effects or an aggregation rule that identifies a clearly defined average effect.

The pooled specification is restrictive in this setting because the common post period begins in 2021, whereas qualifying mandates become effective in 2021, 2023, and 2024. It consequently classifies later adopters as post-policy before their laws take effect and combines genuinely treated observations with not-yet-treated observations inside the adopter group. It also imposes a single coefficient on effects that may differ by adoption cohort and time since implementation. For these reasons, the pooled estimate is interpreted as a benchmark rather than as the principal causal estimand.

A conventional model that respects the state-specific implementation date replaces the pooled interaction with a time-varying treatment indicator:

$$Y_{st} = \alpha_s + \lambda_t + \beta D_{st} + X_{st}^{'}\gamma + \varepsilon_{st} \tag{10}$$

The state-specific model improves timing accuracy by defining D = 1[t G], where G is the first effective treatment year. Its coefficient compares changes associated with actual implementation rather than with a common national break. However, the conventional two-way fixed-effects estimator combines multiple two-group, two-period comparisons. Already-treated states may serve as controls for later adopters, and heterogeneous effects can receive negative or otherwise nontransparent weights (Goodman-Bacon, 2021). The estimator is therefore informative as a conventional benchmark but does not, by itself, isolate a well-defined average effect when treatment responses vary across cohorts or horizons.

The corresponding conventional event-study specification is:

$$Y_{st} = \alpha_s + \lambda_t + \sum_{k \neq -1} \beta_k 1\left(t - G_s = k\right) + X_{st}^{'}\gamma + \varepsilon_{st} \tag{11}$$

The conventional event study replaces the single treatment indicator with leads and lags indexed by k = t-G and uses the year immediately before implementation, k = -1, as the reference period. The leads describe whether treated and comparison states were already evolving differently before implementation, while the lags trace the timing of the post-policy response. Under conditional parallel trends, no anticipation, and cohort-invariant dynamics, measures the outcome difference at event time k relative to the omitted year. With staggered adoption and heterogeneous effects, however, already-treated states can enter the comparison for later cohorts, causing the coefficients to combine cohort-specific responses with nontransparent or negative weights (Sun and Abraham, 2021). The conventional event study is therefore presented as a diagnostic benchmark, while the Callaway and Sant'Anna estimator supplies the main dynamic estimates.

3.4.2 Treatment definition and timing

Treatment begins when a statewide rule requires employers to include a wage or salary range in the job advertisement itself. This definition isolates policies that disclose employer-controlled compensation information before application and bargaining decisions are completed. Laws requiring disclosure only upon request or at a later recruitment stage are excluded because they intervene at a different point in the matching process and therefore do not represent the same treatment.

Table 3.2. Treatment timing under baseline and full-exposure definitions

Jurisdiction	**Enacted/approved**	**Effective date**	**Baseline year**	**Full-exposure year**
Colorado	May 22, 2019	January 1, 2021	2021	2021
California	September 27, 2022	January 1, 2023	2023	2023
Washington	March 30, 2022	January 1, 2023	2023	2023
New York	December 21, 2022	September 17, 2023	2023	2024
Hawaii	July 3, 2023	January 1, 2024	2024	2024
District of Columbia	January 12, 2024	June 30, 2024	Excluded	2025
Maryland	April 25, 2024	October 1, 2024	2024	2025
Illinois	August 11, 2023	January 1, 2025	Outside period	2025
Minnesota	May 17, 2024	January 1, 2025	Outside period	2025
New Jersey	November 18, 2024	June 1, 2025	Outside period	2026
Vermont	June 4, 2024	July 1, 2025	Outside period	2026
Massachusetts	July 31, 2024	October 29, 2025	Outside period	2026

Notes: The baseline analysis covers 2010–2024 and records treatment in the calendar year of implementation. The full-exposure definition identifies the first survey year more plausibly containing complete exposure, given the ACS rolling 12-month income-reference period. Enactment and effective dates are compiled from the official bill histories and labor-agency guidance for Colorado SB19-085, California SB1162, Washington SB5761, New York S9427-A/A10477, Hawaii SB1057, DC Act 25-367, Maryland HB649, Illinois HB3129, Minnesota SF3852, New Jersey P.L. 2024 c.91, Vermont H.704 and the Massachusetts Salary Range Transparency Act.

The ACS reports wage and salary income received during the preceding 12 months. Observations recorded in the calendar year of a midyear or late-year reform may therefore combine pre- and post-policy earnings. This concern is limited for January reforms but potentially important for New York, the District of Columbia and Maryland. The baseline convention is retained for reproducibility and to preserve available post-treatment information. A conservative sensitivity analysis shifts late-year reforms to the first plausible full-exposure year and excludes ambiguous state-years in an additional specification. The baseline coding retains Maryland as treated in 2024 under the calendar-year implementation convention, whereas the District of Columbia is omitted because its June 2024 implementation was not incorporated in the baseline treatment file. This asymmetry reflects the reported baseline coding rather than a claim that Maryland had greater 2024 exposure. The full-exposure sensitivity addresses the issue by requiring a survey year with plausibly complete exposure and therefore treats neither Maryland nor the District of Columbia as fully exposed within the 2010-2024 windows.

3.4.3 Staggered difference-in-differences

The principal design follows Callaway and Sant'Anna (2021). States are grouped into cohorts according to their first effective treatment year, g, and an average treatment effect is estimated for each cohort-year pair. The estimand ATT (g,t) is the average difference, in calendar year t, between the observed outcome for cohort g and the outcome that the same cohort would have experienced without the mandate. This structure allows policy effects to differ across adoption cohorts and across years since implementation.

$$ATT(g,t) = E\left[Y_t(1) - Y_t(0)|G = g\right], t \geq g \tag{12}$$

The counterfactual outcome for an adopting cohort is inferred from jurisdictions that remain untreated in the same calendar year. Both never-treated and not-yet-treated jurisdictions may contribute to this comparison, whereas already-treated states are excluded from the control group. The estimator therefore compares each adopting cohort only with jurisdictions that are still eligible to represent untreated outcomes, removing the principal source of forbidden comparisons in the conventional staggered two-way fixed-effects model.

Cohort-time effects are reorganized by event time k = tg. The event-time effect is:

$$ATT_k = \sum_{g \in \mathcal{G}_k} \omega_{gk} ATT(g, g+k) \tag{13}$$

The cohort-year estimates are aggregated in two complementary ways. The overall ATT summarizes the average post-treatment response across the supported cohort-year comparisons and serves as the main estimate in Chapter Four. Event-time aggregation instead aligns cohorts by years relative to implementation and shows how outcomes evolve before and after the law. Equal-event-time and cohort-size-weighted summaries assign different importance to cohorts and horizons; they therefore answer different questions and are reported separately in Appendix B rather than interpreted as interchangeable estimates.

$$ATT_{0:3} = \frac{1}{4} \sum_{k=0}^{3} ATT_k \tag{14}$$

The event-study coefficients remain the basis for examining dynamics around implementation. Alternative summaries place different weights on post-treatment horizons and treatment cohorts and are therefore not treated as numerically equivalent to the main overall ATT. State-year cells are population weighted in the main analysis, while an equally weighted state-year specification is reported as a sensitivity check.

The composition-adjusted event-study specification can be summarized as:

$$Y_{st} = \alpha_s + \lambda_t + \sum_{k \neq -1} \beta_k D_{st}^k + X_{st}^{'} \gamma + \varepsilon_{st} \tag{15}$$

X_{st} Includes the state-year female, Black, Hispanic and bachelors-degree shares and mean age. These variables account for observable compositional changes and may improve precision. Because composition may itself respond to treatment, the unadjusted model remains the baseline and the adjusted specification is treated as a robustness exercise.

3.4.4 Identification assumptions

The central identifying condition is conditional parallel trends. In the absence of the mandate, adopting states must have experienced the same average outcome changes as jurisdictions eligible to serve as their contemporaneous controls:

$$E\left[Y_{st}(0) - Y_{s,t-1}(0) | G_s = g\right] = E\left[Y_{st}(0) - Y_{s,t-1}(0) | s \in C_{gt}\right] \tag{16}$$

Parallel trends concerns changes rather than equality of levels. Event-study leads provide a diagnostic: systematically nonzero pre-treatment estimates would weaken the design. Individual leads and joint restrictions are therefore reported, although failure to reject does not establish parallel trends and tests may have limited power (Roth et al., 2023).

No anticipation requires employers and workers not to alter relevant behavior before implementation:

$$Y_{st}(0) = Y_{st}(1), t < G_s \tag{17}$$

This assumption may be threatened because laws are announced before they become effective and firms may revise advertisements or salary structures in advance. The year immediately before implementation may therefore be imperfectly clean. Alternative timing and exclusion of ambiguous announcement-to-implementation observations are used to assess this concern.

The stable-unit-treatment assumption requires that one state's mandate does not materially change outcomes in untreated states. Remote work, national advertisements and multistate employers create plausible spillovers. If employers disclose ranges nationally after becoming covered in one jurisdiction, untreated-state workers may receive part of the information benefit, attenuating estimated treated-control differences.

Treatment consistency requires the intervention to be sufficiently comparable across adopters. The statutes share mandatory disclosure in advertisements but differ in employer-size thresholds, remote-work coverage, benefit disclosure, internal notice requirements, enforcement and penalties.

3.4.5 Inference with a limited number of treated states

Treatment varies at the state level. Millions of ACS records improve the precision of state-year outcome construction but do not create millions of independent treatment observations. Standard errors are clustered by state, permitting arbitrary serial correlation within states.

The baseline contains 51 state and District clusters but only six states are treated through 2024: Colorado, California, Washington, New York, Hawaii and Maryland. The District of Columbia is excluded from baseline treatment because its mid-2024 implementation was not incorporated. With few treated clusters, conventional cluster-robust Wald tests may over-reject even when the total number of clusters appears adequate. This baseline count follows the reported estimation coding and should be distinguished from sensitivity or placebo samples that use a broader eventual-adopter definition. Table A5 documents the coding audit, including the omission of the District of Columbia and the partial-year status of Maryland.

Joint inference on four conventional pre-treatment leads is therefore supplemented with a restricted wild-cluster bootstrap using Webb and Rademacher weights. The null hypothesis imposes that the four conventional leads equal zero, and the bootstrap approximates the finite-sample distribution by reweighting cluster-level residual contributions (Cameron et al., 2008; Roodman et al., 2019). This procedure improves finite-sample inference when few clusters are treated; it does not repair a genuine violation of parallel trends.

3.4.6 Sensitivity and robustness analyses

The sensitivity analysis examines whether the substantive conclusion depends on modeling, weighting, timing, or sample choices. Unadjusted estimates are compared with models that account for observed state-year composition. Population-weighted specifications, which emphasize the average person represented by the ACS, are compared with equally weighted specifications, which give each state-year the same influence. Alternative event-time and cohort-weighted aggregations assess whether the result is driven by the weighting of particular cohorts or post-treatment horizons.

Treatment timing is audited using a conservative full-exposure convention because the ACS annual-income measure covers the preceding 12 months. The appendix identifies late-year reforms for which a survey-year treatment indicator implies only partial exposure. A pseudo-treatment placebo assigns adoption six years earlier than the actual date. Table 4.1 also reports estimates excluding 2020 and excluding Colorado, the earliest adopter with the longest post-treatment history.

4. Results and discussion

This chapter presents the empirical results. It begins with the average earnings response and the event-study estimates. It then considers demographic earnings differentials and group-specific earnings before turning to differences by educational attainment. The main identification and robustness checks are discussed with the results, with additional details reported in Appendix A.

4.1 Average earnings response to salary-range disclosure

Table 4.1 reports the average treatment effect for treated jurisdictions under the baseline specification and three sensitivity checks.

Table 4.1. Effect of salary-range disclosure laws on real annual earnings

	(1) Baseline	**(2) Equal state-year weight**	**(3) Exclude 2020**	**(4) Exclude Colorado**
Salary-range disclosure	0.009844**	0.01343*	0.005707***	0.005707***
	(0.004355)	(0.00694)	(0.001442)	(0.001442)
Exact effect (%)	0.99	1.35	0.57	0.57

*Notes: Entries are average treatment effects on treated jurisdictions estimated using the Callaway and Sant'Anna method (CSDID). Standard errors are in parentheses. All specifications use not-yet-treated comparison states. Column (2) gives each state-year equal weight; Column (3) excludes observations from 2020; and Column (4) excludes Colorado, the earliest adopting state. Percentage effects use 100 × [exp(coefficient) 1]. * P < 0.10, ** p < 0.05, *** p < 0.01.*

Column 1: Do salary-range disclosure laws raise real annual earnings? The baseline estimate is 0.009844 and is statistically significant at the 5% level. Applying 100 × [exp(coefficient) - 1] gives an estimated increase of approximately 0.99%. Relative to the estimated counterfactual without the mandate, this suggests a modest improvement in real annual earnings following implementation. The outcome is the state-year mean of log real annual wage and salary earnings, so the estimate should not be interpreted as a change in hourly pay or as an identical gain for every worker. Annual earnings may change through wage rates, hours worked, employment duration, or workforce composition. A causal interpretation remains conditional on the identifying assumptions discussed in Section 3.4.

Column 2: Does the result depend on population weighting? Giving each state-year equal weight produces an estimated increase of 1.35%, compared with 0.99% in the baseline. This specification increases the relative influence of smaller states and therefore estimates a differently weighted average response. The positive association persists, but the larger standard error means that the estimate is significant only at the 10% level. The comparison suggests that the direction of the result is not specific to population weighting, while its magnitude and precision depend on how states contribute to the aggregate estimate.

Column 3: Does excluding 2020 change the estimated earnings response? Removing observations from 2020 produces an estimated increase of 0.57%, significant at the 1% level. The smaller estimate suggests a more modest earnings response under this sample restriction, while the positive association remains. Its greater statistical significance reflects the smaller reported standard error rather than a larger economic effect. Excluding 2020 changes the available pre-treatment comparisons and does not remove pandemic-related influences that may persist into subsequent years.

Column 4: Is the positive estimate driven solely by Colorado? Excluding Colorado, the earliest adopter and the state with the longest post-treatment history, gives an estimated increase of 0.57%, significant at the 1% level. The positive association therefore remains when Colorado is removed. However, this restriction also changes the treatment cohorts and post-treatment periods represented in the estimate, so it should not be read as evidence that later earnings patterns are the same across states.

The estimated effects range from 0.57% to 1.35% and are positive in all four specifications. Their size and statistical precision vary with the weighting and sample restrictions. The sensitivity checks therefore show that the positive association is not limited to the baseline specification, although the pre-treatment and placebo concerns discussed below and in Appendix A remain.

Arnold, Quach and Taska (2025) report wage increases of 1.3–3.6% following salary-posting mandates. The positive direction is consistent with the present estimate, although differences in outcomes, samples and specifications limit direct comparison of magnitudes. Cullen (2024) reviews evidence that information about pay across employers can support job search, wage negotiation and employer competition. These mechanisms provide possible explanations for higher earnings, but the ACS does not observe negotiations or outside offers and therefore cannot test them directly.

4.1.1 When do earnings effects emerge?

Figure 4.1. Event study for average real annual earnings using the main estimator

Notes: Points report estimated percentage changes in real annual earnings and bars show 95% confidence intervals. Later post-treatment horizons have fewer contributing treatment cohorts.

Figure 4.1 suggests that the positive earnings association appears in the implementation year and is larger at later post-treatment horizons. The estimated increase is approximately 0.50% in the implementation year and 0.85% one year later, rising to 5.09% after two years and 4.98% after three years.

Economically, the smaller initial estimates are consistent with gradual adjustment: workers may need time to use disclosed salary information when searching for jobs or negotiating compensation. The ACS annual-income measure may also dilute the initial estimate because earnings reported during the implementation year include income received before the law took effect. These are possible explanations rather than mechanisms directly demonstrated by the analysis.

However, the later estimates require particular caution. The two- and three-year estimates rely mainly on Colorado, the earliest adopter, whereas the earlier horizons include more adopting states. Consequently, the increase across event times may reflect differences between adopting states as well as adjustment over time. Figure 4.1 therefore does not establish that all adopting states experience earnings gains of approximately 5% after two or three years.

4.2 Are the average earnings gains broadly shared?

Following the positive average earnings estimate, this section asks whether earnings gains extend across demographic groups and whether their relative earnings gaps narrow. These are distinct questions: equal proportional gains can leave log earnings gaps unchanged. Table 4.2 examines demographic differentials, while Table 4.3 reports earnings responses for each group separately.

Table 4.2. Raw and residualized effects on demographic annual-earnings differentials

Differential	Raw ATT (1)	SE	Residualized ATT (2)	SE
Men–women	0.0065	0.0051	0.0082***	0.0022
White–Black	0.0045	0.0137	0.0102	0.0158
White–Hispanic	0.0121	0.0084	0.0101	0.0083
White–Asian	0.0244*	0.0145	0.0167	0.0114

*Notes: ATT denotes the average treatment effect on treated jurisdictions estimated using the main method. Standard errors are reported in separate columns. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Raw gaps are differences in weighted group mean log real annual earnings. Residualized gaps are differences in weighted mean residuals constructed using Equations (7) and (8). Positive values indicate an increase in the named reference-group differential. White, Black, and Asian groups are non-Hispanic.*

Table 4.2 provides no consistent evidence that the principal demographic earnings gaps narrow after implementation. It compares changes in raw earnings differentials with changes in residualized differentials, which account for observed differences in age, education, occupation, usual weekly hours and calendar-year effects.

For the raw differentials in Column 1, none of the raw estimates is statistically significant at the 5% level, so the table does not establish systematic changes in the observed demographic earnings gaps. The White–Asian estimate is significant only at the 10% level and is treated as weaker evidence. Statistical insignificance does not establish that the effects are zero.

After accounting for observed worker characteristics in Column 2, the residualized men–women estimate is 0.0082 log points, significant at the 1% level, indicating an increase in the adjusted differential in men's favour. Thus, the positive average earnings association is not accompanied by evidence of greater gender equality in this specification. The residualized racial and ethnic estimates are statistically insignificant. Residualization accounts for the included characteristics but does not isolate discrimination.

The demographic specifications reject their joint pre-treatment tests, weakening a causal interpretation of both the raw and residualized estimates. The findings should therefore be treated as patterns associated with implementation rather than firm evidence that disclosure caused demographic earnings gaps to change. Statistically insignificant estimates also do not establish that the effects are zero.

Appendix Figure D3 compares raw and residualized demographic earnings differentials over 2010–2024, illustrating how adjustment for observed worker characteristics changes their levels and trends.

Table 4.2 therefore provides little evidence of systematic convergence in demographic earnings gaps. Table 4.3 looks at each group's earnings separately to show what lies behind the estimated movements in the gaps.

Table 4.3. Group-specific effects on real annual earnings

Group	Coefficient	Effect (%)	SE
Men	0.01332***	1.34	0.00452
Women	0.00686	0.69	0.00544
White, non-Hispanic	0.00184	0.18	0.00752
Black, non-Hispanic	0.00631	0.63	0.01136
Hispanic	0.01398**	1.41	0.00701
Asian, non-Hispanic	0.02627	2.66	0.01855

*Notes: Effects are exact percentage changes in real annual wage and salary earnings. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.*

Table 4.3 reports an estimated earnings increase of 1.34% for men, significant at the 1% level, and 1.41% for Hispanic workers, significant at the 5% level under the reported inference. These estimates indicate positive annual-earnings associations for both groups. The estimates for women and non-Hispanic White, Black and Asian workers are statistically insignificant.

A significant estimate for one group and an insignificant estimate for another do not establish a significant difference between their responses. Direct tests of the differences are required. Table 4.3 therefore identifies groups with statistically significant earnings associations without establishing a ranking of benefits.

The results do not establish that gains are universal or equal across demographic groups. Together with Table 4.2, they show that positive earnings associations for particular groups need not imply narrower demographic earnings gaps. A causal interpretation remains subject to the study's identifying assumptions and the broader pre-treatment and placebo concerns.

4.3 Do earnings responses differ by educational attainment?

The final stage of the analysis examines whether the average earnings response differs across four education groups.

Table 4.4. Estimated earnings effects by education

Education group	Coefficient	Effect (%)	SE
Less than high school	0.01181	1.17	0.00847
High school/GED	0.01440***	1.45	0.00492
Some college/associate degree	0.00535	0.53	0.00486
Bachelor's degree or higher	0.00825	0.83	0.00768

*Notes: Percentage effects are exact transformations of the log coefficients. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.*

For workers with a high-school diploma or GED, the estimated increase in real annual earnings is 1.45% and is significant at the 1% level. The estimates for the other education groups are not statistically significant. This is the only education group with a statistically significant positive estimate, but the table does not show that its response is significantly different from the responses of the other groups.

The positive estimate is in the same direction as Arnold, Quach and Taska (2025), who report higher wages after salary-posting mandates. Cullen (2024) also discusses how information about pay across employers can affect job search, negotiation, and wage competition. These mechanisms may help explain a positive earnings response, but the present analysis cannot determine why the statistically significant estimate is concentrated among high-school/GED workers.

Figure 4.2. Event studies for real annual earnings by education

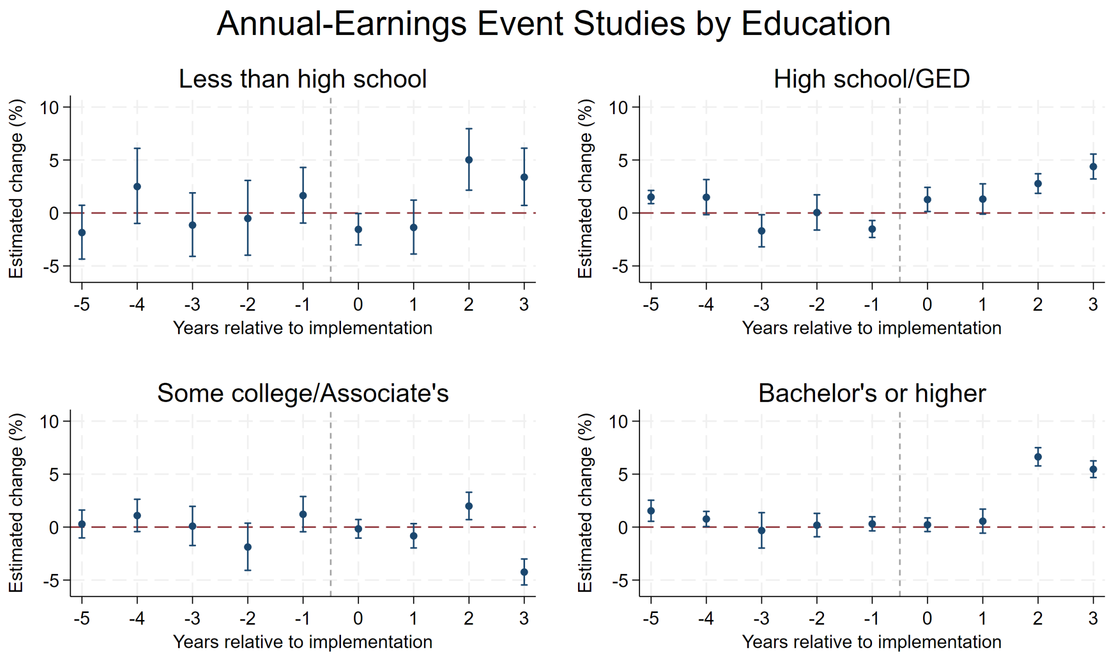

Notes: Panels report event-study estimates from the main estimator by educational attainment. Markers show estimated percentage changes and vertical bars show 95% confidence intervals. Event time is measured relative to policy implementation; event year 1 is the omitted reference period.

4.4 What does the evidence show overall?

The baseline estimate indicates an increase of about 0.99% in real annual earnings following salary-range disclosure. The demographic results are more mixed. The residualized men–women differential increases significantly, while the raw demographic gaps and the residualized racial and ethnic estimates are not significant at the 5% level. At the group level, the estimated increases are significant for men and Hispanic workers, and the education results show a significant increase for high-school/GED workers. These separate significance results should not be interpreted as direct evidence that the groups respond differently from one another.

The combination of higher estimated earnings and limited evidence of demographic earnings convergence is broadly consistent with Arnold, Quach and Taska (2025), who report higher wages without detectable changes in pay dispersion. Their dispersion measures differ from the demographic earnings gaps examined here, so the comparison concerns the broad pattern rather than identical distributional outcomes.

The small number of treated states constrains inference. Restricted wild-cluster bootstrap tests of four conventional pre-treatment leads do not reject their joint null, but the ordinary Wald test, the omnibus pre-treatment test for the Callaway and Sant'Anna estimator, and the six-year-early placebo are less favorable. These diagnostics test different restrictions and do not provide uniformly strong support for the identifying assumptions. The results should therefore be interpreted as earnings patterns associated with implementation, with caution about causal attribution. Appendix A reports the identification and robustness evidence.

5. Summary, conclusion and policy implications

5.1 Summary and conclusion

This study asks whether state salary-range disclosure laws raise real annual earnings and whether any gains are broadly shared across workers. Using the 2010-2024 American Community Survey and staggered policy adoption across states, the analysis applies the Callaway and Sant'Anna staggered difference-in-differences design and follows the average earnings response with distributional and education-based heterogeneity analyses.

The main result is a modest positive association between salary-range disclosure and real annual earnings. The baseline estimate is approximately 0.99 percent. One possible explanation is that salary information available during recruitment gives workers a clearer picture of the pay attached to different job opportunities.

The distributional evidence does not show broad convergence. Raw demographic differentials do not change significantly at the 5% level, while residualization substantially reduces several observed gaps. The residualized men–women estimate is positive and significant, but the residualized White–Black, White–Hispanic, and White–Asian overall effects remain statistically insignificant. Group-specific estimates show significant earnings increases for men and Hispanic workers, but the differences across groups are generally not precise enough to establish systematic narrowing or widening of demographic earnings inequality.

The education results also show no simple pattern across qualification levels. The only statistically significant estimate is a 1.45% increase for workers with a high-school diploma or GED. The estimate for workers with a bachelor's degree or higher is positive but not statistically significant. The estimated response therefore does not increase steadily with educational attainment. In the sample as a whole, disclosure is associated with higher average earnings, but there is no clear evidence that these gains reduce existing earnings inequality.

The identification checks are not uniformly reassuring. In the 3 to +3 event study from the main estimator, the displayed pre-treatment estimates are individually insignificant. In the conventional specification, the restricted wild-cluster bootstrap also fails to reject the joint null for four pre-treatment leads at the 5% level. However, the ordinary Wald test, the Callaway and Sant'Anna omnibus pre-treatment test, and the six-year-early placebo give less favorable results. For this reason, the estimates are interpreted cautiously as evidence on earnings patterns around the introduction of salary-range disclosure laws rather than as unqualified causal effects.

5.2 Policy implications

The results suggest that salary-range disclosure may make the recruitment process more informative and is associated with modestly higher annual earnings. When compensation is shown before negotiation, applicants have more information about the pay attached to a vacancy and can compare job opportunities more directly.

However, the absence of broad changes in gender and racial or ethnic earnings differentials indicates that transparency alone is unlikely to eliminate persistent earnings inequality. Salary information addresses information asymmetry but does not directly address discrimination, occupational segregation, and differences in mobility, or unequal access to high-paying employment. The effectiveness of disclosure laws also depends on compliance and on whether posted salary ranges provide credible information about the compensation available for a position.

5.3 Limitations and areas for further research

Several limitations should be kept in mind. Only a small number of states are treated within the sample period, and the post-treatment window is short, especially at the longer event-study horizons. The ACS is a repeated cross-section, so the same workers and employers cannot be followed over time. It also contains no direct information on job postings, compliance with salary-range requirements, wage negotiations, or outside offers. Finally, the pre-treatment and placebo checks give mixed results, which limits the strength of the causal interpretation.

Further research using longer post-treatment periods, additional adopting states, longitudinal worker records, and employer or job-posting data would provide stronger evidence on the durability and mechanisms of the estimated effects. Such data would also help distinguish changes in wage rates from changes in hours or employment duration and clarify whether the effects operate through job mobility, bargaining, employer competition, or compliance with disclosure requirements.

References

Arnold, D., Quach, S., & Taska, B. (2025). The impact of pay transparency in job postings on the labor market. NBER Working Paper No. 34480. National Bureau of Economic Research.
urlhttps://doi.org/10.3386/w34480

Baker, M., Halberstam, Y., Kroft, K., Mas, A., & Messacar, D. (2023). Pay transparency and the gender gap. American Economic Journal: Applied Economics, 15(2), 157–183.
urlhttps://doi.org/10.1257/app.20210141

Batra, H., Michaud, A., & Mongey, S. (2023). Online job posts contain very little wage information. NBER Working Paper No. 31984. National Bureau of Economic Research.
urlhttps://doi.org/10.3386/w31984

Bennedsen, M., Simintzi, E., Tsoutsoura, M., & Wolfenzon, D. (2022). Do firms respond to gender pay gap transparency? The Journal of Finance, 77(4), 2051–2091.
urlhttps://doi.org/10.1111/jofi.13136

Biasi, B., & Sarsons, H. (2022). Flexible wages, bargaining, and the gender gap. The Quarterly Journal of Economics, 137(1), 215–266.
urlhttps://doi.org/10.1093/qje/qjab026

California Department of Industrial Relations. (2023). California Equal Pay Act: Pay scale requirements in job postings.
urlhttps://www.dir.ca.gov/dlse/california_equal_pay_act.htm

Callaway, B., & Sant'Anna, P. H. C. (2021). Difference-in-differences with multiple time periods. Journal of Econometrics, 225(2), 200–230.
urlhttps://doi.org/10.1016/j.jeconom.2020.12.001

Cameron, A. C., Gelbach, J. B., & Miller, D. L. (2008). Bootstrap-based improvements for inference with clustered errors. The Review of Economics and Statistics, 90(3), 414–427.

urlhttps://doi.org/10.1162/rest.90.3.414

Colorado Department of Labor and Employment. (2024). Equal Pay for Equal Work Act.
urlhttps://cdle.colorado.gov/dlss-home-page/equal-pay-for-equal-work-act

Cullen, Z. B. (2024). Is pay transparency good? Journal of Economic Perspectives, 38(1), 153–180.
urlhttps://doi.org/10.1257/jep.38.1.153

Cullen, Z. B., & Pakzad-Hurson, B. (2023). Equilibrium effects of pay transparency. Econometrica, 91(3), 765–802.
urlhttps://doi.org/10.3982/ECTA19788

Cullen, Z. B., & Perez-Truglia, R. (2023). The salary taboo: Privacy norms and the diffusion of information. Journal of Public Economics, 222, 104890.
urlhttps://doi.org/10.1016/j.jpubeco.2023.104890

District of Columbia Office of the Attorney General. (2024). Wage transparency law guidance.
urlhttps://oag.dc.gov/worker-rights/wage-transparency

European Parliament and Council of the European Union. (2023). Directive (EU) 2023/970 of the European Parliament and of the Council of 10 May 2023 to strengthen the application of the principle of equal pay for equal work or work of equal value between men and women through pay transparency and enforcement mechanisms. Official Journal of the European Union, L 132, 21–44.
urlhttps://eur-lex.europa.eu/eli/dir/2023/970/oj/eng

Feng, K. (2024). Pay transparency and gender differences in job search [Working paper]. SSRN.
urlhttps://doi.org/10.2139/ssrn.4984519

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. Journal of Econometrics, 225(2), 254–277.
urlhttps://doi.org/10.1016/j.jeconom.2021.03.014

Hawai‘i Civil Rights Commission. (2024). Act 203 pay transparency frequently asked questions.
urlhttps://labor.hawaii.gov/hcrc/

IPUMS USA. (2026). IPUMS USA: U.S. Census data for social, economic, and health research. University of Minnesota.
urlhttps://www.ipums.org/projects/ipums-usa

IPUMS USA. (2026). Monetary adjustment feature. University of Minnesota.
urlhttps://usa.ipums.org/usa/adjusted_monetary_values.shtml

Jäger, S., Roth, C., Roussille, N., & Schoefer, B. (2024). Worker beliefs about outside options. The Quarterly Journal of Economics, 139(3), 1505–1556.
urlhttps://doi.org/10.1093/qje/qjae001

International Labour Organization. (2022). Pay transparency legislation: Implications for employers' and workers' organizations. International Labour Organization.
urlhttps://www.ilo.org/publications/pay-transparency-legislation-implications-employers-and-workers

Maryland Department of Labor. (2024). Equal Work for Equal Pay—Wage Range Transparency frequently asked questions.
urlhttps://www.labor.maryland.gov/labor/wages/esswagerangefaq.shtml

New York State Department of Labor. (2023). Pay transparency law.
urlhttps://dol.ny.gov/pay-transparency

Pissarides, C. A. (2000). Equilibrium unemployment theory (2nd ed.). MIT Press.

Roodman, D., Nielsen, M. Ø., MacKinnon, J. G., & Webb, M. D. (2019). Fast and wild: Bootstrap inference in Stata using boottest. The Stata Journal, 19(1), 4–60.
urlhttps://doi.org/10.1177/1536867X19830877

Roth, J., Sant'Anna, P. H. C., Bilinski, A., & Poe, J. (2023). What's trending in difference-in-differences? A synthesis of the recent econometrics literature. Journal of Econometrics, 235(2), 2218–2244.
urlhttps://doi.org/10.1016/j.jeconom.2023.03.008

Sun, L., & Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. Journal of Econometrics, 225(2), 175–199.
urlhttps://doi.org/10.1016/j.jeconom.2020.09.006

U.S. Department of Labor. (2015). Pay transparency nondiscrimination provision. Office of Federal Contract Compliance Programs.
urlhttps://www.dol.gov/agencies/ofccp/faqs/pay-transparency

U.S. Equal Employment Opportunity Commission. (n.d.). The Equal Pay Act of 1963.
urlhttps://www.eeoc.gov/statutes/equal-pay-act-1963

Washington State Department of Labor & Industries. (2023). Equal Pay and Opportunities Act: Pay transparency requirements.
urlhttps://www.lni.wa.gov/workers-rights/wages/equal-pay-opportunities-act

Appendix A. Identification, robustness, and supporting results

This appendix reports identification tests, conventional benchmark estimates, placebo results, and policy-coding checks that support the empirical analysis but are secondary to the central earnings and distributional results in Chapter Four.

Table A1. Conventional two-way fixed-effects benchmarks

Outcome	**Coefficient**	**SE**	**Approx. effect (%)**
Adopter × post-2020 mean earnings	0.0122	0.0131	1.23
Staggered treated mean earnings	0.0274**	0.0130	2.78
Gender earnings differential	0.0028	0.0056	0.28
White-Black differential	0.0255	0.0153	2.58
White-Hispanic differential	0.0001	0.0075	0.01
White-Asian differential	0.0186	0.0205	1.84

*Notes: These specifications are conventional benchmarks rather than the Callaway and Sant'Anna estimator. The pooled model uses eventual-adopter × post-2020 exposure; the staggered TWFE models use state-specific treatment timing. Standard errors are clustered by state. * p < 0.10, ** p < 0.05, *** p < 0.01. For this benchmark table, stars are based on two-sided tests computed from the displayed state-clustered standard errors using 50 residual degrees of freedom.*

The conventional models provide a reference point for the main estimates. Their interpretation is limited under staggered adoption because a common treatment coefficient can combine comparisons across treatment cohorts and post-treatment periods.

Table A2. Extended-window event-study coefficients from the main estimator

Event time	**Coefficient**	**Effect (%)**	**SE**
5	0.01188***	1.20	0.00282
4	0.01346***	1.35	0.00304
3	0.00354	0.35	0.00362
2	0.00536	0.53	0.00517
1	0.00229	0.23	0.00402
0	0.00503**	0.50	0.00200
+1	0.00851***	0.85	0.00213
+2	0.04966***	5.09	0.00449
+3	0.04863***	4.98	0.00428
Average pre-treatment	0.00283	0.28	0.00203
Average post-treatment	0.02796***	2.83	0.00267

*Notes: This table retains the extended 5 to +3 event window using the main estimator. Unlike a conventional event study normalized to 1, the CSDID routine reports pre-treatment pseudo-effects, including an estimate at 1. The +2 and +3 estimates are supported primarily by the earliest cohort. An average pre-treatment estimate does not replace a joint test of all leads. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.*

The main 3 to +3 event study shows positive post-treatment coefficients, particularly at +2 and +3. The displayed 3 and 2 coefficients are individually insignificant and do not indicate a directional pre-trend. The appendix retains the wider-window estimates, whose earlier coefficients provide weaker support for parallel trends. Later post-treatment estimates represent fewer cohorts and therefore identify effects for a narrower composition of adopting states.

Table A3. Parallel-trends and wild-cluster bootstrap tests

Test	**Statistic**	**df**	**Conclusion**
CSDID test of all pre-treatment periods	$^2 = 2.080e+10$***	35	Reject
Displayed-lead conventional Wald	F = 3.314**	4, 50	Reject
Wild-cluster bootstrap, Webb	F = 3.314	4, 50	Do not reject
Wild-cluster bootstrap, Rademacher	F = 3.314	4, 50	Do not reject

*Notes: The tests do not impose the same null hypothesis. The bootstrap tests address the four conventional pre-treatment leads and are useful when the number of treated clusters is small. The wider pre-treatment test for the Callaway and Sant'Anna estimator uses all estimable pre-treatment periods. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.*

The diagnostics differ in both estimand and null hypothesis. The omnibus CSDID test rejects equality of all available pre-treatment pseudo-effects to zero. A separate conventional event-study regression produces an ordinary Wald rejection for four pre-treatment leads, but restricted wild-cluster bootstrap inference does not reject the same four coefficients under Webb or Rademacher weights. The coefficients from the main estimator shown in Figure 4.1 are individually insignificant but are not the coefficients used in the bootstrap exercise. The omnibus rejection and significant early placebo therefore warrant caution in interpreting the estimates, although the results remain informative about earnings patterns around policy adoption.

Table A4. Placebo and additional sensitivity checks

Check	**Coefficient**	**Effect (%)**	**SE**
Six-year-early pseudo-adoption	0.01085**	1.09	0.00429
Composition-adjusted TWFE	0.02595**	2.63	0.01251

Notes: Table 4.1 reports the principal timing and weighting sensitivities. This table retains the early-adoption placebo and a conventional composition-adjusted TWFE benchmark. The placebo uses only eventual adopters before their actual implementation dates and therefore has 11 state clusters. The 11-cluster count belongs to the placebo construction, which uses a broader eventual-adopter sample; it is not the number of jurisdictions treated in the 2010–2024 baseline estimation. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The six-year-early placebo estimate is positive and statistically significant. This result indicates that part of the estimated relationship may reflect earlier differential changes, anticipation, other policy changes, or limitations of the placebo comparison. The composition-adjusted TWFE estimate is also positive, but it is not the main estimand under staggered adoption.

Table A5. Policy-coding audit

Jurisdiction	**FIPS**	**Exact effective date**	**Baseline coding**	**Full-exposure status through 2024**
Colorado	08	January 1, 2021	2021	Treated from 2021
California	06	January 1, 2023	2023	Treated from 2023
Washington	53	January 1, 2023	2023	Treated from 2023
New York	36	September 17, 2023	2023	Treated from 2024
Hawaii	15	January 1, 2024	2024	Treated from 2024
District of Columbia	11	June 30, 2024	Omitted	No fully exposed annual observation
Maryland	24	October 1, 2024	2024	No fully exposed annual observation
Illinois	17	January 1, 2025	Not treated	Outside observed post-period
Minnesota	27	January 1, 2025	Not treated	Outside observed post-period
New Jersey	34	June 1, 2025	Not treated	Outside observed post-period
Vermont	50	July 1, 2025	Not treated	Outside observed post-period
Massachusetts	25	October 29, 2025	Not treated	Outside observed post-period

Notes: *The audit distinguishes the coding used for the reported baseline estimates from a full-exposure convention. The District of Columbia was omitted from the baseline treatment file.*

The ACS annual earnings measure covers the preceding 12 months. Under a full-exposure convention, Colorado, California, Washington, New York, and Hawaii provide the strongest treated observations within the 2010–2024 period. Maryland and the District of Columbia provide no fully exposed post-treatment annual-earnings observation by 2024. This reduces the effective number of treated policy clusters, lowers precision, and limits the generalizability of the estimated effects.

Appendix B. Alternative aggregation and weighting

Alternative aggregation and weighting schemes produce different estimates because they place different weights on treatment cohorts and post-treatment horizons. Table B1 reports these specifications as sensitivity exercises. They answer related but distinct treatment-effect questions and are not treated as numerically identical to the main overall ATT reported in Chapter Four.

Table B1. Sensitivity to alternative aggregation and weighting

Specification	**Estimate**	**Standard error**
Equal event-year average, unadjusted	0.0274***	0.0078
Equal event-year average, adjusted	0.0216**	0.0082
Cohort-size-weighted ATT, unadjusted	0.0097***	0.0016
Cohort-size-weighted ATT, adjusted	0.0099***	0.0026
Equally weighted state-year cells	0.0264**	0.0099
Naive adopter × post-2020 model	0.0122	0.0131

Notes: Alternative weighting changes the estimand. The cohort-size-weighted ATT is not directly equivalent to an equal average of event years 0 through +3. The estimates and standard errors are reproduced from the supplied robustness results. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The sign of the mean-earnings estimate remains positive across the alternative specifications, but its magnitude and precision depend on the aggregation and weighting scheme. Equal event-year averages place equal weight on the selected post-treatment horizons, whereas cohort-size weighting changes the contribution of treatment cohorts. The naive adopter × post-2020 model is not used as the main specification because it assigns a common post period to states that adopted in different years.

The cohort-size-weighted estimates are close in magnitude to the main Chapter Four estimate, while the equal-event-year summaries are larger because later post-treatment horizons receive equal weight. These alternative aggregations identify different weighted averages of cohort- and event-time-specific effects.

Appendix C. Demographic earnings-differential sensitivity

This appendix reports cohort-size-weighted sensitivity estimates for the demographic earnings differentials. The estimates assess whether the main distributional conclusions are sensitive to alternative weighting.

Table C1. Cohort-weighted sensitivity of demographic earnings-differential estimates

Earnings-differential outcome	**Cohort-weighted estimate**	**SE**
Men-women	0.0061	0.0052
White-Black	0.0033	0.0135
White-Hispanic	0.0127*	0.0069

Earnings-differential outcome	Cohort-weighted estimate	SE
White-Asian	0.0246*	0.0133

*Notes: Estimates are differences in mean log annual earnings with state-clustered standard errors. Failure to reject the null does not establish that an earnings differential is exactly zero. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.*

The alternative weighting does not change the principal distributional conclusion. The men-women, White-Black, and White-Hispanic estimates are not statistically significant at the 5 percent level. The White-Asian estimate remains positive but is significant only at the 10 percent level. The available ACS evidence therefore does not show a broad statistically significant narrowing or widening of the main demographic annual-earnings differentials.

The appendix results are subject to the same inferential limitations as the main analysis. Only a small number of states provide post-treatment observations by 2024, and an even smaller subset provides fully exposed annual-income observations. The large ACS sample does not increase the number of independent policy clusters.

Appendix D. Supporting event-study results and descriptive evidence

This appendix reports supporting event-study figures from the main estimator for the raw and residualized demographic annual-earnings differentials, descriptive trends in those differentials, and pre-policy earnings dispersion by education. The main chapter reports the corresponding average effects, while Figure 4.2 presents the education-specific event studies.

Figure D1. Event studies for raw demographic annual-earnings differentials

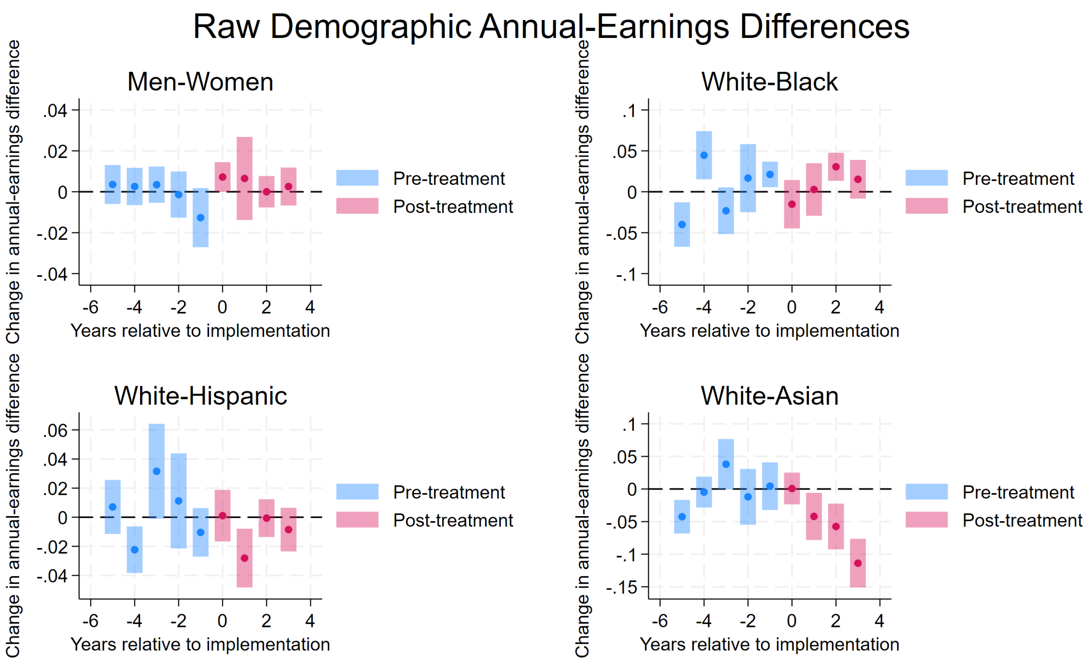

Notes: Panels show estimates from the Callaway and Sant'Anna method (CSDID) for the men-women, White-Black, White-Hispanic, and White-Asian raw log annual-earnings differentials. Bars show 95% confidence intervals. Event time is measured relative to policy implementation.

Figure D2. Event studies for residualized demographic annual-earnings differentials

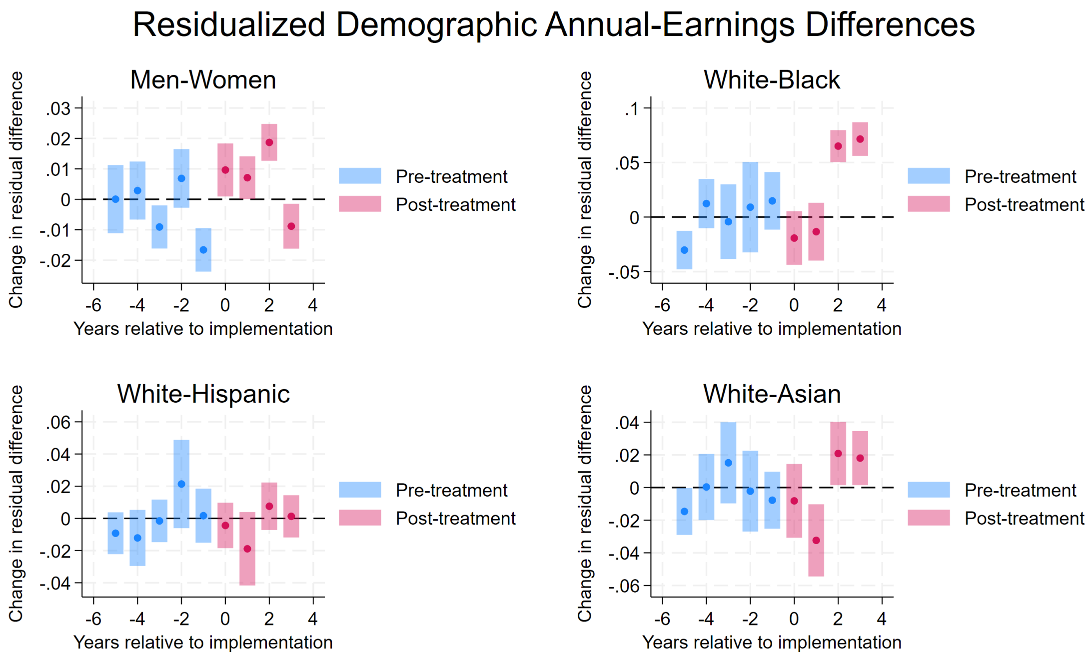

Notes: Panels show estimates from the Callaway and Sant'Anna method (CSDID) for residualized demographic differentials constructed using Equations (7) and (8). Bars show 95% confidence intervals. The joint pre-treatment tests reject for these specifications, so the dynamic estimates should be interpreted with caution because the evidence supporting the underlying identifying assumptions is not strong for these outcomes.

Figure D3. Raw and residualized gender and racial/ethnic annual-earnings differentials, 2010–2024

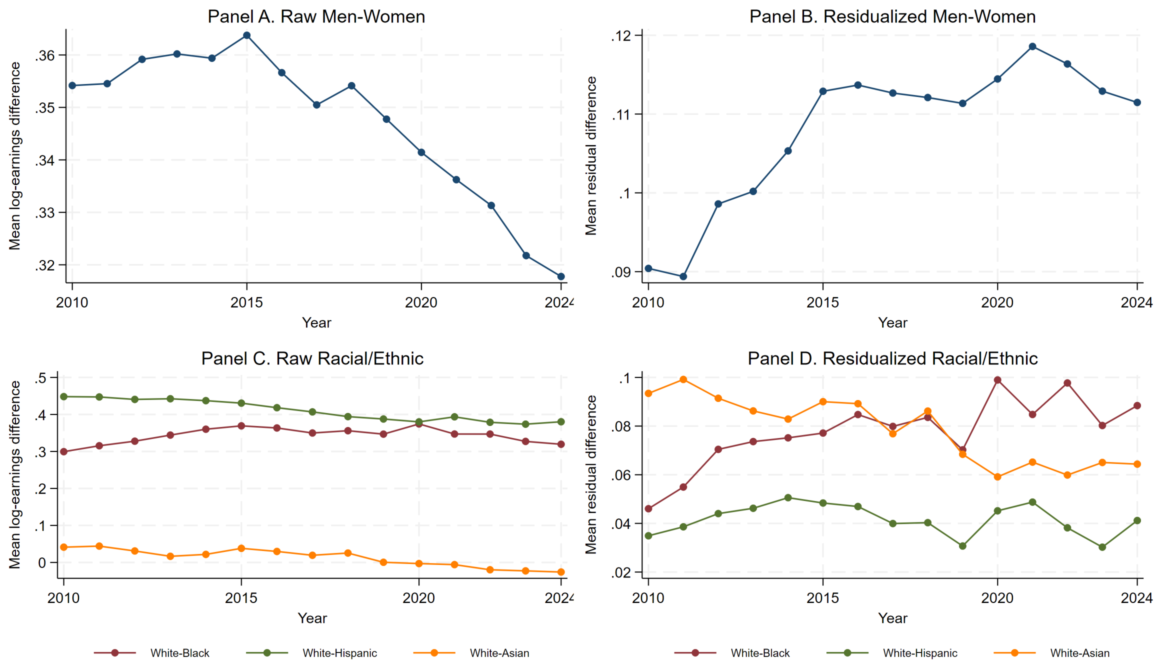

Notes: The upper panels report raw annual-earnings differentials. The lower panels report residualized differentials after controlling for age, age squared, education, occupation, usual weekly hours, and year effects. Positive gender values indicate higher earnings for men; positive racial or ethnic values indicate higher earnings for non-Hispanic White workers.

Residualization materially changes the scale of the demographic gaps. The raw gender differential declines over the sample, while the residualized gender differential is much smaller and does not follow the same downward path. Raw White-Black and White-Hispanic differentials are also substantially larger than their residualized counterparts. The raw White-Asian comparison shows higher Asian earnings, whereas the residualized differential is close to zero for much of the period. These patterns show that observable characteristics explain an important share of the raw gaps, but they do not eliminate all differences. The policy event-study evidence for these outcomes is reported in Figures D1 and D2.

Table D1. Raw and residualized demographic CSDID summary

Outcome	**Raw ATT (SE)**	**Raw post avg. (SE)**	**Residualized ATT (SE)**	**Residualized post avg. (SE)**
Men–women	0.0065 (0.0051)	0.0041 (0.0038)	0.0082*** (0.0022)	0.0066*** (0.0021)
White–Black	0.0045 (0.0137)	0.0084 (0.0089)	0.0102 (0.0158)	0.0260*** (0.0074)
White–Hispanic	0.0121 (0.0084)	0.0090 (0.0063)	0.0101 (0.0083)	0.0036 (0.0060)
White–Asian	0.0244* (0.0145)	0.0531*** (0.0136)	0.0167 (0.0114)	0.0004 (0.0074)

*Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Post avg. denotes the average of the displayed post-treatment event-time coefficients. All demographic joint pre-treatment tests reject the null of no pre-treatment effects; these estimates should therefore be interpreted with caution, although they remain informative about the dynamic patterns around policy implementation.*

Table D2. Pre-policy annual-earnings dispersion by education

Education group	Mean log earnings	P10	P90	P90-P10
Less than high school	10.059	8.815	11.069	2.254
High school/GED	10.338	9.091	11.364	2.272
Some college/associate degree	10.451	9.081	11.523	2.442
Bachelor's degree or higher	11.134	9.963	12.141	2.178

Notes: Statistics use 2010-2020 ACS person weights. P90P10 is the difference between the 90th and 10th percentiles of log real annual earnings.

Appendix E. Policy status by jurisdiction, 2019–2025

Table E1. Policy status by jurisdiction, 2019–2025

JURISDICTION	2019	2020	2021	2022	2023	2024	2025
ALABAMA	N	N	N	N	N	N	N
ALASKA	N	N	N	N	N	N	N
ARIZONA	N	N	N	N	N	N	N
ARKANSAS	N	N	N	N	N	N	N
CALIFORNIA	NYT	NYT	NYT	NYT	T	T	T
COLORADO	NYT	NYT	T	T	T	T	T
CONNECTICUT	N	N	N	N	N	N	N
DELAWARE	N	N	N	N	N	N	N
DISTRICT OF COLUMBIA	NYT	NYT	NYT	NYT	NYT	T	T
FLORIDA	N	N	N	N	N	N	N
GEORGIA	N	N	N	N	N	N	N
HAWAII	NYT	NYT	NYT	NYT	NYT	T	T
IDAHO	N	N	N	N	N	N	N
ILLINOIS	NYT	NYT	NYT	NYT	NYT	NYT	T
INDIANA	N	N	N	N	N	N	N
IOWA	N	N	N	N	N	N	N
KANSAS	N	N	N	N	N	N	N
KENTUCKY	N	N	N	N	N	N	N
LOUISIANA	N	N	N	N	N	N	N
MAINE	N	N	N	N	N	N	N
MARYLAND	NYT	NYT	NYT	NYT	NYT	T	T
MASSACHUSETTS	NYT	NYT	NYT	NYT	NYT	NYT	T
MICHIGAN	N	N	N	N	N	N	N
MINNESOTA	NYT	NYT	NYT	NYT	NYT	NYT	T
MISSISSIPPI	N	N	N	N	N	N	N
MISSOURI	N	N	N	N	N	N	N
MONTANA	N	N	N	N	N	N	N
NEBRASKA	N	N	N	N	N	N	N
NEVADA	N	N	N	N	N	N	N

JURISDICTION	**2019**	**2020**	**2021**	**2022**	**2023**	**2024**	**2025**
NEW HAMPSHIRE	N	N	N	N	N	N	N
NEW JERSEY	NYT	NYT	NYT	NYT	NYT	NYT	T
NEW MEXICO	N	N	N	N	N	N	N
NEW YORK	NYT	NYT	NYT	NYT	T	T	T
NORTH CAROLINA	N	N	N	N	N	N	N
NORTH DAKOTA	N	N	N	N	N	N	N
OHIO	N	N	N	N	N	N	N
OKLAHOMA	N	N	N	N	N	N	N
OREGON	N	N	N	N	N	N	N
PENNSYLVANIA	N	N	N	N	N	N	N
RHODE ISLAND	N	N	N	N	N	N	N
SOUTH CAROLINA	N	N	N	N	N	N	N
SOUTH DAKOTA	N	N	N	N	N	N	N
TENNESSEE	N	N	N	N	N	N	N
TEXAS	N	N	N	N	N	N	N
UTAH	N	N	N	N	N	N	N
VERMONT	NYT	NYT	NYT	NYT	NYT	NYT	T
VIRGINIA	N	N	N	N	N	N	N
WASHINGTON	NYT	NYT	NYT	NYT	T	T	T
WEST VIRGINIA	N	N	N	N	N	N	N
WISCONSIN	N	N	N	N	N	N	N
WYOMING	N	N	N	N	N	N	N

Notes: T = treated; NYT = not yet treated but adopts by 2025; N = no qualifying statewide mandate effective through 2025. DC is classified by its actual June 2024 effective date here but remains excluded from the baseline estimation coding, as documented in Table 3.2 and Appendix A.