Combating Inequality with Transparency? Evidence from Colorado [*]

Sebastian Brown[†], Thomas Fullagar [‡]

May 16, 2025

Abstract

Does increasing wage transparency effectively reduce pay disparities? We address this question by examining the impact of Colorado's Equal Pay for Equal Work Act, a pioneering statewide policy mandating that employers disclose salary ranges in all job postings. Employing a synthetic control method to compare Colorado's experience with a carefully constructed counterfactual, we find no evidence that the law narrowed the gender earnings gap among newly hired workers. In fact, our estimates indicate a widening of approximately 15 percent, a statistically significant increase relative to the synthetic control. We perform additional correlational industry-level analysis, finding suggestive patterns consistent with gender differences in search or bargaining behavior. Our results highlight critical challenges in designing transparency policies and emphasize the necessity of ensuring that information interventions align closely with the behavioral responses of the intended beneficiaries.

JEL Codes: D83, J16, J31, J38, J63, K31

[*]This paper benefited from feedback from Peter Kuhn and Mitchell Hoffman. We would also like to thank Ivan Strahof at IPUMS User Support for helpful information about the CPS.

[†]Department of Economics, UC Santa Barbara. sebastian_brown@ucsb.edu.

[‡]Department of Economics, UC Santa Barbara. fullagar@ucsb.edu.

1 Introduction

Wage transparency has increasingly become a central policy tool for addressing persistent gender pay disparities in labor markets. Advocates argue that greater transparency in salary information empowers workers—particularly those historically disadvantaged—to make informed employment decisions, negotiate better compensation, and challenge wage discrimination. Despite the intuitive appeal of these policies, their effectiveness remains an open empirical question. In this paper, we examine Colorado's *Equal Pay for Equal Work Act* (hereafter, the "Colorado Law"), the first comprehensive statewide mandate requiring employers to include salary ranges in all job postings, regardless of job type or employer size (Colorado Department of Labor & Employment, 2024a). Leveraging this unique policy setting and rich administrative data, we provide new causal evidence on the effectiveness of wage transparency laws in achieving their intended goal: reducing gender-based earnings gaps among newly hired workers.

In recent years, wage transparency laws similar to Colorado's have gained momentum across the United States, driven by ongoing concerns about pay equity. Following Colorado's law—which took effect on January 1, 2021—several other states have implemented comparable measures, including California and Washington (January 1, 2023), New York (September 17, 2023), Hawaii (January 1, 2024), and Illinois (January 1, 2025). Additional states, such as Maryland, Connecticut, Nevada, and Rhode Island, require pay transparency upon request or during the hiring process. Local transparency initiatives have also emerged in Washington D.C. (June 30, 2024) and in jurisdictions within New Jersey, New York, and Ohio, and at the federal level, national transparency legislation was introduced in March 2023 (Arnold et al., 2022; Marfice, 2024; PayAnalytics, 2025). These expanding legislative efforts underscore the urgency of rigorously evaluating the efficacy of transparency policies in achieving meaningful reductions in gender wage disparities.

The Colorado law mandates that employers include salary information in nearly all job postings, representing a substantial departure from previous norms in job advertising.[1] Prior to this legislation, wage disclosure in job advertisements was often rare. For instance, Marinescu and Wolthoff (2020) report that in the first quarter of 2011, only 20% of job postings in Chicago and Washington, D.C., on CareerBuilder.com included salary information. Similarly, outside the United States, Banfi and Villena-Roldán (2019) find that just 13.3% of listings on a major Chilean job board (www.trabajando.com) disclosed pay information.

Given the recent nature of these legislative efforts, there is limited empirical evidence on their effects. A notable exception is the study by Arnold, Quach and Taska (2022), which examines the impact of Colorado's law on job postings. The authors find that the legislation led to a 3.6 percent increase in posted salaries and a 30 percentage point increase in the share of postings that included salary information. However, they acknowledge a key limitation of their analysis: because their data source (Burning Glass Technologies) captures only job advertisements, it does not allow for assessment of downstream outcomes such as realized wages or employment composition.

This paper contributes to the emerging literature on wage transparency by addressing two central questions:

1. Have recent pay transparency laws been effective in achieving their intended goal of reducing gender disparities in worker compensation?

2. If these laws have affected pay gaps, what mechanisms underlie those

[1] According to the Colorado Department of Labor, there are only four exceptions to this requirement: (1) the posting is for a non-competitive promotion, (2) it concerns acting, interim, or temporary roles, (3) it involves a confidential replacement of a current employee, or (4) the position is fully remote, located out of state, and offered by a Colorado-based employer with no physical worksite and fewer than fifteen employees (Colorado Department of Labor & Employment, 2024b).

effects?

To address the first question, we examine how Colorado's pay transparency law affected the earnings gap between newly hired men and women (the "earnings gap"). We focus on Colorado because it offers the earliest and most comprehensive data, which also facilitates the construction of a credible control group for causal inference. In addition to this primary case study, we present an alternative analysis in Appendix C that estimates the average effect of pay transparency laws across all treated states.

Employing the synthetic-control framework of Abadie et al. (2010) (implemented with the `Synth` routine of Abadie et al. (2011)) we find that the Colorado law *widened* rather than narrowed the gender earnings gap among newly hired workers. In the two years after enactment, Colorado's gap rose by roughly $155 relative to its synthetic counterpart, a 15 percent increase over the pre-period average of $1,062. Placebo re-assignments confirm that this divergence is statistically significant at conventional levels.

To understand why increased wage transparency did not narrow—and may have even widened—the earnings gap, we investigate three potential mechanisms. First, we examine whether greater transparency shifted the gender composition of new hires across industries. If women used newly available wage information to pursue higher-paying opportunities, we would expect to see an increase in their share of new hires in industries with higher average wages. Instead, we find no such pattern: changes in gender composition were not positively correlated with industry wage levels, suggesting that women did not systematically sort into higher-paying sectors in response to the policy.[2] Second, within industries, we observe that increased transparency is correlated with a widening of the earnings gap: male earnings grew more substantially in in-

[2]Industries are classified using two-digit North American Industry Classification System (NAICS) codes.

dustries that experienced greater transparency, while female earnings remained largely stagnant. While consistent with differential directed search responses, this pattern may also reflect differences in bargaining behavior. For example, men might be more likely to negotiate salaries near the top of newly disclosed salary ranges, while women might settle nearer the lower end of those same ranges. Third, we find that the gap between advertised and realized wages remained virtually unchanged following the law's implementation, indicating that additional disclosure did not significantly enhance the informativeness or accuracy of salary signals.

Although these analyses are descriptive rather than causal, together they offer suggestive evidence for why the policy did not achieve its intended goal of reducing the gender earnings gap. These findings underscore the importance of aligning informational policies with behavioral responses to effectively achieve distributional and equity objectives.

The rest of this paper is organized as follows. Section 2 gives an overview of the literature. Section 3 discusses our data. Section 4 discusses our empirical method and evaluates the plausibility of its main assumptions. Section 5 gives our main results. Section 6 discusses mechanisms through which the law could have affected our results and gives suggestive evidence related to these mechanisms. Section 7 concludes.

2 Literature Review

This paper contributes to a growing literature examining the effects of wage transparency policies, a topic of increasing interest among economists and policymakers. A central feature of this literature is its focus on distinct forms of transparency—ranging from internal reporting mandates to public disclosure and protections for wage discussion—each targeting different segments of the labor market and operating through different channels. While these studies have offered valuable insights, few have directly examined the causal effect of wage transparency in job postings, especially in the context of U.S. state-level reforms aimed explicitly at reducing gender pay gaps.

Several recent studies focus on internal transparency policies. For example, Böheim and Gust (2021) and Gulyas, Seitz and Sinha (2023) evaluate Austria's 2011 law requiring firms to provide employees with anonymized reports on average wages by gender and occupation group. Despite using different empirical approaches, both papers find that the policy had no significant effect on worker wages or the gender pay gap—perhaps in part because the information was not made public (Gulyas et al., 2023).

Other research examines transparency policies centered on wage-sharing protections and public disclosure. Cullen and Pakzad-Hurson (2023) develop a model of wage bargaining and use U.S. state-level variation to show that laws protecting employees' right to discuss wages lead to a modest decline in overall pay. In contrast, Mas (2017) finds that California's 2010 policy requiring public disclosure of municipal employee salaries reduced total compensation by 7% and substantially increased quit rates among top management. Similarly, Perez-Truglia (2020) documents that the expansion of online access to income tax records in Norway in 2001—effectively revealing individuals' incomes to the broader public—reduced reported life satisfaction, particularly among lower-income individuals, likely due to heightened social comparisons.

One of the few settings in which wage transparency appears to reduce pay disparities is in the public sector. Baker, Halberstam, Kroft, Mas and Messacar (2023) study Canadian laws mandating salary disclosure for public university faculty and find that they reduced the gender wage gap by 20 to 40 percentage points, underscoring the potential for transparency to improve equity in certain institutional contexts.

Our paper advances this literature in two key ways. First, we study a policy that is uniquely broad in scope and directly targets the job search process by requiring salary information to be posted in job advertisements—a feature that distinguishes it from most prior transparency laws. Second, while existing studies often focus on wage levels or disclosure behavior, we directly test the central claim behind many recent reforms: that pay transparency laws can reduce gender pay disparities among newly hired workers. Using rich individual-level data and a transparent identification strategy, we provide one of the first rigorous estimates of the causal impact of public wage posting requirements on the gender pay gap in the U.S. labor market.

3 Data Sources and Summary Statistics

Our primary data source is the Longitudinal Employer-Household Dynamics (LEHD) program, administered by the U.S. Census Bureau (U.S. Census Bureau, 2024). Specifically, we rely on the Quarterly Workforce Indicators (QWI) to measure earnings, employment, and hiring activity, disaggregated by state, industry (classified by two-digit NAICS codes), and worker sex. The QWI are constructed from a combination of administrative datasets and cover more than 95% of private-sector employment in the United States (US Census Bureau, 2022). Earnings information is derived from unemployment insurance (UI) wage records, while establishment-level industry and location data are obtained from the Quarterly Census of Employment and Wages (QCEW). Additional worker characteristics are drawn from other administrative sources. A key advantage of the QWI is that it links workers to specific employers at the job level, enabling detailed analysis of labor market outcomes by demographic group—including age, sex, education, and race/ethnicity.

Our panel spans from the second quarter of 2011 to the fourth quarter of 2023, covering a period of roughly ten years before and three years after the implementation of Colorado's Equal Pay for Equal Work Act. To maintain a strongly balanced panel—required by our empirical strategy—we include only quarters for which complete data are available across all states in the sample.[3] We exclude the first quarter of 2011 to retain Massachusetts in the sample, as it is missing data for that period.

Table 1 summarizes the data sources used in our analysis and the corresponding variables drawn from each. Our primary outcome measure — the earnings gap for newly hired men and women — is sourced from the QWI (in bold). A key advantage of the QWI is that it allows us to focus on newly hired workers, the subgroup most directly affected by wage transparency in job advertisements. Indeed, our main outcome of interest is the earnings gap between newly hired women and men. Using this dataset, we construct quarterly measures of labor market outcomes at the state-by-industry-by-gender level, separately for all workers and for new hires. "New hires" are defined as individuals employed by a given firm in a particular quarter who were not employed by that firm in any of the previous four quarters. Earnings are calculated by aggregating total quarterly wages and dividing by the number of workers in each group. Since individuals can hold multiple jobs within a quarter, the data reflect job-level—rather than person-level—outcomes.

To account for confounding factors and improve the precision of our estimates, we augment the QWI with additional control variables from a variety of publicly available sources, also listed in Table 1. These include demographic and socioeconomic indicators from the American Community Survey (ACS) (Ruggles et al., 2025), macroeconomic variables such as real GDP and personal

[3] At the time of writing, only six states had released QWI data for the first quarter of 2024. Including this quarter would require excluding all other states without comparable data.

income from the Bureau of Economic Analysis (BEA) (U.S. Bureau of Economic Analysis, 2024), and hours worked from the Current Population Survey (CPS) (Flood et al., 2024). We also incorporate job vacancy data from the Job Openings and Labor Turnover Survey (JOLTS) (U.S. Bureau of Labor Statistics, 2024a), labor force statistics from the Local Area Unemployment Statistics (LAUS) (U.S. Bureau of Labor Statistics, 2024b), and COVID-19 health metrics from *The New York Times* (The New York Times, 2021). Monthly and annual variables are aggregated or interpolated to the quarterly level to match the frequency of the QWI.

Table 2 presents summary statistics for key labor market outcomes in Colorado and the donor pool of states, disaggregated by gender and by period (pre- and post-policy implementation). The variables include counts of new hires and total employment, along with earnings for new hires — our outcome variable (in bold). number of new hires, total employment, and average monthly earnings for both newly hired and all workers, as reported in the Quarterly Workforce Indicators (QWI). Standard deviations are shown in brackets. The donor pool includes all states eligible to contribute to the synthetic control, whether or not they ultimately received positive weights.

Across both periods, Colorado's labor market characteristics appear broadly similar to those of the donor pool, suggesting that the control group provides a credible counterfactual. In the post-period, average monthly earnings for newly hired male workers in Colorado were $5,120, compared to $4,236 in the donor pool. For newly hired female workers, the corresponding figures were $3,703 in Colorado and $3,056 in the donor pool. These differences suggest that Colorado generally exhibits higher wages for both genders relative to the control group, consistent with its relatively high cost of living and overall wage levels.

Notably, a persistent gender earnings gap is evident across both time periods. In our outcome of interest, the earnings gap in earnings for newly hired workers, we see that newly hired women earned approximately 72% of what men earned in the pre-period in Colorado women ($2,448 vs. $3,510), and 72.3% in the post-period ($3,703 vs. $5,120), indicating little change over time. A similar pattern holds in the donor pool, where newly hired women earned 67% of newly hired men in the pre-period and 72% in the post-period. While the gap narrowed slightly in the donor pool, the magnitude of the change is small. These descriptive statistics suggest that, at least in aggregate, gender differences in earnings remained relatively stable before and after Colorado's wage transparency law, motivating a more formal analysis of the law's causal effect. As additional controls, we include counts of hires, total employment, and separations. These are also similar for both groups and periods.

Table 3 presents summary statistics for the full set of control variables used in constructing the synthetic control, excluding employment-related variables derived from the QWI. The variables in this table are sourced from the ACS, BEA, CPS, JOLTS, LAUS, and NYT datasets, as detailed in Table 1. The table reports averages and standard deviations (in brackets) for Colorado and the donor pool, separately for the pre- and post-treatment periods. These covariates capture a broad set of economic, demographic, and public health characteristics relevant to earnings dynamics and labor market composition.

Overall, Colorado appears well-matched to the donor pool across most dimensions. In both the pre- and post-periods, differences between Colorado and the donor pool are small in absolute terms and consistent in direction. For instance, Colorado exhibits slightly higher real personal income per capita and consumption levels, and somewhat higher labor force participation rates—features consistent with its above-average economic performance. Racial and educational composition are also broadly comparable: Colorado has a marginally higher share of White and college-educated individuals, and a slightly lower share of Black and Asian residents. COVID-19 case and death rates are low in both Colorado and the donor states during the post-period, and labor supply indicators—such as weekly hours and weeks worked by gender—track closely between the two groups. Collectively, these similarities lend further support to the credibility of the synthetic control as a counterfactual for Colorado in the absence of the policy.

Lastly, we incorporate data from Lightcast to examine employer posting behavior in Colorado, which provides suggestive evidence on potential mechanisms underlying our main findings. Lightcast aggregates job postings scraped from over 65,000 online sources, including employer websites, job boards, and staffing agencies, to construct a comprehensive dataset on job vacancies and their attributes (Lightcast, 2024). These data are aggregated by state, year, and industry, and include key variables such as the total number of postings, the share of postings that include salary information, and median posted salaries.

This dataset allows us to analyze how employer behavior evolved before and after the implementation of the law, and whether changes in transparency were accompanied by changes in advertised pay or job composition. While these analyses are descriptive and not causal, they offer important context for understanding the extent to which employers responded to the policy mandate and how those responses may have interacted with worker behavior. We use these data in Section 6 to assess the correlation between earnings & hirings and job-posting information across industries.

4 Empirical Strategy

4.1 Synthetic Control Estimation Method Description

To evaluate the causal impact of Colorado's Equal Pay for Equal Work Act on the earning gap, we apply the synthetic control method (SCM), a technique introduced by Abadie et al. (2010) and widely used for comparative case studies in policy evaluation (Abadie and Gardeazabal, 2003; Abadie et al., 2012). The central idea is to estimate a counterfactual outcome for the treated unit (Colorado) by constructing a synthetic control—a weighted average of other states ("donor states")—whose pre-treatment outcomes and predictors, or controls, closely resemble those of the treated unit. The resulting synthetic control closely matches the treated unit's outcome before policy enactment and serves as a control group following enactment. Thus, after policy enactment, the difference in outcomes between the treated unit and its synthetic control counterpart reveals the policy's effectiveness.

Formally, Let J denote the number of control units (states), and suppose we observe outcomes over T time periods, with treatment occurring at time T_1. Define $\mathbf{X}_1$ as a $(k \times 1)$ vector of pre-treatment characteristics (including lagged outcomes) for the treated unit, and $\mathbf{X}_0$ as a $(k \times J)$, a matrix of the same characteristics for the control units. The goal is to choose a vector of weights $\mathbf{W} = (w_1, \ldots, w_J)'$ such that:

$$\mathbf{W} \in \mathcal{W} = \left\{ \mathbf{w} \in \mathbb{R}^J : w_j \geq 0 \text{ for all } j, \sum_{j=1}^{J} w_j = 1 \right\},$$

and minimize the distance between the treated unit and the weighted average of control units:

$$\|\mathbf{X}_1 - \mathbf{X}_0\mathbf{W}\|_V = \sqrt{(\mathbf{X}_1 - \mathbf{X}_0\mathbf{W})'\mathbf{V}(\mathbf{X}_1 - \mathbf{X}_0\mathbf{W})},$$

where $\mathbf{V}$ is a positive semi-definite, diagonal matrix assigning importance to each predictor. The resulting optimal weights $\hat{\mathbf{W}}$ are used to construct the synthetic control's outcomes:

$$\hat{Y}_{1t}^{\text{SC}} = \sum_{j=1}^{J} \hat{w}_j Y_{jt} \quad \text{for } t = T_1, \ldots, T,$$

and the estimated treatment effect is:

$$\hat{\tau}_{1t} = Y_{1t} - \hat{Y}_{1t}^{\text{SC}}.$$

To address potential small-sample bias in the original estimator, we implement the bias-corrected synthetic control method developed by Wiltshire (2022), using the `allsynth` Stata package. This method applies a jackknife correction, in which each donor unit j is iteratively left out of the donor pool, and the synthetic gap is recomputed. The bias-corrected estimator is then given by:

$$\tilde{\tau}_{1t} = \hat{\tau}_{1t} - \left(\frac{1}{J} \sum_{j=1}^{J} (\hat{\tau}_{jt}^{(-j)} - \hat{\tau}_{jt}) \right),$$

where $\hat{\tau}_{jt}^{(-j)}$ denotes the synthetic gap for control unit j when it is excluded from the donor pool.

Additionally, we conduct inference through randomization (permutation) inference, a procedure that is standard in synthetic-control applications (Abadie et al., 2010; Abadie and Gardeazabal, 2003; Abadie et al., 2012). The method re-estimates the synthetic control for every donor-pool state *as if that state were treated*. The resulting distribution of post-treatment gaps forms an empirical reference distribution, allowing us to compute non-parametric p-values. If Colorado's estimated effect falls in the extreme tail of this placebo distribution, we interpret it as evidence that the treatment effect is unusually large relative to what would be expected from random reassignment of the treatment.

Finally, we note that recent work in the synthetic control literature has proposed adjustments to the basic synthetic control method to overcome cases where the basic method is unable to build a synthetic control that matches the characteristics of the treated unit well (Abadie and L'Hour, 2021; Ben-Michael et al., 2021). While the package we use offers the functionality to implement these adjustments (Wiltshire, 2022), we do not find that it makes a meaningful difference to our results and thus only present results using the classic method previously described.

5 Synthetic Control, Results

In this section, we describe the composition of the synthetic control—specifically, the weights assigned to control states and predictor variables—as well as the estimated impact of the law on the earnings gap. It is important to note that the donor pool does not include all 50 states but is limited to 41. We exclude Alaska due to sparse QWI data. North Carolina and Michigan are also excluded because their QWI data end prematurely (after 2021 Q3 and 2023 Q1, respectively). In addition, we remove California and Washington, which enacted similar statewide transparency laws, and New Jersey, New York, and Ohio, which implemented localized initiatives affecting only certain jurisdictions. We test the robustness of our findings by iteratively excluding each of these states from the donor pool and re-estimating the synthetic control model. The results remain consistent across these specifications (see Appendix B).

5.1 Similarity of Colorado and the Synthetic Control

The key identifying assumption of the synthetic control method is that the constructed synthetic control closely approximates the trajectory of the treated unit—Colorado—in the absence of treatment. To evaluate this assumption, we compare pre-treatment trends in the earnings gap between Colorado and a synthetic control composed of a weighted combination of other states. The synthetic control is designed to match Colorado on a set of pre-treatment earnings gaps and a rich set of covariates, drawn from sources including the ACS, BEA, LAUS, and the New York Times, as summarized in Table 1.[4] The quality of this match provides evidence supporting the credibility of our identification strategy.

Table 4 displays the states that receive non-negligible weights in constructing the synthetic control. The largest weights are assigned to Texas ($w = 0.196$), Utah ($w = 0.158$), Montana ($w = 0.140$), Nebraska ($w = 0.128$), and Tennessee ($w = 0.127$), with the remaining weights spread across a handful of other states. These weights reflect the degree to which each state's characteristics resemble those of Colorado. The relatively concentrated distribution suggests that a small group of states plays a dominant role in forming the counterfactual. Moreover, the total assigned weight (0.968) indicates that the synthetic control relies on a targeted set of similar states, rather than a diffuse average across many. This focused weighting enhances interpretability and further supports the validity of the synthetic control as a comparison group.

Table 5 reports the balance of predictor variables between Colorado and the synthetic control. A key strength of the synthetic control method lies in its ability to closely match the treated unit on pre-treatment covariates, which improves the credibility of the estimated treatment effect. As shown in the table, the synthetic control replicates Colorado's values with high precision across a wide range of predictors, including demographic characteristics, labor

[4] For the data described in this section, all measures from the ACS and BEA are matched using the year of the current quarter. The NYT measures are from the first month of each quarter, and the LAUS data (average quarterly unemployment rate) averages the monthly rates over each quarter and state.

market indicators, and economic aggregates. This close alignment suggests that the synthetic control provides a plausible counterfactual for what Colorado's earnings gap would have looked like in the absence of the policy.

The weights assigned to each predictor reflect their relative importance in constructing the synthetic control. Variables such as the pre-treatment earnings gap ($w = 0.182$), average weeks worked by women and men ($w = 0.105$ and $w = 0.101$), and unemployment rate ($w = 0.091$) are among the most influential. These variables capture key features of the labor market that are likely predictive of future earnings dynamics. The remaining weights are more dispersed across demographic and macroeconomic variables, such as education levels, racial composition, GDP, and population size. The close correspondence between Colorado and the synthetic control across nearly all of these predictors, both in levels and relative proportions, provides strong evidence of covariate balance, which is critical for ensuring that any post-treatment divergence can be attributed to the policy itself rather than to pre-existing differences.

5.2 Other Assumptions

In addition to the requirement of a good pre-treatment match, the synthetic control method relies on several key assumptions for the estimated treatment effects to be interpreted causally. Two particularly important assumptions are the *no anticipation* assumption and the *Stable Unit Treatment Value Assumption (SUTVA)*.

The *no anticipation* assumption holds that the treated unit—in this case, Colorado—did not change behavior in anticipation of the policy going into effect. If individuals or firms adjusted their actions before the law's formal implementation, the observed pre-treatment outcomes would be contaminated by treatment effects, undermining the identification strategy. Evidence from

Arnold et al. (2022) suggests that employers did not meaningfully adjust their behavior prior to the law's enactment, as there was no significant increase in the fraction of job postings that included salary information before the law took effect. While we are currently exploring ways to validate this assumption on the worker side, we note that it would require workers to be both aware of the forthcoming change and willing to delay job applications or acceptances in anticipation of increased transparency—behavior that seems unlikely in the absence of a visible employer response.

The second assumption, $SUTVA$, requires that the potential outcomes of one unit (e.g., Colorado) are unaffected by the treatment status of other units (i.e., the donor states). Violations of this assumption could arise if, for example, the law induced spillover effects—such as firms relocating postings across state lines to avoid compliance, or workers migrating in or out of Colorado in response to the policy. However, Arnold et al. (2022) find no evidence of a decline in job postings in Colorado following the law's passage, which suggests that employers did not shift job advertisements out of state to circumvent the law. Consequently, we view the risk of substantial spillover effects as limited, and thus consider the SUTVA assumption to be reasonably satisfied in our setting.

5.3 Hours and Earnings

While our primary outcome variable, derived from the QWI, reflects average monthly earnings, it does not allow us to directly observe changes in labor supply—such as hours worked—among the individuals in our sample. That is, the QWI provides information on earnings but not on how those earnings relate to hours worked. If, for instance, the policy led to a narrowing of the gender gap in hourly pay, but women simultaneously chose to reduce their labor supply, then average monthly earnings (our outcome measure) might remain unchanged, obscuring underlying improvements in wage equality. Despite this limitation, we continue to prefer the QWI over survey-based alternatives like the CPS. The QWI offers near-universal worker coverage and, crucially, enables us to distinguish between the earnings of newly hired workers and all workers—a distinction not possible in most survey data.

To address this limitation, we turn to the CPS, which contains data on both weekly hours worked and weeks worked per year. Using this information from the Annual Social and Economic Supplement (ASEC), we estimate average annual hours worked by gender.[5] Figure 1 plots these estimates separately for Colorado and the donor pool.

As the figure shows, average annual hours worked by men consistently exceed those of women in both Colorado and the donor pool. Importantly, we do not observe a dramatic divergence in these trends during the treatment period. In Colorado, men's average hours fell slightly from 1,944 to 1,878 post-treatment, while women's hours actually increased marginally from 1,705 to 1,713. In the donor pool, men's hours also decreased slightly, from 1,917 to 1,901, while women's hours increased from 1,696 to 1,717. Thus, gender gaps in labor supply remained fairly stable before and after the law and moved in similar directions across Colorado and the donor pool.

This suggests that changes in labor supply are unlikely to account for any changes we observe in earnings outcomes. In fact, if anything, the small relative increase in hours worked by women would tend to reduce the observed gender earnings gap in Colorado—meaning that any stability or lack of change in the earnings gap could actually understate potential improvements in hourly pay.

[5] The ASEC is a subsample of the CPS, which reduces the number of available respondents. However, to consistently estimate hours worked for both hourly and salaried workers, we rely on questions from this supplement that are asked of all respondents. We removed top-coded responses and also people that reported 0 hours or weeks. Specifically, we use responses to the questions: "How many weeks did you work last year?" and "How many hours did you usually work per week last year?" We are grateful to the IPUMS support team for their guidance in constructing these estimates.

Finally, to more formally account for the potential role of hours, we conduct a synthetic control analysis using CPS-based measures of average weekly hours. This allows us to impute average hourly wages by dividing QWI monthly earnings by estimated hours. The results of this exercise, presented in Appendix A, closely mirror the findings of our main analysis, lending further support to the conclusion that our results are not driven by gender differences in labor supply.

5.4 Results, Estimated Earnings Gaps

Figure 2 presents our principal estimates. Panel A plots the quarterly gender earnings gap among newly hired workers in Colorado (solid line) alongside the corresponding series for the synthetic control (dashed line). Prior to the law's enactment (marked by the vertical dotted line at 2021 Q1) the two series move almost one-for-one, with average differences close to zero. This close pre-treatment fit supports the credibility of the synthetic control as Colorado's counterfactual.

Panel B highlights the post-treatment divergence by graphing the difference between Colorado's gap and that of the synthetic control. The series oscillates around zero before 2021 but shifts sharply upward thereafter, remaining positive in every quarter. During the first eight quarters after implementation, Colorado's gap exceeds the synthetic benchmark by about \$155 on average, a 15 percent increase over the pre-period mean of \$1,062. The size and persistence of this excess point to a widening, not a narrowing, of gender pay disparities among new hires.

Table 6 corroborates the visual evidence with year-averaged figures. From 2011 to 2020, Colorado's annual gap is statistically indistinguishable from its synthetic counterpart; from 2021 onward, the difference stabilizes in the \$145–\$165 range. These numbers underscore that the post-policy widening is both economically meaningful and sustained.

Section 6 explores why the transparency mandate may have had this unintended effect. In brief, industry-level analyses suggest that greater salary visibility is associated with weaker growth in women's earnings and a declining share of female new hires. Although these correlations are not causal, they are consistent with men reaping larger gains from the policy than women—an outcome at odds with the statute's equity objective.

5.5 Placebo Tests and Inference

Figure 3 displays the randomization-inference exercise. The bold line traces Colorado's post-treatment gap, while each thin gray line shows the gap that would arise if the law were (counter-factually) assigned to a donor-pool state. Following the screening rule in Abadie et al. (2010), we drop placebo states whose pre-treatment mean-squared prediction error (MSPE) exceeds Colorado's by a factor of five; eleven states meet that criterion.[6] The remaining 32 placebo paths cluster tightly around zero before 2021, confirming good pre-period fit, and fan out modestly thereafter. Colorado's gap, by contrast, jumps sharply in 2021 Q1 and stays well above virtually all placebo trajectories throughout the post-period.

Figure 4 summarizes the same information in a single statistic: the ratio of post- to pre-period MSPE. A value greater than one indicates that the treated–synthetic gap grew after policy adoption; larger values signal larger, and/or more persistent,divergence. Colorado's ratio is almost 12, the highest among all 43 placebo states (including those with poor pre-fit). Taken together, the visual evidence and the MSPE ratios imply a p-value well below conventional thresholds, signifying that the observed widening of the gender earnings gap is

[6] Connecticut, Kentucky, Minnesota, Mississippi, Nevada, New Hampshire, New Mexico, North Dakota, West Virginia, and Wyoming. Colorado is of course excluded as well.

highly unlikely to be due to chance assignment of the treatment.

6 Mechanisms

Why might a mandate to post salary ranges fail to shrink gender gaps? Economic theory—and recent empirical evidence—suggests three broad channels: (i) directed search, (ii) bargaining, and (iii) the credibility of pay signals. For directed search, we find that increased transparency did not lead to relatively more women selecting into higher-paying industries. Regarding bargaining, we find that increased wage transparency within an industry was associated with increases rather than decreases in the earnings gap among newly hired workers, indicating that women may not have leveraged the additional information effectively in salary negotiations. Finally, we find no evidence that the mandate substantially improved the accuracy or informativeness of salary signals: the difference between advertised and realized earnings remained largely unchanged following the law's implementation, implying that the posted ranges provided limited additional value in shaping workers' wage expectations. Together, these results illustrate how nuanced worker responses can significantly alter, and even reverse, the intended outcomes of transparency policies. (The data used in the creation of the plots in this section is in Table 10.)

6.1 Directed Search

Greater transparency can, in principle, reduce gender pay disparities by guiding workers toward better-compensated positions. Informative posted salary ranges can help workers—particularly women employed in or searching within lower-paying sectors—redirect their efforts toward higher-paying occupations or firms. If barriers to entry are relatively low, increased applications from women for higher-paying jobs should, in theory, exert downward pressure on wages in these jobs, while simultaneously reducing labor supply in traditionally lower-paying positions, pushing those wages upward and compressing the overall gender earnings gap. However, this beneficial sorting effect relies critically on two conditions: (i) entry barriers, such as required credentials or skills, must be sufficiently low to allow mobility, and (ii) workers must actively observe and respond to the newly available salary information. Empirical support for this directed-search mechanism has been documented by Banfi and Villena-Roldán (2019), who find that jobs posting higher visible salaries receive significantly more applications, while hidden salary postings attract fewer candidates.

Figure 5 provides empirical evidence from the Colorado context to examine whether increased wage transparency led to gender-specific sorting across industries. Specifically, the top panel plots the percentage-point change in the share of newly hired workers who are women against average wages of new hires by industry in the two years before and after the policy. The bottom panel provides analogous information for men. Each observation represents a separate industry, accompanied by a fitted trend line to summarize the overall pattern.

Contrary to the directed-search hypothesis, the top panel shows a negative relationship between industry wages and the change in the share of female new hires. Rather than increasingly sorting into better-paying industries, women's representation as new hires actually tended to decline in industries offering higher average wages. This suggests that increased salary transparency did not significantly improve women's occupational mobility into higher-paying sectors. Conversely, the bottom panel, while exhibiting a subtler positive relationship, indicates that men slightly increased their representation in higher-paying industries following the policy.

Together, these findings suggest that the transparency mandate may not only have been ineffective at promoting beneficial sorting behavior among female workers but also inadvertently benefited male workers more. This outcome underscores the complexities involved in transparency initiatives and highlights the importance of ensuring that such policies are carefully targeted to support their intended beneficiaries.

6.2 Bargaining

Another channel through which salary transparency might narrow the gender earnings gap is by equalizing bargaining power between men and women. Consider a scenario in which firms newly mandated to disclose wage ranges consistently post accurate and informative salary bounds reflecting expected pay for new hires. Assuming no changes to the applicant pool or the underlying value of labor to firms, we would expect greater transparency to reduce the dispersion of realized wages. This would arise because more risk-averse or less confident workers—disproportionately women, according to prior research (see next paragraph)—would feel safer negotiating wages at or above the posted minimum. At the same time, firms could credibly reject demands exceeding their advertised range, limiting the wages paid to the most assertive negotiators.

Substantial experimental and survey evidence indicates that gender disparities in bargaining behavior significantly contribute to gender pay gaps. Women consistently exhibit greater risk aversion, lower competitiveness, and less assertive bargaining compared to men (see Croson and Gneezy 2009 for an extensive review). Field studies and laboratory experiments reinforce these findings, demonstrating women's reluctance to negotiate for higher pay and their tendency to request lower initial salaries relative to similarly-qualified men (e.g., Dohmen and Falk, 2011; Buser et al., 2014; Flory et al., 2014). Surveys from employment platforms like Glassdoor similarly report that women are less likely than men to seek raises (Glassdoor Team, 2021). Moreover, Roussille (2024), using data from an online engineering job marketplace, finds that women initially request salaries on average 2.9% lower than men with comparable qualifications, and employers respond by offering women initial salaries roughly 2.2% lower. These findings suggest that reduced uncertainty in wage expectations could disproportionately benefit female workers, potentially narrowing the pay gap.

To examine whether increased transparency translated into tangible bargaining advantages for women, we explore aggregate correlations between salary visibility and changes in the gender pay gap. If greater transparency significantly improved women's bargaining outcomes relative to men's, we would expect a negative correlation between the share of job postings disclosing salary information and the gender pay gap.[7] That is, increased transparency would lead to bargaining outcomes narrowing the gender gap.

Figure 6 illustrates the relationship between changes in wage transparency and the gender earnings gap across industries following the implementation of Colorado's Equal Pay for Equal Work Act. The horizontal axis displays the percentage point increase in the share of job postings including salary information, comparing averages from the two years before and after the policy. The vertical axis shows the percent change in the gender earnings gap, calculated as the difference between average wages of newly hired women and newly hired men, again comparing averages from the two years before and after policy implementation. Each point represents an industry, with its size proportional to the industry's average number of new hires in the post-treatment period.

The upward-sloping trend line reveals an unexpected pattern: industries that experienced greater increases in wage transparency generally saw expansions—rather than reductions—in the gender pay gap among newly hired work-

[7] We acknowledge that other factors, such as shifts in applicant composition or labor market competitiveness, could also influence earnings changes at the sector level.

ers. This finding contradicts the bargaining hypothesis, which predicts that increased salary transparency would empower women to negotiate relatively higher wages, thereby narrowing the earnings gap. Instead, these results indicate that greater visibility of salary information did not enhance bargaining outcomes for female applicants and may have disproportionately benefited male applicants, exacerbating the gender pay gap.

These industry-level findings and our principal result of an increased overall gender gap may initially appear to contradict existing literature on competitive behavior and pay transparency. While some previous studies have documented null effects of transparency policies (e.g., Böheim and Gust (2021)), we are not aware of prior work demonstrating an *increase* in inequity. Even studies reporting or predicting adverse effects (e.g., Cullen and Pakzad-Hurson (2023), Mas (2017), Perez-Truglia (2020)) typically find lower pay and reduced satisfaction for all or for higher-paid workers, rather than specifically disadvantaging the intended policy beneficiaries. Furthermore, theoretical models such as Cullen and Pakzad-Hurson (2023) still predict greater within-occupation pay compression, despite negative overall wage effects.[8]

However, our context significantly differs from previous research in one crucial respect: the transparency policy we study mandates employers to post salary *intervals*, whereas earlier interventions revealed *single numbers* (either averages or specific individuals' salaries). This increased ambiguity allows workers to employ varying bargaining strategies in response to the revealed information. Consider, for example, a simple scenario in which all workers initially expect unposted jobs to pay \$20 per hour. Then, employers disclose job postings typically paying between \$30 and \$50 per hour. If male applicants, due to greater confidence, request wages nearer \$50, while female applicants request

[8]Importantly, Cullen and Pakzad-Hurson explicitly note that their model does not necessarily predict greater across-occupation pay compression.

wages nearer \$30, the resultant pay gap could widen merely because more positions introduced a larger variance into workers' expectations.[9] Our findings underscore the need for further research into how workers with varying attitudes and attributes interpret and utilize different forms of wage information in bargaining situations.

6.3 Signal Informativeness

For the two previously discussed mechanisms—directed search and bargaining—we assumed that mandated wage disclosures provide informative signals about the distribution of salaries for advertised positions, thus influencing workers' search behavior and negotiation strategies. However, it is plausible that posted salary ranges may not convey meaningful information to applicants for several reasons. Workers might question the credibility of the posted wage information, employers might disclose excessively broad ranges that fail to guide expectations clearly, or posted ranges might systematically misrepresent the actual wages employers anticipate paying. Additionally, workers may already possess accurate salary information from alternative sources or prior experience, reducing the impact of newly disclosed ranges.

Previous analysis by Arnold et al. (2022) provides some insight into the credibility and precision of posted wage ranges in Colorado following the enactment of its transparency law. They find no systematic evidence that newly visible salary ranges became broader or narrower compared to those visible prior to the policy's implementation. This indicates that changes in the breadth of salary disclosures are unlikely to account for substantial shifts in signal informativeness.

While we do not have direct data on workers' perceptions of salary credibility,

[9] Note that in this example, the signal remains informative because all workers initially underestimated the salary range. Thus, this mechanism does not require newly posted ranges to be broader or insincerely posted to generate wider pay differentials.

we can indirectly assess whether salary postings became more indicative of actual realized salaries after the law. Figure 7 compares median annual salaries advertised in job postings (from Lightcast data) with mean annualized monthly earnings realized by workers (from QWI administrative data). Generally, advertised and realized salaries track each other closely over time, though realized earnings exhibit greater short-term volatility.[10]

To further explore this relationship, Figure 8 explicitly graphs the difference between advertised and realized salaries. Due to timing discrepancies between when jobs are posted and filled, we present two measures: one comparing median posted salary in a given period to realized earnings in the same period, and another comparing median posted salary to earnings realized in the subsequent period. A systematic narrowing of these differences post-policy would suggest increased informativeness of posted salaries, while a widening would imply the opposite.

As shown in Figure 8, there is no clear pattern indicating greater convergence between advertised and realized wages following the policy change. The gap fluctuates substantially around zero both before and after the law's implementation without exhibiting any consistent trend toward improved alignment. Together with findings from Arnold et al. (2022), our results do not provide evidence that the transparency law substantially enhanced or diminished the informativeness of salary postings. Consequently, we find no support for changes in wage-signal accuracy as a mechanism by which Colorado's law influenced gender earnings disparities.

[10]Ideally, we would compare median-to-median or mean-to-mean salary figures. However, our datasets differ in their reported measures of central tendency, preventing a direct comparison.

7 Conclusion

In this paper, we evaluated whether Colorado's Equal Pay for Equal Work Act achieved its primary objective of reducing gender pay disparities. Although similar pay transparency initiatives have previously been implemented at various jurisdictional levels, Colorado's law represented a pioneering statewide mandate requiring salary disclosure on all job postings, irrespective of employer size or industry. Given this unique comprehensiveness, we specifically assessed the impact of the law on earnings for newly hired employees using a synthetic control approach.

Contrary to the law's intended outcome, we found robust evidence that the gender earnings gap among newly hired workers widened following the law's implementation. Specifically, our estimates indicate an approximately 15 percent increase in the earnings gap, corresponding to roughly \$155 annually relative to the synthetic control group. These findings were statistically significant under standard inference procedures, including placebo and randomization inference tests. Thus, despite its well-intentioned design, the policy appears not only ineffective but counterproductive in narrowing gender-based wage disparities.

We then explored several theoretical mechanisms—directed search, bargaining, and the informativeness of wage signals—at the industry level to better understand the law's unintended consequences. First, we find no evidence that increased transparency led women to systematically sort into higher-paying industries: changes in the female share of new hires were not positively associated with industry wage levels. Second, we observe a positive correlation between increased transparency and the gender pay gap, driven by rising male earnings in more transparent industries while female earnings remained relatively flat. This pattern is consistent with gender differences in how workers respond to posted wage ranges—whether through job search or bargaining behavior. If men are more likely to anchor on the top of posted ranges and women on the bottom, transparency could inadvertently widen disparities. Finally, we find no meaningful improvement in the accuracy of wage signals: the gap between advertised and realized earnings remained stable before and after the policy. Together, these findings suggest that wage transparency alone—absent attention to how different groups interpret and act on disclosed information—may fall short of, or even work against, equity goals.

Our study underscores the importance of understanding the nuances of how job search behavior and information dissemination interact with transparency policies. Future research should prioritize detailed examinations of worker search processes, particularly how and whether information reaches its intended beneficiaries. Additionally, it may be valuable to investigate complementary policies or targeted information campaigns designed to ensure salary disclosures effectively reach disadvantaged or historically lower-paid worker segments. Ultimately, while transparency remains a potentially powerful tool for equity, our findings highlight the complexity involved in its implementation and the critical importance of targeted policy design to achieve desired equity outcomes.

References

Abadie, Alberto, Alexis Diamond, and Jens Hainmueller, "Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program," *Journal of the American Statistical Association*, 2010, *105* (490), 493–505.

_ , _ , and _ , "Synth: An R Package for Synthetic Control Methods in Comparative Case Studies," *Journal of Statistical Software*, 2011, *42* (13), 1–17.

_ , _ , and _ , "Comparative Politics and the Synthetic Control Method," *American Journal of Political Science*, 06 2012, *59.*

_ and Javier Gardeazabal, "The Economic Costs of Conflict: A Case Study of the Basque Country," *American Economic Review*, March 2003, *93* (1), 113–132.

_ and Jérémy L'Hour, "A Penalized Synthetic Control Estimator for Disaggregated Data," *Journal of the American Statistical Association*, 2021, *116* (536), 1817–1834.

Arnold, David, Simon Quach, and Bledi Taska, "The Impact of Pay Transparency in Job Postings on the Labor Market," *SSRN*, 2022.

Baker, Michael, Yosh Halberstam, Kory Kroft, Alexandre Mas, and Derek Messacar, "Pay Transparency and the Gender Gap," *American Economic Journal: Applied Economics*, April 2023, *15* (2), 157–83.

Banfi, Stefano and Benjamín Villena-Roldán, "Do High-Wage Jobs Attract More Applicants? Directed Search Evidence from the Online Labor Market," *Journal of Labor Economics*, 2019, *37* (3), 715–746.

Ben-Michael, Eli, Avi Feller, and Jesse Rothstein, "The Augmented Synthetic Control Method," Working Paper 28885, National Bureau of Economic Research June 2021.

Buser, Thomas, Muriel Niederle, and Hessel Oosterbeek, " Gender, Competitiveness, and Career Choices *," *The Quarterly Journal of Economics*, 05 2014, *129* (3), 1409–1447.

Böheim, René and Sarah Gust, "The Austrian pay transparency law and the gender wage gap," VfS Annual Conference 2021 (Virtual Conference): Climate Economics 242428, Verein für Socialpolitik / German Economic Association 2021.

Callaway, Brantly and Pedro H.C. Sant'Anna, "Difference-in-Differences with multiple time periods," *Journal of Econometrics*, 2021, *225* (2), 200–230. Themed Issue: Treatment Effect 1.

— **and** — , “Introduction to DiD with Multiple Time Periods,” 2024. https://bcallaway11.github.io/did/articles/multi-period-did.html.

Colorado Department of Labor & Employment, “Equal Pay for Equal Work Act,” 2024. https://cdle.colorado.gov/dlss/labor-laws-by-topic/equal-pay-for-equal-work-act.

— , “Interpretive Notice & Formal Opinion #9A,” 2024. https://cdle.colorado.gov/infos.

Croson, Rachel and Uri Gneezy, “Gender Differences in Preferences,” *Journal of Economic Literature*, June 2009, *47* (2), 448–74.

Cullen, Zoë B. and Bobak Pakzad-Hurson, “Equilibrium Effects of Pay Transparency,” *Econometrica*, 2023, *91* (3), 765–802.

Dohmen, Thomas and Armin Falk, “Performance Pay and Multidimensional Sorting: Productivity, Preferences, and Gender,” *American Economic Review*, April 2011, *101* (2), 556–90.

Flood, Sarah, Miriam King, Renae Rodgers, Steven Ruggles, J. Robert Warren, Daniel Backman, Annie Chen, Grace Cooper, Stephanie Richards, Megan Schouweiler, and Michael Westberry, “IPUMS CPS: Version 12.0 Current Population Survey,” 2024.

Flory, Jeffrey A., Andreas Leibbrandt, and John A. List, “Do Competitive Workplaces Deter Female Workers? A Large-Scale Natural Field Experiment on Job Entry Decisions,” *The Review of Economic Studies*, 10 2014, *82* (1), 122–155.

Glassdoor Team, “Pay During COVID-19: Employed Women 19% Less Likely to Ask for More Money In The Next 12 Months,” Dec 2021. https://www.glassdoor.com/blog/covid-19-pay-survey/.

Gulyas, Andreas, Sebastian Seitz, and Sourav Sinha, “Does Pay Transparency Affect the Gender Wage Gap? Evidence from Austria,” *American Economic Journal: Economic Policy*, May 2023, *15* (2), 236–55.

Lightcast, “Lightcast Data: Basic Overview,” 2024. https://kb.lightcast.io/en/articles/6957498-lightcast-data-basic-overview.

Marfice, Christina, “Pay transparency laws: A state-by-state guide,” *Rippling Blog*, 2024. https://www.rippling.com/blog/pay-transparency-laws-state-by-state-guide.

Marinescu, Ioana and Ronald Wolthoff, “Opening the Black Box of the Matching Function: The Power of Words,” *Journal of Labor Economics*, 2020, *38* (2), 535–568.

Mas, Alexandre, “Does Transparency Lead to Pay Compression?,” *Journal of Political Economy*, 2017, *125* (5), 1683–1721.

PayAnalytics, "US Pay Transparency Laws by State 2025," June 2025. Last updated June 30, 2025.

Perez-Truglia, Ricardo, "The Effects of Income Transparency on Well-Being: Evidence from a Natural Experiment," *American Economic Review*, April 2020, *110* (4), 1019–54.

Roussille, Nina, "The Role of the Ask Gap in Gender Pay Inequality," *The Quarterly Journal of Economics*, 02 2024, *139* (3), 1557–1610.

Ruggles, Steven, Sarah Flood, Matthew Sobek, Daniel Backman, Grace Cooper, Julia A. Rivera Drew, Stephanie Richards, Renae Rodgers, Jonathan Schroeder, and Kari C.W. Williams, "IPUMS USA: Version 16.0 American Community Survey," 2025.

The New York Times, "Coronavirus (COVID-19) Data in the United States," 2021. Accessed April 2025.

U.S. Bureau of Economic Analysis, "Gross Domestic Product (GDP) by State," 2024. Accessed April 2025.

U.S. Bureau of Labor Statistics, "Job Openings and Labor Turnover Survey (JOLTS)," 2024. Accessed April 2025.

_ , "Local Area Unemployment Statistics (LAUS)," 2024. Accessed April 2025.

U.S. Census Bureau, "Longitudinal Employer-Household Dynamics (LEHD)," 2024. Accessed April 2025.

US Census Bureau, "LEHD Data," 2022. https://www.census.gov/programs-surveys/ces/data/restricted-use-data/lehd-data.html.

Wiltshire, Justin C., "allsynth: (Stacked) Synthetic Control Bias-Correction Utilities for Stata," 2022.

8 Figures

Figure 1: Estimated Annual Hours Worked, Pre- and Post-treatment Periods, by Gender

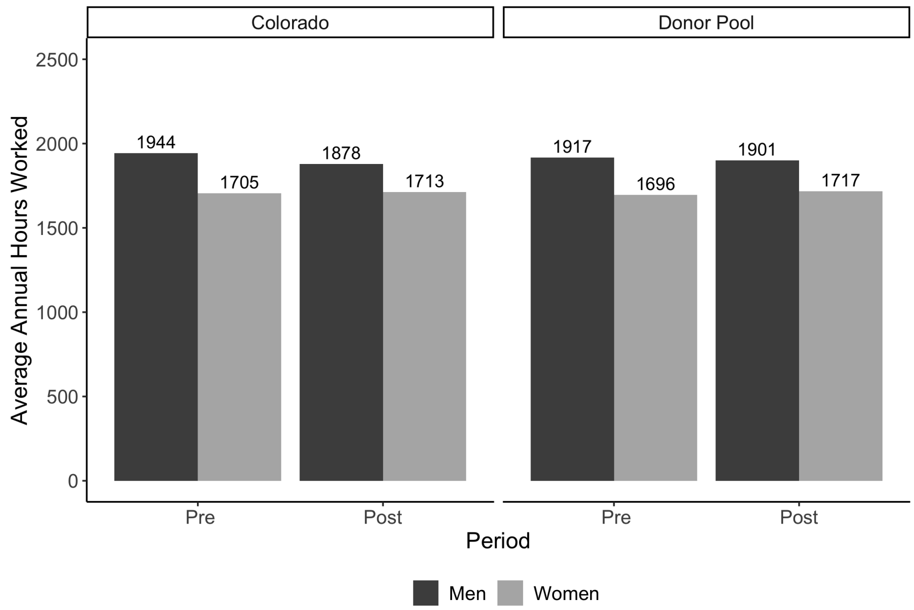

Note: Average hours worked per year by women and men for Colorado and the donor pool. Annual hours are calculated by multiplying number of hours worked per week by number of weeks worked per year. Hours worked per week and weeks worked per year are estimated using the Annual Social and Economic Supplement (ASEC) from the CPS.

Figure 2: Wage gaps for Colorado and the Synthetic Control

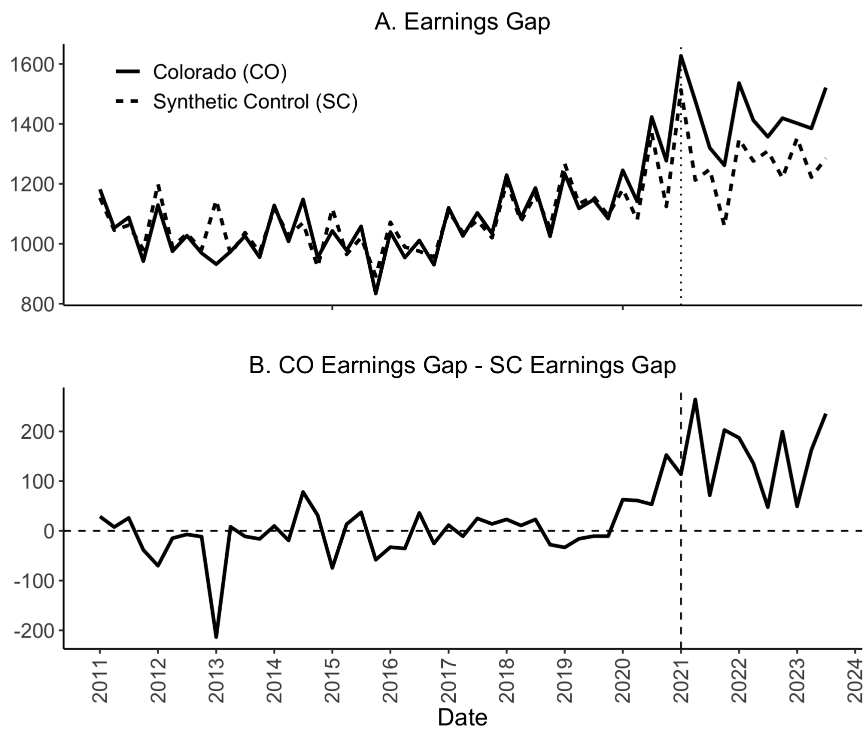

Note: Synthetic control trends. The graphed synthetic control was constructed using data all periods before treatment. The dashed vertical line indicates the first quarter of 2021 when the Colorado law became effective.

Figure 3: Gaps Between the “treated” units and their Synthetic Controls

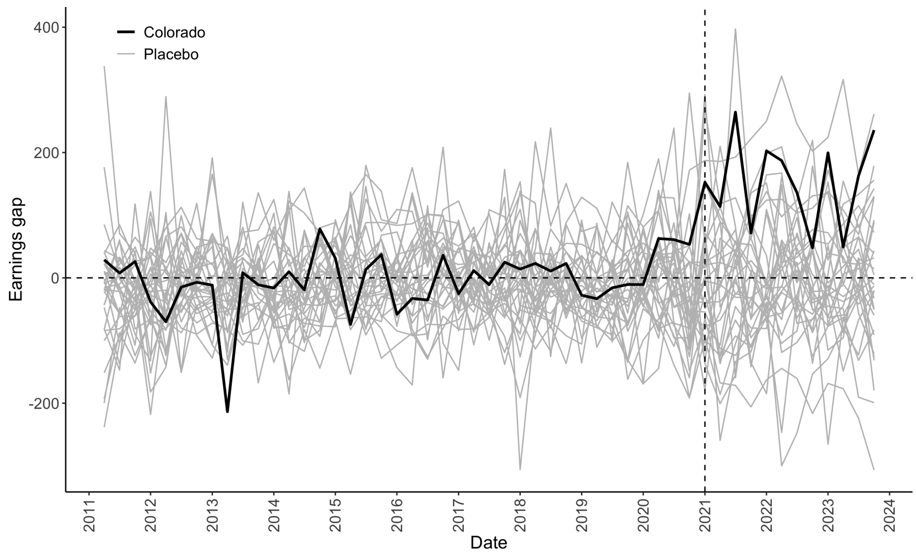

Note: Gaps between the “treated unit” and the synthetic control for placebo states and Colorado. States with a pre-period MSPE fives times that of Colorado were dropped.

Figure 4: Post-period MSPE/pre-period MSPE Ratios

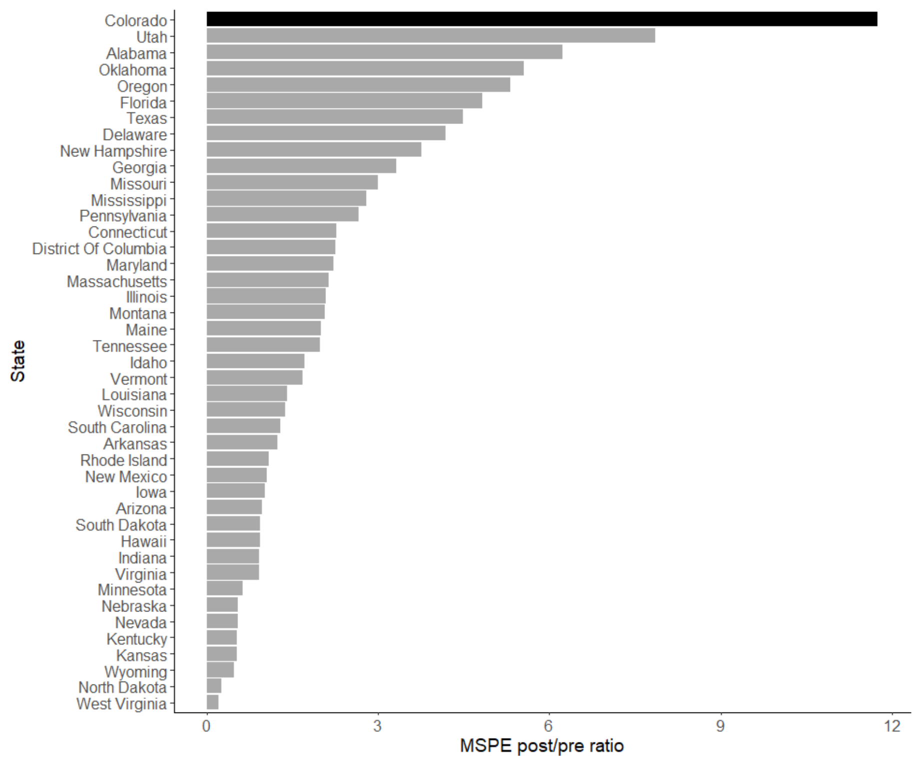

Note: Ratios of post-period MSPEs to pre-period MSPEs for all 43 placebo states and Colorado.

Figure 5: Percentage Point Change in Female and Male Hiring Shares and Average Earnings for New Hires

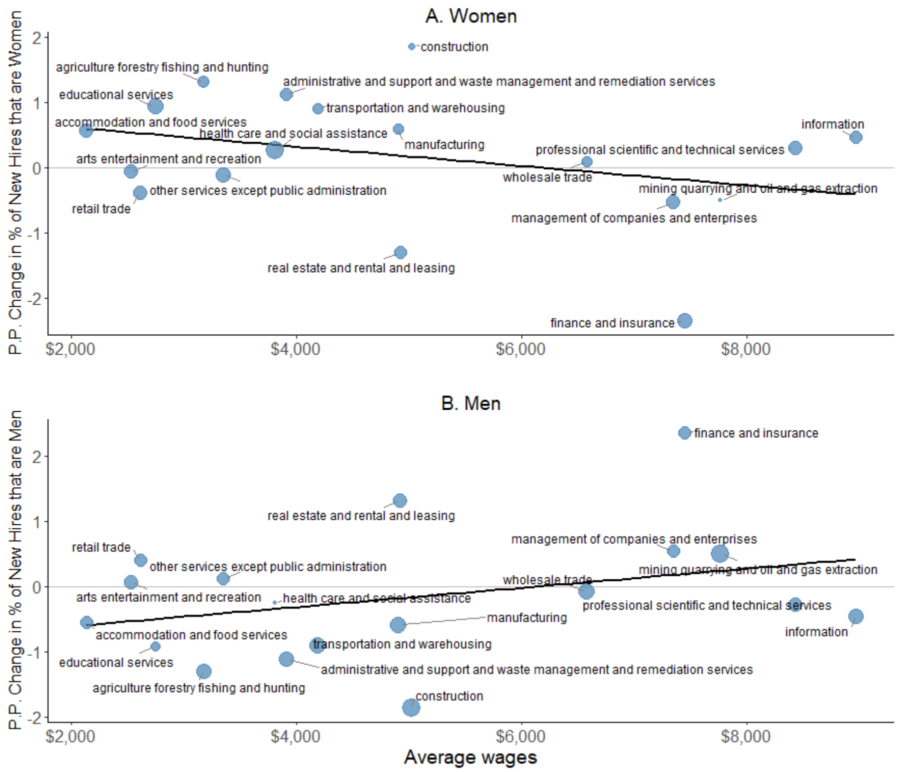

Note: Percentage point changes in the share of new hires who are women (Panel A) or men (Panel B) across industries, comparing the two years before and after the policy's implementation. The horizontal axis indicates average industry wages for newly hired workers in the post-treatment period. The size of each dot reflects the industry's share of new hires who are women (Panel A) or men (Panel B) in the post-treatment period.

Figure 6: Relationship Between Changes In Salary Transparency And The Gender Earnings Gap

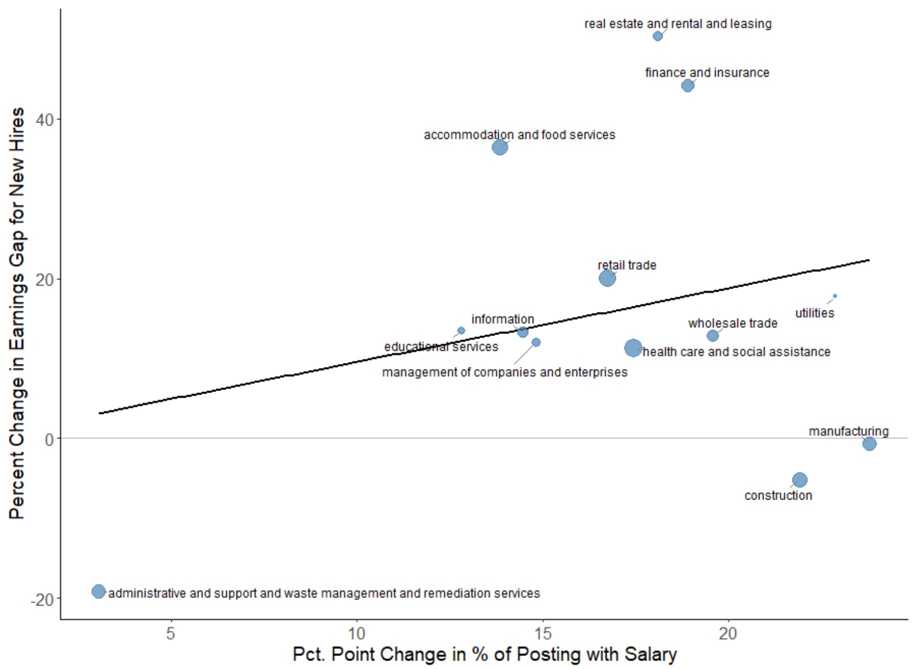

Note: The horizontal axis measures the percentage-point increase in job postings with salary information, comparing averages from the two years before and after the implementation of Colorado's Equal Pay for Equal Work Act. The vertical axis shows the percent change in the gender earnings gap (average wages of newly hired women compared to newly hired men) across the same periods. Each observation represents an industry, with point size proportional to the industry's average number of new hires in the post-treatment period. The upward-sloping trend indicates that industries with greater increases in salary transparency did not experience reductions in the gender pay gap, and instead saw expansions in gender earnings disparities.

Figure 7: Posted and realized salaries for Colorado

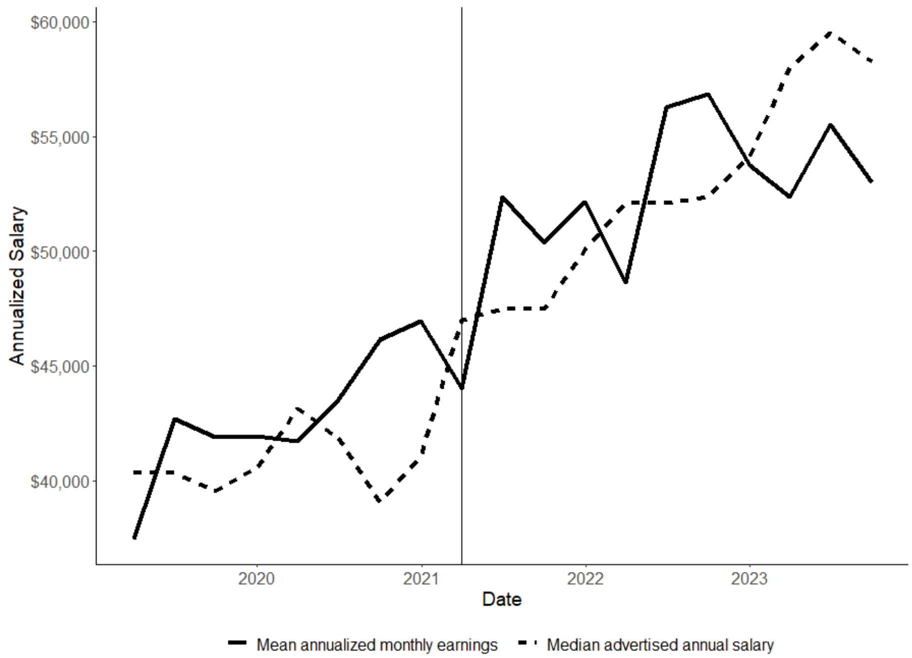

Note: Average realized earnings for new hires from QWI and median posted salaries from Lightcast. For QWI data, average monthly earnings are multiplied by 12 to get annualized amounts. For Lightcast data, these aggregations were generated automatically by limiting the date range of the requested medians to one quarter at a time in Lightcast Analyst.

Figure 8: Difference in posted and realized annualized salaries for Colorado

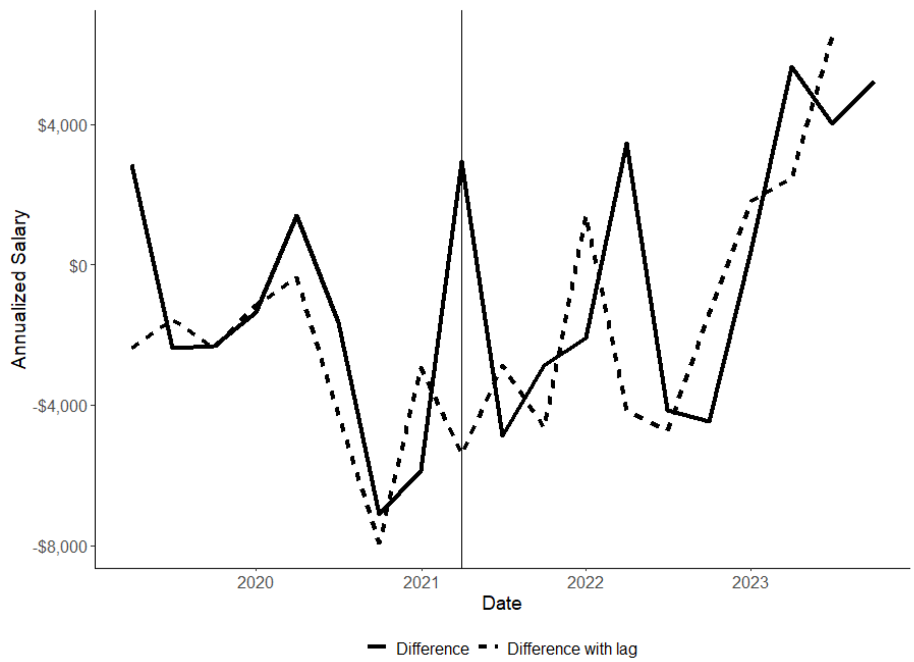

Note: Reported difference is median posted salary from the Lightcast data minus average realized earnings from the QWI data. For QWI data, average monthly earnings are multiplied by 12 to get annualized amounts. For Lightcast data, these aggregations were generated automatically by limiting the date range of the requested medians to one quarter at a time in Lightcast Analyst. "Difference" refers to taking the difference between the two series for the same quarter, whereas "Difference with Lag" refers to taking the difference between posted salaries from one quarter and earnings from the following quarter.

9 Tables

Table 1: Data Sources

Source	**Variable**
American Community Survey (ACS)	
[State; Annual]	Population
	Education
	Sex
	Race
	Marital Status
Bureau of Economic Analysis (BEA)	
[State; Quarterly]	Real GDP
	Real Personal Income
	Expenditure per Capita
Current Population Survey (CPS)	
[State, Sex; Annual]	Average Hours Worked
	Average Weeks Worked
Job Openings and Labor Turnover Survey (JOLTS)	
[State; Monthly]	Job Openings
Local Area Unemployment Statistics (LAUS)	
[State; Monthly]	Unemployment Rate
	Labor Force Participation Rate
Quarterly Wage Indicators (QWI)	
[State, Sex, Industry, Newly Employed; Quarterly]	**Earnings**
	Hires
	Employees
The New York Times (NYT)	
[State; Daily]	Covid Cases per 100k
	Covid Deaths per 100k
Lightcast	
[State, Industry; Monthly]	Median Posted Salary
	Number of Postings
	Percent of Postings with Salary Information

Note: All data is publicly available. All daily and monthly variables were averaged up to the quarterly level. NYT data was reported as 7-day rolling averages before being aggregated to the quarterly level. The outcome of interest is earnings for new hires from the QWI (in bold).

Table 2: Summary Statistics, QWI Variables

	Males		Females	
Variable	Colorado	Donor Pool	Colorado	Donor Pool
A. Post-period				
New Hires	229,937	222,328	207,009	219,528
	[30,346]	[257,326]	[28,444]	[248,466]
Earnings, New Hires	5,120	4,236	3,703	3,056
	[338]	[776]	[281]	[626]
Separations	255,008	243,063	227,150	235,224
	[32,477]	[281,087]	[29,178]	[265,278]
Total Employment	1,260,168	1,295,169	1,089,145	1,221,806
	[47,736]	[1,494,107]	[46,949]	[1,387,074]
Earnings, Total	7,060	6,107	4,719	4,041
	[437]	[1,245]	[297]	[851]
B. Pre-period				
New Hires	202,124	195,246	171,880	178,678
	[30,097]	[226,332]	[28,101]	[201,887]
Earnings, New Hires	3,510	3,181	2,448	2,134
	[415]	[655]	[329]	[501]
Separations	228,453	222,874	191,713	200,625
	[35,550]	[259,065]	[33,010]	[227,456]
Total Employment	1,119,952	1,205,946	967,905	1,122,842
	[87,674]	[1,358,609]	[71,317]	[1,237,246]
Earnings, Total	4,946	4,825	3,316	3,025
	[476]	[1,001]	[355]	[664]

Note: Averages of select QWI variables with standard deviations are reported in brackets. Our outcome of interest in the gap between earnings for new hires between men and women (in bold). Donor pool numbers are for all states that could contribute to the synthetic control, whether or not they actually received positive weights in the synthetic control's construction. Earnings are average monthly earnings of newly hired workers in a quarter using QWI data.

Table 3: Summary Statistics, Control Variables

	Pre-period		Post-period	
Variable	Colorado	Donor Pool	Colorado	Donor Pool
Population	5,843,202	6,608,151	5,478,890	6,372,503
	[28,051]	[7,511,122]	[232,193]	[7,203,889]
Real GDP	420,776	435,614	333,773	372,824
	[14,239]	[570,411]	[35,180]	[468,138]
Real Personal Income per Capita	64,787	58,310	51,958	50,087
	[631]	[6,510]	[4,608]	[6,670]
Real PCE per Capita	50,441	45,908	41,258	40,070
	[1,587]	[4,378]	[2,871]	[4,483]
Number of job openings	216	201	104	107
	[30]	[215]	[31]	[118]
Unemployment rate	3.97	3.82	5.03	5.65
	[1.22]	[1.32]	[2.43]	[2.35]
Labor force participation rate	68.07	62.58	68.06	63.72
	[0.30]	[3.86]	[0.96]	[4.08]
Weekly hours, males	40.41	40.42	40.94	40.75
	[0.51]	[0.82]	[0.65]	[0.87]
Weekly hours, females	36.81	36.54	35.92	36.03
	[0.46]	[0.97]	[0.39]	[1.02]
Weeks worked, males	46.62	46.92	47.47	47.06
	[1.07]	[1.31]	[0.89]	[0.98]
Weeks worked, females	46.03	45.78	46.09	46.09
	[1.34]	[1.43]	[0.69]	[0.91]
Covid cases per 100k	22.79	22.98	1.57	1.66
	[23.44]	[24.85]	[7.72]	[8.05]
Covid deaths per 100k	0.15	0.21	0.02	0.03
	[0.15]	[0.24]	[0.09]	[0.13]
Fraction White	0.70	0.68	0.83	0.75
	[0.00]	[0.15]	[0.04]	[0.14]
Fraction Black	0.04	0.11	0.04	0.11
	[0.00]	[0.10]	[0.00]	[0.11]
Fraction Asian	0.03	0.05	0.03	0.04
	[0.00]	[0.07]	[0.00]	[0.07]
Fraction other races	0.22	0.17	0.10	0.09
	[0.00]	[0.09]	[0.04]	[0.06]
Fraction male	0.51	0.50	0.50	0.49
	[0.00]	[0.01]	[0.00]	[0.01]
Fraction married	0.42	0.40	0.41	0.39
	[0.00]	[0.03]	[0.00]	[0.03]
Fraction some high school	0.23	0.25	0.26	0.27
	[0.01]	[0.02]	[0.01]	[0.03]
Fraction high school	0.24	0.29	0.24	0.29
	[0.00]	[0.03]	[0.00]	[0.03]
Fraction some college	0.18	0.18	0.19	0.19
	[0.00]	[0.02]	[0.01]	[0.02]
Fraction college	0.21	0.16	0.18	0.14
	[0.00]	[0.02]	[0.01]	[0.02]
Fraction post college	0.12	0.10	0.10	0.08
	[0.00]	[0.03]	[0.01]	[0.03]

Note: Averages of control variables with standard deviations reported in brackets. Donor pool numbers are for all states that could contribute to the synthetic control, whether or not they actually received positive weights in the synthetic control's construction. Real GDP is in millions of chained 2017 dollars; job openings is in thousands.

Table 4: SCM Results, State Weights

State	Weight
Texas	0.196
Utah	0.158
Montana	0.140
Nebraska	0.128
Tennessee	0.127
Delaware	0.075
Nevada	0.060
New Mexico	0.046
District Of Columbia	0.037
Total	0.968

Note: States with weights that are at least 0.03.

Table 5: SCM Results, Variable Weights

Variable	Weight (w)	Colorado	Synthetic Control
Earnings Gap (pre-period)	0.182	1,065	1,070
Weeks Worked, Females	0.105	46	46
Weeks Worked, Males	0.101	47	47
Unemployment Rate	0.091	5	5
Some College	0.086	0.19	0.19
Real Personal Consumption Per Capita	0.061	41,351	41,331
Labor Force Participation Rate	0.059	68	67
Asian	0.055	0.03	0.04
Cases Per 100K	0.049	2	2
Other Race	0.042	0.10	0.09
Real GDP	0.031	335,087	344,612
White	0.025	0.83	0.79
Weekly Hours, Females	0.019	36	36
Hires, Male	0.013	235,657	228,898
Separations, Male	0.012	230,492	223,840
Population	0.011	5,488,175	5,964,141
Job Openings	0.009	105	104
Turnover Rate, Female	0.008	0.11	0.10
Total Employment, Female	0.006	970,880	1,018,042
New Hires, Male	0.006	203,910	203,952
Total	0.972		

Note: Average values for Colorado and the synthetic control for variables with weights that are at least 0.005. The "Weight" column is the weight assigned to each variable. The synthetic control is a weighted average computed by taking the values of the variables by state and multiplying them by their associated state weights.

Table 6: SCM Results, Earnings Gap for Colorado and Synthetic Control Group

Year	Avg. Earnings Gap		CO minus SC
	Synthetic Control (SC)	Colorado (CO)	
2011	1,060	1,066	6
2012	1,051	1,025	-26
2013	1,030	972	-58
2014	1,034	1,060	25
2015	999	978	-20
2016	998	984	-14
2017	1,061	1,071	10
2018	1,124	1,132	7
2019	1,163	1,146	-18
2020	1,189	1,272	82
2021	1,258	1,422	163
2022	1,288	1,431	143
2023	1,287	1,436	149

Note: Earnings for Colorado and the synthetic control group, averaged by year.

A Hourly Wage Synthetic Control Analysis

For our main analysis, we directly use the monthly earnings measures for newly hired workers from the QWI. In this section, we instead construct a measure of average hourly wages by dividing each gender group's average monthly earnings measures by the product of its corresponding average weekly hours worked measure from the CPS and 4.345, the average number of weeks in a month.[11] We then use this as our outcome variable and repeat our synthetic control analysis as before.

The synthetic control group composition, predictor balance, and effect estimates are given in tables 7, 8, and 9 below, respectively.

Notably, the selection of states into the synthetic control group is very similar as when using monthly earnings directly, with the largest difference being that North Dakota goes from receiving a small positive weight to receiving 0 weight. Otherwise, the states included remain the same, with Minnesota still receiving a majority of the weight.

Similar to our main analysis, we find no evidence that the law narrowed pay gaps; for each quarter following treatment, the gap between Colorado and its synthetic control is positive, although the difference is never significant at the 5% level.

State	Weight
Minnesota	.548
Utah	.177
Vermont	.143
District of Columbia	.063
Hawaii	.034
Oregon	.034

Table 7: Synthetic control, hourly wages: Composition from donor pool states.

[11] Note that this proxies the average hours worked by new hires with an estimate of the average hours worked by all workers.

Predictor	Colorado	Synthetic Colorado
Gender Wage Gap	3.99	4.03
Real Personal Income Per Capita	52,119.44	51,984.87
Percent White	82.89	80.67
Percent Black	4.10	6.86
Percent Married	40.83	40.37
Percent Male	50.32	49.70
Percent with High School Only	23.90	26.25
Percent with Four Years of College Only	18.13	16.41
Percent with Postgraduate Education	10.20	9.17
Percent in Labor Force	68.02	68.82
Average Quarterly Unemployment Rate	4.94	4.52
COVID-19 Cases Per 100,000	1.61	1.69
COVID-19 Deaths Per 100,000	.02	.02

Table 8: Predictor balance. ACS measures (demographics, education) include general population of all ages.

Quarter	Gap	p
2021 Q1	.15	.9762
2021 Q2	.56	.7619
2021 Q3	1.05	.3571
2021 Q4	.36	.4048
2022 Q1	1.06	.2619
2022 Q2	1.13	.2381
2022 Q3	.21	.2619
2022 Q4	.40	.3095
2023 Q1	.76	.3333
2023 Q2	.26	.3571
2023 Q3	.70	.3333
2023 Q4	1.33	.2143

Table 9: Synthetic control: Hourly wage results. "Gap" refers to the difference in the outcome variable (the gender wage gap) between Colorado and the synthetic control.

B Different Donor Pools

Figure 9: Difference Between Earnings Gaps Between Colorado and the Synthetic Control Group

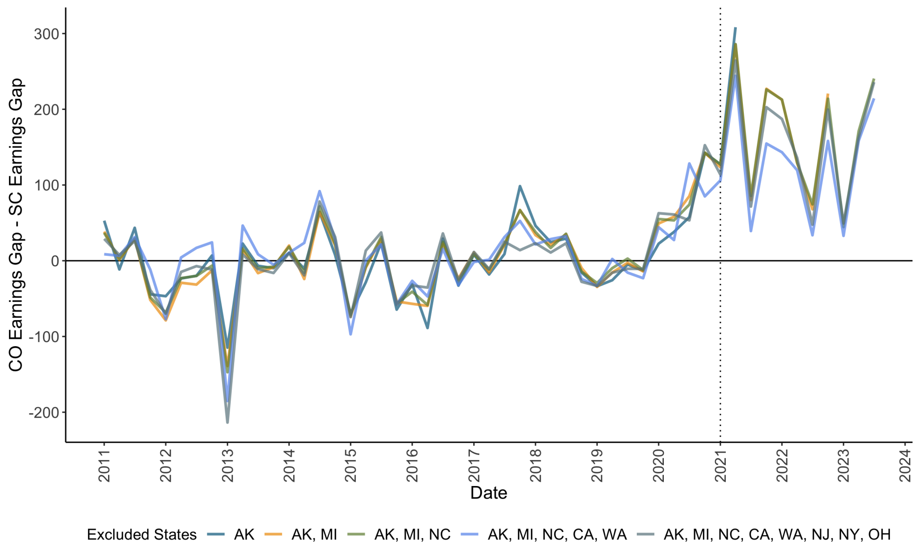

Figure 10: Difference Between Earnings Gaps Between Colorado and the Synthetic Control Group

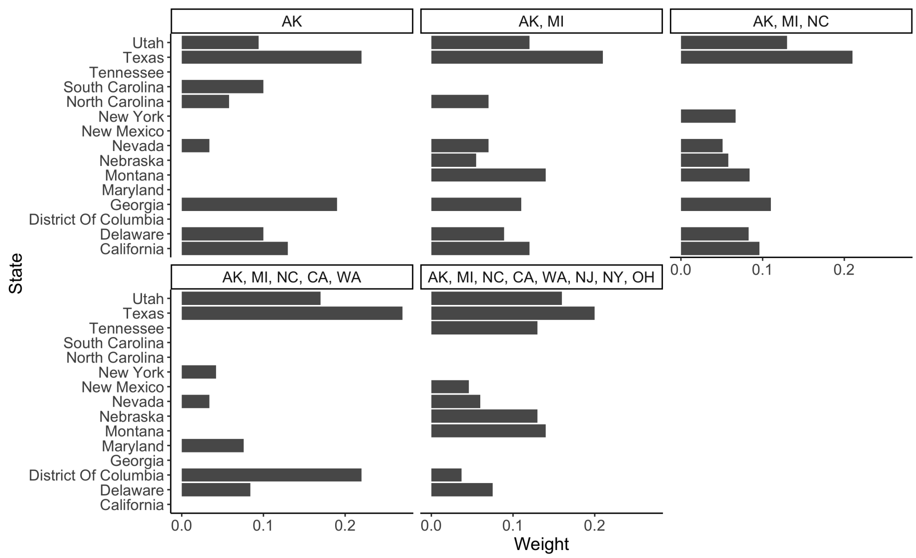

C Multiple State Analysis

For our main analysis, we only examine the effect of the Colorado law. We do this since Colorado's law became effective two years before the law of any other state, and so it has the most available data for analysis. Further, it is most straightforward to construct a synthetic control for a single state with a single time of treatment. In this section, we attempt an additional analysis which aims to get an aggregate average effect of pay transparency laws across all states for which we can get any data. To do this, we use a method for difference-in-differences with multiple treatments introduced by Callaway and Sant'Anna (2021).

C.1 Estimator Description

For this analysis, we use the Callaway and Sant'Anna (2021) staggered difference-in-difference estimator to measure the effect of the state laws on worker outcomes.[12] Specifically, this estimator attempts to estimate *group-time average treatment effects* Callaway and Sant'Anna (2021). That is, for the group of individuals who receive treatment in period g, the average treatment effect at time t is

$$ATT(g,t) = E[Y_t(g) - Y_t(0)|G_g = 1]$$

While these group-time average treatment effects can be reported separately for each treatment group, we instead use two types of aggregation in reporting our main results to ease visualization and interpretation.

First, the "simple" aggregator:

$$\theta_S^O := \sum_{g=2}^{T} \frac{1}{T-g+1} \sum_{t=2}^{T} \mathbf{1}\{g \leq t\} ATT(g,t) P(G=g)$$

where T is the number of time periods in the sample. This aggregation measures the average effect of treatment participation among all ever-treated groups (Callaway and Sant'Anna, 2024). It has the advantage of giving a single number which makes for an easily-interpretable result.

Second, the "dynamic" aggregator:

$$\theta_D(e) := \sum_{g=2}^{T} \mathbf{1}\{g+e \leq T\} ATT(g, g+e) P(G=g|G+e \leq T)$$

which gives the average effect for units treated for e periods (Callaway and Sant'Anna, 2024). This aggregation is used for the coefficients shown in Figure

[12]We make use of the *csdid* command in Stata to implement this estimator.

14. It provides a sense of the lasting impacts of the transparency laws.

C.2 Selection of Treatment and Control Groups

The treated states are those that passed transparency laws which became effective over the sample period. These are Colorado, Washington, and California. Although New York state passed a similar state-wide policy which went into effect in the fall of 2023, it was excluded from the sample because New York City had passed a similar, more local law the year prior. Similarly, New Jersey and Ohio were excluded from both treated and control groups since each experienced partial treatment in the form of more local transparency laws. Following Cullen and Pakzad-Hurson (2023), we use all states never treated prior to 2024 as a control group. Only four states (Alaska, Michigan, Mississippi, and North Carolina) were dropped due to a lack of available data.

C.3 Identifying Assumptions: Difference-in-differences with multiple treatments

The validity of our estimator depends on the following assumptions, the formulas for which are taken directly from Callaway and Sant'Anna (2021).[13] In this section, we describe how each of these apply to our context.

C.3.1 Irreversibility of Treatment

$$D_1 = 0 \text{ almost surely (a.s.)}$$

$$\text{For t} = 2, \ldots, T, D_{t-1} = 1 \implies D_t = 1 \text{ a.s.}$$

[13] Callaway and Sant'Anna also describe an additional assumption needed if using not-yet treated observations in the control group. This is that there must be parallel trends between treated and not-yet treated observations if using not-yet treated observations in the control group. However, we do not currently do this, instead only using never treated units. Therefore, our current analysis does not require this assumption.

This assumption requires that once an observation becomes treated, it remains treated thereafter. Since none of the laws requiring greater transparency were repealed over the course of our sample, this assumption should be satisfied in our sample.

C.3.2 Limited Treatment Anticipation

There exists a known $\delta \geq 0$ such that

$$E[Y_t(g)|X, G_g = 1] = E[Y_t(0)|X, G_g = 1] \text{ a.s.}$$
$$\forall g \in G, t \in 1, \ldots, T \text{ such that } t < g - \delta$$

That is, if treated states anticipate treatment, there must be some known limit to this anticipation. Since we look at worker outcomes, this assumption would be violated if worker behavior changed in an undetectable way in advance of the laws coming into effect (e.g. waiting to apply for jobs for some unknown number of months ahead of the January 2021 in Colorado anticipation of more wages being visible afterwards). It could also be affected by employers changing in anticipation of the policy, since employer behavior could affect worker outcomes.[14]

Arnold et al. (2022) suggests that (in Colorado's case), trends appear fairly parallel for employer posting behavior prior to transparency law implementation with untreated states. We show in the graphs below that trends in posted wages do not differ greatly between treatment and control groups either immediately before the 2021 Colorado law or immediately before the 2023 California and Washington laws.[15] The first graph shows trends in levels, whereas the second

[14]Currently, we are only able to test for changes in employer behavior, but we are planning to add additional analysis explicitly examining changes in worker behavior in a future draft.

[15]These graphs are taken from existing summary methods made available by Lightcast. Because of this, the displayed trends use median advertised wages rather than an estimate of the mean. The median should be sufficient to show a change in the direction of posted

shows trends in percent changes.

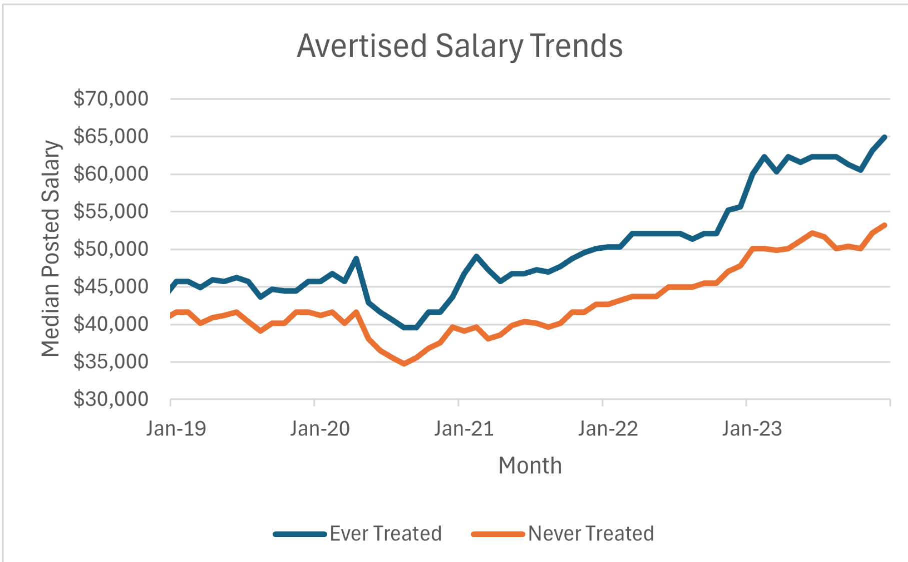

Figure 11: Median advertised annual wages, levels.

wages, but we note that our main results using the QWI are based on average wages, not medians. Because of this, we instead use an estimate of average posted wages when making direct comparisons between Lightcast and QWI data, as in our section on evaluating changes to signal informativeness (see section 6.3).

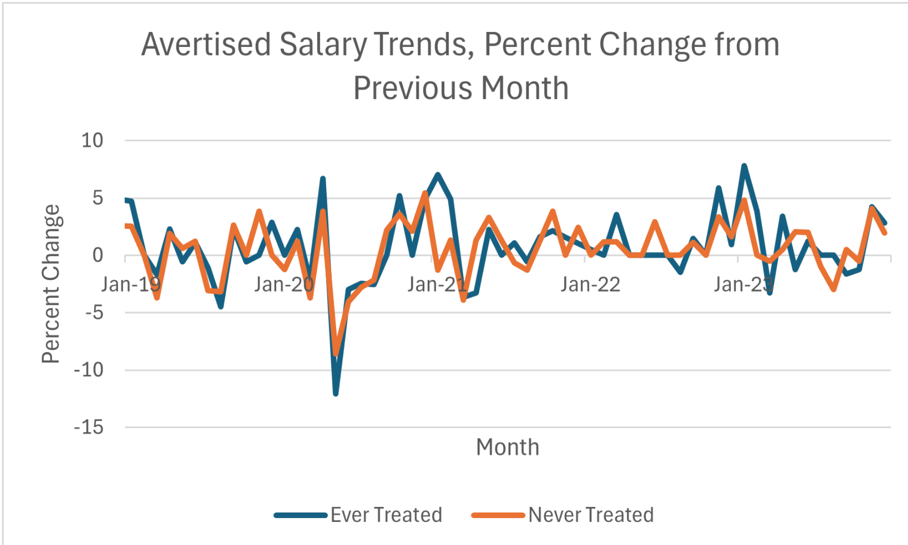

Figure 12: Median advertised annual wages, change from previous month.

C.3.3 Conditional Parallel Trends Assumption

For each $g \in G$ and $t \in 2, \ldots, T$ such that $t \geq g - \delta$,

$$E[Y_t(0) - Y_{t-1}(0)|X, G_g = 1] = E[Y_t(0) - Y_{t-1}(0)|X, C = 1] \text{ a.s.}$$

where C is an indicator equal to 1 if observations are in the never treated group. That is, counterfactual trends must be parallel between treated observations and never treated observations, conditional on included covariates.

The following graph illustrates the trends in our main outcome of interest between treated and control states.

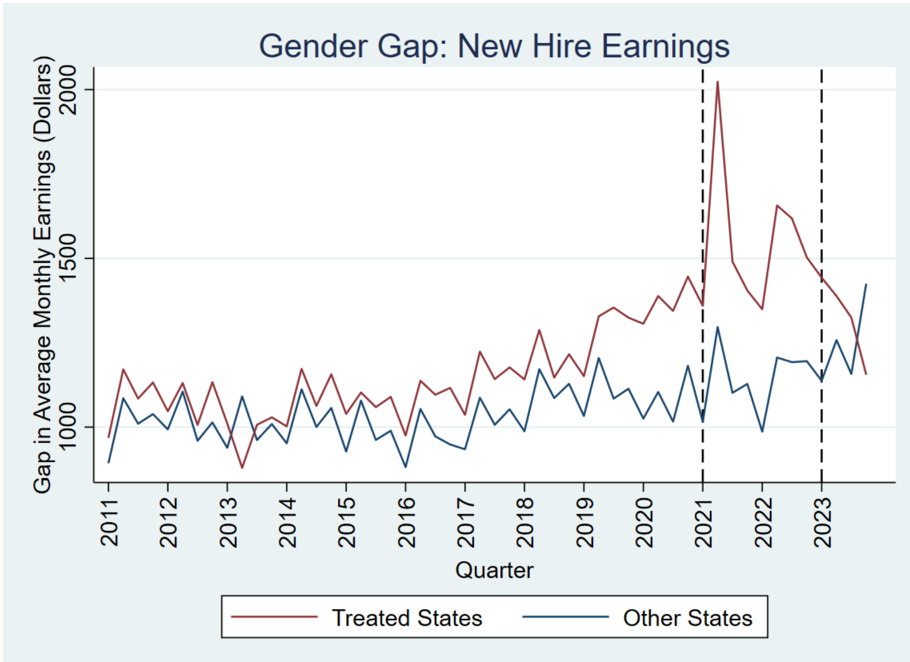

Figure 13: Gender gap in average monthly earnings of newly hired workers. The vertical axis displays male average earnings minus female average earnings. The dashed vertical lines mark the effective dates of the pay transparency laws.

Additionally, we note that Figure 14 presented in the results section displays the estimated coefficients for the aggregate average treatment effect on the treated both before and after treatment. The coefficients do not differ significantly from 0 in the periods before treatment, so we do not have evidence that pre-treatment trends differ between the two groups.

Table 10: Averages by Industry, Post-Treatment Period

Industry	% of Postings with Pay (PwP)	PP change in % of PwP	Employees	New Hires	Percent Female, New Hires (FNH)	PP chg. in % FNH	Wages, All New Hires	Wages, Female New Hires	% chg. in Earnings Gap
Accommodation And FS	29.71	13.86	260691	86032	50.93	0.55	2134	2000	36.41
ASWRMS	35.60	3.06	157422	52801	42.21	1.12	3909	3671	-19.15
Construction	46.04	21.92	177019	30585	19.72	1.87	5021	4356	-5.13
Educational Services	37.43	12.83	42549	6457	64.71	0.94	2747	2492	13.50
Finance And Insurance	33.12	18.90	117693	10482	53.27	-2.35	7454	5935	44.15
Health Care And SA	26.23	17.45	303920	46597	77.25	0.26	3808	3555	11.24
Information	25.59	14.45	79858	7521	40.75	0.47	8972	7672	13.27
MCE	25.44	14.81	47505	5691	49.21	-0.54	7348	5770	12.03
Manufacturing	31.90	23.82	156283	17616	34.33	0.59	4902	4262	-0.73
Real Estate and RL	34.41	18.10	56481	8406	44.46	-1.31	4920	4366	50.33
Retail Trade	29.05	16.74	265216	55926	49.75	-0.40	2610	2236	20.05
Transportation and WH	46.09	13.60	100251	24004	33.16	0.90	4187	3215	87.01
Utilities	51.75	22.89	9050	516	31.45	5.85	8194	7099	17.75
Wholesale Trade	34.42	19.60	114787	13474	33.17	0.09	6580	6068	12.84

Note:
This Table contains the data used in the plots in the mechanisms section. Some industry names were abbreviated for dispaly: ”ASWMRS” is an abbreviation for Administrative And Support And Waste Management And Remediation Services; ”MCE” is Management of Companies and Enterprises; ”Health Care And SA” is Health Care and Social Assistance; ”Real Estate and RL” is Real Estate and Rental And Leasing; ”Transportation and WH” is Transportation and Warehousing. The columns with changes take the changes, in either percentage points or percents, for the two years and before the Colorado Law

C.4 Multiple State Results

In this section, we present our main results. First, the table below gives the simple aggregation results of the average effect of treatment on the treated.

Outcome Variable	ATT	Standard Error	Observations
Gender Gap in New Hire Earnings	-96.5667	97.4635	2,249

Table 11: Simple aggregation results.

That is, we estimate that the gap in earnings in treated states was about $100 smaller on average following the policy than it would have been otherwise, although this effect is not statistically significant. The graph below instead looks at the average treatment effect on the treated by length of exposure to treatment.

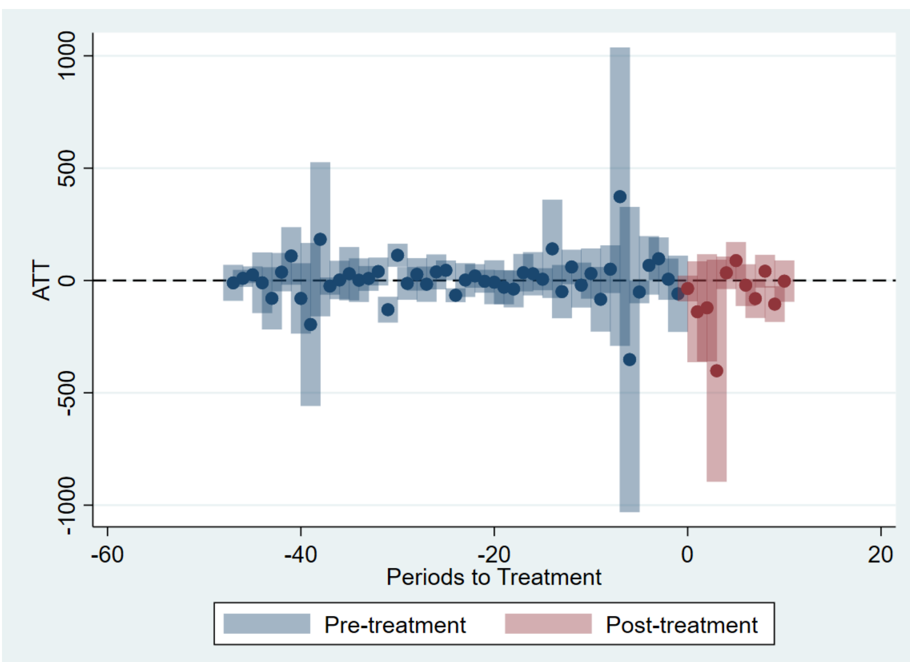

Figure 14: Dynamic effects. Vertical axis represents the average treated effect on the treated for the gender gap in monthly earnings.

Here we see that the treatment effect is most negative in the few periods after treatment. Therefore, we do not have evidence of an enduring impact of this legislation on the gender gap.[16]

[16] Note, however, that only Colorado has enough observations to last for more than the first few quarters after treatment. It is therefore unsurprising that there is no negative effect in the later quarters, since, as we show in the main body of the paper, the gender gap in earnings was higher in Colorado than in other states during the treatment period.