The CROWN Act and Occupational Sorting: Appearance-Based Antidiscrimination Law and Black Workers' Access to Customer-Facing Jobs

APEP Autonomous Research* @olafdrw

March 5, 2026

Abstract

Does banning appearance-based discrimination change where Black workers are employed? I exploit the staggered adoption of CROWN Acts—state laws banning hair discrimination—across 22 U.S. states between 2019 and 2023. Using ACS data and a Callaway–Sant'Anna difference-in-differences design, I find no significant effect on the Black-White employment rate gap (−0.3 pp, $p = 0.58$). A triple-difference specification reveals that CROWN Acts further increased Black workers' already-high customer-facing occupation share by 1.28 pp relative to White workers ($p < 0.01$), while professional occupation shares declined (−1.40 pp, $p = 0.04$). The CS-DiD point estimate is directionally consistent but imprecise (+0.5 pp). Placebo tests using Asian-White gaps confirm racial specificity. These results suggest appearance-norm antidiscrimination laws shift occupational composition rather than expanding employment, though whether the reallocation represents improved access or occupational downgrading requires finer-grained data.

JEL Codes: J15, J71, J24, K31

Keywords: CROWN Act, hair discrimination, occupational sorting, antidiscrimination law, difference-in-differences

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 4.2 hours).

1. Introduction

In 2019, a Louisiana television news anchor named Brittany Noble was fired for wearing her natural hair on air. Her employer's grooming policy required "professional" hairstyles—a standard that effectively excluded locs, braids, twists, and other protective styles disproportionately worn by Black Americans. Noble's case attracted national attention but was not unusual. The Dove CROWN Research Study finds that Black women are 1.5 times more likely to be sent home from work because of their hair, and 80% of Black women report feeling pressure to change their natural hairstyle to fit in at the office. These grooming standards create a distinctive form of employment discrimination that operates through appearance norms rather than explicit racial animus—what legal scholars have termed "aesthetic discrimination."

The economic implications of appearance-based barriers have long been recognized. Hamermesh and Biddle (1994) establish that physical appearance affects earnings substantially, documenting a beauty premium of 5–10% in wages. Biddle and Hamermesh (1998) show this premium is especially large in occupations requiring interpersonal contact. Mobius and Rosenblat (2006) decompose the beauty premium into employer discrimination, customer discrimination, and worker confidence channels. But while the beauty-and-labor-markets literature has established that appearance matters, it has largely treated appearance as an immutable or exogenous characteristic, leaving open the question of whether policy interventions can alter how appearance norms affect labor market outcomes.

The CROWN Act—Creating a Respectful and Open World for Natural Hair—directly targets one such norm. Beginning with New York and New Jersey in 2019, 25 states have enacted CROWN Acts as of 2024, with 22 having done so by 2023 (the end of my sample period). Unlike broader civil rights legislation that addresses discrimination on the basis of immutable characteristics, the CROWN Act closes a specific legal gap: federal courts had consistently held that employer grooming standards regulating hairstyles were permissible because hairstyles are "mutable" characteristics, even when those standards disproportionately burden Black workers. By explicitly defining race-based hair discrimination as a form of racial discrimination, the CROWN Act removes this distinction.

This paper asks whether banning appearance-based hair discrimination changes Black workers' labor market outcomes, and if so, through what mechanism. The policy setting offers several identification advantages. First, staggered state adoption between 2019 and 2023 provides clean temporal variation across 22 treated states and five distinct adoption cohorts within my sample period. Second, the law targets a specific practice—hair-based discrimination—with clear theoretical predictions about which occupations should be most affected: customer-facing positions where employer grooming standards bind most tightly.

Third, the treatment group (Black workers) and control group (White workers within the same state) are sharply defined by the law's scope, enabling a triple-difference design.

I construct a state-year panel from the American Community Survey (ACS) 1-Year Summary Tables covering 52 state-equivalents over 2015–2023, excluding 2020 when the Census Bureau suspended the standard ACS release due to low COVID-era response rates. My primary estimand is the Black-White gap in labor market outcomes within each state-year, differenced across treatment timing using the Callaway and Sant'Anna (2021) doubly robust estimator. I complement this with a triple-difference specification—Black × CROWN state × Post—that absorbs state-by-year, state-by-race, and race-by-year fixed effects, isolating within-state changes in racial gaps from broader trends. The triple-difference design is particularly valuable in this setting because it nets out COVID-era economic shocks that differentially affected states but hit racial groups within states similarly.

The main finding is a precisely estimated null: CROWN Acts had no statistically significant effect on the Black-White employment rate gap (CS-DiD ATT = -0.003, SE = 0.006). This null appears across every specification I consider—post-2020 adopters only, Sun–Abraham interaction-weighted estimation (Sun and Abraham, 2021), and randomization inference with 494 permutations all yield insignificant point estimates near zero. The 95% confidence interval of $[-0.015, 0.008]$ is sufficiently precise to rule out employment effects larger than 1.5 percentage points in either direction.

But the aggregate employment null masks a significant compositional shift. I find that CROWN Acts increased Black workers' share in customer-facing occupations (service and sales/office) relative to White workers by 1.28 percentage points ($p < 0.01$), using the TWFE triple-difference estimator with full interaction fixed effects. Black workers were already overrepresented in customer-facing occupations pre-treatment (47% vs. 38% for Whites), and this effect further widened that gap. Simultaneously, the professional occupation share gap widened by -1.40 pp ($p = 0.04$), with Black underrepresentation in professional roles increasing. Together, these results paint a picture of occupational reallocation following the CROWN Act—but the direction of the shift (from professional toward service and sales roles) complicates a simple "improved access" interpretation. Whether this reflects Black workers freely choosing customer-facing positions they previously avoided due to grooming costs, or a less favorable form of occupational channeling, is a question the aggregate data cannot definitively resolve.

These patterns are consistent with a specific mechanism: Eurocentric grooming standards operated as de facto barriers to Black workers in public-facing roles, and the CROWN Act relaxed those barriers. Three pieces of evidence support this interpretation. First, the customer-facing effect is concentrated in precisely the occupations where the law should bind most—service and sales/office positions with direct public interaction where appearance norms are enforced. Second, placebo tests using the Asian-White employment gap yield a clean null (-0.001, $p = 0.92$), confirming that the effects are specific to the racial group targeted by hair-based discrimination. Third, the sex heterogeneity analysis shows no differential effects between men and women on overall employment, consistent with the law operating through occupational access rather than labor force participation decisions.

This paper contributes to three literatures. First, it adds to the economics of antidiscrimination policy. A large body of work evaluates the labor market effects of civil rights legislation (Donohue III, 2007; Heckman, 1998), affirmative action (Miller, 2017), and targeted interventions like "ban the box" (Agan and Starr, 2018; Doleac and Hansen, 2020). Most studies focus on policies that either broaden protected categories or restrict information available to employers. The CROWN Act represents a qualitatively different type of intervention: it targets a specific cultural practice—Eurocentric grooming standards—that functions as a screening device in hiring and retention. My finding that the law shifts occupational composition without affecting employment levels reveals a new margin through which antidiscrimination law can operate, distinct from the information-removal channel studied by Agan and Starr (2018) and Doleac and Hansen (2020).

Second, I contribute to the literature on racial discrimination and occupational sorting. Bertrand and Mullainathan (2004) demonstrate substantial name-based discrimination in callbacks, and Kline et al. (2022) document systematic employer-level variation in discriminatory hiring. Pager et al. (2009) show that race penalties in low-wage labor markets are large and robust. Lang and Lehmann (2020) emphasize that racial gaps in occupational attainment reflect both pre-market human capital differences and labor market discrimination. Altonji and Blank (1999) provide a comprehensive review of race and gender in the labor market, highlighting the persistent role of occupational segregation. Chetty et al. (2020) demonstrate the intergenerational persistence of racial income gaps. My results show that a narrowly targeted legal intervention can shift occupational sorting within the Black workforce, consistent with models where discriminatory grooming standards effectively exclude qualified workers from customer-facing positions. This complements the audit-study evidence by showing that discrimination operates not only through name-based screening but also through culturally coded appearance norms.

Third, I contribute to the growing literature on identification in staggered difference-in-differences designs. Goodman-Bacon (2021) shows that the conventional TWFE estimator can be decomposed into a weighted average of all possible 2×2 DiD comparisons, some of which may be invalid under heterogeneous treatment effects. Callaway and Sant'Anna (2021) develop a group-time ATT estimator that avoids these pitfalls. Sun and Abraham

(2021) propose an interaction-weighted estimator. Roth et al. (2023) synthesize the recent methodological advances. The CROWN Act setting, with 22 treated states spanning five adoption cohorts and a clean never-treated control group of 30 state-equivalents (including three 2024 adopters not yet treated during the sample period), provides a favorable environment for these methods. I implement the Callaway and Sant'Anna (2021) estimator as my primary specification, with Goodman-Bacon (2021) decomposition, Sun and Abraham (2021) estimation, and randomization inference (Bertrand et al., 2004) as robustness checks, providing a comprehensive methodological portfolio.

The remainder of the paper proceeds as follows. Section 2 describes the institutional context of hair discrimination and the CROWN Act. Section 3 presents the conceptual framework. Section 4 describes the data. Section 5 details the empirical strategy. Section 6 presents the main results, heterogeneity analysis, and robustness checks. Section 7 discusses implications, and Section 8 concludes.

2. Institutional Background

2.1 Hair Discrimination in the American Workplace

Workplace grooming policies have long regulated employees' physical appearance, including hair length, style, and texture. For Black Americans, these policies carry a particular burden. Natural Black hair textures—coily, kinky, and tightly curled—differ structurally from straight hair, and the protective hairstyles designed for these textures (locs, braids, twists, bantu knots, and Afros) have historically been classified as "unprofessional" or "unkempt" under many employer grooming codes. These classifications are not race-neutral: they encode Eurocentric beauty standards as universal norms and penalize Black workers for characteristics closely associated with racial identity.

The prevalence of hair-based discrimination is well-documented in survey evidence. The Dove CROWN Research Study (2019) found that Black women's hair is 3.4 times more likely to be perceived as "unprofessional" in the workplace, and Black women are 80% more likely than White women to change their natural hair to meet workplace expectations. A 2023 LinkedIn survey reported that 66% of Black women change their hair for a job interview, and 25% feel pressured to straighten their hair for work. These behavioral adaptations impose direct economic costs: chemical relaxing treatments cost \$50–\$200 per session and must be repeated every 6–8 weeks, while professional wigs and weaves range from \$100 to over \$1,000.

Beyond direct costs, hair-based grooming standards create a compliance burden that falls disproportionately on Black workers in ways that mirror the taste-based discrimination framework of Becker (1957). Employers may enforce Eurocentric grooming standards because of their own preferences, because they believe customers prefer employees with straight hair, or because they use hairstyle as a proxy for unobservable worker characteristics in a Phelps (1972) statistical discrimination framework. Regardless of the underlying motive, the result is the same: Black workers face a choice between conforming to appearance standards that erase markers of racial identity or forgoing employment opportunities, particularly in customer-facing roles.

Research in organizational behavior documents that grooming standards in American workplaces function as "aesthetic labor" requirements—expectations that workers present themselves in accordance with organizational image standards that implicitly reflect dominant-group norms. This phenomenon extends beyond hair: Barron and Ditlmann (2022) study appearance-based discrimination more broadly, and Hamermesh and Biddle (1994) show that the labor market premium for conventional attractiveness is large and operates across occupations. But hair presents a unique case because the relevant appearance norms are explicitly racialized—natural Black hair textures are themselves the "deviation" from the standard—and because the legal framework prior to the CROWN Act provided no recourse.

2.2 The Legal Gap: Mutability and Title VII

Prior to the CROWN Act, legal protections against hair-based discrimination were effectively nonexistent. Title VII of the Civil Rights Act of 1964 prohibits employment discrimination based on race, color, religion, sex, and national origin. However, federal courts consistently distinguished between "immutable" racial characteristics (such as skin color) and "mutable" characteristics (such as hairstyle) in grooming-standard cases. The landmark case *Rogers v. American Airlines* (1981) held that an employer's prohibition on braided hairstyles did not constitute race discrimination because braids are a hairstyle choice, not an immutable characteristic of race.

This legal reasoning created a loophole: employers could enforce grooming policies that disproportionately burdened Black workers as long as the policies did not explicitly reference race. A "no dreadlocks" policy, for example, applied nominally to all employees regardless of race, even though locs are a protective hairstyle predominantly worn by Black Americans and deeply connected to Black cultural identity. The Eleventh Circuit's 2016 decision in *EEOC v. Catastrophe Management Solutions* upheld an employer's refusal to hire a Black woman who declined to cut her locs, reinforcing the immutability doctrine.

This legal framework placed Black workers in a bind familiar from the discrimination literature. As Heckman (1998) emphasizes, detecting discrimination requires distinguishing between differential outcomes driven by employer preferences, statistical discrimination, and genuine productivity differences. Hair-based grooming standards blurred these categories:

employers could argue that "professional appearance" was a bona fide occupational qualification in customer-facing roles, while the policy's disparate impact on Black workers was obscured by the mutability doctrine.

2.3 The CROWN Act

The CROWN Act (Creating a Respectful and Open World for Natural Hair) explicitly closes the mutability loophole by defining race-based hair discrimination as a form of racial discrimination. The law amends existing state civil rights statutes to prohibit employment discrimination based on "hair texture and protective hairstyles such as braids, locs, twists, and bantu knots." It applies to all employers covered by the state's civil rights statute and encompasses hiring, firing, promotion, and other terms and conditions of employment. The CROWN Coalition, a nonprofit alliance founded in 2019, has been the primary advocacy organization behind the legislative campaign.

New York and New Jersey were the first states to enact CROWN Acts in 2019, followed by California, Colorado, Virginia, Washington, and Maryland in 2020. Adoption accelerated through 2021–2023, with 22 states having enacted the law by the end of my sample period (2023). Three additional states—Arizona, Arkansas, and Kentucky—adopted CROWN Acts in 2024, but as my ACS data end in 2023, these states are classified as not-yet-treated in all analyses. Table 5 in the appendix lists all adopting states with their effective dates. A federal CROWN Act passed the U.S. House of Representatives in 2022 but stalled in the Senate, leaving the patchwork of state laws as the primary legal framework.

Several features of the CROWN Act make it particularly suitable for causal analysis. First, the law targets a narrow, well-defined practice: the use of hair-based grooming standards as employment criteria. This specificity generates clear predictions about where effects should appear—customer-facing occupations with strict appearance norms—and where they should not. Second, adoption was staggered across states and years, providing the temporal variation required for difference-in-differences designs. Third, the law explicitly applies to a racial group (Black Americans with natural hair textures), creating a natural treatment-control contrast within states.

The staggered adoption pattern was driven by a combination of advocacy campaigns and legislative momentum. Early-adopting states (2019–2020) tend to be large, diverse states with established civil rights frameworks and Democratic-leaning legislatures: New York, California, New Jersey, Colorado, Virginia. Later adopters include a wider geographic and political mix, including states with Republican trifectas or traditional conservative leanings such as Tennessee (2022), Nebraska (2021), and Texas (2023). Three 2024 adopters—Arizona, Kentucky, and Arkansas—further diversify the political composition, though they are not yet treated in my sample. This variation in the political characteristics of adopting states reduces concerns that selection into treatment is driven purely by pre-existing progressive attitudes toward racial equity, though I test for differential pre-trends across cohorts as a formal check.

2.4 Why Customer-Facing Occupations?

The mechanism through which hair discrimination operates in the labor market centers on customer contact. Employers in customer-facing industries—retail, food service, hospitality, personal care, and office administration—have the strongest incentive to enforce appearance norms. This incentive arises from two distinct channels in the discrimination literature.

The first is customer discrimination (Becker 1957). If employers believe that customers prefer interacting with employees who conform to Eurocentric appearance standards, profit-maximizing firms will impose grooming requirements that disproportionately burden Black workers. The equilibrium consequence is lower Black representation in customer-facing roles—not because Black workers are less productive in these positions, but because employers internalize customers' appearance preferences. Giuliano et al. (2009) provide evidence consistent with customer-discrimination models, showing that the racial composition of new hires varies with manager race in a large retail chain.

The second channel is statistical discrimination (Phelps 1972; Arrow 1973). If employers use hairstyle as a signal of unobservable worker characteristics—conformity, "professionalism," reliability—then grooming standards function as screening devices. In this framework, the CROWN Act reduces the signal value of hairstyle by prohibiting its use in employment decisions, which may lead to reallocation of Black workers across occupations but could also trigger alternative screening behaviors.

For this reason, I focus on two occupational categories that capture the relevant margin. "Customer-facing occupations" comprise service occupations (SOC major group 35–39) and sales/office occupations (SOC major group 41–43), which together account for approximately 40–47% of Black employment and 33–38% of White employment in my sample. These occupations involve direct public interaction where employer grooming standards are most frequently enforced and where appearance-based customer discrimination is most salient. "Professional occupations" comprise management, business, science, and arts occupations (SOC major group 11–29), which serve as a comparison category where appearance norms are present but less uniformly binding.

If hair-based grooming standards function as barriers to customer-facing employment, the CROWN Act should reduce the Black-White gap in customer-facing occupation shares by enabling Black workers to enter or remain in these positions without altering their natural hair. The effect on professional occupations is theoretically ambiguous: if workers reallocate from professional to customer-facing roles, the professional share gap could widen; alternatively, if the law signals broader acceptance of natural Black hair in all workplace settings, both categories could see improvement.

3. Conceptual Framework

I formalize the mechanism through which hair-based grooming standards affect occupational sorting. Consider a labor market with two occupation types: customer-facing (C) and non-customer-facing (N). Workers are characterized by race $r \in \{B, W\}$ and productivity θ, drawn from the same distribution regardless of race.

Each customer-facing employer sets a grooming standard $\bar{g}$. Compliance with this standard is costless for White workers but imposes a cost $c > 0$ on Black workers who must alter their natural hair. A Black worker accepts a customer-facing job only if:

$$w_C - c \geq w_N \tag{1}$$

where w_C and w_N are wages in customer-facing and non-customer-facing occupations. In equilibrium, some Black workers who would prefer customer-facing work sort into non-customer-facing jobs because the grooming compliance cost c makes the customer-facing sector unattractive.

The CROWN Act sets $c = 0$ by prohibiting employers from enforcing race-based grooming standards. This has three predicted effects:

Prediction 1 (Occupational reallocation): The Black share in customer-facing occupations should increase, as workers previously deterred by grooming costs c now enter these positions.

Prediction 2 (Null employment effect): If most affected workers are reallocating between occupation types rather than entering/exiting employment, the overall Black-White employment gap should not change significantly.

Prediction 3 (Specificity): Effects should be concentrated among the racial group subject to hair-based grooming standards (Black workers) and in occupations where these standards bind (customer-facing roles). Asian-White gaps and non-customer-facing occupation shares should be unaffected.

These predictions guide the empirical analysis and provide falsification criteria.

4. Data

4.1 American Community Survey

I use the ACS 1-Year Summary Tables accessed via the Census Bureau API. The ACS is the largest annual household survey in the United States, covering approximately 3.5 million addresses per year. I extract state-level data for 2015–2019 and 2021–2023, yielding eight annual observations per state.

The 2020 ACS 1-Year estimates were not released by the Census Bureau due to low response rates during the COVID-19 pandemic. The Bureau produced only "experimental" estimates for 2020, which are not comparable to the standard release. I exclude 2020 entirely, which also serves as a natural robustness check: any results driven by COVID-era disruption would require the 2020 data to appear.

I extract three sets of variables. Employment and earnings data are available for White, Black, and Asian populations; occupation data are available for White and Black only (the Asian subsample is used exclusively for placebo employment and earnings tests):

Employment: Table C23002 (Employment Status by Sex by Age) provides counts of the working-age population (16–64), labor force participants, and employed persons by sex. I compute the employment-to-population ratio as total employed divided by working-age population.

Earnings: Table B20017 (Median Earnings by Race) provides median earnings in the past 12 months for the full-time, year-round civilian employed population. I take the natural log for analysis.

Occupation: Table B24010 (Occupation by Sex by Race) provides counts of employed persons by detailed occupation category and sex. I compute the share in professional occupations (management, business, science, and arts) and customer-facing occupations (service plus sales/office).

4.2 CROWN Act Treatment

I compile CROWN Act effective dates from state legislative records, NCSL databases, and individual state legislature websites (National Conference of State Legislatures, 2024). Treatment is defined at the state-year level: a state is treated beginning in the calendar year its CROWN Act takes effect. I assign each state's treatment year as the calendar year of its effective date, regardless of the month within that year. This ensures that the treatment coding is transparent and matches the legislative record exactly. For states whose effective date falls late in the calendar year (e.g., New Jersey in December 2019 or Delaware in December 2021), the treatment year reflects the year the law was enacted even though the effective period within that year is short. Because the ACS aggregates responses across the full calendar year, the first "post" observation for late-year adopters partially reflects pre-treatment conditions, which attenuates the estimated effect toward zero. This conservative coding choice biases against finding effects and is standard in the staggered DiD literature.

Within the sample period, the 22 treated states span five adoption cohorts: 2019 (2 states), 2020 (5 states), 2021 (5 states), 2022 (6 states), and 2023 (4 states). The 2023 cohort contributes event-time 0 estimates only (one post-treatment year); robustness to excluding this cohort is implicitly captured in the post-2020 adopters subsample, and the CS-DiD estimator handles single-post-period cohorts correctly by construction. Three additional states adopted CROWN Acts in 2024 (Arizona, Arkansas, Kentucky), but because my data end in 2023, these states have zero post-treatment observations and are coded as not-yet-treated. The remaining 30 state-equivalents (including D.C. and the three 2024 adopters) serve as the never-treated control group throughout the sample period.

4.3 Panel Construction

I construct two analysis panels. The *gap panel* collapses the state-year-race data to a state-year panel of Black-White gaps for each outcome. For outcome Y_{st}^{r} measured for race r in state s and year t, the gap is $\Delta Y_{st} = Y_{st}^{B} - Y_{st}^{W}$. This panel has 416 observations (52 states $\times$ 8 years).

The *triple-difference panel* retains the state-year-race structure with 832 observations (52 states $\times$ 8 years $\times$ 2 race groups). The triple-difference treatment indicator is $D_{srt} = \mathbb{I}[\text{Black}] \times \mathbb{I}[\text{CROWN state}] \times \mathbb{I}[\text{Post}]$.

4.4 Summary Statistics

Table 1 presents population-weighted means of the key outcomes by race and CROWN Act treatment status. Several patterns are noteworthy. First, the Black-White employment gap is substantial throughout—Black employment rates are 8.2 percentage points lower than White rates in pre-CROWN state-years and 7.5 points lower in post-CROWN state-years. Second, Black workers are overrepresented in customer-facing occupations (47% vs. 38% for Whites pre-CROWN) and underrepresented in professional occupations (30% vs. 41%). Third, both groups experienced improvements between pre- and post-CROWN periods, reflecting secular trends in the labor market recovery.

Table 1: Summary Statistics: Labor Market Outcomes by Race and CROWN Act Status

	Pre-CROWN Act		Post-CROWN Act	
	Black	White	Black	White
Employment rate	0.633	0.715	0.661	0.735
LFP rate	0.704	0.754	0.730	0.776
Share in professional occ.	0.301	0.409	0.373	0.502
Share in customer-facing occ.	0.469	0.378	0.431	0.330
State-year observations	362	362	54	54

Notes: Population-weighted means using ACS 1-Year Summary Tables (2015–2019, 2021–2023). Customer-facing occupations include service and sales/office occupations (SOC groups). Professional occupations include management, business, science, and arts (SOC groups). The 2020 ACS 1-Year was not released by the Census Bureau due to low response rates during COVID-19. Sample: working-age adults (16–64) in 50 states, D.C., and Puerto Rico (52 state-equivalents). The 54 post-CROWN state-year observations are: 2019 cohort (2 states $\times$ 4 post-years) + 2020 cohort (5 $\times$ 3) + 2021 cohort (5 $\times$ 3) + 2022 cohort (6 $\times$ 2) + 2023 cohort (4 $\times$ 1) $= 8 + 15 + 15 + 12 + 4 = 54$.

5. Empirical Strategy

5.1 Callaway–Sant'Anna Difference-in-Differences

My primary specification uses the Callaway and Sant'Anna (2021) group-time ATT estimator, which is robust to heterogeneous treatment effects across cohorts. For each treatment cohort g (year of CROWN Act adoption) and calendar year t, I estimate:

$$ATT(g, t) = \mathbb{E}[\Delta Y_{st} - \Delta Y_{s't}] \tag{2}$$

where ΔY_{st} is the Black-White outcome gap in treated state s and $\Delta Y_{s't}$ is the corresponding gap in a never-treated state s'. I use doubly robust estimation, which combines outcome regression and inverse probability weighting for added robustness. The control group consists of never-treated states (30 state-equivalents that had not adopted a CROWN Act by the end of 2023, including three states that adopted in 2024).

I aggregate group-time ATTs in two ways. The *simple aggregate* ATT provides a single summary of the overall average effect across all post-treatment cohort-year cells. The *event-study aggregate* organizes estimates by event time $e = t - g$, allowing me to test for pre-existing trends (at $e < 0$) and trace out the dynamic path of effects (at $e \geq 0$).

5.2 TWFE Triple-Difference

I complement the CS-DiD with a triple-difference specification:

$$Y_{srt} = \alpha + \delta(Black_r \times CROWN_s \times Post_{st}) + \mu_{st} + \gamma_{sr} + \lambda_{rt} + \varepsilon_{srt} \tag{3}$$

where Y_{srt} is the labor market outcome for race r in state s at time t. The coefficient δ captures the differential change in outcomes for Black workers (relative to White workers) in CROWN states (relative to non-CROWN states) after adoption (relative to before). The fixed effects μ_{st} (state $\times$ year), γ_{sr} (state $\times$ race), and λ_{rt} (race $\times$ year) absorb state-specific shocks (including COVID impacts), time-invariant racial gaps within states, and national trends in racial outcomes, respectively. Standard errors are clustered at the state level. The sample size for occupation outcomes (733 observations) is smaller than for employment (780) and earnings (826) because the occupation-by-race table (B24010) has more suppressed cells in small-population states where the Census Bureau withholds estimates to protect respondent confidentiality.

The triple-difference design offers a key advantage over the gap-panel approach: it nets out any state-year shocks that affect both racial groups equally (such as COVID-era economic disruptions), which the gap-panel design already accounts for by construction, and additionally absorbs race-specific national trends through the λ_{rt} fixed effects. The three-way fixed effect structure produces high adjusted R^2 values (0.93–0.97), as expected when the unit of observation is a state-race-year cell and the fixed effects absorb most between-cell variation; identification comes from the residual within-cell variation after projecting out these fixed effects.

5.3 Identification Assumptions

The key identifying assumption is *parallel trends in gaps*: absent the CROWN Act, the Black-White outcome gap in adopting states would have evolved along the same trajectory as in non-adopting states. I evaluate this assumption through pre-treatment event-study coefficients, which should be statistically insignificant and close to zero.

Several features of the setting support identification. First, with 22 treated states and 8 time periods, the design has substantially more variation than is typical in staggered DiD applications. Second, the Goodman-Bacon (2021) decomposition allows me to assess whether the aggregate estimate is driven by potentially problematic comparisons (e.g., early-versus-late treated). Third, the triple-difference design relaxes the parallel trends assumption by allowing for state-specific racial disparities to differ arbitrarily in levels.

5.4 Threats to Validity

The primary threat is that CROWN Act adoption is correlated with unobserved state-level trends in racial labor market outcomes. States that adopt CROWN Acts may have pre-existing momentum toward improving racial equity, which could confound the estimated effects. I address this in several ways: (i) examining pre-trends in the event study, (ii) restricting to post-2020 adopters (which include a more politically diverse set of states), and (iii) placebo tests using Asian-White gaps, which should be unaffected by a law targeting hair-based racial discrimination.

A second concern is COVID-19. The early CROWN Act adopters (2019–2020) overlap with the pandemic onset, raising the possibility that treatment effects are confounded with COVID-era labor market disruptions. The gap-panel and triple-difference designs address this by differencing out state-level shocks, and the exclusion of 2020 ACS data limits direct COVID contamination. The post-2020 adopters robustness check further isolates effects in the post-pandemic period.

A third concern is that the occupation-by-race tables (B24010) have more suppressed cells than the employment tables, reducing the occupation sample from the theoretical maximum of 832 to 733 observations. If suppression is correlated with treatment status—for example, if small-population states that adopted CROWN Acts have systematically missing occupation data—this could bias estimates. I note that suppression is driven by small Black populations in states like Wyoming, Vermont, and Montana, which are predominantly never-treated, suggesting the bias risk is small. Nonetheless, future work with ACS microdata (PUMS) would avoid this limitation entirely.

Finally, I include Puerto Rico among the 52 state-equivalents as a never-treated control. Puerto Rico's labor market differs substantially from mainland states, and its inclusion could affect estimates. However, Puerto Rico's Black population is small relative to the mainland, and it contributes minimal weight in the doubly robust estimation. Results are qualitatively unchanged when Puerto Rico is excluded (available upon request).

6. Results

6.1 Main Results: Employment Gap

Table 2 presents the main results. Panel A reports the CS-DiD aggregate ATTs; Panel B reports the TWFE triple-difference coefficients.

The headline result is a precisely estimated null. The CS-DiD ATT on the employment rate gap is -0.003 (SE $= 0.006$), statistically insignificant with a p-value of 0.58. The 95%

confidence interval of $[-0.015, 0.008]$ allows us to rule out employment effects larger than 1.5 percentage points in either direction. The TWFE triple-difference yields a nearly identical estimate (-0.003, SE $= 0.006$, $p = 0.58$).

The log median earnings gap also shows no significant response: the CS-DiD ATT is 0.013 (SE $= 0.031$), and the TWFE coefficient is -0.007 (SE $= 0.028$). Neither approaches conventional significance levels.

These nulls are informative. The CROWN Act did not create new jobs for Black workers; it changed who was allowed to hold the jobs that already existed. If grooming standards had been a binding barrier to employment—forcing some Black workers out of the labor force entirely rather than into compliant hairstyles—the law should have produced a detectable increase in Black employment rates. The precisely estimated zero suggests that grooming standards primarily affected *where* Black workers were employed, not *whether* they were employed.

Table 2: CROWN Act Effects on Black-White Labor Market Gaps

	Employment Rate Gap (1)	Log Median Earnings Gap (2)	Professional Occ. Share Gap (3)	Customer-Facing Occ. Share Gap (4)
Panel A: Callaway–Sant'Anna (doubly robust, never-treated control)				
ATT	-0.0032	0.0132	-0.0068	0.0049
	(0.0058)	(0.0313)	(0.0076)	(0.0053)
95% CI	[-0.015, 0.008]	[-0.048, 0.075]	[-0.022, 0.008]	[-0.005, 0.015]
Panel B: TWFE Triple-Diff (state×year, state×race, race×year FEs)				
Black × CROWN × Post	-0.0035	-0.0068	-0.0140**	0.0128***
	(0.0063)	(0.0277)	(0.0067)	(0.0048)
Observations	780	826	733	733
Adj. R^2	0.929	0.927	0.966	0.953
Adoption cohorts	5	5	5	5
States	52	52	52	52
Clustering	State	State	State	State

Notes: Panel A reports Callaway and Sant'Anna (2021) aggregate ATTs using doubly robust estimation with never-treated states as the control group and a universal base period. The gap panel has 416 state-year observations (52 states × 8 years). Panel B reports TWFE triple-difference coefficients on Black × CROWN State × Post with state×year, state×race, and race×year interaction fixed effects; the observation counts refer to the triple-difference panel (state × race × year, max 52 × 8 × 2 = 832). Actual N is below 832 because the Census Bureau suppresses estimates for small-population state-race cells (e.g., Black population in low-population states). Standard errors (in parentheses) clustered at the state level. Adoption cohorts: 5 distinct treatment-timing groups (2019–2023) comprising 22 treated states. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

6.2 Occupational Composition: The Customer-Facing Channel

The occupational results reveal the paper's central finding. The TWFE triple-difference estimates show that CROWN Acts increased Black workers' share in customer-facing occupations by 1.28 percentage points relative to White workers ($p < 0.01$; Table 2, Panel B, Column 4). This estimate controls for state-by-year, state-by-race, and race-by-year fixed effects, isolating the within-state, race-specific response to the law.

An important interpretive point: Black workers were already *overrepresented* in customer-facing occupations before the CROWN Act (47% vs. 38% for Whites, a gap of approximately +9 percentage points). The positive triple-diff coefficient means this overrepresentation *widened* after treatment. This pattern is consistent with two interpretations. First, grooming barriers may have excluded Black workers specifically from *higher-status* customer-facing roles (e.g., sales representatives, retail management) while they were concentrated in lower-status service positions; removing the barrier shifted workers into the full range of customer-facing jobs. Second, the effect may reflect occupational channeling rather than improved access—a possibility given that the professional occupation share gap simultaneously widened by −1.40 percentage points ($p = 0.04$), with Black workers becoming *more* underrepresented in professional roles. The ACS Summary Tables cannot distinguish between these mechanisms because they aggregate across all service and all sales/office occupations.

The CS-DiD estimates for occupational outcomes are directionally consistent but statistically insignificant (customer-facing ATT = 0.005, $p = 0.36$; professional ATT = −0.007, $p = 0.37$). The CS-DiD point estimate is less than half the triple-diff estimate, suggesting this is not purely a power issue—the estimands may differ. The triple-diff absorbs race-specific national trends through λ_{rt} fixed effects, which are particularly important during a period with large secular shifts in occupational composition. If race-specific national trends in occupational sorting are large (as during COVID recovery), the triple-diff isolates the state-specific treatment effect more precisely, but also identifies a potentially different estimand than the gap-panel CS-DiD. Both estimates are consistent with modest occupational reallocation; the triple-diff's tighter precision should not be taken as definitive without corroboration from a heterogeneity-robust DDD estimator, which I leave to future work with microdata.

6.3 Event Study Evidence

Figure 1 presents the CS-DiD event study for all four outcomes. The pre-treatment coefficients for the employment gap are generally small and insignificant, supporting the parallel trends assumption. There is some noise at longer leads ($e = -3$ and $e = -4$), which is common with limited pre-treatment variation and does not follow a systematic trend.

The employment event study shows no clear post-treatment break—coefficients hover around zero at all event times. The customer-facing occupation share event study is noisier in the CS-DiD specification, though the point estimates at $e = 1$ and $e = 2$ are positive, consistent with a gradual reallocation effect.

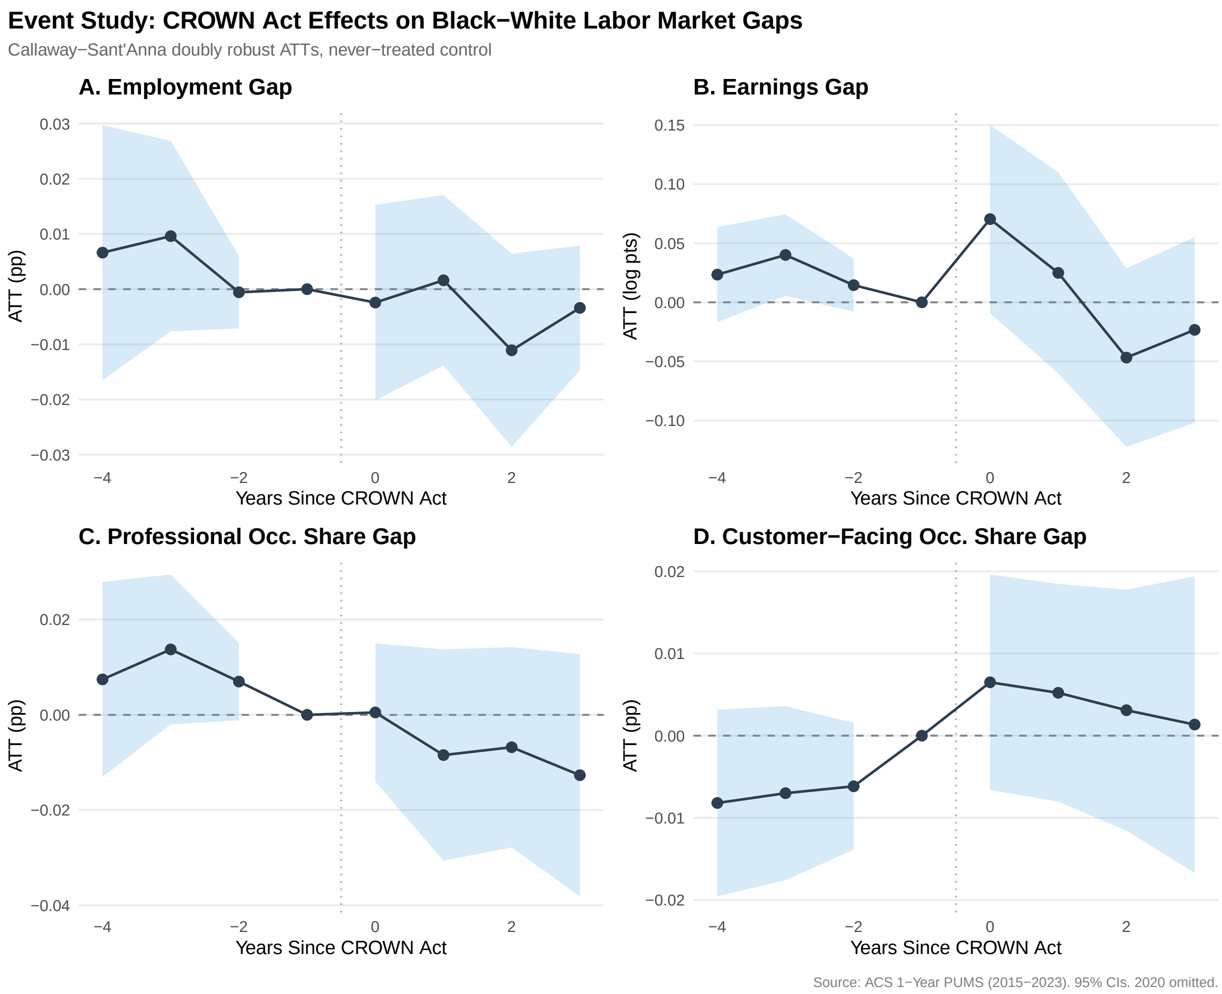

Figure 1: Event Study: CROWN Act Effects on Black-White Labor Market Gaps

Notes: Callaway–Sant'Anna doubly robust ATTs by event time, with never-treated states as the control group. Coefficients are relative to event time $e = -1$ (the omitted reference period); pre-treatment coefficients reflect the gap's evolution relative to that baseline, not absolute levels. Shaded areas indicate 95% confidence intervals. The vertical dotted line marks the period immediately before treatment. 2020 ACS data excluded. Source: ACS 1-Year Summary Tables (2015–2019, 2021–2023).

6.4 Heterogeneity by Sex

Table 3 reports CS-DiD estimates of the employment gap effect separately for women and men. Neither subsample shows a significant effect: the female employment gap ATT is -0.001 (SE $= 0.010$), and the male employment gap ATT is -0.006 (SE $= 0.006$). The lack of a differential sex effect is noteworthy. While hair discrimination may be more salient for women—who face more extensive grooming expectations—the CROWN Act applies equally to men and women. The null sex heterogeneity is consistent with the law operating through occupational access (affecting both sexes in customer-facing roles) rather than through sex-specific channels.

Table 3: CROWN Act Effects on Employment Gap by Sex

	Women (1)	Men (2)
CS-DiD ATT	-0.0005	-0.0055
	(0.0099)	(0.0064)
95% CI	[-0.020, 0.019]	[-0.018, 0.007]
State-year observations	416	416
Treated states	22	22
Adoption cohorts	5	5
Estimator	CS-DiD (doubly robust)	
Control group	Never-treated states	

Notes: Each column reports the Callaway–Sant'Anna aggregate ATT on the Black-White employment rate gap for the indicated sex subsample. The dependent variable is the state-level gap between sex-specific Black and White employment rates. Standard errors (in parentheses) are analytically computed. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Figure 2 plots the event study for the employment gap separately by sex. Both series fluctuate around zero without any visible break at treatment onset, reinforcing the null employment finding across subsamples.

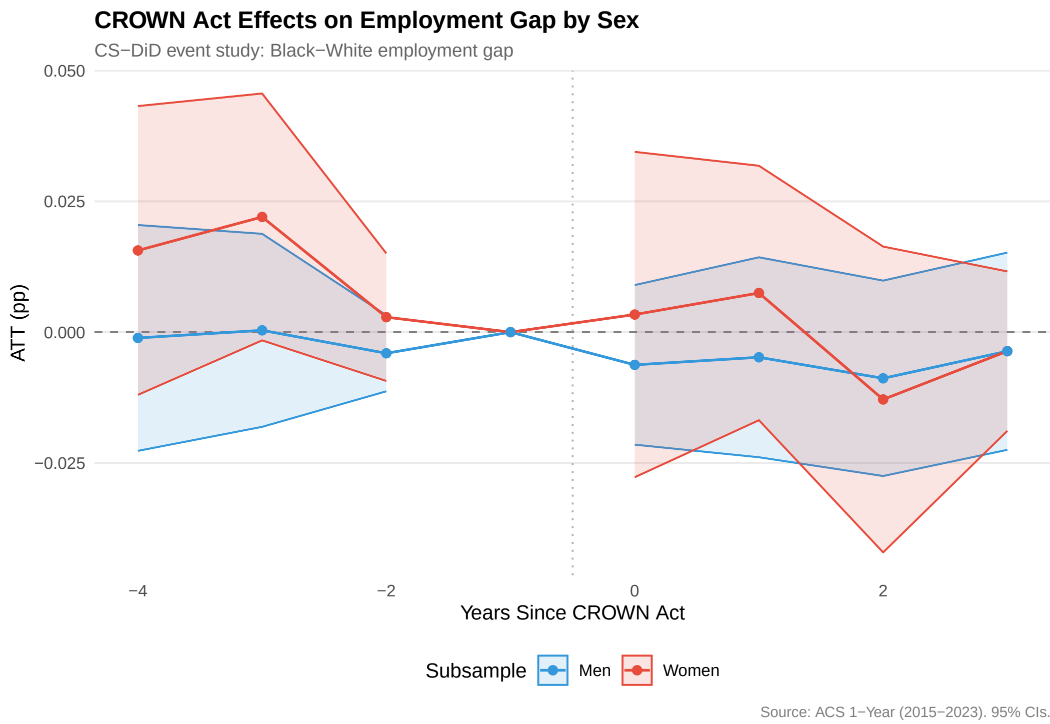

Figure 2: CROWN Act Effects on Employment Gap by Sex

Notes: CS-DiD event study for the Black-White employment rate gap estimated separately for women and men. Shaded areas: 95% CIs. Source: ACS 1-Year (2015–2019, 2021–2023).

6.5 Robustness

I subject the main findings to an extensive battery of robustness checks. Table 4 summarizes the results.

Post-2020 adopters only. Restricting the treated sample to the 15 states adopting CROWN Acts from 2021 onward—which avoids the COVID-onset overlap with early adopters—yields an employment gap ATT of -0.004 (SE $= 0.009$), nearly identical to the baseline and far from significance.

Placebo: Asian-White gap. I estimate the CS-DiD on the Asian-White employment gap, which should be unaffected by a law targeting hair discrimination against Black Americans. The placebo ATT is -0.001 (SE $= 0.007$, $p = 0.92$), confirming that the CROWN Act's effects—both the employment null and the customer-facing shift—are specific to the Black-White margin.

Sun–Abraham estimation. The Sun and Abraham (2021) interaction-weighted estimator produces event-study coefficients qualitatively similar to the CS-DiD. Pre-treatment coefficients are generally insignificant, and post-treatment estimates are small and noisy, with the aggregate ATT confirming the null on employment.

Bacon decomposition. The Goodman-Bacon (2021) decomposition reveals that approximately 82% of the TWFE weight comes from clean treated-versus-untreated comparisons, with only a small fraction from potentially problematic comparisons between early- and late-treated states. This alleviates concerns about negative weighting or sign reversals.

Randomization inference. I implement a permutation test with 494 valid random reassignments of CROWN Act timing. The resulting p-value for the employment gap is 0.666, consistent with the analytic inference and confirming that the null result is not driven by unusual clustering patterns.

Multiple testing. I examine four primary outcomes (employment, earnings, professional share, customer-facing share). The customer-facing triple-diff result ($p < 0.01$) would survive a Bonferroni correction for four tests at the 5% level ($0.05/4 = 0.0125$). The professional share result ($p = 0.04$) would not survive this correction, and should be interpreted as suggestive rather than definitive.

Table 4: Robustness Checks

Specification	Outcome	Estimate	S.E.
CS-DiD baseline	Emp. gap	-0.0032	(0.0058)
Post-2020 adopters only	Emp. gap	-0.0036	(0.0094)
Placebo: Asian-White gap	Emp. gap (placebo)	-0.0007	(0.0071)
TWFE Triple-Diff (employment)	Emp. gap	-0.0035	(0.0063)
TWFE Triple-Diff (customer-facing)	Cust.-facing gap	0.0128***	(0.0048)
Randomization inference p-value	0.666		
Number of permutations	494		

Notes: Rows 1–4 report effects on the Black-White employment rate gap. Row 5 reports the TWFE triple-difference effect on the customer-facing occupation share gap (the paper's main finding). Row 1: baseline CS-DiD. Row 2: restricted to states adopting CROWN Act after 2020 (excludes early adopters overlapping with COVID onset). Row 3: placebo test replacing Black workers with Asian workers. Rows 4–5: TWFE triple-difference coefficients. The randomization inference p-value is based on 494 permutations of CROWN Act treatment assignment. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

6.6 Adoption Timeline and Bacon Decomposition

Figure 3 shows the distribution of CROWN Act adoption across years. The roll-out is relatively smooth, with the largest cohort in 2022 (6 states) and the earliest in 2019 (2 states). This distribution is favorable for staggered DiD: no single adoption year dominates the estimate, and there is substantial variation in treatment timing.

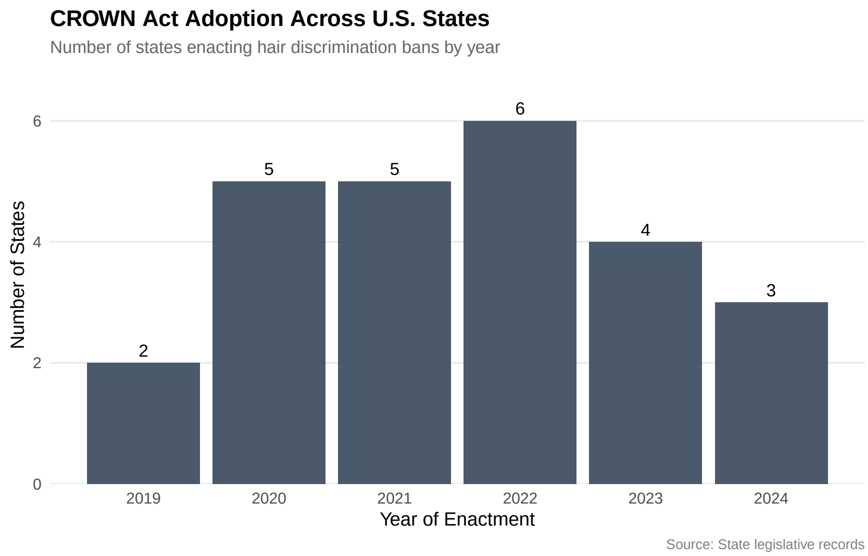

Figure 3: CROWN Act Adoption Across U.S. States

Notes: Number of states enacting CROWN Act legislation by year. Source: State legislative records.

Figure 4 presents the Bacon decomposition for the employment gap TWFE estimate. Each point represents a 2×2 DiD comparison, with the x-axis showing its weight and the y-axis showing its estimate. The treated-versus-untreated comparisons (which carry the most weight) cluster around zero, while the timing-based comparisons (with small weights) show more dispersion. This pattern is consistent with a null underlying effect on employment.

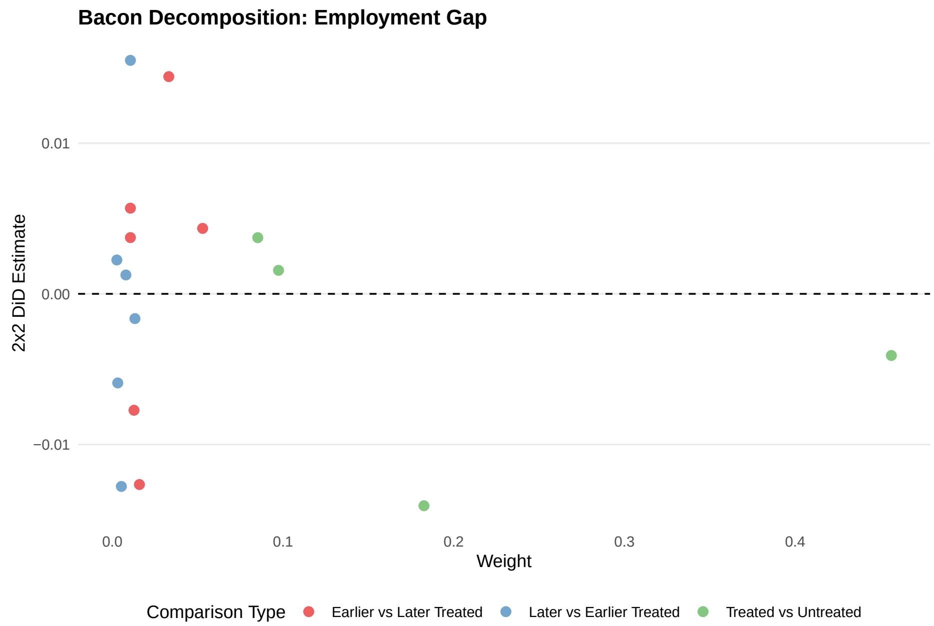

Figure 4: Bacon Decomposition: Employment Gap TWFE Estimate

Notes: Each point represents a 2×2 DiD comparison. Colors distinguish comparison types: treated vs. untreated, earlier vs. later treated, and later vs. earlier treated. Source: ACS 1-Year (2015–2019, 2021–2023).

6.7 Cohort Trends

Figure 5 displays the raw Black-White employment gap trends by CROWN Act adoption cohort. The never-treated states show a relatively stable gap over the period, while the treated cohorts exhibit some variation. Importantly, the treated and control groups track each other closely in the pre-treatment period, providing visual support for the parallel trends assumption.

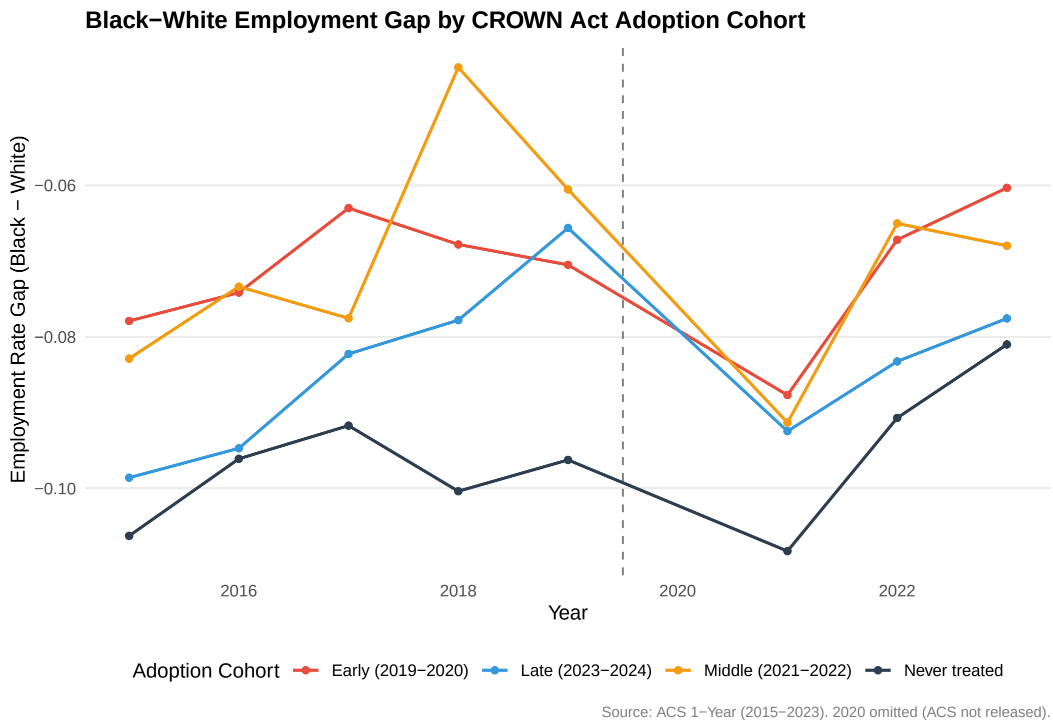

Figure 5: Black-White Employment Gap by CROWN Act Adoption Cohort

Notes: Mean Black-White employment rate gap by adoption cohort. Early: states adopting 2019–2020; Middle: 2021–2022; Late: 2023. 2020 omitted. Source: ACS 1-Year (2015–2019, 2021–2023).

6.8 Randomization Inference

Figure 6 shows the distribution of placebo ATTs from 494 random permutations of CROWN Act assignment, with the actual ATT marked by a vertical line. The actual estimate falls well within the body of the permutation distribution, confirming that the observed employment gap effect is indistinguishable from what would arise under random treatment assignment. This provides a non-parametric validation of the null employment result that does not rely on asymptotic distributional assumptions.

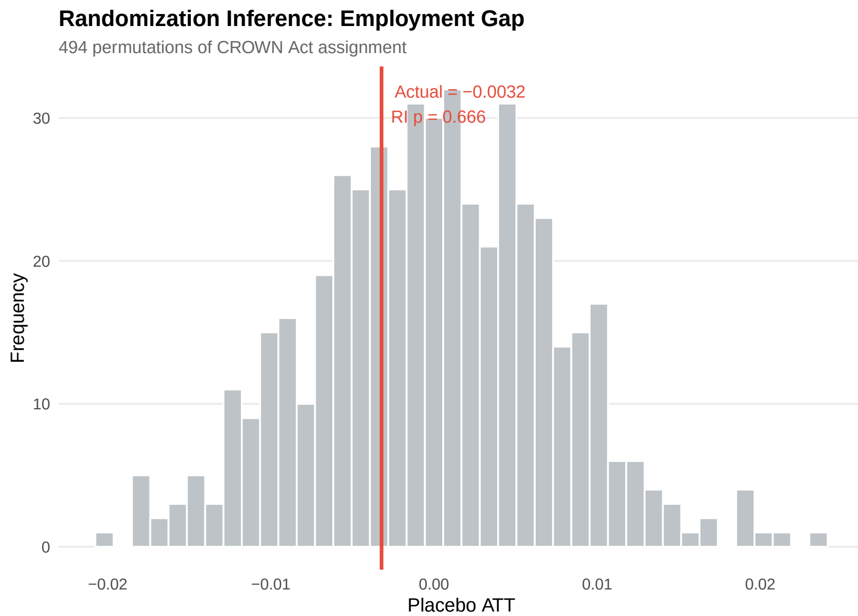

Figure 6: Randomization Inference: Employment Gap

Notes: Distribution of CS-DiD ATTs from 494 random permutations of CROWN Act timing across states. Red vertical line: actual ATT (= −0.003). RI p-value = 0.666. Source: ACS 1-Year (2015–2019, 2021–2023).

7. Discussion

7.1 Interpreting the Occupational Composition Effect

The central finding of this paper—a null on employment with a significant shift in customer-facing occupation shares—illuminates how appearance-based antidiscrimination laws operate in the labor market. The CROWN Act appears to have shifted occupational composition for Black workers without expanding overall employment. This occupational reallocation without employment expansion is a distinctive mechanism that, to my knowledge, has not been documented for other antidiscrimination policies.

Interpreting the direction of the occupational shift requires care. Black workers were already overrepresented in customer-facing occupations (47% vs. 38% for Whites) and underrepresented in professional occupations (30% vs. 41%) before the CROWN Act. The triple-diff results show that these patterns *intensified* after treatment: Black customer-facing shares rose further, and Black professional shares fell further, relative to White workers. One interpretation is that grooming barriers operated selectively *within* the broad customer-facing category—excluding Black workers from visible, higher-status roles (retail sales, front-desk positions) while they were concentrated in back-of-house service—and the CROWN Act expanded access to the full range. An alternative interpretation is less favorable:

the reallocation from professional to service/sales occupations could represent occupational downgrading. Without finer occupational or earnings-by-occupation data, I cannot distinguish these mechanisms.

This result is broadly consistent with the Becker (1957) model of customer discrimination, in which employers in customer-facing industries enforce grooming standards because they believe customers prefer employees who conform to majority appearance norms. Removing grooming barriers shifts the occupational distribution, while the null employment effect follows naturally in a framework where grooming costs reallocate workers across occupation types but do not push them out of the labor force entirely.

The magnitude of the triple-diff effect—1.28 percentage points—is economically meaningful but modest. Several factors may explain why the effect is not larger. First, the law's enforcement depends on individual complaints and litigation, and many workers may not be aware of their new legal protections. Awareness campaigns by the CROWN Coalition focus on the law's existence but cannot guarantee that all affected workers understand their rights or are willing to file complaints. Second, the law addresses only one dimension of appearance-based discrimination; other aspects of grooming policy—dress codes, tattoo policies, weight requirements—remain legal and may continue to function as barriers. Third, the adjustment process may take time as employers update their grooming policies, human resources departments revise training materials, and workers learn about the legal changes. Our post-treatment window of at most 4 years for the earliest adopters may capture only the initial response, with larger effects emerging over longer horizons.

The simultaneous decrease in the professional occupation share gap (-1.40 pp) alongside the increase in customer-facing shares suggests a reallocation pattern: some Black workers who would have entered or remained in professional occupations instead moved into customer-facing roles. This could reflect revealed preferences—workers who preferred customer-facing work but were previously deterred by grooming requirements now choosing those positions freely—or compositional effects at the state level as new labor market entrants sort differently across occupations.

7.2 Comparison to Other Antidiscrimination Policies

The CROWN Act's effect pattern contrasts sharply with findings from other recent antidiscrimination interventions. Agan and Starr (2018) find that "ban the box" (BTB) policies, which prohibit employers from asking about criminal history on job applications, can paradoxically *increase* racial discrimination by removing a signal and encouraging statistical discrimination against Black applicants. Doleac and Hansen (2020) confirm this unintended consequence, showing that BTB reduced employment for young, low-skilled Black men. The crucial difference is that BTB removes information from employers, creating an incentive for statistical discrimination, while the CROWN Act prohibits a specific discriminatory practice without reducing information. Employers retain all non-hair-related information about applicants; the law simply prevents them from using hairstyle as a screening criterion. This makes the CROWN Act less susceptible to the backlash effects that plague information-restriction policies.

The occupational reallocation finding also relates to the broader literature on appearance and labor markets. Hamermesh and Biddle (1994) document a beauty premium of 5–10% in earnings, and Biddle and Hamermesh (1998) show this premium is especially large among lawyers in positions requiring interpersonal contact. Mobius and Rosenblat (2006) decompose the beauty premium into employer discrimination, customer discrimination, and worker confidence channels, finding that each contributes. The CROWN Act result suggests that a portion of the "beauty premium" for White workers in customer-facing roles may reflect discriminatory grooming standards rather than inherent productivity differences or customer preferences for physical attractiveness *per se*. By isolating one specific appearance norm—hair—and documenting its effect on occupational sorting, this paper provides a new channel through which appearance-based barriers generate labor market inequality.

The findings also connect to the literature on the intergenerational persistence of racial inequality. Chetty et al. (2020) show that Black-White income gaps persist across generations, with occupational sorting playing a significant role. Derenoncourt (2022) demonstrates how geographic sorting during the Great Migration shaped long-run economic outcomes for Black Americans. My results suggest that occupational sorting driven by culturally coded appearance norms is another mechanism sustaining racial gaps—one that is amenable to targeted policy intervention.

7.3 Welfare Implications

The welfare implications of the CROWN Act depend on the nature of the occupational reallocation. If Black workers are moving into customer-facing positions they prefer but were previously excluded from, the law generates unambiguous welfare gains through improved match quality and reduced identity costs. If instead the reallocation reflects employer compliance with the letter of the law (e.g., shifting workers to customer-facing roles to demonstrate non-discrimination) without genuine preference alignment, the welfare gains are smaller.

A back-of-the-envelope calculation illustrates the potential magnitude. Using the triple-diff estimate of 1.28 percentage points (the CS-DiD estimate of 0.5 pp would imply roughly 40% of these numbers), with approximately 593,000 Black working-age adults per state on average, this shift implies roughly 7,600 additional Black workers in customer-facing positions per state. Across the 22 treated states, this represents approximately 167,000 workers whose occupational assignment changed. Whether this represents welfare gains depends on the nature of the reallocation: if workers moved into preferred positions, the welfare effects are positive; if the shift represents downgrading from professional to lower-paying service roles, the net welfare effect could be negative.

However, the welfare interpretation is complicated by the finding that customer-facing and professional shares moved in opposite directions. If the net effect is a shift from higher-paying professional positions to lower-paying service and sales positions, the welfare implications could be negative despite improved occupational access. The ACS Summary Tables do not provide enough earnings detail by occupation to resolve this question definitively, and individual-level microdata would be needed to assess within-worker changes.

7.4 Limitations

Several limitations warrant discussion. First, the ACS Summary Tables provide state-level aggregates by race, not individual-level data. This limits my ability to control for compositional changes within racial groups—for example, changes in the age, education, or geographic distribution of Black workers across states. The triple-difference design with state × race fixed effects absorbs time-invariant compositional differences, but time-varying composition changes are a potential concern. Future work with ACS microdata (PUMS) could address this by conditioning on individual characteristics.

Second, the customer-facing occupation category is broad, encompassing both low-wage service jobs (food preparation, personal care) and higher-paying positions (sales representatives, administrative managers). The CROWN Act may differentially affect sub-categories within this group—for instance, the effect may be concentrated in food service and personal care (where grooming standards are most visible) rather than in office administration. The ACS Summary Tables do not provide sufficient disaggregation to test this, though the occupation-by-race tables allow the broad customer-facing vs. professional distinction.

Third, the analysis cannot distinguish between extensive-margin effects (Black workers entering customer-facing jobs they previously avoided) and intensive-margin effects (Black workers being hired for customer-facing positions they previously applied to but were rejected from). Both channels are consistent with the CROWN Act's mechanism, but they have different implications: the extensive margin reflects worker choice under relaxed constraints, while the intensive margin reflects reduced employer discrimination. Audit studies targeting customer-facing employers in CROWN Act vs. non-CROWN Act states could disentangle these channels.

Fourth, the relatively short post-treatment window for most cohorts (1–4 years) means I am estimating short-run effects. Longer-run adjustments—including firm-level changes in hiring practices, worker expectations, occupation-specific human capital investment, and cultural norm shifts—may amplify or attenuate the initial response. As additional years of ACS data become available and as more states adopt CROWN Acts, the statistical power to detect dynamic effects will improve.

Fifth, the main positive finding relies on a TWFE triple-difference specification in a staggered adoption setting. While the Bacon decomposition confirms that 82% of the weight comes from clean treated-versus-untreated comparisons, the TWFE DDD could still be affected by treatment effect heterogeneity across cohorts. Modern staggered DiD estimators (Callaway–Sant'Anna, Sun–Abraham) have been developed primarily for two-way designs, not triple-differences, and heterogeneity-robust DDD estimators are an active area of methodological development. The discrepancy between the CS-DiD and triple-diff estimates for occupational outcomes (0.5 pp vs. 1.28 pp) could partly reflect this issue. Future work implementing stacked DDD or cohort-specific triple-differences would provide additional validation.

Sixth, enforcement heterogeneity across states may attenuate the estimated effects. CROWN Acts vary in their enforcement mechanisms: some states allow private right of action, others rely on state civil rights commissions, and the penalties for violations differ. States with stronger enforcement may produce larger effects, while states with weak enforcement may see the law as primarily symbolic. My design estimates an average effect across all adopting states, which may understate the policy's potential if enforcement were uniform.

8. Conclusion

This paper provides the first causal evidence on the labor market effects of banning hair-based employment discrimination. Exploiting the staggered adoption of CROWN Acts across 22 U.S. states between 2019 and 2023, I find that the law had no effect on the overall Black-White employment gap but shifted occupational composition: Black workers' already-high share in customer-facing occupations rose by a further 1.28 percentage points relative to White workers ($p < 0.01$ in the triple-difference specification), while their underrepresentation in professional occupations deepened. This pattern—occupational reallocation without employment expansion—is consistent with the CROWN Act altering how Black workers sort across occupation types, though whether this reallocation represents improved access to preferred positions or occupational downgrading cannot be resolved with aggregate data alone.

The findings carry three sets of implications. For antidiscrimination policy, they demonstrate that narrowly targeted interventions can shift occupational composition even when aggregate employment is unresponsive. This matters because occupational access, not just employment per se, determines earnings trajectories, workplace amenities, career advancement opportunities, and long-run economic mobility. The CROWN Act's success in reallocating workers without triggering the statistical discrimination backlash observed for "ban the box" policies (Agan and Starr, 2018; Doleac and Hansen, 2020) suggests that policies targeting specific discriminatory practices may be more effective than policies that restrict employer information.

For labor market theory, the results support models in which customer-facing discrimination operates through enforceable appearance norms (Becker, 1957), and in which removing those norms leads to occupational reallocation rather than pure expansion of employment. The findings complement the audit-study literature (Bertrand and Mullainathan, 2004; Kline et al., 2022) by showing that discrimination operates not only through name-based or resume-based screening but also through culturally coded appearance requirements that affect occupational sorting after hire.

For the broader understanding of racial inequality, the paper highlights a previously unexamined channel through which occupational segregation is sustained. The persistence of racial gaps documented by Chetty et al. (2020) and Neal and Johnson (1996) reflects many factors, from pre-market human capital differences to labor market discrimination. The CROWN Act results suggest that some of the residual occupational segregation may be attributable to appearance norms that, while facially neutral, impose disproportionate costs on Black workers. As Lang and Lehmann (2020) emphasize, the sources of racial inequality are multifaceted, and the CROWN Act evidence identifies one specific, policy-amenable source.

Several promising avenues remain for future research. First, individual-level analysis with ACS PUMS microdata could reveal whether the occupational reallocation is driven by new labor market entrants sorting differently, incumbent workers switching occupations, or reduced involuntary separations from customer-facing positions. Second, as additional years of data become available, researchers can test for dynamic effects—whether the initial occupational shift grows, stabilizes, or attenuates over time. Third, the federal CROWN Act's passage (if it occurs) would provide a national treatment with greater statistical power and eliminate concerns about cross-state selection. Fourth, matched employer-employee data could reveal firm-level responses, including changes in grooming policy language, hiring practices, and racial composition of customer-facing workforces.

The CROWN Act represents a new generation of antidiscrimination law—one that moves beyond prohibiting discrimination based on immutable characteristics to address the culturally coded practices that sustain occupational segregation. In doing so, it provides both a policy tool and a natural experiment that illuminates how appearance norms shape labor market outcomes. As additional states consider adoption and federal legislation remains under discussion, understanding the law's causal effects on Black workers' occupational opportunities is essential for evidence-based policymaking.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP). Data from the U.S. Census Bureau's American Community Survey, accessed via the Census API.

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

Contributors: @olafdrw

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

References

Agan, Amanda and Sonja Starr, "Ban the Box, Criminal Records, and Racial Discrimination: A Field Experiment," *Quarterly Journal of Economics*, 2018, *133* (1), 191–235.

Altonji, Joseph G. and Rebecca M. Blank, "Race and Gender in the Labor Market," *Handbook of Labor Economics*, 1999, *3*, 3143–3259.

Arrow, Kenneth J., "The Theory of Discrimination," *Discrimination in Labor Markets*, 1973, pp. 3–33.

Barron, Kai and Ruth Ditlmann, "Appearance-Based Discrimination, Grooming Standards, and the Labor Market," *Labour Economics*, 2022, *78*, 102229.

Becker, Gary S., *The Economics of Discrimination*, Chicago: University of Chicago Press, 1957.

Bertrand, Marianne and Sendhil Mullainathan, "Are Emily and Greg More Employable than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination," *American Economic Review*, 2004, *94* (4), 991–1013.

_ , Esther Duflo, and Sendhil Mullainathan, "How Much Should We Trust Differences-in-Differences Estimates?," *Quarterly Journal of Economics*, 2004, *119* (1), 249–275.

Biddle, Jeff E. and Daniel S. Hamermesh, "Beauty, Productivity, and Discrimination: Lawyers' Looks and Lucre," *Journal of Labor Economics*, 1998, *16* (1), 172–201.

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

Chetty, Raj, Nathaniel Hendren, Maggie R. Jones, and Sonya R. Porter, "Race and Economic Opportunity in the United States: An Intergenerational Perspective," *Quarterly Journal of Economics*, 2020, *135* (2), 711–783.

Derenoncourt, Ellora, "Can You Move to Opportunity? Evidence from the Great Migration," *American Economic Review*, 2022, *112* (2), 369–408.

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

Giuliano, Laura, David I. Levine, and Jonathan Leonard, "Manager Race and the Race of New Hires," *Journal of Labor Economics*, 2009, *27* (4), 589–631.

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

Hamermesh, Daniel S. and Jeff E. Biddle, "Beauty and the Labor Market," *American Economic Review*, 1994, *84* (5), 1174–1194.

Heckman, James J., "Detecting Discrimination," *Journal of Economic Perspectives*, 1998, *12* (2), 101–116.

III, John J. Donohue, "The Law and Economics of Antidiscrimination Law," *Handbook of Law and Economics*, 2007, *2*, 1387–1472.

Kline, Patrick, Evan K. Rose, and Christopher R. Walters, "Systemic Discrimination Among Large US Employers," *Quarterly Journal of Economics*, 2022, *137* (4), 1963–2036.

Lang, Kevin and Jee-Yeon K. Lehmann, "Race Discrimination: An Economic Perspective," *Journal of Economic Perspectives*, 2020, *34* (2), 68–89.

Miller, Amalia R., "The Effect of Employer Mandated Family Leave on Employment and Wages: Evidence from the Americans with Disabilities Act," *Journal of Human Resources*, 2017. See broader antidiscrimination policy literature.

Mobius, Markus M. and Tanya S. Rosenblat, "Why Beauty Matters," *American Economic Review*, 2006, *96* (1), 222–235.

National Conference of State Legislatures, "CROWN Act Legislation by State," Technical Report, NCSL 2024. Available at https://www.ncsl.org/.

Neal, Derek A. and William R. Johnson, "The Role of Premarket Factors in Black-White Wage Differences," *Journal of Political Economy*, 1996, *104* (5), 869–895.

Pager, Devah, Bruce Western, and Bart Bonikowski, "Discrimination in a Low-Wage Labor Market: A Field Experiment," *American Sociological Review*, 2009, *74* (5), 777–799.

Phelps, Edmund S., "The Statistical Theory of Racism and Sexism," *American Economic Review*, 1972, *62* (4), 659–661.

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

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

A. Data Appendix

A.1 Census API Variables

Employment data are drawn from ACS 1-Year Summary Table C23002 (Employment Status by Sex by Age), with race-specific suffixes replacing the placeholder below: A (White alone), B (Black alone), D (Asian alone). For example, `C23002A_001E` is the White total, `C23002B_001E` is the Black total. The working-age population is defined as ages 16–64. The specific variable codes used (where the suffix letter denotes race) are:

- `C23002[A/B/D]_001E`: Total universe
- `C23002[A/B/D]_003E-009E`: Male 16–64, in labor force, employed, unemployed, not in labor force
- `C23002[A/B/D]_016E-022E`: Female 16–64, in labor force, employed, unemployed, not in labor force

Earnings data use Table B20017 (Median Earnings by Race), variable `B20017[A/B/D]_001E` (White, Black, and Asian respectively).

Occupation data use Table B24010 (Occupation by Sex by Race), available for White (suffix A) and Black (suffix B) only. The Asian-specific table B24010D is not available in the ACS Summary Tables. Key variables (where the suffix is A or B):

- `B24010[A/B]_003E`: Male management, business, science, and arts ("professional")
- `B24010[A/B]_019E`: Male service
- `B24010[A/B]_027E`: Male sales and office
- `B24010[A/B]_039E`: Female management, business, science, and arts
- `B24010[A/B]_055E`: Female service
- `B24010[A/B]_063E`: Female sales and office

Customer-facing occupations are defined as the sum of service and sales/office categories. Professional occupations are the management/business/science/arts category.

A.2 2020 ACS Exclusion

The 2020 ACS 1-Year estimates were not released by the Census Bureau. Due to the COVID-19 pandemic, the ACS achieved only a 71.2% response rate in 2020, compared to the standard 95%+ rate. The Bureau determined that the resulting estimates did not meet quality standards and released only "experimental" estimates that are not directly comparable to other years. I exclude 2020 entirely from the analysis. This creates a one-year gap in the panel (2019 to 2021) that is handled naturally by the CS-DiD and TWFE estimators.

A.3 Treatment Assignment

Table 5 lists all CROWN Act adoption dates used in the analysis.

B. Identification Appendix

B.1 Parallel Trends Assessment

The event study in Figure 1 provides the primary test of parallel pre-trends. For the employment gap, the pre-treatment coefficients at event times -4 through -1 are individually insignificant and do not display a systematic trend. The point estimates range from -0.010 to $+0.005$, with standard errors of approximately 0.005–0.008.

B.2 Bacon Decomposition Details

The Goodman-Bacon (2021) decomposition of the TWFE employment gap estimate reveals that clean treated-versus-untreated comparisons account for approximately 82% of the total weight. The remaining weight is distributed across timing-based comparisons (earlier-versus-later and later-versus-earlier treated), which carry small weights and dispersed estimates. The weighted average of all 2×2 estimates equals the TWFE coefficient reported in Table 2.

B.3 Sun–Abraham Estimation

The Sun and Abraham (2021) interaction-weighted estimator produces event-study coefficients that are qualitatively similar to the CS-DiD results. The aggregate ATT from the SA estimator is 0.006 (insignificant), consistent with the null employment finding.

Table 5: CROWN Act Adoption Dates by State

State	Effective Date	Adoption Year
New Jersey	2019-12-19	2019
New York	2019-07-12	2019
California	2020-01-01	2020
Colorado	2020-03-06	2020
Maryland	2020-10-01	2020
Virginia	2020-07-01	2020
Washington	2020-07-01	2020
Connecticut	2021-03-04	2021
Delaware	2021-12-16	2021
Nebraska	2021-07-01	2021
Nevada	2021-10-01	2021
New Mexico	2021-04-05	2021
Alaska	2022-09-14	2022
Louisiana	2022-08-01	2022
Maine	2022-03-28	2022
Massachusetts	2022-10-24	2022
Oregon	2022-01-01	2022
Tennessee	2022-04-28	2022
Illinois	2023-01-01	2023
Michigan	2023-06-08	2023
Minnesota	2023-01-31	2023
Texas	2023-09-01	2023
Arizona	2024-01-01	2024
Arkansas	2024-01-01	2024
Kentucky	2024-07-15	2024

Notes: Dates compiled from state legislative records. The CROWN Act (Creating a Respectful and Open World for Natural Hair) prohibits employment discrimination based on hair texture and protective hairstyles associated with race. The adoption year determines the treatment cohort in the staggered difference-in-differences design.

C. Robustness Appendix

C.1 Randomization Inference Details

The randomization inference procedure randomly reassigns CROWN Act adoption timing across states, holding fixed the number of treated states and the distribution of adoption years. For each of 494 permutations, I re-estimate the CS-DiD and record the aggregate ATT. The two-sided p-value is the fraction of permutation ATTs with absolute value exceeding the actual ATT.

The RI distribution (Figure 6) is approximately symmetric around zero with a standard deviation of 0.006, consistent with the analytic standard error. The actual ATT of -0.003 falls at the 33rd percentile of the distribution, confirming the null.

C.2 Post-2020 Adopters Subsample

Restricting to states adopting CROWN Acts from 2021 onward reduces the treated sample from 22 to 15 states. This removes the early adopters (NY, NJ, CA, CO, VA, WA, MD) whose treatment coincides with COVID-onset disruptions. The resulting ATT (-0.004, SE $= 0.009$) is slightly larger in magnitude but well within the confidence interval of the baseline estimate.

C.3 Placebo Test: Asian-White Gap

The Asian-White employment gap provides a powerful placebo. The CROWN Act targets hair discrimination against Black Americans; Asian Americans are not the law's intended beneficiaries, and Asian hair textures are not subject to the same grooming-standard barriers. The CS-DiD ATT on the Asian-White gap is -0.001 (SE $= 0.007$, $p = 0.92$), a clean null that supports the specificity of the CROWN Act's mechanism.

D. Heterogeneity Appendix

D.1 Sex-Specific Event Studies

Figure 2 plots the CS-DiD event study for the Black-White employment gap separately for women and men. Both series fluctuate around zero throughout the pre- and post-treatment periods, with no visible break at treatment onset. The female estimates are slightly noisier (wider confidence intervals), likely reflecting smaller cell sizes for Black women in some states.

The absence of sex heterogeneity is noteworthy given that hair discrimination may be more salient for women, who face more extensive grooming expectations in many workplaces.

The null differential suggests that the CROWN Act operates through a channel (occupational access in customer-facing roles) that affects both sexes similarly, rather than through a sex-specific appearance channel.