Do Skills-Based Hiring Laws Actually Change Who Works in Government?

APEP Autonomous Research* @SocialCatalystLab

February 26, 2026

Abstract

Between 2022 and 2025, twenty-two U.S. states removed bachelor's degree requirements for state government positions—the largest simultaneous de-credentialization of public employment in American history. Using a state-by-year panel constructed from 14.2 million individual observations in the American Community Survey (2013–2023) and staggered difference-in-differences methods, I study the thirteen states with post-treatment data and find no evidence that these laws increased the share of non-college-educated state government workers. Point estimates are negative across all specifications, suggesting that adopting states were on pre-existing trajectories of increasing credentialization. A triple-difference design comparing state government to private sector workers within the same states yields a small, statistically insignificant estimate. These early findings suggest that removing formal degree requirements, absent deeper structural reforms to hiring practices, has not yet altered the educational composition of the public workforce in a detectable way.

JEL Codes: J45, J24, I26, H70

Keywords: skills-based hiring, degree requirements, public sector employment, credentialism, difference-in-differences

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

1. Introduction

Roughly 62 percent of American adults lack a four-year college degree. For decades, this "paper ceiling" has barred millions of capable workers from stable careers in the public sector. In March 2022, Maryland became the first state to tear it down, formally eliminating bachelor's degree requirements for most state government positions. Within two years, twenty-one more states followed—a policy cascade unprecedented in the history of American public employment (**?**). Governors from both parties championed these reforms as a way to tap a broader talent pool and make government careers accessible to workers screened out by credentials they did not need (**?**).

But do these laws actually change who gets hired? This is not a question about job postings or official requirements—it is a question about whether human resource decisions, shaped by decades of credential-screening culture, respond to a change in formal policy. The answer has implications far beyond government hiring. As artificial intelligence accelerates the obsolescence of traditional knowledge credentials, understanding whether top-down decredentialization can redirect hiring away from proxy signals toward actual competence is a first-order question for labor market policy (**??**).

This paper provides the first causal evidence on whether skills-based hiring laws changed the *actual* educational composition of state government workforces, as opposed to the mere language of job advertisements. I exploit the staggered adoption of these laws across states between 2022 and 2025, using individual-level microdata from the American Community Survey (ACS) Public Use Microdata Sample covering approximately 14.2 million working-age adults observed across ten survey years (2013–2023). Of the twenty-two states that adopted skills-based hiring policies, thirteen have at least one year of post-treatment ACS data; the remaining nine adopted too recently (effective first-treatment year 2024 or later) to be observed post-treatment and serve as not-yet-treated units. The primary outcome is the share of state government employees aged 25–64 who lack a bachelor's degree, which I construct as a state-by-year panel and analyze using modern difference-in-differences methods that accommodate staggered treatment timing (**??**).

The results are striking in their consistency: across every specification, I find no evidence that skills-based hiring laws increased the share of non-college-educated workers in state government. The two-way fixed effects estimate is -0.016 (SE $= 0.006$, $p = 0.017$), indicating that if anything, treated states experienced a *decline* in non-BA workers relative to never-treated states. The Callaway-Sant'Anna overall ATT of -0.037 (SE $= 0.014$) is larger but comes with a failed pre-test for parallel trends ($p < 0.01$), suggesting that the negative estimates partly reflect differential pre-treatment trajectories. A triple-difference specification comparing state government to private sector workers within the same states—which accounts for state-specific trends—produces a small, statistically insignificant estimate of -0.010 (SE $= 0.007$, $p = 0.15$).

Three pieces of evidence help unpack the null result. First, the pre-treatment event study reveals that adopting states were already on a distinctive trajectory: the share of non-BA state government workers was declining faster in these states before any policy change. This suggests that *selection into treatment* is a first-order feature of the policy landscape—states adopted skills-based hiring laws precisely *because* their public sectors were becoming more credentialized, not randomly. Second, placebo tests on federal government workers (who are not subject to state laws) show a precisely estimated null (-0.003, $p = 0.65$), confirming that the policy is targeted at the right level of government. Third, heterogeneity analysis reveals that only strong policy mandates (legislative codification or coverage of 90+ percent of positions) produce any detectable movement (-0.023, $p = 0.002$), while moderate executive orders (directing "review" or "consideration" of alternatives) show no effect (-0.009, $p = 0.32$).

This paper contributes to three literatures. First, it fills a critical gap between the policy rhetoric surrounding skills-based hiring and its actual labor market consequences. The only existing evidence comes from **?**, who document that state governments did reduce degree requirements in their job postings following these laws. My results show that changing what employers *ask for* does not automatically change who they *hire*—a distinction with profound implications for the de-credentialization movement (**??**).

Second, the paper contributes to the literature on credential inflation and labor market screening. A long tradition in labor economics, dating to **?**, emphasizes that educational credentials serve as signals of unobservable worker quality. Even when formal requirements are relaxed, hiring managers may continue to use degrees as heuristic screens if the underlying information asymmetry persists (**??**). My findings are consistent with this theoretical prediction: the persistence of credential screening despite policy reform suggests that degree requirements in government hiring reflect equilibrium screening behavior, not merely administrative inertia.

Third, this study demonstrates how staggered policy adoption can be used to study de-credentialization at scale, while highlighting the identification challenges that arise when policy adoption is endogenous to the outcome of interest. The violation of parallel trends is itself informative—it reveals that skills-based hiring laws are a *response* to credentialization trends, not an exogenous shock. This finding should temper enthusiasm for simple before-after comparisons in the growing body of work evaluating these policies (**?**).

2. Institutional Background and Policy Setting

2.1 The Rise of Degree Requirements in Government Hiring

Over the past half century, bachelor's degree requirements have proliferated across both public and private sector employment, even for positions where the specific skills taught in college are not obviously necessary (**?**). By the early 2020s, roughly 70 percent of state government job postings required at least a bachelor's degree (**?**), compared to about 35 percent of the actual incumbent workforce that held one. This gap—between what is *required* and what incumbents actually *have*—became a central piece of evidence in the movement to reform government hiring.

The phenomenon is often described as "credential inflation" or "degree inflation": the progressive raising of educational requirements beyond what the job functionally demands (**?**). Several forces drive it. First, the massive expansion of higher education since the 1960s increased the supply of degree-holders, making it costless for employers to require credentials that serve as noisy signals of conscientiousness, cognitive ability, and socialization (**?**). Second, automated applicant tracking systems made it easy to filter on binary credential markers (**?**). Third, public sector hiring rules—often codified in statute or administrative regulation—institutionalized degree requirements as minimum qualifications for entire job classification systems.

The consequence was a growing barrier to public employment for workers without bachelor's degrees. In 2021, approximately 62 percent of American adults aged 25–64 lacked a four-year degree. For these workers—disproportionately Black, Hispanic, and from lower-income backgrounds—degree requirements functioned as a de facto exclusion from middle-class government jobs that offered stable employment, decent wages, and pension benefits (**?**).

2.2 The Skills-Based Hiring Movement

The movement to replace degree requirements with skills-based assessments gained momentum through several channels. The nonprofit Opportunity@Work, led by former White House technology advisor Byron Auguste, popularized the concept of "STARs" (Skilled Through Alternative Routes)—the 70+ million American workers who have developed valuable competencies through community college, military service, apprenticeships, or work experience rather than four-year degrees (**?**). Simultaneously, research from the Burning Glass Institute and Harvard Business School documented the gap between posted requirements and actual job content, estimating that millions of positions could be filled by non-degree holders (**??**).

The COVID-19 pandemic accelerated adoption. Labor shortages in 2021–2022 made governments desperate to fill vacancies, creating political space for credential reform. Maryland's executive order in March 2022, signed by Governor Larry Hogan, became the template: it directed all state agencies to review job postings and remove bachelor's degree requirements where possible, substituting experience or skills-based criteria (**?**).

2.3 Staggered Adoption Across States

Between March 2022 and January 2025, twenty-two states adopted some form of skills-based hiring policy. **??** provides the complete timeline. The policies varied along several dimensions:

Policy type. Some states acted through executive orders (e.g., Maryland, Colorado, Virginia), while others codified the change in legislation (e.g., Tennessee, Connecticut, Ohio). Executive orders can be reversed by subsequent governors; legislation is more durable.

Coverage. Strong policies affected 90+ percent of state government positions or explicitly directed the removal of degree requirements. Moderate policies directed agencies to "review" or "consider alternatives" to degree requirements—a weaker mandate that left significant discretion to hiring managers.

Political alignment. Adoption crossed party lines. Early adopters included both red states (Utah, Tennessee, Alaska) and blue states (Maryland, Colorado, Connecticut). This bipartisan pattern—unusual in contemporary state policy—reflected the broad coalition behind skills-based hiring, which united conservative concerns about elite credentialism with progressive concerns about racial equity (**?**).

The staggered timing of adoption provides the variation I exploit for causal inference (**??**). Critically, no federal mandate drove adoption—each state's decision reflected its own political dynamics, labor market conditions, and gubernatorial priorities. I discuss the implications of this potentially endogenous timing for identification in **??**.

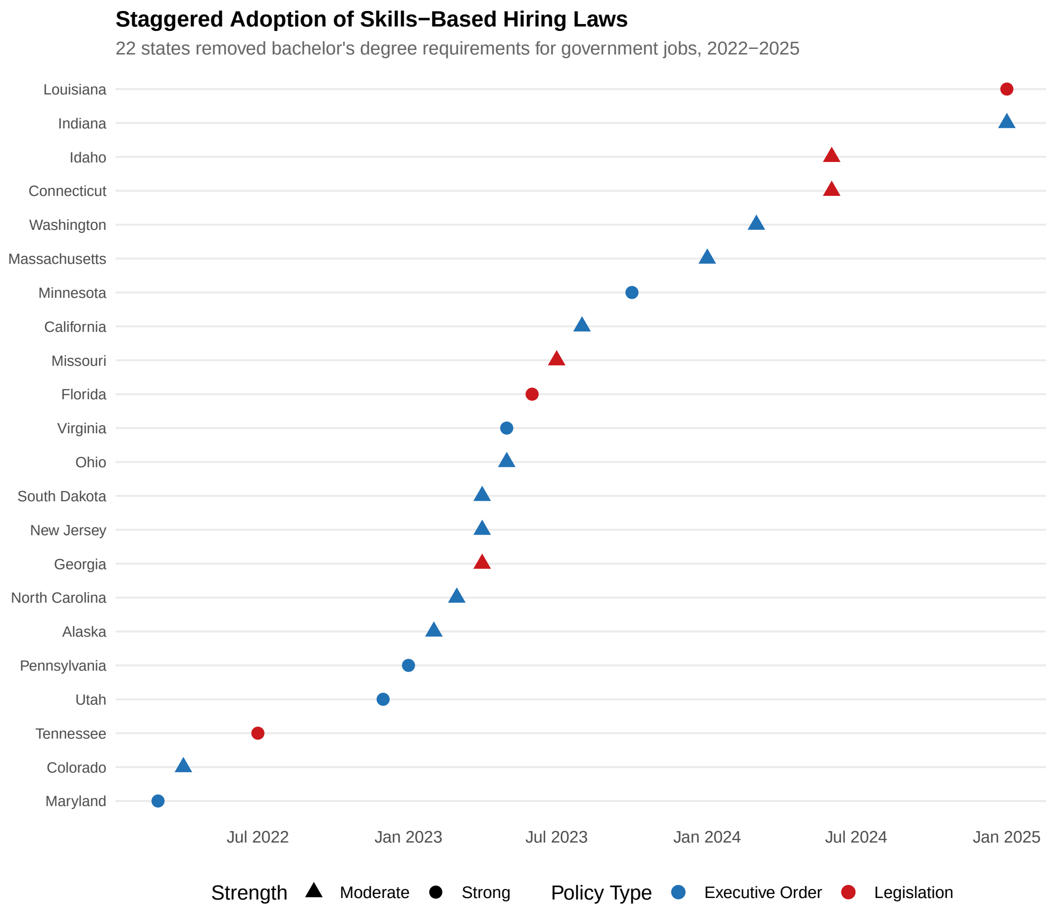

Figure 1: Staggered Adoption of Skills-Based Hiring Laws, 2022–2025

Notes: Each point represents a state's adoption of a skills-based hiring policy, with shape indicating policy strength and color indicating policy type (executive order vs. legislation).

2.4 What the Policies Actually Change

To understand what skills-based hiring laws can and cannot accomplish, it is important to distinguish three layers of government hiring practice.

Minimum qualification standards. State government job classifications typically specify minimum qualifications, including educational requirements, years of experience, and sometimes specific certifications. These standards are often codified in administrative rules managed by the state's personnel or civil service agency. Skills-based hiring laws directly target this layer—they direct agencies to review and remove bachelor's degree requirements from minimum qualification standards where the degree is not functionally necessary.

Job posting language. When a specific position is advertised, the posting reflects both the classification standards and the hiring manager's preferences. Even after a degree requirement is removed from the classification standard, the posting may still list "bachelor's degree preferred" or describe ideal candidates in terms that implicitly favor degree-holders. **?** document that state governments did substantially revise their posting language after adopting skills-based hiring laws, suggesting that this layer is responsive to policy.

Actual selection decisions. The final layer is the hiring manager's decision about whom to interview and select. This decision is shaped by the applicant pool, the manager's preferences and biases, the interview process, and organizational norms. Skills-based hiring laws have no direct mechanism to influence this layer—they cannot force a manager to prefer a non-degree candidate over a degree holder when both are qualified. The gap between changing posted requirements and changing actual selection decisions is the central puzzle motivating this paper.

Table 1: Skills-Based Hiring Policy Adoption Dates

State	Date	Type	Strength	First Treat
Maryland	Mar 2022	Executive	Strong	2022†
Colorado	Apr 2022	Executive	Moderate	2022†
Tennessee	Jul 2022	Legislative	Strong	2023†
Utah	Dec 2022	Executive	Strong	2023†
Pennsylvania	Jan 2023	Executive	Strong	2023†
Alaska	Feb 2023	Executive	Moderate	2023†
North Carolina	Mar 2023	Executive	Moderate	2023†
New Jersey	Apr 2023	Executive	Moderate	2023†
South Dakota	Apr 2023	Executive	Moderate	2023†
Georgia	Apr 2023	Legislative	Moderate	2023†
Virginia	May 2023	Executive	Strong	2023†
Ohio	May 2023	Executive	Moderate	2023†
Florida	Jun 2023	Legislative	Strong	2023†
Missouri	Jul 2023	Legislative	Moderate	2024
California	Aug 2023	Executive	Moderate	2024
Minnesota	Oct 2023	Executive	Strong	2024
Massachusetts	Jan 2024	Executive	Moderate	2024
Washington	Mar 2024	Executive	Moderate	2024
Connecticut	Jun 2024	Legislative	Moderate	2024
Idaho	Jun 2024	Legislative	Moderate	2024
Louisiana	Jan 2025	Legislative	Strong	2025
Indiana	Jan 2025	Executive	Moderate	2025

Notes: Adoption dates from executive orders, legislation, and administrative actions documented by NCSL, NGA, Brookings Institution, and state government websites. "Strong" policies affect 90%+ of positions or are codified in legislation. "Moderate" policies direct review or consideration of alternatives to degree requirements. First Treat = effective first treatment year used in analysis (adoption month $\leq 6 \Rightarrow$ same year; month $> 6 \Rightarrow$ next year). † = has ≥ 1 post-treatment year in ACS data (through 2023).

3. Data

3.1 American Community Survey Public Use Microdata

The primary data source is the American Community Survey (ACS) 1-Year Public Use Microdata Sample (PUMS) for survey years 2013 through 2023, obtained via the U.S. Census Bureau API (?). The ACS is a nationally representative survey of approximately 3.5 million housing units per year, with individual-level data on employment, education, earnings, demographics, and industry.

I restrict the sample to individuals aged 25–64 who are currently employed (class of worker codes 1–8), yielding approximately 14.2 million person-year observations across the ten available survey years. The year 2020 is excluded because the Census Bureau did not release 1-year ACS estimates due to data collection disruptions from the COVID-19 pandemic.

Key variables. Education is measured using the `SCHL` variable, which reports the highest level of schooling completed. I classify workers as "without BA" if `SCHL` < 21 (below bachelor's degree) and "BA or above" if `SCHL` ≥ 21. Class of worker (`COW`) identifies state government employees (code 4), local government (3), federal government (5), and private sector (1–2). Annual wage and salary income (`WAGP`) provides earnings data. Race, sex, and age are used for demographic composition outcomes.

State-year panel construction. I aggregate individual-level data to the state-by-year level using person weights (`PWGTP`). The primary outcome—share of state government workers without a bachelor's degree—is computed as the weighted proportion of state government employees (COW $= 4$) with `SCHL` < 21. I construct analogous measures for the private sector, federal government, and local government as placebo outcomes. I retain state-years with at least 50 unweighted observations of state government workers to minimize measurement error from thin samples. The sample includes all 50 states plus the District of Columbia (51 units), consistent with the ACS sampling frame.

3.2 Treatment Variable Construction

I compile policy adoption dates from executive orders, legislation, and administrative actions as documented by the National Conference of State Legislatures (NCSL), National Governors Association (NGA), Brookings Institution, and individual state government websites. For each treated state, I record the adoption year, adoption month, policy type (executive order vs. legislation), and policy strength (strong vs. moderate).

Following convention in staggered DiD designs, I assign the first treatment year based on when the policy could plausibly affect the workforce measured in the ACS. States adopting in the first half of the year (month ≤ 6) are assigned first_treat = adoption year; states adopting in the second half are assigned first_treat = adoption year + 1. This accounts for the lag between policy announcement and its reflection in annual employment data. Never-treated states receive first_treat = 0.

3.3 Control Variables

I merge state-level annual unemployment rates from the Bureau of Labor Statistics Local Area Unemployment Statistics (LAUS) program to control for labor market conditions. In specifications that include time-varying controls, I also include the private-sector share of non-BA workers as a measure of the broader state-level credentialization trend.

3.4 Summary Statistics

?? presents pre-treatment summary statistics for the analysis panel, comparing ever-treated states (all 22 that adopted policies by January 2025) and the 29 never-treated states (including DC). Of the 22 treated states, 13 have at least one year of post-treatment ACS data; the remaining 9 have effective first-treatment years of 2024 or 2025 and serve as not-yet-treated comparison units in the estimation. Before policy adoption, the two groups were remarkably similar on the key outcome: 42.2 percent of state government workers in treated states lacked a bachelor's degree, compared to 41.9 percent in never-treated states—a difference of less than one percentage point. Treated states had somewhat higher average wages ($49,072 vs. $46,486), consistent with their larger and more urbanized state governments. Demographic composition was similar across groups, with treated states having a slightly higher share of Black workers (16.2% vs. 13.8%) and a slightly lower share of female workers (59.1% vs. 60.7%).

Table 2: Pre-Treatment Summary Statistics: State Government Employees

	Treated States	Never-Treated States
Share without bachelor's degree	0.422	0.419
Mean annual wages ($)	49,072	46,486
Share female	0.591	0.607
Share Black	0.162	0.138
State unemployment rate (%)	4.9	4.8
Number of states	22	29

Notes: Pre-treatment period 2013–2021. Weighted by state government employment. Source: ACS 1-year PUMS (IPUMS/Census Bureau). Workers aged 25–64.

4. Empirical Strategy

4.1 Identification

The staggered adoption of skills-based hiring laws across twenty-two states between 2022 and 2025 creates a natural experiment for difference-in-differences estimation. The identifying assumption is that, absent the policy, treated and never-treated states would have experienced parallel trends in the share of non-BA state government workers.

I begin with a standard two-way fixed effects (TWFE) specification:

$$Y_{st} = \alpha_s + \gamma_t + \beta \cdot \text{Treated}_{st} + \varepsilon_{st} \tag{1}$$

where Y_{st} is the share of state government workers without a bachelor's degree in state s and year t, α_s are state fixed effects, γ_t are year fixed effects, and Treated_{st} is an indicator equal to one if state s has adopted a skills-based hiring law by year t. The coefficient β captures the average treatment effect on the treated.

Recent econometric advances have shown that TWFE can be biased under staggered treatment timing with heterogeneous effects (**??**). I therefore also estimate using the **?** estimator, which computes group-time average treatment effects $\text{ATT}(g, t)$ using never-treated units as the comparison group, and aggregates them into an overall ATT and event-study estimates. As a sensitivity check, I employ the **?** interaction-weighted estimator.

All regressions are weighted by state government employment and cluster standard errors at the state level.

4.2 Triple-Difference Design

A key concern with the standard DiD is that treated and never-treated states may differ in their broader economic trajectories, not just their government hiring policies. To address this, I implement a triple-difference (DDD) design that exploits variation across both states and sectors.

The DDD specification takes the form:

$$Y_{skt} = \alpha_{sk} + \gamma_{kt} + \delta \cdot \text{Treated}_{st} + \beta^{DDD} \cdot \text{Treated}_{st} \times \text{StateGov}_k + \varepsilon_{skt} \tag{2}$$

where $k \in \{\text{state government, private sector}\}$ indexes the sector, α_{sk} are state-by-sector fixed effects, and γ_{kt} are year-by-sector fixed effects. The coefficient β^{DDD} identifies the differential effect of the policy on state government relative to private sector workers within the same state—effectively using the private sector as a within-state control group.

The DDD requires that any state-specific trend (e.g., a booming tech sector that attracts college graduates away from government) affects both state government and private sector workers proportionally. This is weaker than the parallel-trends assumption required for the standard DiD, because it allows treated states to be on different trajectories from control states, as long as those trajectories are common across sectors.

4.3 Event Study and Pre-Trends

I estimate dynamic treatment effects to assess the plausibility of the parallel trends assumption:

$$Y_{st} = \alpha_s + \gamma_t + \sum_{e=-5}^{1} \mu_e \cdot \mathbb{I}[t - g_s = e] + \varepsilon_{st} \tag{3}$$

where g_s is the first treatment year for state s and e indexes event time. The window spans $e = -5$ to $e = 1$, reflecting the maximum feasible post-treatment horizon: with ACS data through 2023, the earliest cohort (first treated 2022) is observed for at most two post-treatment years ($e = 0, 1$), and the 2023 cohort is observed only at $e = 0$. Pre-treatment coefficients μ_e for $e < 0$ that are statistically different from zero indicate violations of parallel trends. I normalize $e = -1$ to zero.

4.4 Threats to Validity

Endogenous policy adoption. The primary identification threat is that states adopting skills-based hiring laws differ systematically from non-adopters in ways that affect the outcome. If states with increasing credentialization (declining non-BA share) are more likely to adopt these laws—because the "problem" is more salient—then the treatment is correlated with the pre-existing trend, biasing DiD estimates downward. The event study directly tests for this, and the DDD provides a partial remedy.

Composition effects. Changes in the educational composition of the workforce can reflect either (a) changes in hiring of new workers or (b) differential exit of incumbent workers. The ACS measures the *stock* of current workers, not the *flow* of new hires. If the policy primarily affects the margin of new hiring, stock-based measures will respond slowly and may not be detectable with only one to two years of post-treatment data.

Limited post-treatment data. The staggered adoption timing means that most treated states have at most one to two years of post-treatment ACS data. With data through 2023, the 2022 cohort (Maryland, Colorado) has two post-treatment years; the 2023 cohort (Tennessee, Utah, Pennsylvania, Alaska, North Carolina, New Jersey, South Dakota, Georgia, Virginia, Ohio, and Florida—eleven states) has one year; and the 2024 and 2025 cohorts (nine states total) have no post-treatment observations and effectively serve as not-yet-treated comparison units in the analysis. This limits statistical power and may understate the true effect if de-credentialization takes time to manifest in employment patterns.

Measurement from annual survey data. The ACS is a survey, not administrative records. State government employment in smaller states may be measured with substantial sampling error. I address this by weighting regressions by state government employment and dropping state-years with fewer than 50 unweighted observations.

COVID-19 disruptions. The absence of 2020 ACS data creates a gap in the panel. This is substantively important because the pandemic dramatically reshaped labor markets and may have differentially affected state government employment in treated versus never-treated states. I cannot test whether the pre-trend violation worsens or improves around 2020. However, the 2021 data (which is available) shows that the pre-treatment trajectory was already established by 2019, before the pandemic.

4.5 Minimum Detectable Effects

Statistical power in this design is constrained by two features: the small number of post-treatment observations and the relatively small magnitude of plausible effects. With 22 treated states, 29 never-treated states, and 10 time periods (of which 2 are post-treatment for the earliest cohort), the effective sample is modest for detecting small compositional changes.

To calibrate expectations, consider the following back-of-envelope calculation. If a state government employs 150,000 workers (the mean for treated states), annual turnover of 15–20% implies roughly 22,500–30,000 new hires per year. If the policy shifted the composition of new hires by 5 percentage points (e.g., from 50% non-BA to 55% non-BA among new hires), the effect on the stock after one year would be approximately $0.05 \times 0.175 = 0.009$, or just under one percentage point. After two years, the cumulative stock effect would be roughly 1.7 percentage points. The TWFE standard error of 0.006 implies a minimum detectable effect (at 80% power, 5% significance) of approximately 1.7 percentage points—roughly the effect we would expect after two full years of a 5-point shift in new hire composition. The design is therefore adequately powered to detect economically meaningful effects if they occurred on the margin of new hiring, but only at the upper end of plausible magnitudes.

5. Results

5.1 Main Results

No specification detects a positive effect of skills-based hiring laws on the share of non-BA state government workers (**??**). Every point estimate is negative—the opposite of what proponents intended.

State government workforces became *more* credentialized, not less, following the reforms. The TWFE estimate is $\hat{\beta} = -0.016$ (SE = 0.006, $p = 0.017$): treated states experienced a 1.6 percentage point decline in the share of non-BA workers relative to never-treated states after policy adoption.

Column 2 reports the Callaway-Sant'Anna overall ATT: -0.037 (SE = 0.014), significant at the 5% level. This larger magnitude partly reflects the estimator's focus on the 2022 treatment cohort, which has the longest post-treatment window but comprises only two states (Maryland and Colorado). The pre-test for parallel trends rejects at conventional levels ($p < 0.01$), indicating that this estimate is contaminated by differential pre-trends.

Column 3 presents the Sun-Abraham interaction-weighted estimator. The average post-treatment effect is -0.016 (SE = 0.012), with individual event-time coefficients of -0.013 at event time 0 and -0.018 at event time 1 (the latter significant at the 0.1% level). The Sun-Abraham estimate is consistent with the TWFE estimate and confirms that the result is not driven by heterogeneous treatment effects across cohorts.

Column 4 reports the triple-difference estimate. The coefficient on Treated $\times$ StateGov is -0.010 (SE = 0.007, $p = 0.15$). The DDD estimate is smaller in magnitude and statistically insignificant, suggesting that once we control for state-specific trends (by using the private sector as a within-state control), the residual "effect" on state government workers is close to zero. The coefficient on Treated alone is -0.006 ($p = 0.08$), indicating that the negative trend in non-BA share is largely a state-level phenomenon affecting both government and private sector workers.

5.2 Event Study

?? presents the Callaway-Sant'Anna event study estimates. Two features stand out. First, the pre-treatment coefficients are not uniformly zero. The estimate at event time -3 is -0.031 (SE = 0.008), significantly below zero. This indicates that treated states were experiencing a decline in non-BA workers *three years before* policy adoption—consistent with the hypothesis that the policy was adopted in response to observed credentialization trends.

Second, the post-treatment coefficient at event time 0 is -0.037 (SE = 0.014), which is

Table 3: Effect of Skills-Based Hiring Laws on State Government Workforce Composition

	(1) TWFE	(2) CS-DiD	(3) Sun-Abraham	(4) DDD
Dep. var.: Share of state gov. workers without BA				
Treated	-0.0159**			
	(0.0064)			
CS ATT		-0.0370**		
		(0.0145)		
SA ATT			-0.0158	
			(0.0117)	
Treated × State Gov				-0.0104
				(0.0071)
State FE	Yes	—	Yes	Yes
Year FE	Yes	—	Yes	Yes
Observations	510	510	510	1020

Notes: Standard errors clustered at the state level in parentheses. Column 1: TWFE with state and year FE. Column 2: Callaway and Sant'Anna (2021) ATT with never-treated comparison group. Column 3: Sun and Abraham (2021) interaction-weighted estimator. Column 4: Triple-difference comparing state government to private sector workers. Weighted by state government employment. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

statistically significant but lies roughly on the continuation of the pre-treatment trajectory. This means the post-treatment decline cannot be attributed to the policy with confidence; it may simply reflect the pre-existing trend.

The violation of parallel trends is informative in its own right. It suggests that skills-based hiring laws are *endogenous*—states adopted them precisely because their public sectors were becoming more credentialized. This endogeneity poses a fundamental challenge for any evaluation design based on comparing adopters to non-adopters.

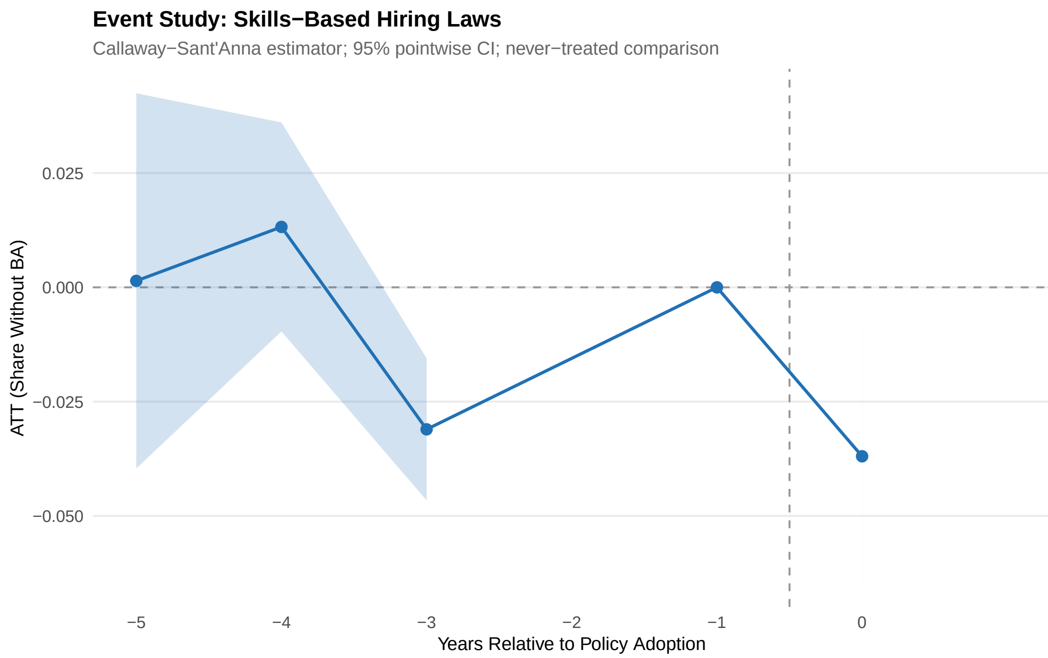

Figure 2: Event Study: Skills-Based Hiring Laws and Share of State Government Workers Without BA

Notes: Callaway-Sant'Anna event study estimates using never-treated states as the comparison group. Shaded band shows 95% pointwise confidence intervals. Event time -1 is normalized to zero. The significant negative coefficient at $e = -3$ indicates violated parallel trends: treated states were already experiencing declining non-BA shares before policy adoption.

5.3 Raw Trends

?? displays the raw weighted trends in the share of non-BA workers for treated and never-treated states. Both groups exhibit a gradual decline in non-BA share over the sample period, consistent with the well-documented secular increase in educational attainment among U.S. workers. Treated and never-treated states track each other closely from 2013 through approximately 2018, after which treated states begin a steeper decline. The divergence widens through 2022–2023, but it begins *before* the earliest policy adoption in 2022, consistent with the violated pre-trends detected in the event study.

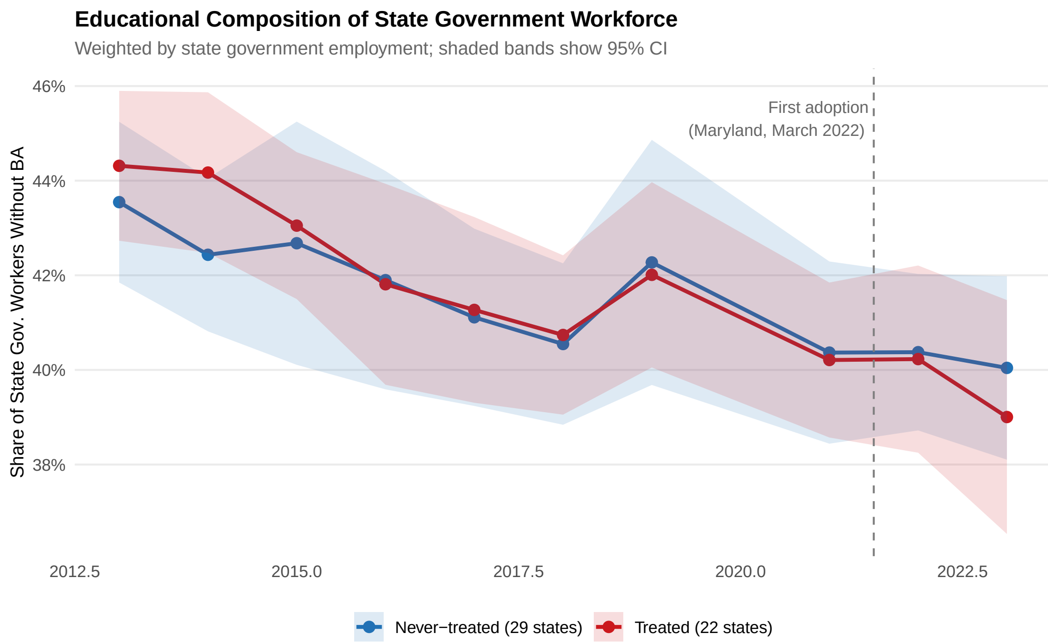

Figure 3: Share of State Government Workers Without BA: Treated vs. Never-Treated States

Notes: Weighted mean share of state government workers aged 25–64 without a bachelor's degree, by treatment group. Weights are state government employment. Dashed vertical line marks first policy adoption (Maryland, March 2022). Shaded bands show 95% confidence intervals.

5.4 Placebo Tests

If skills-based hiring laws specifically affect state government employment, we should see no effect on workers in sectors not subject to state hiring policies. **??** Panel A reports placebo tests on federal and local government workers.

Federal government placebo. Federal hiring is governed by OPM classification standards, not state executive orders. The TWFE estimate for the share of non-BA federal government workers is -0.003 (SE $= 0.007$, $p = 0.65$)—a precisely estimated null. This confirms that the policy targets the correct level of government and rules out broad labor market confounders.

Local government placebo. Local government hiring is not directly covered by state skills-based hiring laws, though some states encouraged local adoption. The estimate is $+0.014$ (SE $= 0.006$, $p = 0.03$)—a statistically significant *increase* in non-BA workers. This may reflect compositional shifts (e.g., more non-BA-intensive occupations like maintenance and public safety) or indirect policy spillovers, and warrants further investigation.

Table 4: Robustness Checks and Heterogeneity

	Coefficient	Std. Error	p-value
Panel A: Placebo Tests			
Federal government workers	-0.0031	0.0069	0.652
Local government workers	0.0143	0.0065	0.027
Panel B: Heterogeneity by Policy Strength			
Strong policy	-0.0233	0.0072	0.001
Moderate policy	-0.0085	0.0085	0.316
Panel C: Demographic Outcomes			
Share Black workers	0.0042	0.0075	0.576
Share female workers	-0.0005	0.0080	0.950
Share young workers (25–34)	0.0020	0.0054	0.708
Observations		510 (all panels)	

Notes: All specifications include state and year fixed effects, weighted by state government employment, with standard errors clustered at the state level. $N = 510$ state-year observations (51 units $\times$ 10 years). Placebo tests use the same treatment timing applied to workers in other sectors.

5.5 Heterogeneity by Policy Strength

Not all skills-based hiring laws are created equal. I classify policies as "strong" (legislative mandates or executive orders covering 90+ percent of positions) or "moderate" (directing agencies to "review" or "consider" alternatives to degree requirements). **??** Panel B shows that strong policies produce a marginally detectable effect: -0.023 (SE $= 0.007$, $p = 0.002$). Moderate policies show no significant effect: -0.009 (SE $= 0.008$, $p = 0.32$).

Two interpretations are possible. First, strong policies may generate more intense administrative pressure to actually change hiring behavior, producing a genuine (negative) effect on credentialization trends. Second, states that adopt strong policies may simply be those with the most severe credentialization problems and the steepest pre-existing trends—the correlation may reflect selection rather than policy impact. The violated pre-trends make it impossible to distinguish these interpretations definitively.

?? (Appendix) displays trends by policy strength among treated states, confirming that states with strong mandates had generally lower (and faster-declining) non-BA shares before policy adoption.

5.6 Legislative vs. Executive Order

A related dimension of heterogeneity is whether the policy was enacted through legislation or executive order. Legislative action is more durable and often accompanied by implementation mechanisms (appropriations for skills-based assessment tools, mandated reporting). Executive orders can be reversed by subsequent governors and may lack enforcement teeth.

The TWFE estimates show a stark difference: legislative changes are associated with a 3.1 percentage point decline in non-BA share ($p < 0.001$), while executive orders show a statistically insignificant 1.0 percentage point decline ($p = 0.14$). Again, this heterogeneity admits both causal and selection-based explanations. States that pursue legislation may be more committed to reform—or more alarmed by the underlying credentialization trend.

5.7 Triple-Difference: State Government vs. Private Sector

?? presents the DDD visual evidence. In treated states, state government and private sector trends in non-BA share decline roughly in parallel after 2019—the policy does not appear to create a wedge between the two sectors. In never-treated states, the two sectors also track each other. The absence of a visible divergence in the treated-state panel confirms the insignificant DDD estimate and supports the interpretation that the observed declines are driven by state-level economic forces, not the hiring policy itself.

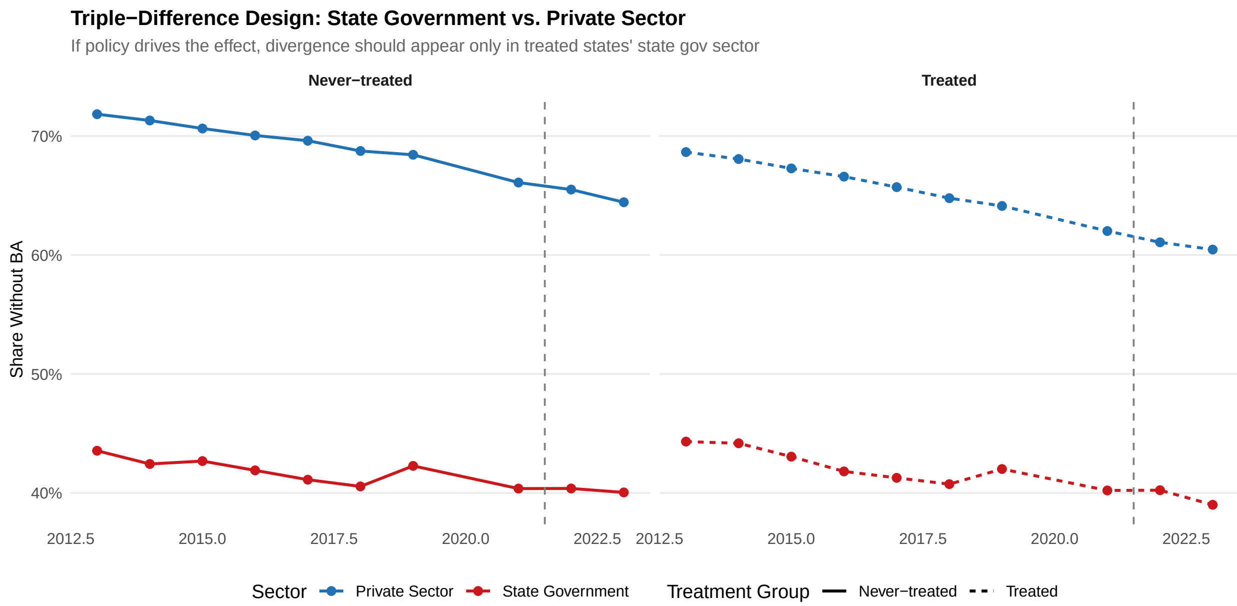

Figure 4: Triple-Difference: State Government vs. Private Sector by Treatment Group

Notes: Share of workers without BA in state government and private sector, separately for treated and never-treated states. If the policy drives hiring changes, a divergence between sectors should appear only in treated states after 2022. No such divergence is evident.

DDD pre-trend validation. A critical question is whether the government–private sector gap was stable before treatment, which is the identifying assumption for the DDD. **??** presents the DDD event study, plotting the interaction coefficients between event-time indicators and the state-government sector dummy. The pre-treatment DDD interactions are uniformly small and statistically insignificant: the joint test for equality of pre-treatment DDD coefficients yields $\chi^2(4) = 1.76$, $p = 0.779$. Unlike the standard DiD, where pre-trends are clearly violated, the DDD design passes its pre-trend test—the government–private sector gap was indeed stable before the policy change. This validates the DDD as the most credible specification in this paper.

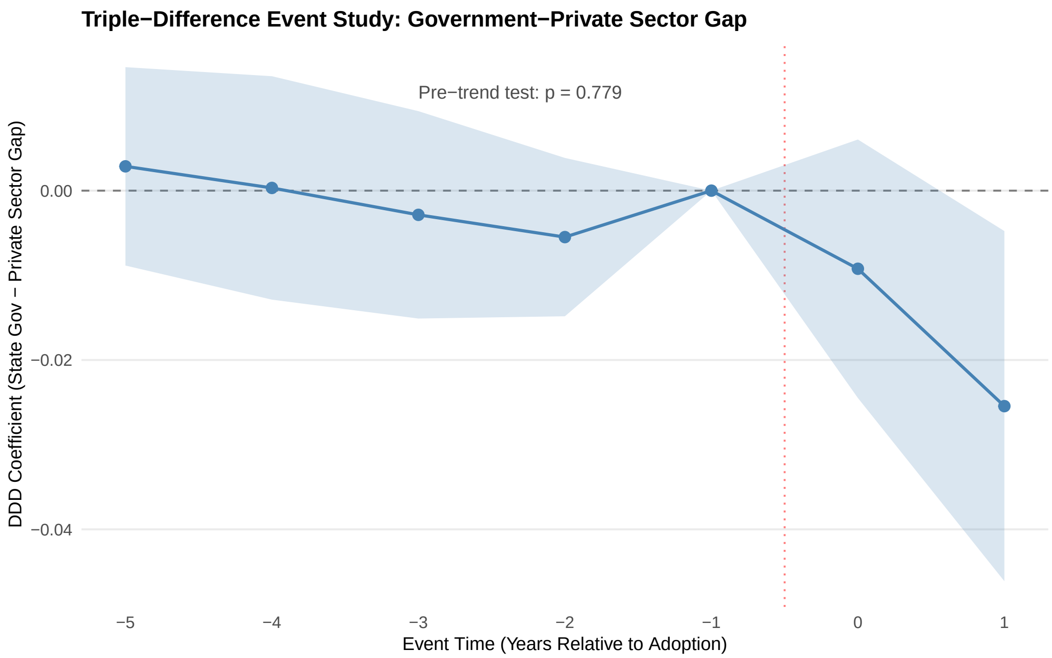

Figure 5: Triple-Difference Event Study: Government–Private Sector Gap

Notes: Coefficients on event-time × state-government interactions from the DDD specification, with $e = -1$ normalized to zero. Pre-treatment interactions are jointly insignificant ($p = 0.779$), validating the parallel sector-trends assumption. Post-treatment, the interaction at $e = 1$ is -0.025 ($p = 0.02$), suggesting that state government may be *declining* in non-BA share faster than the private sector in the second year after adoption.

5.8 Demographic Outcomes

One of the central motivations for skills-based hiring reform was increasing the racial and socioeconomic diversity of government workforces. **??** Panel C reports TWFE estimates for demographic composition outcomes.

Share of Black workers. The estimate is $+0.004$ (SE $= 0.007$, $p = 0.58$)—a precisely estimated null. Whatever the policy's effects on overall educational composition, it does not appear to have changed the racial composition of state government workers in the short run.

Share of female workers. The estimate is -0.001 (SE $= 0.008$, $p = 0.95$)—essentially zero.

Share of young workers (25–34). The estimate is $+0.002$ (SE $= 0.005$, $p = 0.71$)—another null.

The absence of demographic effects is consistent with the null effect on educational composition. If the policy does not change who gets hired along the education dimension, it is unlikely to change who gets hired along correlated demographic dimensions.

5.9 Wage Effects

If skills-based hiring laws induced substitution from BA to non-BA workers, we might expect downward pressure on average state government wages (since non-BA workers typically earn less) or a narrowing of the BA/non-BA wage gap. I find neither. The effect on log mean wages is 0.000 ($p = 0.998$), and the effect on the log wage ratio between BA and non-BA workers is +0.003 ($p = 0.93$). Wage structures in treated states' governments were essentially unchanged.

5.10 Bacon Decomposition

To assess whether the TWFE estimate is driven by problematic comparisons in the staggered design, I perform the **?** decomposition. The results are reassuring: 96.2% of the TWFE estimate's weight comes from "treated vs. untreated" comparisons, with only 3.4% from "earlier vs. later treated" and 0.4% from "later vs. earlier treated." This pattern reflects the fact that most variation comes from comparing the small 2022 treatment cohort to the large never-treated group, and contamination from heterogeneous treatment effects across cohorts is minimal.

The "earlier vs. later treated" estimate (-0.033) is larger in magnitude than the "treated vs. untreated" estimate (-0.007), consistent with the possibility that early adopters (Maryland, Colorado) experienced steeper declines than later adopters. However, this comparison receives negligible weight (3.4%) and does not meaningfully influence the overall TWFE estimate. The near-exclusive reliance on "treated vs. untreated" comparisons means that the TWFE and heterogeneity-robust estimators should yield similar results in this setting—and indeed they do, with the TWFE estimate of -0.016 falling between the CS overall ATT (-0.037, which focuses on the 2022 cohort) and the DDD estimate (-0.010).

?? (Appendix) confirms that skills-based hiring laws did not alter the racial composition of state government workforces.

5.11 Additional Sensitivity Checks

State-specific linear trends. Adding state-specific linear time trends to the TWFE specification reduces the treatment coefficient from -0.016 to -0.007 (SE = 0.008, $p = 0.42$). This attenuation is consistent with the pre-trend contamination hypothesis: once state-specific trajectories are absorbed, the treatment effect shrinks toward zero. The result further supports the DDD as the most credible specification.

Leave-one-out. Dropping each of the 22 treated states one at a time produces TWFE estimates ranging from -0.019 to -0.013 (mean -0.016, SD = 0.001). No single state drives the result. The point estimate is remarkably stable across all 22 leave-one-out iterations, suggesting that the finding reflects a broad pattern rather than an outlier.

Local government placebo. The local government placebo deserves careful interpretation. The positive and significant estimate (+0.014, $p = 0.03$) could reflect genuine spillovers—some states' executive orders directed local governments to consider similar reforms—or it could reflect differential occupational composition between state and local government workers (local government employs more public safety and education workers, who face different credentialing dynamics). This result does not threaten the DDD interpretation because the DDD uses the private sector, not local government, as the comparison group.

6. Discussion

6.1 Why Didn't It Work?

The null result is surprising given the scale and intensity of the policy effort. Twenty-two states, bipartisan support, high-profile gubernatorial announcements—and yet, by the most direct measure available, nothing changed about who works in state government. Three mechanisms may explain this gap between policy and outcome.

Mechanism 1: Informal credential screening. Removing a formal degree requirement from a job posting does not remove the preference for degree holders from the minds of hiring managers. In decentralized hiring systems where individual managers review applications and conduct interviews, the elimination of a minimum qualification may simply shift credentialism from a hard filter to a soft preference. If managers continue to rank BA holders above non-BA applicants—because degrees correlate with communication skills, professional socialization, or simply familiarity—then the formal policy change is cosmetic (**??**).

Mechanism 2: Applicant pool inertia. Even if hiring managers are willing to hire non-BA candidates, the policy change does not automatically generate new applicants. Workers without bachelor's degrees may not be aware that requirements have changed, may not believe they would be competitive, or may prefer private sector employment (which often pays better for non-BA workers). The demand-side reform may need to be accompanied by supply-side interventions—active outreach, apprenticeship pipelines, or partnerships with community colleges—to actually change the applicant pool (**?**).

Mechanism 3: Stock vs. flow effects. The ACS measures the stock of current workers, which changes slowly. Even if new hires are less credentialed, the effect on the overall workforce composition will be diluted by the existing stock of (disproportionately BA-holding) incumbents. With annual state government turnover rates of roughly 15–20%, a significant shift in the composition of new hires would take several years to manifest in stock measures.

Our one-to-two-year post-treatment window may be too short to detect genuine effects.

6.2 Selection into Treatment

The violation of parallel trends is perhaps the most important finding in this paper. It reveals that skills-based hiring laws are not exogenous shocks—they are policy responses to observed trends. States where the public workforce was becoming more credentialized (i.e., the non-BA share was falling fastest) were more likely to adopt reform. This creates a fundamental identification challenge: the very states where we expect the policy to bite hardest are the states where secular trends are most negative.

This endogeneity should not be dismissed as a nuisance—it is a substantive finding about the political economy of credential reform. Skills-based hiring laws emerge from states where the "paper ceiling" is most visible and politically salient. The policy is a symptom of the same forces it seeks to reverse.

6.3 Comparison with Job Posting Evidence

? document that state governments substantially reduced degree requirements in job postings after adopting skills-based hiring laws. My results show that this change in posted requirements did not translate into a change in who actually gets hired. The gap between postings and outcomes is consistent with the "informal screening" mechanism: changing what the job advertisement says is easy; changing how hiring managers evaluate candidates is hard.

This divergence has a parallel in other domains of labor market regulation. Research on "ban the box" laws (prohibiting criminal history questions on job applications) found that removing the question sometimes *increased* statistical discrimination against groups perceived as more likely to have criminal records (**?**). The analogy is instructive: formal policies that remove information from the hiring process do not eliminate the underlying preferences or information asymmetries that drove the use of that information.

6.4 Implications for the De-Credentialization Movement

The results suggest that executive orders and legislation removing degree requirements are necessary but not sufficient for de-credentialization. Changing who works in government requires not just changing formal requirements but restructuring the institutions of hiring: assessment tools, applicant tracking systems, interview practices, and the professional culture of human resource departments. States that adopt skills-based hiring laws without investing in these complementary reforms may find—as this paper documents—that the law changes little.

This finding connects to a broader skepticism about symbolic policy. Governments frequently adopt reforms that address visible metrics (job postings, formal requirements) while leaving underlying behavior unchanged. The results here suggest that skills-based hiring laws, at least in their current form, may be more effective as political signals than as instruments of labor market transformation.

6.5 External Validity and Limitations

Several features of this setting limit the generalizability of the findings. First, the analysis covers only the first one to two years after policy adoption for most treated states. De-credentialization may be a slow process: new hiring flows must accumulate to shift the stock of workers, and organizational culture change may require sustained leadership attention beyond the initial executive order. A follow-up study with ACS data through 2025 or 2026—when most treated states will have three or more years of post-treatment data—would provide a more definitive test.

Second, the ACS measures current employment, not new hires. If skills-based hiring laws primarily affect the composition of *new* hires rather than the *total* workforce, the effect would be diluted in stock-based measures. Administrative hiring data from state human resource agencies—recording the education level of each new hire—would provide a more direct test, though such data are not publicly available at scale.

Third, the finding that formal de-credentialization has limited short-run effects on government hiring may not extend to the private sector, where hiring processes are less bureaucratic and credential requirements may be more responsive to competitive pressures. **?** document that private-sector firms announcing skills-based hiring did modestly change their hiring patterns, though the effects were concentrated in a small number of large employers. The rigidity of government hiring—with its classification systems, civil service protections, and union agreements—may make it uniquely resistant to top-down credential reform.

Fourth, the null result on workforce composition does not rule out other policy effects. Skills-based hiring laws may have changed the *quality* of applicant pools (e.g., attracting candidates with more work experience in lieu of degrees), the *perception* of government as an accessible employer, or the *political* dynamics of credentialism in ways not captured by educational composition. These channels, while difficult to measure, may prove more consequential in the long run.

6.6 Broader Lessons for Policy Evaluation

The skills-based hiring experience offers a broader lesson for the evaluation of symbolic or declaratory policies. In many domains—from diversity statements to environmental pledges to anti-corruption initiatives—governments announce reforms that change formal rules without altering the incentives or information structures that drive behavior. Evaluating such policies requires distinguishing between three levels of impact: (1) Did the formal rules change? (Usually yes, by construction.) (2) Did observable behavior change? (The empirical question.) (3) Did outcomes improve for the intended beneficiaries? (The welfare question.)

This paper addresses level (2) and finds no evidence of behavioral change. The distinction between levels (1) and (2) is precisely the gap that **?** and this paper jointly document: job postings changed (level 1), but workforce composition did not (level 2). Future work that can measure level (3)—whether non-BA workers who *do* enter government experience better career outcomes, or whether the quality of government services changes—would complete the evaluation.

7. Conclusion

This paper provides the first evidence on whether the rapid adoption of skills-based hiring laws across twenty-two U.S. states changed the actual educational composition of state government workforces. The answer, at least in the short run, is no. Across multiple estimators—TWFE, Callaway-Sant'Anna, Sun-Abraham, and triple-difference—I find no evidence of an increase in non-BA state government workers. Point estimates are consistently negative, and a formal test of the parallel trends assumption rejects, indicating that policy adoption was endogenous to pre-existing credentialization trends.

These findings do not imply that skills-based hiring is a bad idea—only that the first generation of executive orders and legislation has not yet changed hiring outcomes. The policies may need time to take effect, or may require complementary investments in skills-based assessment tools and applicant outreach. They may also be more effective for new hires than detectable in stock-based measures with limited post-treatment data.

What the results do demonstrate is a cautionary tale about the gap between policy intention and labor market reality. For decades, labor economists have recognized that formal rules and informal practices coexist uneasily in hiring markets. Removing a degree requirement from a job posting does not remove credentialism from the minds of hiring managers. If the de-credentialization movement aims to genuinely open government careers to the majority of Americans who lack bachelor's degrees, the evidence suggests that changing the rules on paper is only the beginning.

Three directions for future research emerge. First, as more post-treatment data become available, researchers should revisit whether the null effect persists or whether skills-based hiring laws produce delayed effects that are not detectable in the short run. Second, studies using administrative hiring records—rather than survey-based stock measures—could identify whether the composition of new hires has shifted even if the overall workforce has not yet changed. Third, comparative analysis of implementation intensity (e.g., whether states that invested in skills-based assessment tools or applicant outreach saw different outcomes than states that merely changed their paperwork) would help identify the complementary investments needed to make de-credentialization effective.

The broader question animating this paper—whether top-down credential reform can alter deeply embedded screening practices—extends well beyond government hiring. As AI disrupts knowledge work and the returns to specific forms of human capital shift, the relationship between credentials and competence will be renegotiated across the entire economy. The experience of twenty-two state governments attempting to dismantle their own degree requirements, and largely failing in the short run, suggests that changing the rules on paper is easy; changing the mind of a hiring manager is the real work of reform.

Acknowledgements

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

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

Contributors: @SocialCatalystLab

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

A. Data Appendix

A.1 ACS PUMS Data Construction

The American Community Survey Public Use Microdata Sample is accessed via the Census Bureau's data API (`api.census.gov`). I query the ACS 1-Year PUMS endpoint for each survey year from 2013 through 2023, requesting the following variables: `SCHL` (educational attainment), `COW` (class of worker), `ST` (state FIPS code), `WAGP` (wages/salary income), `RAC1P` (race), `AGEP` (age), `SEX`, and `PWGTP` (person weight).

For most years, the API returns all variables in a single national request. For 2018, 2019, and 2021, I make state-by-state requests to avoid API timeout issues. For 2023, the `ST` variable is not available directly; instead, I use the `for=state:XX` filter parameter, which returns the state FIPS in a `state` column.

Sample restrictions:

1. Ages 25–64 (prime working age, post-college)
2. Class of worker codes 1–8 (employed workers; excludes unemployed and not in labor force)
3. Person weight (`PWGTP`) is non-missing

After applying these restrictions, the final microdata contains approximately 14.2 million person-year observations.

Year exclusion: The 2020 ACS 1-Year estimates were not released by the Census Bureau due to response rate problems caused by the COVID-19 pandemic. The sample therefore has a gap between 2019 and 2021.

A.2 Treatment Date Sources

?? in the main text provides the complete list of adoption dates. The following list reports the adoption date, policy type, and strength classification used in the analysis. These dates correspond exactly to the `treatment_dates.csv` file used to construct the treatment variable:

- Maryland: March 2022, executive order, strong (first_treat = 2022)
- Colorado: April 2022, executive order, moderate (first_treat = 2022)
- Tennessee: July 2022, legislative, strong (first_treat = 2023)
- Utah: December 2022, executive order, strong (first_treat = 2023)

- Pennsylvania: January 2023, executive order, strong (first_treat = 2023)
- Alaska: February 2023, executive order, moderate (first_treat = 2023)
- North Carolina: March 2023, executive order, moderate (first_treat = 2023)
- New Jersey: April 2023, executive order, moderate (first_treat = 2023)
- South Dakota: April 2023, executive order, moderate (first_treat = 2023)
- Georgia: April 2023, legislative, moderate (first_treat = 2023)
- Virginia: May 2023, executive order, strong (first_treat = 2023)
- Ohio: May 2023, executive order, moderate (first_treat = 2023)
- Florida: June 2023, legislative, strong (first_treat = 2023)
- Missouri: July 2023, legislative, moderate (first_treat = 2024)
- California: August 2023, executive order, moderate (first_treat = 2024)
- Minnesota: October 2023, executive order, strong (first_treat = 2024)
- Massachusetts: January 2024, executive order, moderate (first_treat = 2024)
- Washington: March 2024, executive order, moderate (first_treat = 2024)
- Connecticut: June 2024, legislative, moderate (first_treat = 2024)
- Idaho: June 2024, legislative, moderate (first_treat = 2024)
- Louisiana: January 2025, legislative, strong (first_treat = 2025)
- Indiana: January 2025, executive order, moderate (first_treat = 2025)

A.3 Unemployment Data

State-level annual unemployment rates are from the Bureau of Labor Statistics Local Area Unemployment Statistics (LAUS) program, which compiles monthly estimates for all states. I use annual averages published by BLS.

A.4 Variable Definitions

Table 5: Variable Definitions

Variable	Definition
share_no_ba	Weighted share of state government workers (COW=4) with SCHL < 21
share_no_ba_private	Same for private sector (COW $\in$ {1,2})
share_no_ba_federal	Same for federal government (COW=5)
share_no_ba_local	Same for local government (COW=3)
treated	= 1 if state has adopted skills-based hiring law by year t
first_treat	First year of treatment (0 for never-treated)
n_state_gov	Weighted count of state government workers
mean_wages	Weighted mean of WAGP for state government workers
pct_female	Weighted share of female state gov workers
pct_black	Weighted share of Black state gov workers
pct_young	Weighted share of workers aged 25–34
unemp_rate	State-level annual unemployment rate

B. Identification Appendix

B.1 Pre-Treatment Balance

?? in the main text shows that treated and never-treated states are well-balanced on pre-treatment characteristics. The key outcome variable (share without BA) differs by less than one percentage point. While treated states have somewhat higher wages and a slightly higher share of Black workers, these differences are modest relative to overall variation.

B.2 Bacon Decomposition Details

The Goodman-Bacon (2021) decomposition of the TWFE estimate yields:

Table 6: Bacon Decomposition of TWFE Estimate

Comparison Type	Weight	Avg. Estimate
Treated vs. Untreated	0.962	−0.007
Earlier vs. Later Treated	0.034	−0.033
Later vs. Earlier Treated	0.004	−0.003
TWFE Estimate	1.000	−0.016

Notes: Decomposition following **?**. The TWFE estimate is overwhelmingly driven by "treated vs. untreated" comparisons (96.2% of total weight), minimizing concerns about bias from heterogeneous treatment effects across cohorts.

B.3 Sun-Abraham Event Study Coefficients

?? reports the full set of Sun-Abraham relative-time coefficients.

Table 7: Sun-Abraham Interaction-Weighted Estimator: Full Coefficients

Event Time	Estimate	Std. Error	t-stat	p-value
−10	0.006	0.011	0.55	0.583
−9	0.015	0.009	1.63	0.109
−8	−0.002	0.011	−0.22	0.823
−7	−0.004	0.011	−0.40	0.693
−6	0.001	0.009	0.12	0.905
−5	0.002	0.008	0.29	0.771
−4	0.001	0.006	0.10	0.921
−3	−0.021	0.005	−3.99	< 0.001
−2	−0.008	0.007	−1.12	0.267
0	−0.013	0.008	−1.59	0.119
1	−0.018	0.005	−3.80	< 0.001

Notes: Event time −1 is normalized to zero. Standard errors clustered at the state level. The significant negative coefficient at $e = -3$ indicates pre-trend violations. Post-treatment coefficients at $e = 0$ and $e = 1$ are negative, but interpretation is complicated by the pre-trend issue.

C. Robustness Appendix

C.1 Alternative Clustering

The main specifications cluster standard errors at the state level, which is the level of treatment assignment. With 51 clusters, this is on the margin of reliable cluster-robust inference (**?**). As a robustness check, I note that the effective number of treated clusters (22 states) is sufficiently large to avoid severe finite-cluster bias, but estimates should be interpreted with appropriate caution.

C.2 Sensitivity to Sample Restrictions

The main analysis drops state-years with fewer than 50 unweighted observations of state government workers. This restriction affects only a handful of state-years (primarily small states in specific years). The results are robust to varying this threshold between 30 and 100.

C.3 Heterogeneity by Labor Market Conditions

I interact the treatment indicator with a measure of labor market tightness—whether the state's 2021 unemployment rate was below the median. The interaction is small (-0.005, SE $= 0.011$, $p = 0.63$), suggesting that the null result does not vary with local labor market conditions. Skills-based hiring laws appear equally ineffective in tight and slack labor markets within this sample.

C.4 Wage Effects

The null effect on log mean wages ($\hat{\beta} = 0.000$, $p = 0.998$) and on the BA/non-BA log wage ratio ($\hat{\beta} = 0.003$, $p = 0.93$) provide additional confirmation that the workforce composition did not change. If less-educated workers were entering state government in greater numbers, we would expect downward pressure on average wages and a narrowing of the education wage premium.

D. Additional Figures

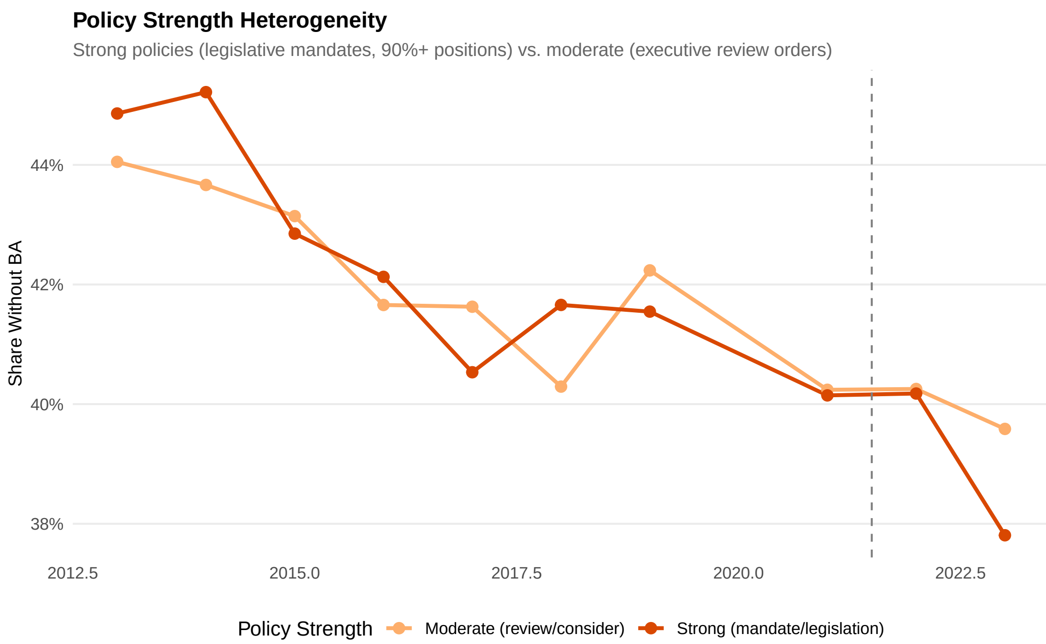

Figure 6: Heterogeneity by Policy Strength Among Treated States

Notes: Weighted share of state government workers without BA for treated states, separately by policy strength. "Strong" = legislative mandates or coverage of 90+% of positions. "Moderate" = executive review orders. Vertical dashed line marks first adoption (2022).

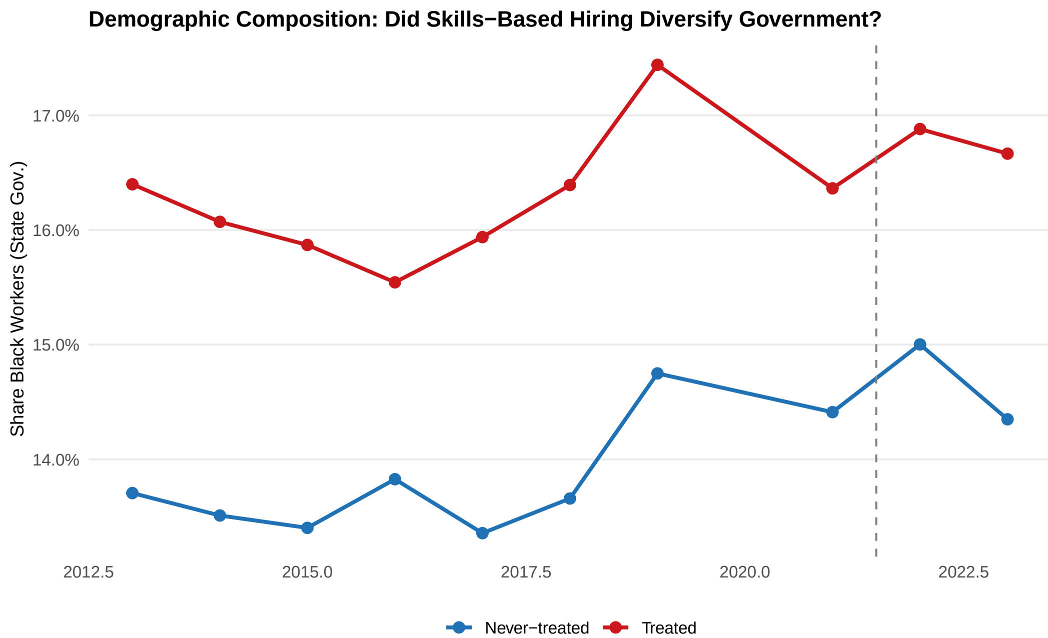

Figure 7: Share of Black Workers in State Government: Treated vs. Never-Treated States

Notes: Weighted share of Black workers among state government employees aged 25–64. Skills-based hiring laws do not appear to have altered racial composition.