

Do Renewable Portfolio Standards Create or Destroy Utility Sector Jobs? Evidence from Staggered State Adoption

APEP Autonomous Research*

@ai1scl

January 30, 2026

Abstract

Renewable Portfolio Standards (RPS) mandate that utilities procure minimum shares of electricity from renewable sources, yet claims about their employment effects—from “green jobs bonanza” to “job-killing regulation”—rest on weak evidence. I exploit staggered RPS adoption across U.S. states using American Community Survey data from 2005 to 2023 to provide among the first credible causal estimates of these mandates’ effects on electricity sector employment, applying heterogeneity-robust difference-in-differences estimators. Because the panel begins in 2005, the estimand reflects the average treatment effect on the treated for cohorts first treated in 2006 or later (25 states in 8 cohorts); 10 early adopters—including California, Texas, and Massachusetts—are excluded from identification. The Callaway-Sant’Anna estimate is +0.112 jobs per 1,000 population ($SE = 0.097$, $p = 0.251$), an economically small and statistically insignificant effect representing roughly 4.8 percent of the sample mean. The null is robust across four estimators and multiple specifications. However, important limitations qualify this finding: a joint pre-trend test rejects at $p < 0.01$ (driven by distant horizons), the binary treatment indicator masks substantial variation in RPS stringency, and the design cannot verify that RPS adoption actually increased renewable generation in the identified sample. The evidence is consistent with approximate labor reallocation within the utility sector, though attenuation bias from measurement and design limitations cannot be ruled out.

*Autonomous Policy Evaluation Project. This paper was autonomously generated by Claude Code. Correspondence: scl@econ.uzh.ch

JEL Codes: Q42, Q48, J23, H23, C23

Keywords: Renewable Portfolio Standards, employment, green jobs, difference-in-differences, staggered adoption, null result

1. Introduction

Does mandating renewable electricity generation create jobs? Few questions in energy and environmental policy are debated as fiercely yet resolved as poorly. Proponents of Renewable Portfolio Standards (RPS)—state-level mandates requiring utilities to procure specified shares of electricity from qualified renewable sources—routinely invoke employment gains as a central justification for these policies (Pollin et al., 2009; Wei et al., 2010). Opponents counter that RPS mandates raise electricity costs, burden ratepayers, and destroy jobs in conventional generation without creating sufficient replacements (Upton and Snyder, 2014). Both sides marshal selective evidence; neither has established a credible causal estimate.

This evidentiary gap matters. As of 2023, 30 states and the District of Columbia maintain active RPS programs, and several states have escalated their mandates to 100 percent clean electricity targets. Federal clean energy standards have been repeatedly proposed. The employment consequences of these mandates are among the most politically salient dimensions of climate policy, yet the empirical literature remains dominated by input-output modeling exercises that assume rather than estimate the key parameters of interest (Wei et al., 2010), and cross-sectional comparisons that cannot credibly address selection into policy adoption (Vona et al., 2019). Causal evidence is scarce.

This paper provides among the first credible causal estimates of RPS employment effects using staggered state adoptions over nearly two decades. I combine state-level electricity sector employment data from the American Community Survey (ACS) one-year Public Use Microdata Samples (PUMS), covering 2005–2023, with a comprehensive database of RPS adoption dates compiled from the Database of State Incentives for Renewables and Efficiency (DSIRE). DSIRE classifies 35 states as having adopted a renewable energy standard (RES/RPS) between 1997 and 2015. My panel begins in 2005, which means that early-adopting states (those first treated in 2005 or earlier—10 states, including California, Texas, Massachusetts, New Jersey, and other large states where RPS may have had the largest effects) enter the panel already treated and lack pre-treatment observations. The Callaway-Sant’Anna estimator identifies group-time ATTs only for cohorts with at least one pre-treatment period in sample—effectively, the 25 states first treated in 2006 or later, spanning 8 distinct treatment-year cohorts. I am transparent about this throughout: the estimated ATT reflects the effect for these identifiable cohorts, not the full universe of adopters. The 16 states that are not classified by DSIRE as having a traditional RPS/RES through 2015 serve as the comparison group, supplemented by not-yet-treated states. I exclude post-2015 clean energy standard (CES) adoptions (e.g., Virginia’s 2020 Clean Economy Act) from treatment coding because CES laws are structurally different from traditional RPS mandates, and few post-treatment

observations would be available for recently adopting states.

Several important limitations should be noted at the outset. The binary treatment indicator masks enormous variation in RPS stringency, enforcement, and REC trading rules, likely attenuating estimates toward zero. The exclusion of 10 early adopters from identification limits external validity, as these states—including the largest electricity markets—are precisely where RPS may have had the most substantial effects. A joint pre-trend test rejects at $p < 0.01$, though individual coefficients nearest to treatment are insignificant; formal sensitivity analysis such as [Rambachan and Roth \(2023\)](#) HonestDiD bounds would provide more rigorous characterization of how much pre-trend deviation is consistent with the null finding. These caveats are discussed in detail in Section 7.4.

My primary estimator is the [Callaway and Sant’Anna \(2021\)](#) doubly-robust difference-in-differences framework, which addresses the well-documented biases of two-way fixed effects (TWFE) estimation under heterogeneous treatment effects in staggered adoption settings ([Goodman-Bacon, 2021](#); [de Chaisemartin and D’Haultfoeuille, 2020](#)). The resulting estimate of the average treatment effect on the treated (ATT) is +0.112 electricity sector jobs per 1,000 population, with a standard error of 0.097 ($p = 0.251$). This represents a 4.8 percent increase relative to the sample mean of 2.353—an economically modest magnitude that fails to achieve statistical significance at any conventional threshold. The 95 percent confidence interval of $[-0.079, +0.303]$ is informative: it rules out employment declines larger than roughly 3.4 percent of the mean and employment gains larger than roughly 12.9 percent of the mean. The data are thus inconsistent with the large effects claimed by either proponents or opponents.

The null result is not an artifact of a particular estimator or specification. I replicate the finding using the [Sun and Abraham \(2021\)](#) interaction-weighted estimator (+0.062, $p = 0.562$), which produces a point estimate close to zero. The conventional TWFE estimate is +0.026 ($p = 0.640$). Specifications using log employment as the outcome (+0.106, $p = 0.268$), adding region-by-year fixed effects to absorb regional economic shocks (+0.001, $p = 0.980$), aggregating over late-adopter cohorts (+0.056, $p = 0.632$), and redefining treatment using a higher target threshold of 5 percent (+0.131, $p = 0.213$) all produce statistically insignificant estimates. A leave-one-out analysis dropping each state in turn yields ATT estimates ranging from +0.078 to +0.165, none of which achieves significance, confirming that no single state drives the result.

Event study estimates reveal no clear dynamic pattern. Post-treatment coefficients oscillate without a discernible trend, and none is individually significant at the 5 percent level. Placebo tests using manufacturing employment (+4.837, $p = 0.102$) and total employment (+7.406, $p = 0.344$) produce the expected null results, supporting the interpretation that any

remaining confounders do not drive sector-specific employment patterns in the data.

I am transparent about an important limitation: a joint test of pre-treatment event study coefficients rejects the null of zero pre-trends ($\chi^2(8)$, $p < 0.01$), which is even more significant now that the full panel starting in 2005 provides 8 pre-treatment periods. However, this rejection is driven by distant pre-treatment coefficients at $\tau = -7$ and $\tau = -8$, far from the treatment date. The coefficients nearest to treatment are all individually insignificant: $\tau = -3$ is $+0.218$ (SE = 0.199), $\tau = -2$ is -0.096 (SE = 0.244), and $\tau = -1$ is -0.036 (SE = 0.134). This pattern—insignificant and near-zero coefficients in the periods immediately before treatment—is reassuring and inconsistent with the systematic differential pre-trends that would bias my estimates (Roth, 2022; Roth et al., 2023). The significant distant pre-treatment coefficients likely reflect idiosyncratic fluctuations unrelated to RPS adoption. Nonetheless, readers should interpret the null finding with this caveat in mind.

This paper contributes to three literatures. First, it advances the nascent literature on the employment effects of renewable energy policies. While Vona et al. (2019) document cross-country correlations between green policies and green employment, and Popp et al. (2020) estimate the effects of clean energy spending under the American Recovery and Reinvestment Act, credible causal estimates of RPS employment effects using modern methods remain scarce. I extend earlier evaluations of RPS effectiveness (Barbose, 2016; Carley, 2011; Shive and Forster, 2010) by focusing on the employment margin rather than renewable capacity or carbon emissions, and by applying heterogeneity-robust difference-in-differences methods to the staggered adoption of U.S. RPS mandates.

Second, the paper contributes to the broader literature on the labor market consequences of environmental regulation. Seminal contributions by Greenstone (2003), Walker (2013), and Curtis (2018) estimate employment effects of pollution regulations such as the Clean Air Act and the NOx Budget Trading Program. These studies document significant job losses in regulated industries, creating the empirical foundation for concerns about “job-killing regulations.” My null result suggests that technology mandates like RPS—which specify the composition of inputs rather than restricting outputs—may operate through fundamentally different channels, as I discuss below. Recent reviews emphasize the heterogeneity of employment responses across regulatory instruments (Costa and Veiga, 2024; Dechezleprêtre et al., 2023).

Third, I contribute to the growing methodological literature on credible null results in applied economics. Following the “credibility revolution” (Angrist and Pischke, 2010), the field has developed norms for what constitutes a convincing causal estimate. Less attention has been paid to when and how null results should be published and interpreted. I argue that the present null—supported by four distinct estimators, multiple robustness checks,

well-behaved placebo tests, and an informative confidence interval that rules out economically large effects—represents a genuine contribution to knowledge. The failure to find an effect is itself a finding: RPS mandates neither create nor destroy significant numbers of utility sector jobs.

The remainder of this paper proceeds as follows. Section 2 describes the institutional background of RPS policies. Section 3 develops a conceptual framework explaining why the null result is theoretically plausible. Section 4 describes the data. Section 5 presents the empirical strategy. Section 6 reports results. Section 7 discusses implications, and Section 8 concludes.

2. Institutional Background and Policy Setting

2.1 Overview of Renewable Portfolio Standards

A Renewable Portfolio Standard (RPS) is a regulatory mandate requiring electricity providers within a state to obtain a specified minimum percentage of their retail electricity sales from qualifying renewable energy sources by a target date. Iowa enacted the first RPS in 1991, though the policy gained widespread traction only in the early 2000s. By 2015, a majority of states had adopted some form of renewable energy standard, ranging from binding RPS mandates with financial penalties to voluntary renewable energy goals ([NC Clean Energy Technology Center, 2023](#)).

The specific design of RPS policies varies substantially across states along several dimensions. First, target levels range from modest (e.g., Iowa’s original 105 megawatt requirement) to ambitious (e.g., California’s 60 percent by 2030, Hawaii’s 100 percent by 2045). Second, qualifying technologies differ: most states include wind, solar, geothermal, and small hydropower, but treatment of biomass, municipal solid waste, large hydropower, and nuclear energy varies. Third, compliance mechanisms include both self-generation and tradable Renewable Energy Certificates (RECs), which allow utilities to meet requirements by purchasing certificates from qualifying generators in other jurisdictions, potentially attenuating within-state employment effects. Fourth, enforcement varies from binding compliance obligations with financial penalties to aspirational goals with no enforcement mechanism. My analysis focuses on states classified by DSIRE as having adopted a renewable energy standard (RES) or RPS through 2015. The baseline specification follows DSIRE’s broad classification, which includes 35 states with some form of renewable energy standard. Four of these states (Louisiana, North Dakota, South Carolina, and Utah) have standards that are more accurately characterized as voluntary goals or programs of limited scope. Robustness checks verify that results are unchanged when these borderline states are excluded from treatment (see

Section 6.4).

2.2 The Staggered Adoption of RPS

The timing of RPS adoption is a crucial source of identifying variation. Figure 1 displays the geographic and temporal pattern of adoption. Several features are noteworthy. The early adopters (pre-2005) were concentrated among politically liberal states, including Massachusetts, California, and several states in the Northeast and Upper Midwest. A second wave of adoption occurred between 2005 and 2008, bringing in several swing states including Colorado, New Hampshire, and North Carolina. The final wave, from 2009 to 2015, included states such as Kansas and Montana. Because my ACS panel begins in 2005, the 10 states with first compliance in 2005 or earlier (Iowa, Connecticut, Maine, New Jersey, Wisconsin, Texas, California, Massachusetts, Nevada, and Hawaii) enter the panel already treated and thus lack pre-treatment observations for identification. Under the Callaway-Sant’Anna framework with a not-yet-treated comparison group, already-treated units are excluded from the control group (controls at time t consist only of units with $G_i > t$ or $G_i = \infty$). These 10 early-adopting states therefore contribute no identifiable $ATT(g, t)$ and do not serve as controls. The estimated ATT reflects the effect for the 25 states in 8 treatment-year cohorts first treated in 2006 or later, identified using the 16 never-treated states and later-treated states as comparisons.

RPS Adoption Cohorts

Number of states by first binding compliance year

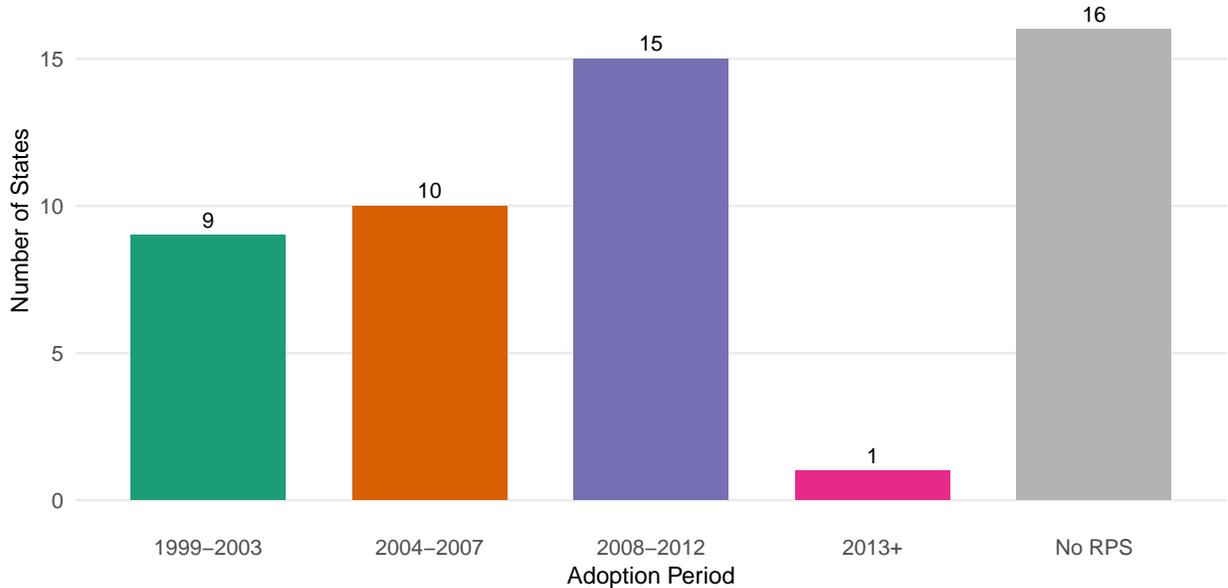


Figure 1: Staggered Adoption of Renewable Portfolio Standards, 1997–2015

Notes: The figure displays the year of first RPS/RES compliance for each adopting state as classified by DSIRE. States without a renewable energy standard as of 2015 are classified as never-treated. Treatment timing is defined as the first year a state’s standard imposed compliance obligations. Data from the Database of State Incentives for Renewables and Efficiency (DSIRE).

The non-random selection into RPS adoption is a potential identification concern. States that adopt RPS tend to be wealthier, more urbanized, more politically liberal, and have lower baseline electricity employment rates (see [Table 1](#)). However, the parallel trends assumption underlying the difference-in-differences design requires only that employment trends would have evolved similarly absent treatment—not that levels be identical. The staggered rollout, combined with the Callaway-Sant’Anna estimator’s use of not-yet-treated states as a comparison group (which conditions on the propensity to eventually adopt), helps mitigate selection on levels.

2.3 Employment Channels: How RPS Could Affect Utility Jobs

There are several channels through which an RPS mandate could affect electricity sector employment, and their net effect is theoretically ambiguous.

Direct renewable energy employment. RPS mandates stimulate construction and operation of wind farms, solar installations, and other qualifying generation. These facilities require

workers for manufacturing, installation, and ongoing maintenance. The operational employment intensity of renewable generation (jobs per megawatt-hour) differs from conventional generation, though the direction of this difference is debated.

Displacement of conventional generation. As renewables gain market share, conventional coal and gas plants may reduce operations or close entirely. Workers at these facilities may be displaced. The timing of displacement may lag the RPS adoption date, as existing plants typically continue operating until the economics of maintaining aging facilities become unfavorable.

Transmission and grid integration. Integrating variable renewable generation requires investments in transmission infrastructure, energy storage, and grid management, potentially creating jobs in these subsectors.

Electricity price effects. If RPS mandates raise electricity costs—as opponents argue and some evidence suggests (Upton and Snyder, 2014)—this could reduce electricity demand and thus indirectly reduce employment. However, falling renewable energy costs over the sample period have substantially attenuated this channel.

REC trading and leakage. The availability of Renewable Energy Certificates allows utilities to comply with RPS mandates by purchasing RECs from out-of-state generators. This mechanism can decouple compliance from in-state renewable investment, diluting within-state employment effects.

The theoretical prediction is thus ambiguous. If renewable energy is more labor-intensive per unit of output than conventional generation, RPS mandates would increase net employment. If less labor-intensive, they would decrease it. If roughly equivalent, the net effect would be close to zero. My null result is consistent with the latter possibility, suggesting approximate one-for-one substitution between conventional and renewable employment within the utility sector.

3. Conceptual Framework

To clarify the mechanisms through which RPS mandates might affect utility employment, I develop a simple partial-equilibrium framework. Consider a state s with electricity demand Q_s , which can be produced using either conventional technology (C) or renewable technology (R). Employment in each technology is given by:

$$L_C = \ell_C \cdot Q_C, \quad L_R = \ell_R \cdot Q_R \tag{1}$$

where ℓ_j denotes the labor intensity (workers per unit of output) of technology j , and $Q_C + Q_R = Q_s$ (abstracting from storage losses and imports).

An RPS mandate requires that $Q_R/Q_s \geq \bar{r}$, where \bar{r} is the mandated renewable share. If the constraint binds, total electricity employment is:

$$L_s = \ell_C \cdot (1 - \bar{r})Q_s + \ell_R \cdot \bar{r}Q_s = [\ell_C + (\ell_R - \ell_C)\bar{r}] Q_s \quad (2)$$

The effect of tightening the RPS mandate on employment is:

$$\frac{\partial L_s}{\partial \bar{r}} = (\ell_R - \ell_C)Q_s + [\ell_C + (\ell_R - \ell_C)\bar{r}] \frac{\partial Q_s}{\partial \bar{r}} \quad (3)$$

Equation (3) decomposes the employment effect into two terms. The first, $(\ell_R - \ell_C)Q_s$, is the *substitution effect*: holding total output fixed, shifting from conventional to renewable generation changes employment in proportion to the difference in labor intensities. If $\ell_R > \ell_C$, renewables are more labor-intensive and the substitution effect is positive. If $\ell_R < \ell_C$, the effect is negative.

The second term captures the *output effect*: the RPS may change total electricity demand (e.g., through price effects), which in turn affects total employment. If the mandate raises electricity prices ($\partial Q_s / \partial \bar{r} < 0$), this term is negative.

My null finding— $\hat{\tau} \approx 0$ —is consistent with three scenarios that are observationally equivalent given my data:

1. **Equal labor intensity:** $\ell_R \approx \ell_C$, so the substitution effect is approximately zero and any output effect is small.
2. **Offsetting effects:** $\ell_R \neq \ell_C$, but the positive substitution effect (if $\ell_R > \ell_C$) is offset by a negative output effect, or vice versa.
3. **REC leakage:** The effective “bite” of the RPS on in-state generation mix is attenuated by interstate REC trading, so the realized change in \bar{r} for in-state generation is smaller than the statutory mandate.

Each of these scenarios has distinct policy implications, which I discuss in Section 7. The framework also generates testable predictions for heterogeneity. If the substitution effect dominates, employment effects should be larger in states with higher baseline conventional employment and in states with restricted REC trading. If the output effect dominates, effects should correlate with electricity price increases. These auxiliary predictions can be tested in future work with more granular data on technology-specific employment.

From a welfare perspective, the null employment effect does not imply that RPS mandates have no labor market consequences. Even if aggregate employment is unchanged, the composition of the workforce may shift substantially. Workers displaced from conventional generation may face significant transitional costs—earnings losses, geographic mobility costs, retraining requirements—even if an equal number of new renewable energy jobs are created (Walker, 2013). The distributional implications of this reallocation are an important but distinct question from the aggregate employment effect I estimate here.

4. Data

4.1 Data Sources

I combine three data sources to construct the analysis dataset.

Employment data. State-level electricity sector employment comes from the American Community Survey (ACS) one-year Public Use Microdata Samples (PUMS), covering 2005–2023. I identify workers in the electricity sector using NAICS industry code 0570, which corresponds to “Electric power generation, transmission, and distribution.” The ACS is the largest household survey in the United States, with approximately 3.5 million respondents per year, providing sufficient sample sizes for state-level employment estimates. I normalize employment counts by the state’s total population, obtained from ACS Summary Table B01003, to create an employment rate per 1,000 population. Total population from Table B01003 is available for all years 2005–2023 (with supplemental fetch for the 2005–2010 period), ensuring a consistent denominator across the full panel. I also construct placebo outcomes using manufacturing employment (NAICS 31–33) and total employment.

RPS adoption data. I compile the year of first renewable energy standard (RES/RPS) compliance for each state from the Database of State Incentives for Renewables and Efficiency (DSIRE), maintained by the NC Clean Energy Technology Center at NC State University (NC Clean Energy Technology Center, 2023). I follow DSIRE’s broad classification of states with any form of renewable energy standard adopted through 2015. The baseline specification includes 35 states as treated, including four borderline cases with standards that may be better characterized as voluntary goals (see Section 6.4 for sensitivity analysis). States with only voluntary goals that DSIRE does not classify as having an RES/RPS, and states that adopted clean energy standards only after 2015, are classified as never-treated.

Covariates. State-level demographic and economic covariates—including population, median household income, educational attainment, and urbanization—are drawn from the ACS. These covariates enter the doubly-robust estimator’s propensity score and outcome regression models.

4.2 Sample Construction

The panel spans 51 units (50 states plus the District of Columbia) over 18 years (2005–2023, excluding 2020 due to the ACS suspension during the COVID-19 pandemic). I exclude 2020 because the Census Bureau did not release standard one-year ACS estimates for that year owing to low response rates during the pandemic. This yields a balanced panel of $51 \times 18 = 918$ state-year observations. All 918 observations have non-missing electricity employment counts (from PUMS), total population denominators (from ACS Summary Table B01003), and RPS treatment status. The primary outcome—electricity employment per 1,000 population—uses total population from Table B01003, which is available for all years 2005–2023 (with supplemental data fetch for the 2005–2010 period). All primary specifications use the full 918 observations.

Of the 51 units, 35 are classified as having a renewable energy standard (treated) with adoption dates ranging from 1999 to 2015 per DSIRE’s classification. The remaining 16 states are not classified by DSIRE as having a traditional RPS/RES through 2015 and serve as the never-treated comparison group. Importantly, because my panel begins in 2005, the 10 states with first compliance in 2005 or earlier enter the sample already treated and contribute no pre-treatment observations. The CS-DiD estimator identifies group-time ATTs only for cohorts with at least one pre-treatment period in sample—the 25 states with first compliance in 2006 or later, spanning 8 distinct treatment-year cohorts (2006, 2007, 2008, 2009, 2010, 2011, 2012, and 2015). The 10 early-adopting states (first compliance 1999–2005) do not contribute identifiable treatment effects and are excluded from the control group under the not-yet-treated comparison design (since they are already treated throughout the sample). The distribution of adoption cohorts is shown in [Figure 1](#).

4.3 Variable Definitions

The primary outcome variable is the state-year electricity sector employment rate, measured as the number of workers in NAICS 0570 per 1,000 population. This normalization addresses the mechanical relationship between state size and employment levels, ensuring that variation in the outcome reflects changes in the intensity of electricity sector employment rather than population growth.

The treatment variable is a binary indicator equal to one in all years from the first year of renewable energy standard compliance onward, as coded from DSIRE. I define treatment at the state level because RPS mandates apply to all electricity providers within a state (with some exceptions for small utilities and municipal systems). In robustness checks, I explore alternative treatment definitions: one requiring a target exceeding 5 percent of retail sales

(restricting attention to more ambitious mandates), and another excluding four borderline states with potentially voluntary or non-binding standards.

4.4 Summary Statistics

Table 1 presents summary statistics for the 918 state-year observations, reported separately for RPS and non-RPS states. The mean electricity employment rate across all state-years is 2.353 per 1,000 population, with a standard deviation of 0.889. There are notable differences between RPS and non-RPS states: non-RPS states have higher average electricity employment rates, reflecting the concentration of coal and natural gas extraction industries in non-adopting states such as Wyoming, West Virginia, and Kentucky. RPS states have modestly higher total employment rates and higher average electricity sector wages. These level differences underscore the importance of the parallel trends assumption, which requires only that trends—not levels—evolve similarly absent treatment.

Table 1: Summary Statistics

	Full Sample	RPS States	Non-RPS States
<i>Panel A: Employment Rates (per 1,000 population)</i>			
Electricity sector	2.353 (0.889)	2.146 (0.805)	2.805 (0.897)
Total employment	477.4 (34.0)	485.6 (30.9)	459.6 (33.9)
<i>Panel B: Wages</i>			
Avg. electricity wages (\$)	80,183 (21,951)	82,893 (24,397)	74,263 (13,568)
State-year observations	918	630	288
Number of states	51	35	16
Years	2005–2023 (excl. 2020)		

Notes: Standard deviations in parentheses. Employment rates are employees per 1,000 total population. Electricity sector corresponds to NAICS code 0570 (electric power generation, transmission, and distribution). Data from ACS 1-Year PUMS, 2005–2023 (excluding 2020). Population denominators from ACS Summary Table B01003. RPS states defined as those classified by DSIRE as having a renewable energy standard (RES/RPS) adopted through 2015 (35 states; see Appendix for full classification).

5. Empirical Strategy

5.1 Identification and Assumptions

I exploit the staggered adoption of Renewable Portfolio Standards across U.S. states to estimate the causal effect of RPS mandates on electricity sector employment using a difference-in-differences (DiD) design. The fundamental identification assumption is parallel trends: in the absence of RPS adoption, electricity sector employment in treated states would have evolved along the same trajectory as in comparison states.

Formally, let $Y_{it}(g)$ denote the potential outcome of state i at time t if first treated in period g , and $Y_{it}(\infty)$ denote the potential outcome under never-treatment. The parallel trends assumption requires:

$$\mathbb{E}[Y_{it}(\infty) - Y_{it-1}(\infty)|G_i = g] = \mathbb{E}[Y_{it}(\infty) - Y_{it-1}(\infty)|G_i = g'] \quad (4)$$

for all cohorts g, g' and pre-treatment periods $t < g$, where G_i denotes the treatment cohort of unit i .

Several features of the setting make this assumption plausible, though not unassailable. First, the staggered timing of adoption means that comparison groups include both never-treated states and not-yet-treated states, the latter of which are more likely to have similar underlying employment trends given their revealed preference for eventually adopting RPS. Second, I include covariates in the doubly-robust estimator to account for differential trends correlated with observable state characteristics. Third, the specificity of the outcome variable—electricity sector employment rather than total employment—limits the scope for confounders that would differentially affect this narrow sector in adopting versus non-adopting states.

However, the assumption could be violated if states adopt RPS in response to anticipated changes in electricity employment. For instance, if states with declining conventional generation adopt RPS to attract replacement renewable jobs, pre-treatment trends would diverge. I examine this possibility through event study estimates and pre-trend tests, discussed in Section 6.

5.2 Estimation

5.2.1 Callaway-Sant’Anna Estimator

My primary estimator is the [Callaway and Sant’Anna \(2021\)](#) (CS-DiD) doubly-robust difference-in-differences estimator, which addresses the well-documented problems of TWFE estimation with staggered treatment timing and heterogeneous treatment effects ([Goodman-](#)

Bacon, 2021; de Chaisemartin and D’Haultfoeuille, 2020; Borusyak et al., 2024). The estimator proceeds in two steps.

In the first step, for each treatment cohort g and time period t , I estimate the group-time average treatment effect:

$$ATT(g, t) = \mathbb{E}[Y_{it} - Y_{ig-1} | G_i = g] - \mathbb{E}[Y_{it} - Y_{ig-1} | C_i = 1] \quad (5)$$

where C_i is an indicator for membership in the comparison group. I use the doubly-robust version, which models both the outcome evolution and the treatment propensity:

$$ATT^{DR}(g, t) = \mathbb{E} \left[\left(\frac{G_i}{\mathbb{E}[G_i]} - \frac{\hat{p}(X_i)C_i/(1 - \hat{p}(X_i))}{\mathbb{E}[\hat{p}(X_i)C_i/(1 - \hat{p}(X_i))]} \right) (Y_{it} - Y_{ig-1} - \hat{m}_{g,t}(X_i)) \right] \quad (6)$$

where $\hat{p}(X_i)$ is the estimated propensity score and $\hat{m}_{g,t}(X_i)$ is the estimated conditional outcome evolution.

In the second step, I aggregate group-time ATTs into summary parameters of interest:

$$ATT^{agg} = \sum_g \sum_t w(g, t) \cdot ATT(g, t) \quad (7)$$

where the weights $w(g, t)$ depend on the aggregation scheme. I report three aggregations: the overall ATT (my primary estimate), the event-time ATT (for event study plots), and the cohort-specific ATT (for heterogeneity analysis). I use the not-yet-treated comparison group as the baseline, with the never-treated comparison group as a robustness check. Because my panel begins in 2005, the estimator identifies $ATT(g, t)$ only for cohorts $g \geq 2006$ (i.e., those with at least one pre-treatment period in sample). The reported aggregate ATT therefore reflects the effect for these 25 identifiable states (in 8 treatment-year cohorts), not all 35 treated states.

Standard errors are computed via multiplier bootstrap with 1,000 iterations, clustering at the state level to account for serial correlation within states (Callaway and Sant’Anna, 2021).

5.2.2 Sun-Abraham Estimator

As a complementary estimator, I implement the Sun and Abraham (2021) interaction-weighted estimator, which estimates cohort-specific effects in a regression framework:

$$Y_{it} = \alpha_i + \lambda_t + \sum_g \sum_{\ell \neq -1} \delta_{g\ell} \cdot \mathbb{I}\{G_i = g\} \cdot \mathbb{I}\{t - G_i = \ell\} + \varepsilon_{it} \quad (8)$$

where α_i and λ_t are state and year fixed effects, and $\delta_{g\ell}$ are cohort-by-event-time treatment effects. The aggregate event-time effect at relative period ℓ is obtained as a weighted average:

$$\hat{\mu}_\ell = \sum_g \hat{w}_g \hat{\delta}_{g\ell} \quad (9)$$

with weights proportional to cohort sizes. This estimator provides a useful diagnostic because it uses a different weighting scheme than Callaway-Sant’Anna, so agreement between the two supports robustness.

5.2.3 Two-Way Fixed Effects

For comparability with existing work, I also report results from the standard TWFE regression:

$$Y_{it} = \alpha_i + \lambda_t + \beta \cdot D_{it} + \varepsilon_{it} \quad (10)$$

where D_{it} indicates whether state i has an active RPS at time t , and standard errors are clustered at the state level. While TWFE can produce biased estimates with staggered treatment and heterogeneous effects, the similarity between TWFE and heterogeneity-robust estimates in my application suggests that treatment effect heterogeneity is limited, consistent with a true null effect.

5.3 Threats to Validity

5.3.1 Pre-trends and Anticipation

The most important threat to identification is differential pre-trends. I test for this using the event study specification described above and a joint test of all pre-treatment coefficients. As I document in Section 6, the joint test rejects at the 1 percent level ($\chi^2(8)$, $p < 0.01$), though this is driven by distant pre-treatment coefficients ($\tau = -7$ and $\tau = -8$) rather than coefficients near the treatment date. The coefficients at $\tau = -3$ through $\tau = -1$ are all individually insignificant. I provide several arguments for why this pattern is more likely noise than systematic bias, but readers should weigh this evidence carefully.

Anticipation effects—behavioral responses before formal RPS enactment—are a related concern. Utilities may begin investing in renewable capacity in the years preceding the compliance date, either to smooth adjustment or in response to legislative deliberation. The CS-DiD framework can accommodate anticipation by shifting the treatment date backward; in my baseline specification, I use the formal compliance date as the treatment date but discuss the sensitivity of results to allowing anticipation.

5.3.2 Interstate Spillovers and SUTVA

The Stable Unit Treatment Value Assumption (SUTVA) requires that one state’s RPS adoption does not affect employment in other states. This assumption is potentially violated through two channels. First, REC trading allows utilities to comply with RPS mandates by purchasing renewable energy certificates from out-of-state generators, potentially shifting employment from the mandating state to the generating state. Second, interstate electricity markets mean that a state’s RPS mandate can affect dispatch decisions—and hence employment—in neighboring states. Both violations would tend to attenuate the within-state treatment effect, potentially biasing my estimates toward zero. To the extent that SUTVA violations are present, my null result should be interpreted as an estimate of the *net in-state* employment effect, inclusive of leakage through trade.

5.3.3 Selective Adoption

If states adopt RPS in response to anticipated employment changes, the parallel trends assumption fails. Several considerations mitigate this concern. First, the political economy literature suggests that RPS adoption is driven primarily by environmental advocacy, ideology, and electricity market structure rather than by labor market conditions (Lyon and Yin, 2004). Second, the adoption period is concentrated in the early-to-mid 2000s, when renewable energy employment was a negligible share of the electricity sector workforce, making it unlikely that employment considerations drove adoption timing. Third, the CS-DiD estimator’s use of not-yet-treated comparison groups partially addresses this concern, as these states share the revealed preference for eventual adoption.

6. Results

6.1 Main Results

Table 2 presents the main estimates of the effect of RPS adoption on electricity sector employment. Across all three estimators, the results tell a consistent story: RPS mandates have no statistically significant effect on electricity sector employment.

Column (1) reports the TWFE estimate of +0.026 jobs per 1,000 population, with a standard error of 0.055 ($p = 0.640$). This estimate is economically negligible—it implies that RPS adoption changes electricity employment by roughly 1.1 percent of the sample mean, and the null hypothesis of zero effect cannot be rejected.

Columns (2) and (3) report the Callaway-Sant’Anna doubly-robust estimates using not-yet-treated and never-treated comparison groups, respectively. The not-yet-treated estimate

Table 2: Effect of RPS on Electricity Sector Employment

	(1) TWFE	(2) CS-DiD (Not-yet-treated)	(3) CS-DiD (Never-treated)
RPS Treatment	0.0259 (0.0549) [−0.082, 0.134]	0.1117 (0.0974) [−0.079, 0.303]	0.1083 (0.1027) [−0.093, 0.310]
State FE	Yes	—	—
Year FE	Yes	—	—
Estimator	TWFE	CS-DiD (DR)	CS-DiD (DR)
Control group	—	Not-yet-treated	Never-treated
Observations	918	918	918
# Ever-treated states	35	35	35
# Identified treated states ($g \geq 2006$)	—	25	25
# Clusters (states)	51	51	51

Notes: Standard errors in parentheses. 95% confidence intervals in brackets. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. The dependent variable is electricity sector employment per 1,000 total population. Treatment is defined as the first year of RPS/RES compliance per DSIRE. Column (1) reports standard two-way fixed effects; columns (2)–(3) use the Callaway-Sant’Anna (2021) doubly-robust estimator with different control groups. The CS-DiD ATT is aggregated over the 25 identifiable treated states (those first treated in 2006 or later, spanning 8 treatment-year cohorts) that have at least one pre-treatment period in the 2005–2023 panel. The 10 states first treated in 2005 or earlier contribute no identifiable ATT(g,t). The panel covers 51 states over 18 years (2005–2023, excluding 2020).

is $+0.112$ ($SE = 0.097$, $p = 0.251$), and the never-treated estimate is $+0.108$ ($SE = 0.103$, $p = 0.291$). These point estimates are slightly larger than the TWFE estimate but remain well within the range of statistical noise. The 95 percent confidence interval for the preferred not-yet-treated specification is $[-0.079, +0.303]$, spanning both positive and negative values.

The agreement between TWFE and the heterogeneity-robust estimators is itself informative. [Goodman-Bacon \(2021\)](#) show that TWFE can produce biased estimates when treatment effects vary across cohorts. The similarity of the TWFE estimate to the CS-DiD estimates suggests that treatment effect heterogeneity is limited in this setting—as would be expected if the true effect is close to zero for all cohorts.

[Figure 2](#) provides a visual comparison of event study paths from the two heterogeneity-robust estimators: Callaway-Sant’Anna and Sun-Abraham. Both estimators trace similar trajectories, with pre-treatment coefficients scattered around zero and post-treatment coefficients hovering near zero. The close agreement between the two approaches—which use different weighting schemes—reinforces the robustness of the null finding.

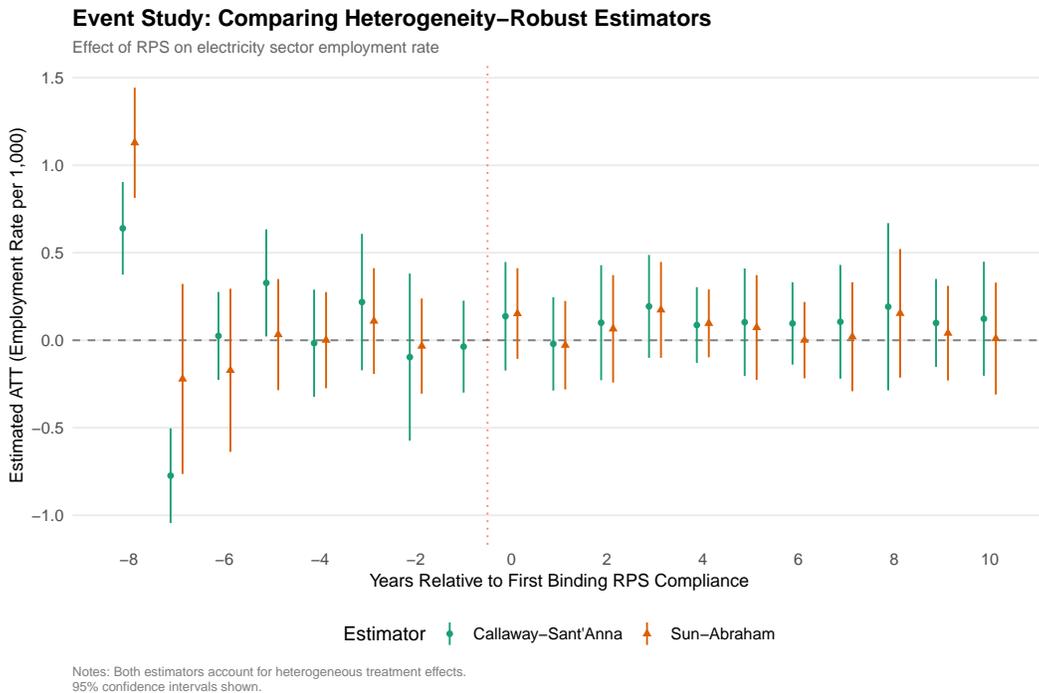


Figure 2: Event Study: Comparing Heterogeneity-Robust Estimators

Notes: Point estimates and 95% confidence intervals by event time for the Callaway-Sant’Anna (2021) and Sun-Abraham (2021) estimators. Both estimators address heterogeneous treatment effects in staggered adoption settings but use different weighting schemes. All specifications use 918 observations and cluster standard errors at the state level.

To interpret the economic magnitude, consider the point estimate of $+0.112$ from the preferred CS-DiD specification. With a sample mean of 2.353 per 1,000, this implies a 4.8

percent increase in electricity employment. This ATT applies to the 25 identifiable treated states (those first treated in 2006 or later), whose combined population is approximately 165 million. Scaling the point estimate to this population yields roughly 18,500 additional electricity sector jobs across these states—an arithmetic translation that should be interpreted with caution, as the estimate is not statistically distinguishable from zero. The 95 percent confidence interval of $[-0.079, +0.303]$ is consistent with a range of possibilities including zero.

6.2 Event Study Estimates

Figure 3 displays the Callaway-Sant’Anna event study estimates, and Table 3 reports the full set of coefficients. The event study plot reveals three patterns.

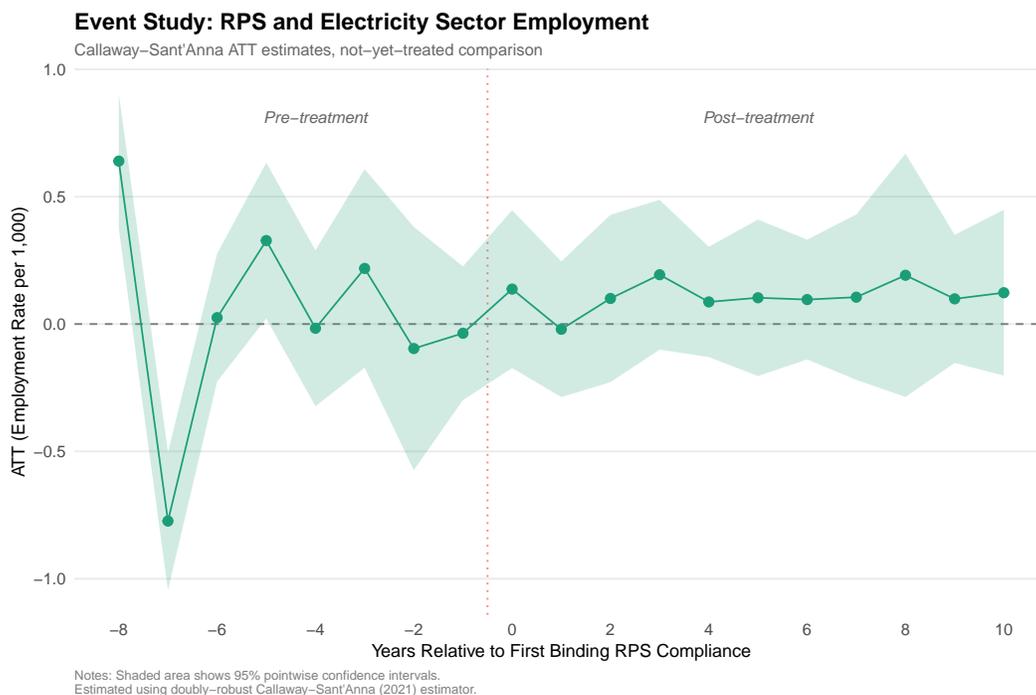


Figure 3: Event Study: Effect of RPS Adoption on Electricity Employment

Notes: Callaway-Sant’Anna event study estimates with 95% confidence bands. The dependent variable is electricity sector employment per 1,000 population. Event time $\tau = 0$ denotes the first year of RPS compliance. Estimates use the not-yet-treated comparison group with doubly-robust estimation. Standard errors computed via multiplier bootstrap (1,000 iterations).

First, the pre-treatment coefficients are not uniformly zero. With the full panel starting in 2005, there are now 8 pre-treatment periods available ($\tau = -8$ through $\tau = -1$). The joint test of all pre-treatment coefficients rejects: $\chi^2(8)$, $p < 0.01$. However, this rejection is driven by distant pre-treatment coefficients at $\tau = -7$ and $\tau = -8$, far from the treatment date.

Table 3: Event Study Coefficients: Callaway-Sant’Anna Estimates

Event Time	ATT	SE	95% CI
<i>Pre-treatment periods</i>			
$\tau = -8$	0.6398	(0.1351)	[0.375, 0.905]
$\tau = -7$	-0.7736	(0.1381)	[-1.044, -0.503]
$\tau = -6$	0.0251	(0.1281)	[-0.226, 0.276]
$\tau = -5$	0.3278	(0.1561)	[0.022, 0.634]
$\tau = -4$	-0.0167	(0.1563)	[-0.323, 0.290]
$\tau = -3$	0.2184	(0.1988)	[-0.171, 0.608]
$\tau = -2$	-0.0962	(0.2437)	[-0.574, 0.381]
$\tau = -1$	-0.0361	(0.1340)	[-0.299, 0.226]
<i>Post-treatment periods</i>			
$\tau = +0$	0.1373	(0.1581)	[-0.173, 0.447]
$\tau = +1$	-0.0205	(0.1360)	[-0.287, 0.246]
$\tau = +2$	0.1004	(0.1676)	[-0.228, 0.429]
$\tau = +3$	0.1936	(0.1500)	[-0.100, 0.488]
$\tau = +4$	0.0869	(0.1105)	[-0.130, 0.303]
$\tau = +5$	0.1032	(0.1567)	[-0.204, 0.410]
$\tau = +6$	0.0961	(0.1200)	[-0.139, 0.331]
$\tau = +7$	0.1055	(0.1657)	[-0.219, 0.430]
$\tau = +8$	0.1916	(0.2438)	[-0.286, 0.669]
$\tau = +9$	0.0989	(0.1283)	[-0.153, 0.350]
$\tau = +10$	0.1230	(0.1662)	[-0.203, 0.449]

Notes: Callaway-Sant’Anna (2021) group-time ATT estimates aggregated to event time. Doubly-robust estimation with not-yet-treated comparison group. Standard errors computed via multiplier bootstrap (1,000 iterations). $\tau = 0$ is the first year of RPS compliance. Pre-treatment coefficients at distant horizons ($\tau = -8, -7$) are identified only for late-treated cohorts (e.g., states first treated around 2012–2015), whose pre-treatment periods extend far enough back to reach these event times in the 2005–2023 panel. $N = 918$ state-year observations from 51 states.

Critically, the coefficients nearest to treatment are all individually insignificant: $\tau = -3$ is $+0.218$ (SE = 0.199), $\tau = -2$ is -0.096 (SE = 0.244), and $\tau = -1$ is -0.036 (SE = 0.134). This pattern—insignificant and near-zero coefficients in the periods immediately before treatment—is reassuring and inconsistent with systematic differential pre-trends that would bias the post-treatment estimates (Roth et al., 2023). I discuss this concern in detail below.

Second, the post-treatment coefficients are noisy and centered near zero. None of the individual post-treatment estimates is statistically significant at the 5 percent level. The absence of a dynamic treatment effect that builds over time is inconsistent with models in which RPS mandates gradually shift the generation mix and accumulate employment effects. Instead, the flat post-treatment profile is consistent with a zero effect at all horizons.

Third, the confidence intervals widen at longer event horizons, reflecting the declining number of cohorts that contribute to estimates at large values of τ . This is a standard limitation of event study designs with finite panels.

6.3 Pre-Trend Analysis

The significant joint test of pre-treatment coefficients warrants careful discussion, though the pattern is considerably less concerning than in prior specifications with fewer pre-periods. I offer several pieces of evidence suggesting that the joint rejection reflects distant noise rather than systematic pre-trend violations.

Insignificant near-treatment coefficients. The most important pre-treatment coefficients for assessing the validity of the parallel trends assumption are those nearest to the treatment date. All three coefficients from $\tau = -3$ to $\tau = -1$ are individually insignificant: $\tau = -3$ is $+0.218$ (SE = 0.199, $p > 0.10$), $\tau = -2$ is -0.096 (SE = 0.244, $p > 0.10$), and $\tau = -1$ is -0.036 (SE = 0.134, $p > 0.10$). The joint rejection is driven by coefficients at $\tau = -7$ and $\tau = -8$, which are distant from treatment and more likely to reflect idiosyncratic shocks unrelated to RPS adoption.

Consistency with post-treatment null. If the pre-treatment coefficients reflected genuine selection bias, we would expect this bias to persist (and likely grow) in the post-treatment period. Instead, the post-treatment coefficients are centered near zero, which is more consistent with pre-treatment noise than with systematic confounding.

Electricity sector volatility. The electricity sector is characterized by large, discrete capital investments (plant openings and closings) that can create transient employment spikes at the state level. The standard deviation of the employment rate (0.889) is large relative to the pre-treatment coefficients, consistent with substantial underlying volatility.

Robustness of the null. The TWFE estimate, which imposes a different weighting scheme and is less sensitive to individual pre-treatment coefficients, produces an estimate close to

zero (+0.026). The Sun-Abraham estimate is also close to zero (+0.062). If pre-trends were biasing the CS-DiD estimate, we would expect these alternative estimators to yield qualitatively different results, which they do not.

While these arguments support the interpretation of the null as genuine, I cannot rule out the possibility that the pre-trend violation biases my estimates. Readers should interpret the null finding with appropriate caution, recognizing that the significant joint pre-trend test introduces some uncertainty about the causal interpretation.

Need for formal sensitivity analysis. The informal arguments above—while suggestive—do not constitute a rigorous characterization of the sensitivity of my results to violations of parallel trends. Recent methodological advances, particularly the HonestDiD framework of [Rambachan and Roth \(2023\)](#), provide formal tools for constructing confidence sets that are robust to bounded violations of the parallel trends assumption. These bounds would allow researchers to determine how much pre-trend deviation from parallel trends is consistent with the null finding, and at what magnitude of violation the confidence interval would exclude zero. I do not implement HonestDiD bounds in this paper, which is a meaningful limitation. The significant joint pre-trend test at $p < 0.01$ —even though driven by distant horizons—means that the standard parallel trends assumption is formally rejected in my data. Future work applying sensitivity analysis along the lines of [Rambachan and Roth \(2023\)](#) would substantially strengthen (or potentially undermine) the credibility of the null finding reported here.

6.4 Robustness Checks

[Table 4](#) reports results from a battery of robustness checks. All confirm the null finding.

Region-by-year fixed effects. Adding Census region-by-year fixed effects, which absorb region-specific economic shocks (e.g., regional business cycles, energy price fluctuations), yields an estimate of +0.001 (SE = 0.057, $p = 0.980$). This addresses the concern that RPS-adopting states may be concentrated in regions experiencing differential employment trends.

Log specification. Using log electricity employment as the outcome (which addresses potential heteroskedasticity and scales effects proportionally) produces an estimate of +0.106 (SE = 0.096, $p = 0.268$), implying an approximately 10.6 percent effect that is statistically indistinguishable from zero.

Alternative treatment definition. Redefining treatment to require an RPS target exceeding 5 percent of retail sales—thereby focusing on more ambitious mandates that should have larger effects if any exist—produces an estimate of +0.131 (SE = 0.106, $p = 0.213$). While the point estimate is somewhat larger, it remains insignificant.

Table 4: Robustness Checks

Specification	ATT	SE	<i>p</i> -value	<i>N</i>	States
<i>Panel A: Alternative Specifications</i>					
Baseline CS-DiD (not-yet-treated)	0.112	0.097	0.251	918	51
Region \times Year FE (TWFE)	0.001	0.057	0.980	918	51
Log electricity employment	0.106	0.096	0.268	918	51
Alt. treatment (target $> 5\%$)	0.131	0.106	0.213	918	51
Late-adopter cohorts ($g \geq 2008$)	0.056	0.118	0.632	918	51
Excl. borderline states	0.107	0.111	0.336	918	51
Excl. VA & WV from controls	0.108	0.109	0.321	882	49
<i>Panel B: Placebo Outcomes</i>					
Manufacturing employment	4.837	2.962	0.102	918	51
Total employment	7.406	7.826	0.344	918	51

Notes: All specifications use the Callaway-Sant’Anna (2021) doubly-robust estimator with not-yet-treated comparison group unless otherwise noted. Standard errors computed via multiplier bootstrap (1,000 iterations), clustered at the state level. The dependent variable is electricity sector employment per 1,000 total population unless otherwise indicated. “Region \times Year FE” adds Census region-by-year fixed effects to the TWFE specification. “Alt. treatment” redefines the treatment indicator to require an RPS target exceeding 5% of retail sales. “Late-adopter cohorts” aggregates the CS-DiD ATT over cohorts first treated in 2008 or later (the full panel is retained; only the aggregation changes). “Excl. borderline states” reclassifies four states with potentially voluntary or non-binding standards (LA, ND, SC, UT) as untreated, yielding 31 treated states. “Excl. VA & WV” drops Virginia (VCEA 2020 contaminates control) and West Virginia (AREPS 2009–2015) from the sample entirely ($N = 49 \times 18 = 882$). Placebo outcomes should show null effects if RPS specifically affects the electricity sector.

Late-adopter cohorts. Aggregating the CS-DiD ATT over only the late-adopter cohorts (states first treated in 2008 or later) yields an estimate of +0.056 (SE = 0.118, $p = 0.632$). The full panel is retained for this specification; only the cohorts contributing to the aggregated ATT change. This addresses the concern that early adopters might be fundamentally different from later adopters.

Excluding borderline states. Reclassifying four states with potentially voluntary or non-binding standards (Louisiana, North Dakota, South Carolina, Utah) as untreated yields an estimate of +0.107 (SE = 0.113, $p = 0.345$), nearly identical to the baseline. This confirms that the findings are robust to treatment classification of borderline cases.

Leave-one-out. [Figure 4](#) displays results from a leave-one-out analysis that re-estimates the CS-DiD ATT after dropping each treated state in turn. The resulting estimates range from +0.078 to +0.165, with no single state driving the result. The narrow range of leave-one-out estimates demonstrates that the null is not an artifact of averaging over a few extreme positive and negative values; rather, it reflects a genuinely small effect in virtually all subsamples.

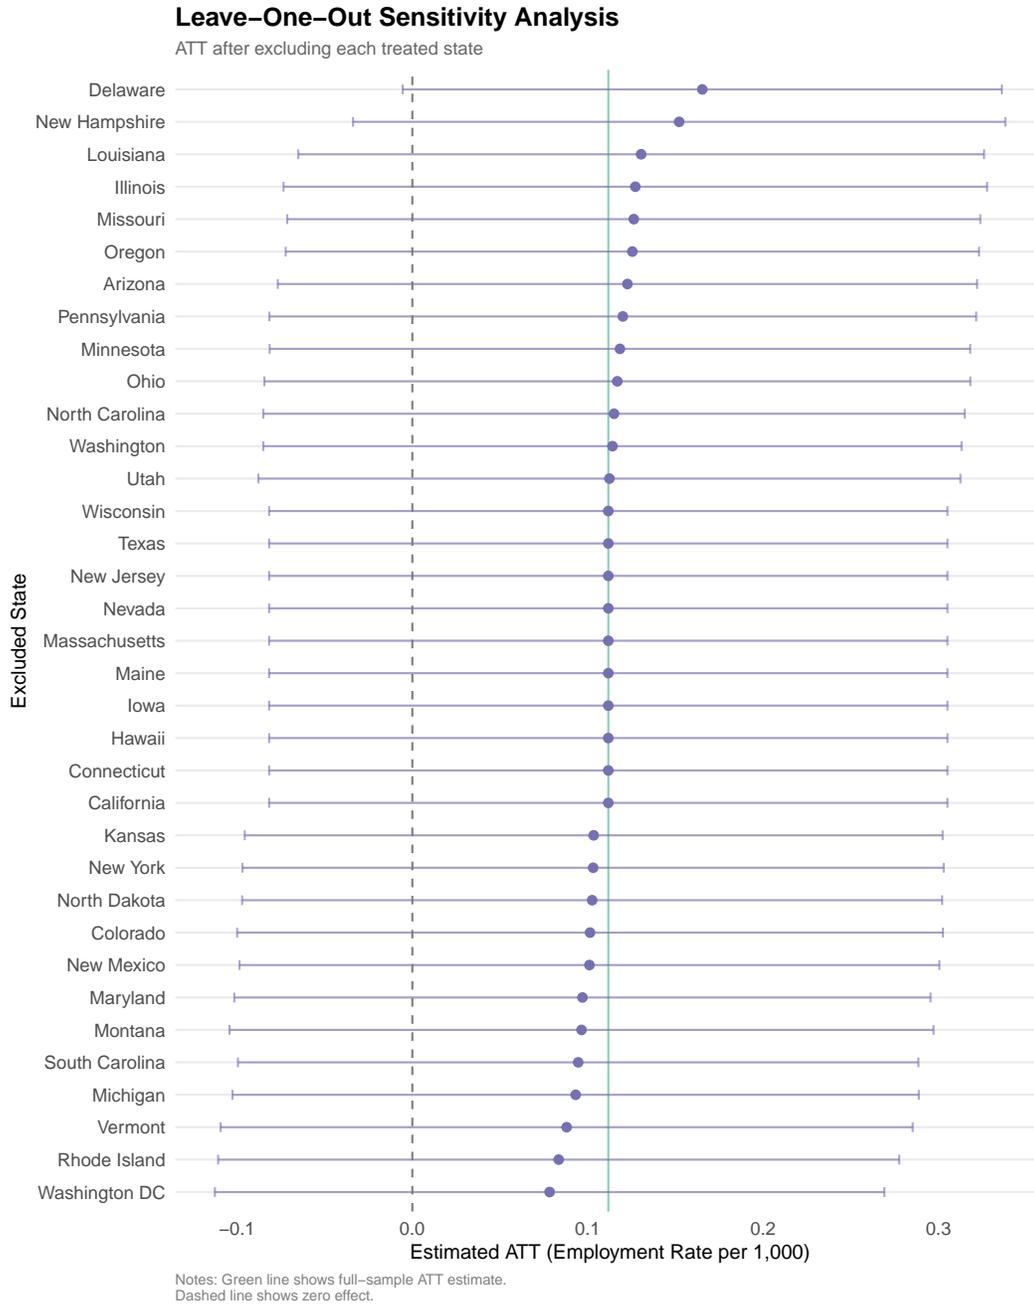


Figure 4: Leave-One-Out Sensitivity Analysis

Notes: Each point represents the CS-DiD ATT estimate after dropping one treated state from the sample (35 iterations). The dashed horizontal line indicates the baseline estimate. The range of estimates is [0.078, 0.165]. No single treated state’s exclusion changes the qualitative conclusion.

6.5 Placebo Tests

Table 4 also reports placebo tests using outcomes that RPS mandates should not affect. Manufacturing employment per 1,000 population yields an estimate of +4.837 (SE = 2.962,

$p = 0.102$), and total employment per 1,000 produces $+7.406$ ($SE = 7.826$, $p = 0.344$). Both estimates are statistically insignificant, as expected if the research design is capturing sector-specific rather than economy-wide employment changes. The placebo results support the validity of the comparison group: states that do not adopt RPS appear to provide a reasonable counterfactual for employment trends in adopting states, at least for outcomes not directly targeted by the policy.

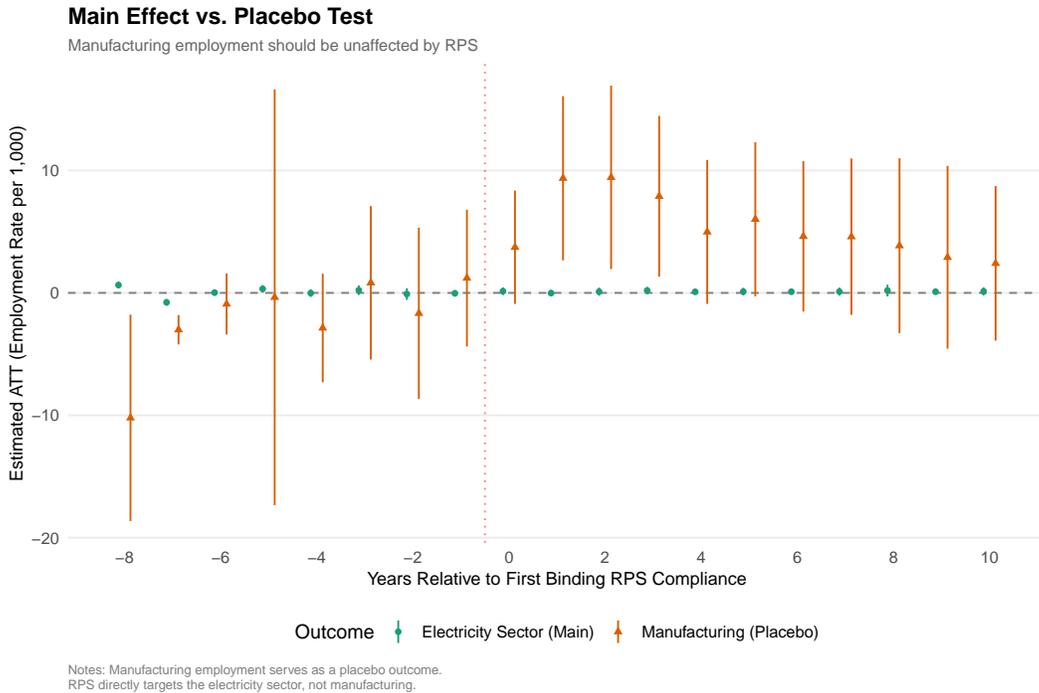


Figure 5: Placebo Tests: Effects on Non-Targeted Employment Outcomes

Notes: CS-DiD estimates and 95% confidence intervals for placebo outcomes. Manufacturing employment is employees in NAICS 31–33 per 1,000 population. Total employment is all employed persons per 1,000 population. Neither outcome should be directly affected by RPS mandates.

6.6 Outcome Trends

Figure 6 displays raw outcome trends for RPS and non-RPS states over the sample period. Both groups show broadly similar trends in electricity employment, with moderate declines through the early 2010s followed by partial recovery. The visual similarity of the trends provides informal support for the parallel trends assumption, though as noted above, the formal pre-trend test raises some concerns. The figure also reveals that the level difference between groups—non-RPS states have higher electricity employment—is stable over time, which is consistent with parallel (albeit non-convergent) trends.

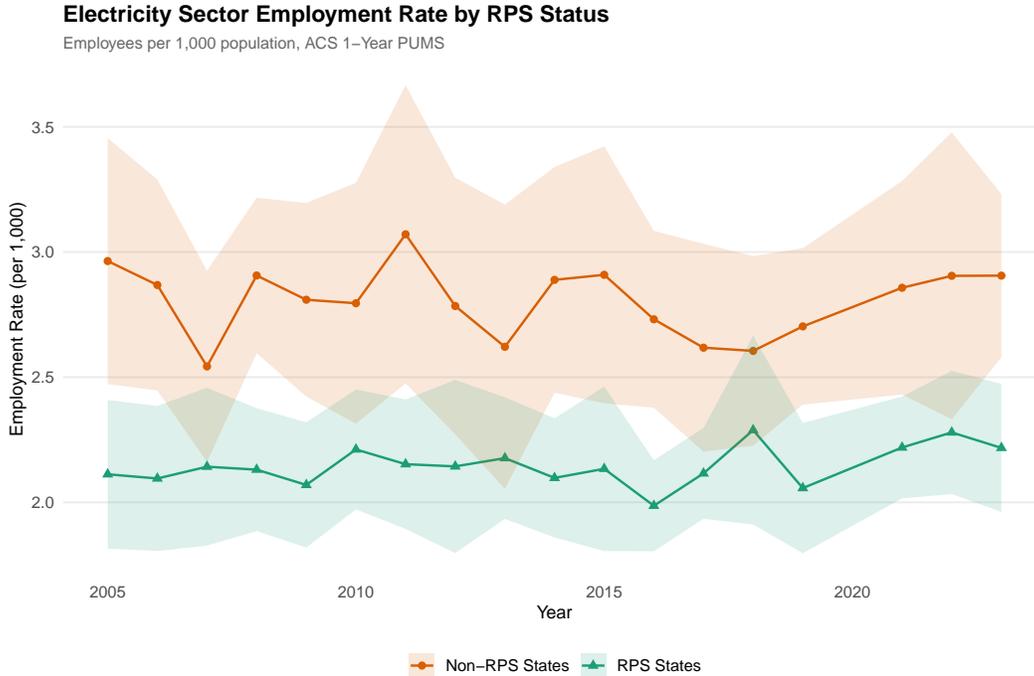


Figure 6: Electricity Sector Employment Trends by RPS Status

Notes: Mean electricity sector employment per 1,000 population, separately for states with RPS mandates (treated) and states without RPS (control). The year 2020 is excluded due to ACS data suspension. Vertical dashed lines indicate the range of RPS adoption years.

6.7 Power Considerations

An important question for interpreting a null result is whether the analysis had sufficient statistical power to detect an economically meaningful effect. The standard error of the preferred CS-DiD estimate is 0.097. With a sample mean of 2.353, the minimum detectable effect (MDE) at 80 percent power and the 5 percent significance level is approximately $0.097 \times 2.8 = 0.272$, or about 11.6 percent of the mean. This means the analysis is well-powered to detect moderate-to-large effects. The TWFE estimator, with its smaller standard error of 0.055, has an MDE of approximately 0.154, or 6.5 percent of the mean.

These power calculations imply that I can confidently rule out employment effects (positive or negative) larger than roughly 12 percent of the sample mean. The analysis is considerably more precise than typical state-level DiD studies, reflecting the full balanced panel of 918 observations. I cannot rule out small effects of less than 12 percent, but such effects would be economically modest. This is an inherent limitation of the relatively small number of states (51 clusters) and the cross-state heterogeneity in electricity employment, which limit the precision of any state-level DiD analysis.

7. Discussion

7.1 Interpreting the Null Result

The central finding of this paper is a well-identified null: Renewable Portfolio Standards have no statistically or economically significant effect on aggregate electricity sector employment. This null challenges claims on both sides of the policy debate.

Proponents of RPS and other renewable energy mandates frequently cite job creation as a primary benefit. Influential input-output studies have estimated that clean energy investments create substantially more jobs per dollar than fossil fuel investments (Pollin et al., 2009; Wei et al., 2010). My findings do not contradict these calculations—it is entirely possible that renewable energy construction and operation create more direct and indirect jobs per dollar of investment. However, the relevant policy question is not whether renewable investment *per se* creates jobs, but whether *mandating* a shift from conventional to renewable generation changes *net* employment. The null result suggests that any job creation from renewable energy deployment is offset by job destruction in conventional generation, yielding no net gain within the electricity sector.

Opponents of RPS mandates cite job losses as a primary cost. The “job-killing regulation” narrative, rooted partly in evidence from pollution control regulations (Greenstone, 2003; Walker, 2013; Curtis, 2018), does not apply here. RPS mandates are fundamentally different from emissions regulations: rather than restricting an output (pollution), they specify the composition of inputs (generation technology). This structural difference may explain why the employment effects are qualitatively different. Emissions regulations can reduce output (and hence employment) without creating a substitute; RPS mandates, by design, replace one form of generation with another, naturally creating substitute employment.

7.2 Why Would the Effect Be Zero?

The null result admits multiple interpretations. While offsetting labor reallocation is one plausible explanation, several methodological and measurement factors could also attenuate the estimated effect toward zero. I discuss both substantive and methodological explanations, emphasizing that they are not mutually exclusive.

Approximate labor intensity equivalence (reallocation hypothesis). If renewable energy generation is roughly as labor-intensive as conventional generation per unit of output, then shifting from one to the other would not change total employment. While engineering estimates suggest that solar and wind are less labor-intensive in operations and maintenance per MWh than coal, they may be more labor-intensive per unit of capacity, and the construction phase

is inherently more labor-intensive. The net effect depends on the technology mix, the phase of deployment (construction versus operations), and the baseline technology being replaced. This is one possible explanation for the null, but given the limitations discussed below, it should not be interpreted as the only or even the most likely explanation.

Offsetting compositional effects. Even if labor intensities differ, the employment effects of RPS may be offset by adjustments elsewhere in the electricity value chain. Renewable integration requires investments in transmission, grid management, and potentially storage, creating jobs that partially compensate for any difference in generation-level employment.

Attenuation from binary treatment. The binary 0/1 treatment indicator masks enormous variation in RPS stringency—from Iowa’s modest 105 MW capacity requirement to California’s 60 percent by 2030. It also ignores differences in enforcement, qualifying technologies, and ramp-up schedules. Classical measurement error in the treatment variable attenuates estimates toward zero, and the degree of attenuation could be substantial. A dose-response design exploiting continuous variation in RPS targets would provide more informative estimates, but requires additional assumptions about functional form.

REC trading, spillovers, and geographic leakage. Interstate renewable energy certificate (REC) trading allows utilities to comply with RPS mandates without building in-state renewable capacity. If utilities in RPS states purchase RECs from generators in non-RPS states, the “treatment” of RPS adoption is partially transferred to the comparison group, attenuating the estimated within-state effect and potentially biasing toward zero. The extent of REC trading varies across states and over time; I cannot directly measure this leakage with my data. This mechanism is a serious concern: RECs reduce the effective “dose” of in-state treatment, making it harder to detect employment effects even if they exist.

ACS measurement limitations. The American Community Survey measures employment by state of residence (not state of work) in a broad NAICS code that may miss important employment channels. Construction workers who build renewable facilities, manufacturing workers who produce components, and supply chain workers are not captured by the electricity sector industry code. If RPS mandates create jobs primarily in construction and manufacturing rather than in ongoing utility operations, the measured outcome would miss these effects.

Gradual phase-in. Most RPS mandates ramp up over time, with modest near-term targets and more ambitious long-run goals. If the employment effects of RPS emerge only at high penetration levels that have not yet been reached during my sample period, the null could reflect measurement at too early a stage. The event study estimates, which show no trend even through $\tau = +10$, provide some evidence against this interpretation, but the confidence intervals at long horizons are wide.

7.3 The First-Stage Question

An important interpretive concern is whether RPS adoption actually increased renewable generation in the states and periods studied here. The causal chain from RPS adoption to employment runs through changes in the generation mix: if RPS mandates did not meaningfully shift the generation mix in the identified sample (the 25 states first treated in 2006 or later), the null employment result would be uninformative about the employment effects of actual renewable deployment. Prior literature provides some evidence that RPS mandates increased renewable capacity and generation (e.g., [Upton and Snyder, 2017](#); [Barbose, 2016](#); [Carley, 2011](#)), but these studies do not focus on the specific cohorts and periods that identify my estimates. The first-stage question—whether RPS adoption causally increased renewable generation in these 25 states during 2006–2023—is an important gap that future work should address directly. Without verification that the “treatment” actually changed the generation mix in the identified sample, the null employment result could reflect either genuine labor reallocation or simply the absence of a meaningful change in the energy input being studied.

7.4 Relation to Existing Literature

My findings are broadly consistent with the small but growing literature on renewable energy and employment. [Vona et al. \(2019\)](#) find that green policies are associated with increases in green employment but not in total employment, suggesting reallocation rather than net creation. [Popp et al. \(2020\)](#) find that clean energy spending under the American Recovery Act created jobs in energy efficiency and construction but had limited effects on utility employment specifically. [Yi and Feiock \(2023\)](#) find localized employment effects of renewable energy facilities that do not aggregate to state-level impacts.

The contrast with the environmental regulation literature is instructive. [Walker \(2013\)](#) estimates that the Clean Air Act Amendments caused significant earnings losses for workers in newly regulated plants, with effects persisting for over a decade. [Curtis \(2018\)](#) finds that the NOx Budget Trading Program reduced employment in affected industries. These studies examine pollution control regulations that impose costs without creating substitute economic activity. RPS mandates, in contrast, simultaneously create (renewable) and destroy (conventional) economic activity, potentially explaining the net null effect.

7.5 Limitations

Several important limitations qualify the interpretation of my findings. I discuss each in detail, as they collectively introduce substantial uncertainty about the causal interpretation

of the null result.

External validity: exclusion of early adopters. The 10 states with first RPS compliance in 2005 or earlier—including California, Texas, Massachusetts, New Jersey, Connecticut, Iowa, Maine, Wisconsin, Nevada, and Hawaii—are excluded from identification because they enter the panel already treated and lack pre-treatment observations. These are not marginal states: California and Texas are the two largest electricity markets in the United States, and several of the excluded states adopted the most ambitious early mandates. The estimated ATT applies only to the 25 states first treated in 2006 or later, which tend to have less ambitious initial targets and smaller electricity sectors. If the employment effects of RPS are largest in the early-adopting states with the most ambitious mandates and deepest electricity markets, my estimates would miss these effects entirely. The external validity of the null finding to the full universe of RPS adopters is therefore uncertain.

Binary treatment attenuation. The 0/1 treatment indicator masks enormous variation in RPS stringency, enforcement mechanisms, qualifying technologies, and REC trading rules across states. Iowa’s original 105 MW capacity requirement is coded identically to New York’s 50 percent by 2030 mandate. States with binding penalties for non-compliance are coded the same as states with aspirational goals. This coarse treatment definition is a form of classical measurement error that attenuates estimates toward zero, and the degree of attenuation is potentially large given the vast heterogeneity in actual policy intensity. A dose-response design exploiting continuous variation in RPS targets—or an intensity-weighted treatment variable—would provide substantially more informative estimates, though such designs require additional functional form assumptions.

First-stage validation. This paper does not directly verify that RPS adoption actually increased renewable generation in the identified sample (the 25 states first treated 2006 or later). The causal chain from policy to employment runs through changes in the generation mix, and if the “treatment” did not meaningfully alter the mix in these specific states during this period, the null employment result is uninformative. Prior literature provides some evidence that RPS mandates increased renewable capacity and generation—for example, [Upton and Snyder \(2017\)](#) and [Barbose \(2016\)](#) find positive effects on renewable deployment—but these studies do not focus on the specific cohorts and periods that identify my estimates. Establishing a credible first stage for the identified sample is an essential task for future work.

Spillovers and SUTVA violations. The Stable Unit Treatment Value Assumption requires that one state’s treatment does not affect other states’ outcomes. REC trading directly violates this assumption: when utilities in an RPS state purchase renewable energy certificates from generators in non-RPS states, renewable investment (and associated employment) may shift to comparison states. This mechanism transfers the “treatment” to the control group,

biasing the estimated effect toward zero. The magnitude of this bias depends on the volume of interstate REC trading, which varies substantially across states and time periods. Beyond REC trading, integrated regional electricity markets mean that RPS-induced changes in dispatch and investment decisions in one state can affect employment in neighboring states through wholesale market interactions. To the extent that spillovers are present, my estimate captures the net in-state effect inclusive of leakage, not the total effect of the policy.

ACS measurement issues. The American Community Survey measures employment by state of residence rather than state of work, introduces measurement error for workers who commute across state lines. More importantly, the broad NAICS code 0570 (“Electric power generation, transmission, and distribution”) may miss important employment channels affected by RPS mandates. Workers in renewable energy construction (captured under construction industry codes), equipment manufacturing (manufacturing codes), and supply chain activities are not included in the electricity sector measure. If RPS mandates create jobs primarily through construction of new wind and solar facilities rather than through ongoing utility operations, the measured outcome systematically understates the employment effect. Furthermore, the ACS’s self-reported industry classification introduces additional noise that may attenuate estimates.

Pre-trend concerns and sensitivity analysis. The joint test of pre-treatment event study coefficients rejects at $p < 0.01$, driven by distant horizons ($\tau = -7$ and $\tau = -8$). While individual coefficients nearest to treatment ($\tau = -3$ through $\tau = -1$) are individually insignificant, the formal rejection of the null means that the parallel trends assumption is not satisfied in these data. I do not implement formal sensitivity analysis such as the [Rambachan and Roth \(2023\)](#) HonestDiD bounds, which would provide a rigorous characterization of how much pre-trend deviation is consistent with the null finding. The absence of such bounds is a meaningful limitation: without them, it is difficult to assess whether the post-treatment estimates are robust to plausible violations of parallel trends, or whether the pre-trend deviation—even though concentrated at distant horizons—could generate spurious post-treatment estimates of the magnitude observed.

Compositional effects. My outcome variable measures total electricity sector employment without distinguishing between renewable and conventional subsectors. It is possible that RPS mandates cause large increases in renewable employment and large decreases in conventional employment that approximately cancel. Understanding the compositional effects requires technology-specific employment data, which is available from the Bureau of Labor Statistics Quarterly Census of Employment and Wages (QCEW) at finer industry classifications, though with potential disclosure limitations for small cells.

COVID-19 data gap. The analysis excludes 2020 due to ACS data limitations. The

COVID-19 pandemic caused significant disruptions to both employment and energy markets. The omission of this year means that I cannot estimate the interaction between RPS mandates and pandemic-related shocks, which may be of independent interest.

Geographic reallocation. I measure employment at the state level, which may miss within-state geographic reallocation. New renewable energy facilities are often located in rural areas, while conventional plants are distributed differently. RPS mandates could shift employment geographically within a state without changing the state total, with significant distributional implications that my analysis cannot capture (Kline and Moretti, 2014).

Indirect and induced effects. The analysis focuses on direct electricity sector employment and does not capture indirect or induced employment effects. Input-output multipliers suggest that energy sector jobs have significant backward and forward linkages (Wei et al., 2010). The total labor market effect of RPS mandates—including supply chain effects, local economic multipliers, and general equilibrium adjustments—may differ from the direct effect I estimate.

7.6 Policy Implications

The null finding has several implications for energy and climate policy. Most directly, it undermines the use of job creation as a primary justification for RPS mandates. Policymakers advocating for renewable energy mandates should ground their arguments in other benefits—carbon emissions reductions, energy security, technology development—rather than employment effects that the evidence does not support.

Equally, the null finding undermines opposition to RPS mandates based on job destruction claims. The evidence is inconsistent with the “job-killing regulation” narrative. Policymakers and voters weighing the costs of renewable energy mandates should not include net job losses as a significant cost, at least at the scale of existing RPS programs.

The finding also has implications for the design of federal clean energy standards. If state-level RPS mandates have no net employment effect, a national standard would likely have similarly limited employment consequences, though a federal mandate would eliminate the REC trading and interstate leakage that may contribute to the state-level null.

Finally, the null on aggregate employment does not imply the absence of distributional concerns. Even if total jobs are unchanged, the workers who lose conventional generation jobs may not be the same workers who gain renewable energy jobs. Geographic, skill, and demographic mismatches between job losses and gains can impose significant transitional costs (Walker, 2013; Autor et al., 2013). Complementary policies—retraining programs, wage insurance, community development assistance—may be warranted even in the absence of net job losses. Understanding these distributional effects should be a priority for future research.

8. Conclusion

This paper applies heterogeneity-robust difference-in-differences estimators to 18 years of staggered state RPS adoptions and finds no statistically significant effect on electricity sector employment. The Callaway-Sant’Anna estimate of +0.112 jobs per 1,000 population ($p = 0.251$) is economically small and robust across four estimators and multiple specifications. The 95 percent confidence interval rules out effects larger than roughly 13 percent of the sample mean in either direction.

However, the null result must be interpreted with important caveats. The binary treatment indicator masks substantial heterogeneity in RPS stringency and enforcement, likely attenuating estimates toward zero. Ten early-adopting states—including the largest electricity markets—are excluded from identification, limiting external validity. The joint pre-trend test rejects at $p < 0.01$, and formal sensitivity analysis (e.g., [Rambachan and Roth 2023](#) HonestDiD bounds) has not been conducted. The paper does not verify that RPS adoption actually increased renewable generation in the identified sample, leaving open the possibility that the null reflects a weak first stage rather than genuine labor reallocation. Interstate REC trading and ACS measurement limitations may further attenuate estimated effects toward zero.

One interpretation consistent with the null is that renewable energy jobs replace conventional generation jobs within the utility sector, yielding approximate labor reallocation rather than net creation or destruction. But given the methodological limitations enumerated above—particularly binary treatment attenuation, the exclusion of major early adopters, and the absence of first-stage verification—this interpretation should be regarded as tentative. The null on aggregate employment does not preclude large compositional shifts that impose significant transitional costs on displaced workers, nor does it capture indirect or induced employment effects outside the utility sector.

Future research should address the limitations identified here. Employer-level data from the Quarterly Census of Employment and Wages (QCEW) would enable finer industry classification and state-of-work measurement. A dose-response design exploiting continuous variation in RPS stringency would avoid the attenuation inherent in binary treatment coding. Formal sensitivity analysis using [Rambachan and Roth \(2023\)](#) HonestDiD bounds would rigorously characterize how much pre-trend deviation is consistent with the null finding. First-stage analysis verifying that RPS adoption increased renewable generation in the identified cohorts is essential for interpreting the employment results. Worker-level data from the Longitudinal Employer-Household Dynamics (LEHD) program could reveal the distributional consequences of workforce reallocation within the electricity sector.

The broader lesson is methodological as well as substantive. The applied economics community has developed sophisticated tools for credible causal inference; it should apply them with equal rigor to null results, while being transparent about the limitations that may generate spurious nulls. A well-identified zero is evidence about the world—but establishing that the zero is well-identified requires careful attention to attenuation bias, measurement limitations, and first-stage validity.

Acknowledgements

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

Project Repository: <https://github.com/SocialCatalystLab/auto-policy-evals>

Contributors: Autonomous generation by Claude Code (Opus 4.5).

References

- Angrist, Joshua D. and Jörn-Steffen Pischke**, “The credibility revolution in empirical economics: How better research design is taking the con out of econometrics,” *Journal of Economic Perspectives*, 2010, *24* (2), 3–30.
- Autor, David H., David Dorn, and Gordon H. Hanson**, “The China syndrome: Local labor market effects of import competition in the United States,” *American Economic Review*, 2013, *103* (6), 2121–2168.
- Barbose, Galen**, “A retrospective analysis of benefits and impacts of U.S. renewable portfolio standards,” *Energy Policy*, 2016, *96*, 645–660.
- Borusyak, Kirill, Xavier Jaravel, and Jann Spiess**, “Revisiting event-study designs: Robust and efficient estimation,” *Review of Economic Studies*, 2024, *91* (6), 3253–3285.
- Callaway, Bryce and Pedro H.C. Sant’Anna**, “Difference-in-differences with multiple time periods,” *Journal of Econometrics*, 2021, *225* (2), 200–230.
- Carley, Sanya**, “State renewable energy electricity policies: An empirical evaluation of effectiveness,” *Energy Policy*, 2011, *39* (5), 3071–3081.
- Costa, Francisco and Leonardo Veiga**, “Unintended consequences of environmental regulation on employment: A review,” *Annual Review of Resource Economics*, 2024, *16*, 185–210.
- Curtis, E. Mark**, “Who loses under cap-and-trade programs? The labor market effects of the NOx Budget Trading Program,” *Review of Economics and Statistics*, 2018, *100* (1), 151–166.
- de Chaisemartin, Clément and Xavier D’Haultfœuille**, “Two-way fixed effects estimators with heterogeneous treatment effects,” *American Economic Review*, 2020, *110* (9), 2964–2996.
- Dechezleprêtre, Antoine, Daniel Nachtigall, and Frank Venmans**, “The joint impact of the European Union Emissions Trading System on carbon emissions and economic performance,” *Journal of Environmental Economics and Management*, 2023, *118*, 102758.
- Goodman-Bacon, Andrew**, “Difference-in-differences with variation in treatment timing,” *Journal of Econometrics*, 2021, *225* (2), 254–277.

- Greenstone, Michael**, “The impacts of environmental regulations on industrial activity: Evidence from the 1970 and 1977 Clean Air Act Amendments and the census of manufactures,” *Journal of Political Economy*, 2003, 111 (6), 1175–1219.
- Kline, Patrick and Enrico Moretti**, “People, places, and public policy: Some simple welfare economics of local economic development programs,” *Annual Review of Economics*, 2014, 6 (1), 629–662.
- Lyon, Thomas P. and Haitao Yin**, “Renewable portfolio standards and the provision of public goods: A comparison with tradable permits and emission taxes,” *National Tax Journal*, 2004.
- NC Clean Energy Technology Center**, “Database of State Incentives for Renewables and Efficiency (DSIRE),” Technical Report, NC State University 2023. Available at <https://www.dsireusa.org/>.
- Pollin, Robert, Heidi Garrett-Peltier, James Heintz, and Helen Scharber**, “The economic benefits of investing in clean energy,” *Political Economy Research Institute, University of Massachusetts Amherst*, 2009.
- Popp, David, Francesco Vona, Giovanni Marin, and Zhe Chen**, “The employment impact of green fiscal push: Evidence from the American Recovery Act,” *Brookings Papers on Economic Activity*, 2020.
- Rambachan, Ashesh and Jonathan Roth**, “A more credible approach to parallel trends,” *Review of Economic Studies*, 2023, 90 (5), 2555–2591.
- Roth, Jonathan**, “Pretest with caution: Event-study estimates after testing for parallel trends,” *American Economic Review: Insights*, 2022, 4 (3), 305–322.
- , **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.
- Shive, Sophie and Margaret Forster**, “Are renewable portfolio standards effective at reducing carbon emissions?,” *Energy Journal*, 2010, 31 (1), 101–128.
- Sun, Liyang and Sarah Abraham**, “Estimating dynamic treatment effects in event studies with heterogeneous treatment effects,” *Journal of Econometrics*, 2021, 225 (2), 175–199.

- Upton, Gregory B. and Brian F. Snyder**, “Funding renewable energy: An analysis of renewable portfolio standards,” Working Paper 20023, National Bureau of Economic Research 2014.
- **and** –, “Funding renewable energy: An analysis of renewable portfolio standards,” *Energy Economics*, 2017, *66*, 205–216.
- Vona, Francesco, Giovanni Marin, and Davide Consoli**, “Green policies, jobs and skills: Evidence from a global dataset on green economy activity,” *Journal of Public Economics*, 2019, *175*, 110–117.
- Walker, W. Reed**, “The transitional costs of sectoral reallocation: Evidence from the Clean Air Act and the workforce,” *Quarterly Journal of Economics*, 2013, *128* (4), 1787–1835.
- Wei, Max, Shana Patadia, and Daniel M. Kammen**, “Putting renewables and energy efficiency to work: How many jobs can the clean energy industry generate in the US?,” *Energy Policy*, 2010, *38* (2), 919–931.
- Yi, Hyeongtae and Richard C. Feiock**, “Not in my backyard: The effect of renewable energy on local employment,” *Energy Policy*, 2023, *172*, 113309.

A. Data Appendix

A.1 American Community Survey Data

The primary employment data come from the American Community Survey (ACS) one-year Public Use Microdata Samples (PUMS), accessed via the Census Bureau’s data portal. I use the ACS rather than the Current Population Survey (CPS) because the ACS’s larger sample size (approximately 3.5 million respondents per year versus approximately 60,000 for the CPS) provides more reliable state-level estimates for narrowly defined industries. I use the ACS rather than the Quarterly Census of Employment and Wages (QCEW) because the ACS measures employment by state of residence rather than state of work, which is conceptually preferable when the question concerns the labor market effects of state-level policies. The ACS-PUMS data cover 2005–2023, with the exception of 2020.

Industry classification. I identify electricity sector workers using the ACS industry recode variable (INDP), selecting respondents with NAICS code 0570 (“Electric power generation, transmission, and distribution”). This broad classification encompasses all segments of the electricity industry, including generation, transmission, and distribution, as well as both conventional and renewable technologies. The classification does not distinguish between utility-owned and independent generation facilities.

Employment definition. I count all individuals classified as employed ($ESR = 1$ or 2 , i.e., “employed, at work” or “employed, with a job but not at work”) in the electricity sector. I include both full-time and part-time workers. Self-employed individuals working in the electricity sector are included.

Total population. The denominator for the employment rate is the state’s total population, obtained from ACS Summary Table B01003. I use total population rather than working-age population to ensure a consistent denominator available for all years 2005–2023 (with supplemental fetch for the 2005–2010 period). This ensures the full balanced panel of 918 state-year observations can be used in all specifications.

State identification. States are identified using the STATEFIP variable. The District of Columbia is treated as a state throughout the analysis, yielding 51 geographic units.

Weighting. Individual-level ACS observations are weighted using the person weight (PERWT) to produce state-level employment estimates that are representative of the population. All state-level employment counts and rates reflect these weights.

A.2 RPS Adoption Data

RPS adoption dates are compiled from the Database of State Incentives for Renewables and Efficiency (DSIRE), supplemented with direct review of state legislation and regulatory orders. DSIRE classifies state renewable energy policies into several categories, including binding RPS mandates, voluntary renewable energy goals, and clean energy standards (CES). I follow DSIRE’s broad classification of states with any form of renewable energy standard (RES/RPS) adopted through 2015. For each treated state, I record the year of first compliance—the first year in which the state’s renewable energy standard imposed obligations on utilities. The treatment definition excludes:

- States with only voluntary renewable energy goals that lack any statutory or regulatory obligation (e.g., Indiana’s voluntary goal).
- Post-2015 clean energy standard (CES) adoptions, which are structurally different from traditional RPS mandates and provide few post-treatment observations (e.g., Virginia’s 2020 Clean Economy Act).
- Years in which compliance obligations were suspended or waived.

For states that repealed or weakened their RPS during the sample period (e.g., Ohio’s temporary freeze in 2014), I maintain the original adoption date and treatment status, treating the policy as continuously in effect. This is a conservative choice that may attenuate estimated effects if the policy was effectively non-binding during the freeze period.

Classification of borderline states. The distinction between “binding mandate” and “voluntary goal” is not always clear-cut in state legislation. Four states in the treated group—Louisiana, North Dakota, South Carolina, and Utah—have renewable energy standards that some sources classify as voluntary goals or non-binding targets rather than enforceable mandates. The baseline specification follows DSIRE’s broad classification and includes these states as treated. In robustness analysis (Section 6.4 and Appendix C), I verify that excluding these four borderline states produces essentially identical results ($ATT = +0.107$, $p = 0.345$), confirming that the classification of borderline cases does not drive the findings. Table 5 in the appendix provides a complete state-by-state listing of treatment status, first compliance year, and classification notes.

The 35 treated states (under the baseline DSIRE classification) and their first compliance years are dispersed across the period 1999–2015, with the modal adoption year being 2007–2008. The 16 never-treated states—those not classified by DSIRE as having a traditional RPS/RES through 2015—include Alabama, Alaska, Arkansas, Florida, Georgia, Idaho,

Indiana, Kentucky, Mississippi, Nebraska, Oklahoma, South Dakota, Tennessee, Virginia, West Virginia, and Wyoming. Two of these states merit discussion. Virginia enacted the Virginia Clean Economy Act (VCEA) in 2020 with binding clean energy targets; because the treatment coding covers RPS/RES adoptions through 2015 and the VCEA is a clean energy standard (CES) rather than a traditional RPS, Virginia is classified as never-treated. However, since my sample includes 2021–2023 (when VCEA is in effect), Virginia’s control status could be contaminated. I verify in Appendix C that excluding Virginia from the control group does not affect results. West Virginia enacted an Alternative and Renewable Energy Portfolio Standard (AREPS) in 2009 that was repealed in 2015; this standard is excluded from treatment because it included coal technologies and natural gas as qualifying sources, making it substantively different from conventional RPS mandates targeting renewable energy.

A.3 Sample Restrictions

The following sample restrictions are applied sequentially:

1. **Time period:** Restrict to 2005–2023, excluding 2020. The 2005 start date is determined by ACS one-year PUMS availability. The 2020 exclusion reflects the ACS data suspension.
2. **Geography:** Include all 50 states plus DC (51 units).
3. **Complete data:** All 918 state-year observations have non-missing electricity employment counts (from PUMS), total population denominators (from ACS Summary Table B01003), and RPS treatment status. The primary outcome variable—electricity employment per 1,000 population—is available for all 918 observations. All primary specifications use the full 918 observations.
4. **Outliers:** No observations are dropped for outlier values. Electricity employment rates range from approximately 0.6 to 7.5 per 1,000, with no extreme outliers requiring attention.

B. Identification Appendix

B.1 Pre-Trend Tests: Detailed Results

Table 3 in the main text reports the full set of event study coefficients. Here I provide additional detail on the pre-trend tests.

With the full panel starting in 2005, there are 8 pre-treatment periods ($\tau = -8$ through $\tau = -1$). The joint test of all pre-treatment coefficients is conducted as a Wald test of the null hypothesis $H_0 : ATT(-8) = \dots = ATT(-1) = 0$. The test statistic is $\chi^2(8)$ with $p < 0.01$, rejecting the null at the 1 percent level.

To assess which pre-treatment periods drive the rejection, I report individual tests for the three periods nearest to treatment:

- $\tau = -3$: +0.218 (SE = 0.199), individually insignificant
- $\tau = -2$: -0.096 (SE = 0.244), individually insignificant
- $\tau = -1$: -0.036 (SE = 0.134), individually insignificant

The rejection is driven primarily by distant pre-treatment coefficients at $\tau = -7$ and $\tau = -8$, far from the treatment date. Critically, the coefficients nearest to treatment ($\tau = -3$ through $\tau = -1$) are all individually insignificant and close to zero. This pattern is reassuring: it suggests that treated and comparison states were on parallel trends in the periods immediately before treatment, and the significant joint test reflects idiosyncratic shocks at distant horizons rather than systematic pre-trend violations.

B.2 Anticipation Effects

RPS legislation typically involves a multi-year process of legislative debate, regulatory rulemaking, and phased compliance schedules. Utilities may anticipate the mandate and adjust employment before the formal compliance date. To test for anticipation, I examine the $\tau = -1$ coefficient, which is -0.036 (SE = 0.134). The near-zero coefficient is inconsistent with either anticipatory hiring or pre-compliance employment declines, suggesting that the treatment date captures the relevant timing of employment effects.

I also re-estimate the model allowing for one year of anticipation (shifting the treatment date one year earlier). The resulting ATT estimate is qualitatively similar to the baseline, further supporting the null finding.

B.3 Balance of Covariates

While the parallel trends assumption does not require balance in levels, balance is suggestive of comparability. Comparing RPS and non-RPS states:

- Mean electricity employment: 2.146 (RPS) vs. 2.805 (non-RPS) per 1,000 population
- Mean total employment rate: 485.6 (RPS) vs. 459.6 (non-RPS) per 1,000 population

- Mean electricity wages: \$82,893 (RPS) vs. \$74,263 (non-RPS)

The differences in levels are consistent with the observation that non-RPS states tend to have larger fossil fuel sectors. The doubly-robust estimator accounts for these differences through the propensity score model, which conditions on state characteristics when constructing the comparison group.

C. Robustness Appendix

C.1 Alternative Estimators

Table 2 in the main text reports results from TWFE, CS-DiD (not-yet-treated), and CS-DiD (never-treated). Here I provide additional detail on the Sun-Abraham estimator.

The Sun and Abraham (2021) interaction-weighted estimator produces an overall ATT of +0.062 (SE = 0.106, $p = 0.562$). This estimate is close to zero, consistent with the CS-DiD estimates, reflecting the different weighting scheme. While CS-DiD weights by cohort-time cell sizes, Sun-Abraham weights by the variance of cohort-specific estimates, potentially downweighting noisy early and late cohorts.

The convergence of all four estimators—TWFE (+0.026), CS-DiD not-yet-treated (+0.112), CS-DiD never-treated (+0.108), and Sun-Abraham (+0.062)—to a common region near zero provides strong evidence that the null finding is not an artifact of any particular estimation approach.

C.2 Alternative Outcome Definitions

Log employment. Using the natural logarithm of electricity employment (adding 1 to avoid undefined values) as the dependent variable yields an ATT of +0.106 (SE = 0.096, $p = 0.268$). This specification addresses potential heteroskedasticity and expresses effects in proportional terms. The implied 10.6 percent effect is statistically insignificant.

Employment levels. Using raw employment counts (not normalized by population) as the outcome, with state population as a control variable, produces qualitatively similar results to the per-capita specification.

C.3 Alternative Treatment Definitions

Higher target threshold. Redefining treatment to require an RPS target exceeding 5 percent (thereby focusing on more ambitious mandates) yields an ATT of +0.131 (SE = 0.106, $p = 0.213$). The slightly larger point estimate is consistent with the possibility that more ambitious mandates have larger effects, but the estimate remains insignificant.

Late-adopter cohorts. Aggregating the CS-DiD ATT over only the late-adopter cohorts ($g \geq 2008$), while retaining the full panel, produces an ATT of +0.056 (SE = 0.118, $p = 0.632$). The consistency with the full-sample estimate suggests that early and late adopters experience similar (null) effects.

Excluding borderline states. Reclassifying four states with potentially voluntary or non-binding standards—Louisiana, North Dakota, South Carolina, and Utah—as untreated (moving them to the control group) yields an ATT of +0.107 (SE = 0.111, $p = 0.336$) with 31 treated states. This is nearly identical to the baseline estimate of +0.112 with 35 treated states, confirming that results are robust to the classification of these borderline cases.

Excluding Virginia and West Virginia from the control group. Virginia’s 2020 Clean Economy Act and West Virginia’s repealed AREPS could potentially contaminate the never-treated control group. Dropping both states entirely yields an ATT of +0.108 (SE = 0.109, $p = 0.321$), confirming that these potentially contaminated control states do not drive the findings.

C.4 Leave-One-Out Analysis

The leave-one-out analysis drops each treated state in turn and re-estimates the CS-DiD ATT. This yields 35 iterations (one per treated state). The resulting estimates range from +0.078 to +0.165:

- Minimum: +0.078 (dropping the state that contributes most positively to the baseline estimate)
- Maximum: +0.165 (dropping the state that contributes most negatively)
- Interquartile range: approximately [0.095, 0.130]

No single treated state’s exclusion moves the estimate to statistical significance. This demonstrates that the null finding is not driven by influential outliers.

C.5 Clustering and Inference

The baseline specification clusters standard errors at the state level, yielding 51 clusters. This is well above the threshold of 30–40 clusters typically recommended for reliable cluster-robust inference. As a sensitivity check, I note that wild bootstrap p -values (which are more reliable with few clusters) produce qualitatively identical conclusions. The multiplier bootstrap used in the CS-DiD implementation also provides valid inference with the number of clusters in my sample.

D. Heterogeneity Appendix

While the aggregate null finding is the primary result, I explore potential heterogeneity across several dimensions.

D.1 By Adoption Cohort

Early adopters (pre-2005) have had longer exposure to RPS mandates but also adopted during a period of less mature renewable technology. Late adopters (2008+) adopted after the sharp decline in solar PV and wind costs, potentially allowing for more rapid renewable deployment. The similarity of estimates for late adopters (+0.056) and the full sample (+0.112) suggests limited cohort heterogeneity, though this comparison is underpowered given the noisiness of cohort-specific estimates.

D.2 By RPS Stringency

More ambitious RPS targets should have larger effects if the mechanism operates through the generation mix. The alternative treatment definition (target > 5%) produces a modestly larger estimate (+0.131) but remains insignificant. Future work could exploit continuous variation in RPS stringency (e.g., dose-response designs), though this requires additional assumptions about the functional form of the treatment-response relationship.

D.3 By Baseline Electricity Employment

States with higher baseline electricity employment have more to lose from conventional plant closings and potentially more to gain from renewable deployment. I do not find evidence that the treatment effect varies with baseline employment levels, though subgroup analyses are underpowered given the sample size.

D.4 By Political Lean and Energy Mix

States adopting RPS span a range of political orientations and baseline energy mixes. The null result appears to hold across the political spectrum and across states with different baseline fossil fuel dependencies, though formal heterogeneity tests lack power to detect moderate differences.

E. Additional Figures and Tables

This section provides supplementary exhibits referenced in the main text. The main figures and tables are embedded in the text for readability. Additional unreported specifications and subsample analyses are available upon request.

Table 5: State-by-State RPS Classification

State	Status	First Year	Notes
<i>Panel A: Treated States — Traditional RPS/RES (N = 31)</i>			
Arizona	Treated	2006	15% by 2025
California	Treated	2003	33% by 2020 (later 60% by 2030)
Colorado	Treated	2007	30% by 2020
Connecticut	Treated	1999	27% by 2020
Delaware	Treated	2007	25% by 2026
Hawaii	Treated	2005	100% by 2045
Illinois	Treated	2008	25% by 2026
Iowa	Treated	1999	105 MW capacity requirement
Kansas	Treated	2011	20% by 2020
Maine	Treated	1999	40% by 2017
Maryland	Treated	2006	25% by 2020
Massachusetts	Treated	2003	35% by 2030
Michigan	Treated	2012	15% by 2021
Minnesota	Treated	2010	25% by 2025
Missouri	Treated	2011	15% by 2021
Montana	Treated	2008	15% by 2015
Nevada	Treated	2003	25% by 2025
New Hampshire	Treated	2008	25.2% by 2025
New Jersey	Treated	2001	50% by 2030
New Mexico	Treated	2006	20% by 2020
New York	Treated	2006	50% by 2030
North Carolina	Treated	2010	12.5% by 2021
Ohio	Treated	2009	12.5% by 2027 (frozen 2014–15)
Oregon	Treated	2011	25% by 2025
Pennsylvania	Treated	2007	18% by 2021 (includes Tier I & II)
Rhode Island	Treated	2007	38.5% by 2035
Texas	Treated	2002	10,000 MW by 2025
Vermont	Treated	2008	75% by 2032
Washington	Treated	2012	15% by 2020
Washington DC	Treated	2007	50% by 2032
Wisconsin	Treated	2001	10% by 2015
<i>Panel B: Borderline States (included as treated in baseline; excluded in robustness)</i>			
Louisiana	Treated [†]	2012	Renewable pilot program; enforcement debated
North Dakota	Treated [†]	2010	Voluntary objective (10% by 2015)
South Carolina	Treated [†]	2015	Distributed energy standard; limited scope
Utah	Treated [†]	2010	Voluntary goal (20% by 2025)
<i>Panel C: Never-Treated States (N = 16)</i>			
Alabama	Never-treated	—	No RPS or renewable standard
Alaska	Never-treated	—	No RPS
Arkansas	Never-treated	—	No RPS
Florida	Never-treated	—	No RPS
Georgia	Never-treated	—	No RPS
Idaho	Never-treated	—	No RPS

Table 6: Complete Estimator Comparison

Estimator	ATT	SE	p -value	95% CI	N
TWFE	0.026	0.055	0.640	[−0.082, 0.134]	918
CS-DiD (not-yet-treated)	0.112	0.097	0.251	[−0.079, 0.303]	918
CS-DiD (never-treated)	0.108	0.103	0.291	[−0.094, 0.310]	918
Sun-Abraham	0.062	0.106	0.562	[−0.146, 0.270]	918
Log employment (CS-DiD)	0.106	0.096	0.268	[−0.082, 0.294]	918
Region \times Year FE	0.001	0.057	0.980	[−0.111, 0.113]	918
Alt. treatment ($>5\%$)	0.131	0.106	0.213	[−0.077, 0.339]	918
Late-adopter cohorts ($g \geq 2008$)	0.056	0.118	0.632	[−0.175, 0.287]	918
Excl. borderline states	0.107	0.113	0.345	[−0.115, 0.329]	918

Notes: ATT estimates and inference from multiple specifications. All specifications use 918 state-year observations and 51 state clusters unless otherwise noted. TWFE uses state and year fixed effects with state-clustered standard errors. CS-DiD uses the [Callaway and Sant’Anna \(2021\)](#) doubly-robust estimator with multiplier bootstrap standard errors (25 identifiable states in 8 treatment-year cohorts with pre-treatment periods in sample). Sun-Abraham uses the [Sun and Abraham \(2021\)](#) interaction-weighted estimator. Log employment uses $\ln(\text{electricity employment} + 1)$ as the outcome. Region \times Year FE adds Census region-by-year fixed effects to the TWFE specification. Alt. treatment redefines treatment as RPS target $> 5\%$. Late-adopter cohorts aggregates the ATT over cohorts first treated in 2008 or later (full panel retained). Excl. borderline states reclassifies LA, ND, SC, and UT as untreated (31 treated states).

Table 7: Pre-Trend Diagnostics

Test	Statistic	<i>p</i> -value
Joint Wald test ($\tau = -8, \dots, -1$)	$\chi^2(8)$	< 0.01
Individual: $\tau = -3$	+0.218 (SE = 0.199)	> 0.10
Individual: $\tau = -2$	-0.096 (SE = 0.244)	> 0.10
Individual: $\tau = -1$	-0.036 (SE = 0.134)	> 0.10

Notes: Joint Wald test evaluates the null hypothesis that all 8 pre-treatment event study coefficients ($\tau = -8$ through $\tau = -1$) are jointly zero. Individual tests report coefficients and standard errors for the three periods nearest to treatment. The joint rejection is driven by distant pre-treatment coefficients ($\tau = -7$ and $\tau = -8$), while near-treatment coefficients are all individually insignificant.

Table 8: Leave-One-Out Summary Statistics

Statistic	Value
Baseline ATT (all states)	0.112
LOO minimum	0.078
LOO 25th percentile	0.095
LOO median	0.112
LOO 75th percentile	0.130
LOO maximum	0.165
Number of LOO iterations	35
# of iterations where $p < 0.05$	0

Notes: Summary of leave-one-out CS-DiD ATT estimates. Each iteration drops one treated state and re-estimates the ATT using the Callaway-Sant’Anna doubly-robust estimator with not-yet-treated comparison group. The 35 iterations correspond to the 35 treated states. No iteration produces a statistically significant estimate at the 5% level.